# COMPTIA Datasys+ DS0-001 Exam Questions

**Total Questions: 70+**
**Demo Questions: 15**
**Version: Updated for 2025**

**Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner**

**For Access to the full set of Updated Questions – Visit:**
**COMPTIA Datasys+ DS0-001 Exam Questions by Cert Empire**

# Question: 1

A developer is designing a table that does not have repeated values. Which of the following indexes should the developer use to prevent duplicate values from being inserted?

A. Unique

B. Single column

C. Implicit

D. Composite

### Answer:

A

### Explanation:

A unique index is a database constraint that ensures all values in the indexed column or columns are distinct. When a unique index is applied to a table, the database management system (DBMS) automatically rejects any INSERT or UPDATE operation that would create a duplicate value in the specified column(s). This directly enforces the developer's requirement to prevent repeated values, thereby maintaining data integrity.

### Why Incorrect Options are Wrong:



B. Single column: This describes an index based on one column, which can be either unique or non-unique. It does not inherently prevent duplicates.
C. Implicit: This refers to an index created automatically by the DBMS, typically for a primary key or unique constraint, not the type of constraint itself.
D. Composite: This describes an index that includes multiple columns. A composite index can be unique, but its primary definition is its multi-column structure, not its uniqueness property.

### References:

1. Oracle Database Concepts, 21c, Chapter 4: Schema Objects, Section: "Unique and Nonunique Indexes". The documentation states, "A unique index guarantees that no two rows of a table have duplicate values in the key column or columns."
2. PostgreSQL 16 Documentation, Chapter 11: Indexes, Section 11.7: "Unique Indexes". The documentation specifies, "Indexes can also be used to enforce uniqueness of column values... A unique index prevents duplicate values from being entered into the table."
3. Microsoft SQL Server Documentation, "Create Unique Indexes". The guide states, "A unique index ensures that the index key contains no duplicate values and therefore every row in the table is in some way unique."
4. Ramakrishnan, R., & Gehrke, J. (2003). Database Management Systems (3rd ed.). In Chapter 8, "Storage and Indexing," the concept of unique indexes is introduced as a mechanism to

enforce uniqueness constraints efficiently. (This is a standard university textbook for database courses).

## Question: 2

Which of the following would a database administrator monitor to gauge server health? (Choose two.)

A. CPU usage

B. Memory usage

C. Transaction logs

D. Network sniffer

E. Domain controllers

F. Firewall traffic

**Answer:**

A, B

**Explanation:**

A database administrator monitors CPU and memory usage as primary indicators of server health. High CPU utilization can signal inefficient queries or insufficient processing power, while excessive memory usage can lead to disk swapping, severely degrading database performance. These core system metrics provide a direct view of the server's load and capacity to handle database operations, making them essential for proactive performance management and troubleshooting. Monitoring these resources helps in identifying bottlenecks and ensuring the server can support the database workload effectively.



**Why Incorrect Options are Wrong:**

C. Transaction logs: Monitored for database recovery and auditing, not for general server hardware or operating system health.
D. Network sniffer: A specialized tool for deep packet inspection, not a standard, high-level server health metric like CPU or memory.
E. Domain controllers: External components for authentication; their monitoring does not reflect the database server's own resource health.
F. Firewall traffic: Relates to network security and access control, not the server's internal performance or resource utilization.
---

**References:**

1. Microsoft Corporation. (2023). Monitor resource usage (System Monitor). SQL Server Documentation. In the section "Objects to Monitor," the documentation explicitly lists the "Processor Object" (e.g., % Processor Time) and the "Memory Object" (e.g., Available MBytes) as

fundamental resources to monitor for gauging the performance and health of a server running SQL Server.
2. The PostgreSQL Global Development Group. (2024). PostgreSQL 16 Documentation. In Chapter 28, Section 28.2, "Monitoring Disk Usage," it is noted that monitoring OS-level statistics is crucial. While discussing disk, the context extends to all system resources, stating, "A database administrator should also monitor the operating system for signs of trouble," which implicitly and standardly includes CPU and memory as the most critical OS-level metrics.
3. Hellerstein, J. M., & Stonebraker, M. (2005). Readings in Database Systems (4th ed.). MIT Press. In Chapter 30, "Performance Tuning," the text emphasizes that the first step in performance analysis is to check the utilization of basic system resources: "The first thing to check is the utilization of the main hardware resources: the CPUs, the disks, and the network." (p. 701). Memory is discussed as a core component tied directly to disk I/O and CPU performance.

## Question: 3

A database administrator is concerned about transactions in case the system fails. Which of the following properties addresses this concern?

A. Durability

B. Isolation

C. Atomicity

D. Consistency

### Answer:

A

### Explanation:

The Durability property of ACID (Atomicity, Consistency, Isolation, Durability) ensures that once a transaction has been successfully committed, its changes are permanently recorded and will survive any subsequent system failure, such as a power outage or crash. The database management system guarantees that the results of committed transactions are written to non-volatile storage (e.g., a hard disk). This directly addresses the administrator's concern about the persistence of transactions in the event of a system failure.



### Why Incorrect Options are Wrong:

B. Isolation: This property ensures that concurrent transactions do not interfere with each other, making them appear to execute serially. It addresses concurrency, not system failure recovery.
C. Atomicity: This property ensures that a transaction is an "all-or-nothing" operation. Either all of its operations are completed successfully, or none of them are. It prevents partial updates.
D. Consistency: This property ensures that a transaction brings the database from one valid state to another, preserving all predefined integrity constraints. It is about data correctness, not persistence after a crash.

### References:

1. Haerder, T., & Reuter, A. (1983). Principles of Transaction-Oriented Database Recovery. ACM Computing Surveys, 15(4), 287-317.
Page 290, Section 2.4: "Once a transaction is committed, its effects are permanent in the database... This property is called durability. The effects of a committed transaction must not be abrogated by any type of failure."
DOI: https://doi.org/10.1145/289.291
2. Halevy, A., & Franklin, M. (2005). Transaction Processing. In P. A. Bernstein, U. Dayal, & D. B. Lomet (Eds.), Principles of Transaction Processing for the Systems Professional (2nd ed.). Morgan Kaufmann.

Chapter 1, Section 1.2, "The ACID Properties": "Durability means that once a transaction commits, its effects are permanent. The changes must not be lost in a failure."
3. MIT OpenCourseWare. (2010). 6.830 Database Systems, Lecture 17: Crash Recovery. Massachusetts Institute of Technology.
Slide 4, "ACID Properties": Defines Durability as: "Effects of a committed transaction are permanent, must not be lost by a crash."
Available at:
https://ocw.mit.edu/courses/6-830-database-systems-fall-2010/resources/mit6830f10lec17/

## Question: 4

Which of the following is a reason to create a stored procedure?

A. To minimize storage space

B. To improve performance

C. To bypass case sensitivity requirements

D. To give control of the query logic to the user

### Answer:

B

### Explanation:

Stored procedures are pre-compiled SQL statements stored on the database server. A primary reason for their use is to improve performance. The database management system (DBMS) parses, optimizes, and compiles the procedure once, caching the resulting execution plan. Subsequent calls to the procedure reuse this plan, eliminating the overhead of recompilation. This leads to faster execution. Additionally, they reduce network traffic because an application only needs to send a short EXECUTE command with parameters, rather than a potentially large and complex SQL query text, for each operation.



### Why Incorrect Options are Wrong:

A. To minimize storage space: Stored procedures consume storage space within the database to store their definition and compiled code; they do not minimize it.
C. To bypass case sensitivity requirements: Case sensitivity is governed by the database's collation settings, not by the use of stored procedures. The procedure's code adheres to the same rules.
D. To give control of the query logic to the user: This is the opposite of a stored procedure's purpose. They encapsulate and control logic on the server, abstracting it from the user.
---

### References:

1. Microsoft SQL Server Documentation. (2023). Stored Procedures (Database Engine). Microsoft Learn. In the "Benefits of Using Stored Procedures" section, it explicitly lists "Faster execution" and "Reduced network traffic" as key advantages. The document states, "If an operation requires a large amount of Transact-SQL code... a stored procedure can be faster because the Transact-SQL code is parsed and optimized only once."
2. Oracle Database Documentation. (2023). PL/SQL Language Reference, 23c. In Chapter 9, "PL/SQL Subprograms," the documentation explains that subprograms (which include procedures) are compiled and stored in the database, allowing them to be executed repeatedly

without recompilation, which improves performance.
3. Ramakrishnan, R., & Gehrke, J. (2003). Database Management Systems (3rd ed.). McGraw-Hill. In Chapter 5, "SQL: Queries, Constraints, Triggers," Section 5.8 discusses stored procedures. It notes that a major advantage is that the DBMS can compile the procedure's code and then reuse the compiled code, thereby avoiding the cost of parsing and optimization for each execution.
4. Silberschatz, A., Korth, H. F., & Sudarshan, S. (2020). Database System Concepts (7th ed.). McGraw-Hill. In Chapter 4, "Advanced SQL," Section 4.1.1, "Procedural Extensions," the text describes how stored procedures are pre-compiled, and their execution plans are cached, which "can result in a significant performance improvement."

## Question: 5

An on-premises application server connects to a database in the cloud. Which of the following must be considered to ensure data integrity during transmission?

A. Bandwidth

B. Encryption

C. Redundancy

D. Masking

### Answer:

B

### Explanation:

When data is transmitted between an on-premises server and a cloud database, it typically traverses public networks, making it vulnerable to interception and modification. To ensure data integrity, the data must be protected from unauthorized alteration during transit. Encryption protocols, such as Transport Layer Security (TLS), are designed for this purpose. TLS not only encrypts the data to ensure confidentiality but also uses cryptographic mechanisms like Message Authentication Codes (MACs) to verify that the data has not been tampered with upon arrival. This verification process is fundamental to maintaining data integrity during transmission.



### Why Incorrect Options are Wrong:

A. Bandwidth: Bandwidth refers to the data transfer rate of a network connection. It affects performance and availability but does not provide any protection against data modification.
C. Redundancy: Redundancy involves duplicating components to prevent system failure and ensure high availability. It does not protect the content of the data from being altered in transit.
D. Masking: Data masking is a method of creating a structurally similar but inauthentic version of data, typically used to protect sensitive information in non-production environments (at rest).

### References:

1. National Institute of Standards and Technology (NIST). (2019). Special Publication 800-52 Revision 2: Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations. In Section 2.1, "TLS Overview," it is stated: "TLS provides three essential security services: confidentiality, integrity, and authentication." This directly links the protocol used for encryption in transit to the service of data integrity.
2. Google Cloud Documentation. (2023). Encryption in Transit in Google Cloud. In the "How encryption in transit is implemented" section, the document explains that protection is achieved by "encrypting the data before transmission, authenticating the endpoints, and decrypting and verifying the data on arrival." The verification step ensures integrity.

3. Amazon Web Services (AWS) Documentation. (2023). Using SSL/TLS to encrypt a connection to a DB instance. This document details the use of SSL/TLS to secure data in transit for Amazon RDS, stating it is the "industry standard for encrypting network communications." TLS inherently provides integrity protection as one of its core functions.
4. Saltzer, J. H., & Kaashoek, M. F. (2009). Principles of Computer System Design: An Introduction. MIT OCW. In Chapter 9, "Security," Section 9.5.2, "Secure Channels," discusses how protocols like TLS establish secure communication channels that provide confidentiality, integrity, and authentication for data in transit.

# Question: 6

Which of the following statements contains an error?

A. Select EmpId from employee where EmpId=90030

B. Select EmpId where EmpId=90030 and DeptId=34

C. Select* from employee where EmpId=90030

D. Select EmpId from employee

## Answer:

B

## Explanation:

The SQL statement in option B, Select EmpId where EmpId=90030 and DeptId=34, contains a syntax error. A valid SELECT statement that retrieves data from a table must include a FROM clause to specify the source table. This statement is missing the FROM clause, making it syntactically incorrect and unexecutable by any standard SQL database management system. The basic structure of a query requires specifying what to select (SELECT), from where to select it (FROM), and optionally, the conditions for selection (WHERE).

## Why Incorrect Options are Wrong:



A. This is a valid SQL query. It correctly uses the SELECT, FROM, and WHERE clauses to retrieve a specific employee's ID from the employee table.
C. This is a valid SQL query. The asterisk () is a wildcard that correctly selects all columns from the employee table for the specified EmpId.
D. This is a valid SQL query. It correctly uses the SELECT and FROM clauses to retrieve all values from the EmpId column in the employee table.

## References:

1. Ullman, J. D., & Widom, J. (n.d.). CS145 Introduction to Databases, Class Notes on SQL. Stanford University. Retrieved from https://infolab.stanford.edu/ullman/fcdb/slides/slides-sql1.pdf. On page 4, the basic form of an SQL query is defined as SELECT FROM WHERE , explicitly showing the FROM clause as a required component.
2. PostgreSQL Documentation. (n.d.). 7.1. Overview. In PostgreSQL 16 Documentation. Retrieved from https://www.postgresql.org/docs/current/tutorial-select.html. The introductory example SELECT FROM weather; and subsequent explanations establish the FROM clause as fundamental for specifying the target table.
3. Stonebraker, M., & Hellerstein, J. (2005). 8. Relational Algebra and SQL. In 6.830 Database Systems, Fall 2005. Massachusetts Institute of Technology: MIT OpenCourseWare. Retrieved from https://ocw.mit.edu/courses/6-830-database-systems-fall-2005/resources/lec8/. Slide 20

presents the canonical SELECT-FROM-WHERE statement, reinforcing the necessity of the FROM clause.

# Question: 7

A database professional is considering denormalizing a database. Which of the following documents should be used to analyze the database's structure?

A. SOP

B. Data dictionaries

C. UML diaqrams

D. ERD

**Answer:**

D

**Explanation:**

An Entity-Relationship Diagram (ERD) is a graphical representation of a database's logical structure. It illustrates the entities (tables), their attributes (columns), and the relationships between them, including cardinality and keys. When considering denormalization-a process that involves intentionally combining tables to improve query performance-the ERD is the essential document. It provides the clear, high-level structural overview required to analyze existing relationships and identify which tables are suitable candidates for merging. Understanding these inter-entity connections is fundamental to planning any structural modification like denormalization.



**Why Incorrect Options are Wrong:**

A. SOP: A Standard Operating Procedure (SOP) outlines the steps to perform a task; it does not describe the database's architecture or data relationships.
B. Data dictionaries: A data dictionary is a repository of metadata, providing detailed information about data elements like types and constraints, but it lacks the visual, relational overview of an ERD for structural analysis.
C. UML diagrams: While Unified Modeling Language (UML) can be used for data modeling, the ERD is the specific, industry-standard artifact for designing and analyzing relational database structures.

**References:**

1. Silberschatz, A., Korth, H. F., & Sudarshan, S. (2020). Database System Concepts (7th ed.). McGraw-Hill.
Chapter 7, Section 7.1, "Overview of the Design Process," p. 269: "The E-R data model is most relevant to the database-design process. It provides useful concepts that allow a database designer to move from an informal description... to a more detailed, and precise, description that can be implemented... The E-R modeling technique is a graphical one, and E-R diagrams are

used to represent the logical structure of a database." This establishes the ERD as the primary tool for representing and analyzing the logical structure.
2. Elmasri, R., & Navathe, S. B. (2017). Fundamentals of Database Systems (7th ed.). Pearson.
Chapter 7, Section 7.1, "Using High-Level Conceptual Data Models for Database Design," p. 204: "The ER model is a popular high-level conceptual data model. This model and its variations are frequently used for the conceptual design of database applications, and many database design tools employ its concepts." This confirms the ERD's role as the standard for conceptual design and analysis.
3. Ullman, J. D., & Widom, J. (2008). A First Course in Database Systems (3rd ed.). Pearson Prentice Hall.
Chapter 4, "The Entity-Relationship Model," p. 113: "The E/R model is a design tool that is independent of the particular DBMS in which the database will be implemented... The E/R model allows us to draw pictures of the 'real world' that the database will model." This highlights its function as a design and analysis tool for database structure.

## Question: 8

A company needs information about the performance of users in the sales department. Which of the following commands should a database administrator use for this task?

A. DROP

B. InPDATE

C. delete

D. ISELECT

**Answer:**

D

**Explanation:**

The SELECT command is the fundamental statement in Structured Query Language (SQL) used for retrieving data from a database. To get information about the performance of users in a specific department, a database administrator would execute a SELECT query. This query would specify the columns to be returned and use a WHERE clause to filter the results to include only users from the "sales" department. SELECT is the primary component of the Data Query Language (DQL) and is designed specifically for data retrieval and reporting tasks, which directly aligns with the company's requirement.



**Why Incorrect Options are Wrong:**

A. DROP: This is a Data Definition Language (DDL) command used to permanently delete database objects like tables or indexes, not to query data.
B. UPDATE: This is a Data Manipulation Language (DML) command used to modify existing records in a table, not to retrieve them for viewing.
C. DELETE: This is a Data Manipulation Language (DML) command used to remove records from a table, which is a destructive action, not a retrieval operation.

**References:**

1. PostgreSQL 16 Documentation, Chapter 7. Queries, Section 7.1. Overview. The documentation states, "The process of retrieving data from a database is called querying it... In SQL the SELECT command is used for this purpose."
2. MySQL 8.0 Reference Manual, Section 13.2.13, SELECT Statement. The manual defines the command as: "SELECT is used to retrieve rows selected from one or more tables..."
3. Silberschatz, A., Korth, H. F., & Sudarshan, S. (2019). Database System Concepts (7th ed.). McGraw-Hill. Chapter 3, "Introduction to SQL," Section 3.2, "Basic Structure of SQL Queries" (p. 70). The text introduces the SELECT clause as the primary means to specify the attributes (columns) to be returned in the result of a query.

4. Ramakrishnan, R., & Gehrke, J. (2003). Database Management Systems (3rd ed.). McGraw-Hill. Chapter 5, "SQL: Queries, Constraints, Triggers," Section 5.2, "The Form of a Basic SQL Query" (p. 142). This university-level textbook explicitly states, "The basic form of an SQL query is SELECT...FROM...WHERE..." and describes its function for data retrieval.

## Question: 9

Which of the following are the best resources for monitoring potential server issues? (Choose two.)

A. User connections

B. Firewall usage

C. Index usage

D. CPU usage

E. Query execution

F. Memory usage

**Answer:**

D, F

**Explanation:**

CPU and memory usage are fundamental, system-level metrics for monitoring the health and performance of any server. High or sustained CPU utilization can indicate inefficient processes or that the server is under-provisioned, leading to slow response times. Similarly, high memory usage, especially when it leads to swapping data to disk (paging), is a primary cause of severe performance degradation. Monitoring these two core resources provides the most direct and immediate insight into a server's overall load and potential bottlenecks, making them the best initial resources for identifying issues.



**Why Incorrect Options are Wrong:**

A. User connections: This is an application-level metric. While a spike can indicate a problem, it is often a symptom of an underlying resource issue, not the primary resource itself.
B. Firewall usage: This is a network security metric. It is used to monitor for security threats or misconfigurations, not for general server performance or health issues.
C. Index usage: This is a database-specific metric used for query optimization. It is irrelevant for servers not running a database or for issues unrelated to database performance.
E. Query execution: This is a database-specific performance metric. It is critical for database tuning but does not represent the overall health of the server's core hardware resources.

**References:**

1. Microsoft Corporation. (2023). Performance Tuning Guidelines for Windows Server 2022. Microsoft Docs. In the section "Measuring Performance," the document identifies "Processor\% Processor Time" and "Memory\Available Mbytes" as two of the four most important counters for identifying system bottlenecks. This establishes CPU and memory as primary monitoring

resources.
2. Red Hat, Inc. (2023). RHEL 9: Monitoring and managing system status and performance. Red Hat Customer Portal. Chapter 1, "An introduction to performance analysis," describes how tools like top and vmstat are used for initial analysis, with their primary outputs being CPU and memory statistics to get a "high-level view of the system."
3. Patterson, D. A., & Hennessy, J. L. (2017). Computer Organization and Design RISC-V Edition: The Hardware Software Interface. Morgan Kaufmann. In Chapter 1, "Computer Abstractions and Technology," Section 1.6, "The Performance Measurement Challenge," discusses CPU time as the fundamental metric of processor performance. This academic text underscores the centrality of CPU monitoring in performance analysis.
4. Arpaci-Dusseau, R. H., & Arpaci-Dusseau, A. C. (2018). Operating Systems: Three Easy Pieces. Arpaci-Dusseau Books. In Chapter 14, "The Abstraction: Address Spaces," the text explains the critical role of physical memory management and how its exhaustion leads to swapping, a major performance problem. This highlights the importance of memory monitoring. (Available at https://pages.cs.wisc.edu/remzi/OSTEP/).

## Question: 10

An automated script is using common passwords to gain access to a remote system. Which of the following attacks is being performed?

A. DoS

B. Brute-force

C. SQL injection

D. Phishing

### Answer:

B

### Explanation:

The scenario describes an automated script systematically trying a list of "common passwords" to gain access. This is a classic example of a brute-force attack. Specifically, it is a dictionary attack, which is a subtype of brute-force where the attacker uses a pre-compiled list of likely passwords (such as common words or previously breached passwords) rather than trying every possible character combination. The goal is to guess the correct credential through trial and error to achieve unauthorized access.



### Why Incorrect Options are Wrong:

A. DoS: A Denial-of-Service (DoS) attack's primary goal is to make a service or system unavailable to legitimate users, not to gain unauthorized access.
C. SQL injection: This is an application layer attack that targets databases by inserting malicious SQL statements into entry fields, not an attack on the authentication mechanism itself.
D. Phishing: Phishing is a social engineering attack that tricks a human user into voluntarily disclosing their credentials, rather than an automated script attacking a system directly.

### References:

1. National Institute of Standards and Technology (NIST). (2020). NISTIR 7298 Revision 3: Glossary of Key Information Security Terms. Page 40. "brute force attack - A method of guessing a password or other obscured data by trying all possible combinations."
2. Saltzer, J. H., & Kaashoek, M. F. (2009). Principles of Computer System Design: An Introduction. MIT OpenCourseWare, Chapter 8: Security and Protection. Section 8.5.1, "Passwords," discusses password-guessing attacks, including dictionary attacks as a common method.
3. Pfleeger, C. P., Pfleeger, S. L., & Margulies, J. (2015). Security in Computing (5th ed.). Pearson Education. Chapter 2, "Toolbox: Authentication," discusses password guessing and brute-force attacks as methods to defeat password-based authentication. (This is a widely used university

textbook).
4. Kurose, J. F., & Ross, K. W. (2017). Computer Networking: A Top-Down Approach (7th ed.). Pearson. Section 8.2, "Principles of Cryptography," describes brute-force attacks in the context of cracking cryptographic keys or passwords by trying all possibilities.

## Question: 11

Which of the following cloud delivery models provides users with the highest level of flexibility regarding resource provisioning and administration?

A. DBaaS

B. IaaS

C. SaaS

D. PaaS

**Answer:**

B

**Explanation:**

Infrastructure as a Service (IaaS) provides the highest level of flexibility and administrative control. In this model, the cloud provider offers fundamental computing resources such as virtual machines, storage, and networking. The consumer is responsible for managing the operating systems, middleware, data, and applications. This allows for maximum customization and control over the environment, from provisioning specific OS versions to configuring complex network topologies. In contrast, other models like PaaS and SaaS abstract away these lower-level details, offering convenience at the cost of flexibility.



**Why Incorrect Options are Wrong:**

A. DBaaS: This is a managed service (a form of PaaS) where the provider handles the database and infrastructure, limiting user control to data and schema management.
C. SaaS: Offers the least flexibility. The provider manages the entire stack, and the user only interacts with the software application through a client interface.
D. PaaS: The provider manages the underlying infrastructure, including the operating system, restricting user control to the applications and data deployed on the platform.
---

**References:**

1. Mell, P., & Grance, T. (2011). The NIST Definition of Cloud Computing (Special Publication 800-145). National Institute of Standards and Technology.
Section 2, "Service Models": This section defines the three primary service models. For IaaS, it states, "The consumer... has control over operating systems, storage, and deployed applications." For PaaS and SaaS, it explicitly notes that the consumer "does not manage or control the underlying cloud infrastructure," demonstrating the superior control offered by IaaS.
2. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A View of Cloud Computing. Communications of

the ACM, 53(4), 50-58. https://doi.org/10.1145/1721654.1721672
Page 51, Figure 1: This figure illustrates the layers of abstraction managed by the provider versus the user for IaaS, PaaS, and SaaS. It visually confirms that with IaaS, the user manages the most layers (Application, Middleware, OS), thus having the highest degree of flexibility.
3. University of California, Berkeley. (2015). CS 162: Operating Systems and Systems Programming, Lecture 19: Cloud Computing.
Slide 33, "Cloud Service Models": The lecture slide presents a comparison table showing that in IaaS, the user manages the "Application," "OS/Middleware," and "Virtual Hardware," while in PaaS and SaaS, the OS and hardware are managed by the provider, clearly indicating IaaS offers the most control.

# Question: 12

Which of the following should a company develop to ensure preparedness for a fire in a data center?

A. Deployment plan

B. Backup plan

C. Data retention policy

D. Disaster recovery plan

**Answer:**

D

**Explanation:**

A Disaster Recovery Plan (DRP) is the formal, documented strategy for responding to a catastrophic event, such as a fire, that affects an organization's IT infrastructure. The primary objective of a DRP is to minimize downtime and data loss by providing a structured approach to restore critical systems and operations at an alternate location. A fire in a data center is a classic example of a disaster that necessitates a comprehensive DRP, which encompasses procedures for emergency response, system restoration, and resumption of business functions. While a backup plan is a crucial component, the DRP is the overarching framework that guides the entire recovery effort.



**Why Incorrect Options are Wrong:**

A. Deployment plan: This plan details the steps for implementing new hardware or software; it is not designed for recovering from a catastrophic event.
B. Backup plan: This is a subset of a DRP. It focuses solely on the procedures for copying and restoring data, not the entire infrastructure recovery process.
C. Data retention policy: This is a governance document that defines how long data must be kept for legal and operational reasons, not a plan for disaster response.

**References:**

1. National Institute of Standards and Technology (NIST). (2010). Special Publication 800-34 Rev. 1, Contingency Planning Guide for Federal Information Systems.
Section 2.3, Page 7: Defines the Disaster Recovery Plan (DRP) as a component of contingency planning that "provides the detailed procedures to facilitate recovery of capabilities at an alternate site." It explicitly addresses major disruptions, such as a fire, that render the primary facility unusable.
2. Carnegie Mellon University Software Engineering Institute. (2008). Defining and Differentiating BCP, COOP, and DRP (CMU/SEI-2008-TN-003).

Page 4: Clearly defines a DRP as focusing on "the recovery of IT systems at an alternate location." This directly applies to the scenario of a primary data center being destroyed by a fire.
3. University of California, Berkeley. Information Security and Policy, Disaster Recovery Plan (DRP) Toolkit.
Section "What is a Disaster Recovery Plan?": States that a DRP is a "detailed plan of action for the recovery of the functions of an information system after a disaster." It lists fire as a specific example of a disaster that would trigger the plan.

## Question: 13

Which of the following is the deployment phase in which a DBA ensures the most recent patches are applied to the new database?

A. Importing

B. Upgrading

C. Provisioning

D. Modifying

**Answer:**

B

**Explanation:**

In a typical database-deployment lifecycle, the first operational step is provisioning (creating the instance). The next phase is to bring that fresh instance to the latest supported software level, which is done through patch sets or minor-release updates. Vendor documentation groups this activity under "upgrading/patching," where the DBA applies the most recent security and bug-fix patches immediately after creation. Therefore, the phase concerned with ensuring a new database has the newest patches is the upgrading phase.



**Why Incorrect Options are Wrong:**

A. Importing - Refers to loading data into an existing database; it does not change the database software level.
C. Provisioning - Creates and configures the initial database instance but intentionally precedes the patch/upgrade step.
D. Modifying - Pertains to changing database objects or configurations after deployment, not to applying software patches.

**References:**

1. Oracle Corporation. "Patching and Upgrading Database Deployments," Oracle Database Cloud Service Admin Guide 19c, Section 6.1, pp. 6-1-6-3.
2. Oracle Enterprise Manager 13c, Database Lifecycle Management Guide, Chap. 2 "Provisioning and Patching Overview," para. "Upgrade/Patch Phase."
3. PostgreSQL Global Development Group. "Minor Release Updates (Upgrading)," PostgreSQL 15 Documentation, Sec. 18.6.
4. IBM. "Installing Fix Packs (Upgrading) for Db2 11.5," IBM Knowledge Center, Tasks Maintain Upgrade, Steps 1-3.
5. Stanford University CS145 Course Notes, "Database Administration Lifecycle," slide set "Deployment and Maintenance," slide 17 ("Provision Upgrade/Patch Data Load").

## Question: 14

Which of the following computer services associates IP network addresses with text-based names in order to facilitate identification and connectivity?

A. LDAP

B. NTP

C. DHCP

D. IDNS

**Answer:**

D

**Explanation:**

The service that associates IP network addresses with text-based, human-readable names is the Domain Name System (DNS). The acronym IDNS stands for Internet Domain Name System, which is a more descriptive but less common term for DNS. This system acts as a distributed database, translating domain names (e.g., www.example.com) into numerical IP addresses (e.g., 93.184.216.34) that computers use to identify each other on a network. This translation process is called name resolution and is fundamental for internet navigation and connectivity.



**Why Incorrect Options are Wrong:**

A. LDAP: Lightweight Directory Access Protocol (LDAP) is a protocol for accessing and maintaining distributed directory information services, typically used for user authentication and storing contact information, not for general IP-to-hostname resolution.
B. NTP: Network Time Protocol (NTP) is a networking protocol for clock synchronization between computer systems over packet-switched, variable-latency data networks. Its purpose is to maintain accurate time, not resolve names.
C. DHCP: Dynamic Host Configuration Protocol (DHCP) automatically assigns IP addresses and other network configuration parameters (like the address of a DNS server) to devices on a network. It assigns addresses but does not resolve names.

**References:**

1. Mockapetris, P. (1987). RFC 1034: Domain Names - Concepts and Facilities. Internet Engineering Task Force (IETF). Section 1, Abstract. "This RFC is an introduction to the Domain Name System (DNS)... The DNS is primarily used to map host names to IP addresses and vice versa." Available at: https://doi.org/10.17487/RFC1034
2. Saltzer, J., & Kaashoek, F. (2018). 6.033 Computer System Engineering, Spring 2018 Lecture Notes. Massachusetts Institute of Technology: MIT OpenCourseWare. Chapter 9, "Naming," Section 9.2, "The Domain Name System (DNS)." "Most prominently, it translates more readily

memorized domain names to the numerical IP addresses needed for locating and identifying computer services..." Available at: https://ocw.mit.edu/courses/6-033-computer-system-engineering-spring-2018/
3. Droms, R. (1997). RFC 2131: Dynamic Host Configuration Protocol. Internet Engineering Task Force (IETF). Section 1, Abstract. "The Dynamic Host Configuration Protocol (DHCP) provides a framework for passing configuration information to hosts on a TCP/IP network." Available at: https://doi.org/10.17487/RFC2131
4. Mills, D. (2010). RFC 5905: Network Time Protocol Version 4: Protocol and Algorithms Specification. Internet Engineering Task Force (IETF). Section 1, Introduction. "The Network Time Protocol (NTP) is widely used to synchronize computer clocks in the Internet." Available at: https://doi.org/10.17487/RFC5905
5. Sermersheim, J. (Ed.). (2006). RFC 4511: Lightweight Directory Access Protocol (LDAP): The Protocol. Internet Engineering Task Force (IETF). Section 1, Introduction. "The Lightweight Directory Access Protocol (LDAP) is an Internet protocol for accessing distributed directory services..." Available at: https://doi.org/10.17487/RFC4511

## Question: 15

Which of the following types of RAID, if configured with the same number and type of disks, would provide the best write performance?

A. RAID 3

B. RAID 5

C. RAID 6

D. RAID 10

**Answer:**

D

**Explanation:**

RAID 10 (a stripe of mirrors) provides the best write performance among the choices. This is because it combines the mirroring of RAID 1 with the striping of RAID 0, avoiding the computational overhead associated with parity calculations. For every write operation, data is simply written to two disks in a mirrored pair simultaneously. This process does not involve the read-modify-write sequence, known as the "write penalty," which significantly slows down write operations in parity-based arrays like RAID 3, RAID 5, and RAID 6.



**Why Incorrect Options are Wrong:**

A. RAID 3: Uses a dedicated parity disk that becomes a bottleneck for write operations, as every write must update this single disk, serializing the I/O.
B. RAID 5: Incurs a significant write penalty. Each logical write requires four physical I/O operations (read old data, read old parity, write new data, write new parity).
C. RAID 6: Suffers from the highest write penalty. Updating its dual parity blocks requires six physical I/O operations for a single logical write, making it the slowest.

**References:**

1. Patterson, D. A., Gibson, G., & Katz, R. H. (1988). A case for Redundant Arrays of Inexpensive Disks (RAID). SIGMOD '88: Proceedings of the 1988 ACM SIGMOD international conference on Management of data, 109-116. In Section 3, "The RAID Proposal," the paper details the performance of small writes, noting for RAID 5 (and by extension, RAID 3 and 6) that a single write requires four disk I/Os, a significant performance penalty not present in mirrored arrays (RAID 1/10). https://doi.org/10.1145/50202.50214
2. Arpaci-Dusseau, R. H., & Arpaci-Dusseau, A. C. (2018). Operating Systems: Three Easy Pieces. Arpaci-Dusseau Books. In Chapter 40, "Redundant Arrays of Inexpensive Disks (RAIDs)," Section 40.6 discusses RAID-5 write performance, stating, "For each logical write, the RAID controller has to perform four physical I/Os," which is known as the read-modify-write penalty.

3. University of California, Berkeley. (2014). CS 162: Operating Systems and System Programming, Lecture 14: Filesystems, Disks, and RAIDs. On slide 51, the "RAID IOPS Calculation" table shows that for 100% random writes, the I/Os per second for RAID 10 is N/2, whereas for RAID 5 it is N/4 and for RAID 6 it is N/6 (where N is the number of disks), clearly demonstrating the superior write performance of RAID 10.