



Up-to-date Practice Test with Latest Questions and Answers covering latest syllabus and topics of the exam. Makes you ready to face actual exam.



- C1000-122 Practice Questions
- C1000-122 Practice Test
- C1000-122 Practice Exam
- C1000-122 Exam Questions
- C1000-122 Study Guide



[killexams.com](http://killexams.com)

**IBM**

# C1000-122

*Db2 12 for z/OS DBA Fundamentals*

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/C1000-122>



### Question: 906

On reviewing the DDF configuration, a performance analyst sees delays in 'queue time of connection request awaiting DBAT.' What should be reviewed to resolve bottleneck?

- A. Assign all connections to SECPORT
- B. Increase MAXDBAT and CONDBAT
- C. Change APPLCOMPAT to highest level
- D. Enable EF and EA on catalog

**Answer:** B

Explanation: If queue time rises, resource or thread starvation is likely. Increasing MAXDBAT and CONDBAT relieves pressure and improves throughput for distributed connections.

### Question: 907

In Db2 12 for z/OS package versioning, PKG\_MULTI bound with VERSION('FL505') at APPLCOMPAT(V12R1M505) for LISTAGG, rebound to VERSION('FL508') for UTS sep (PH04424). What REBIND VERSION('FL508') creates new copy, and what APRETAINDUP(YES) retains dups if identical?

- A. FREE PACKAGE VERSION('FL505'); PKGCOPYMAX(4)
- B. BIND PACKAGE(MULTI\_COLL.PKG\_MULTI) COPY VERSION('FL505'); EXTENTCP(2)
- C. REBIND PACKAGE(MULTI\_COLL.PKG\_MULTI) VERSION('FL508') APRETAINDUP(YES); RETAINEDCP(3)
- D. ALTER PACKAGE VERSION PRUNE; RETAINEDCP(4)

**Answer:** C

Explanation: Versioning tags for mgmt; REBIND VERSION creates tagged copy. APRETAINDUP(YES) keeps if same (avoid bloat). RETAINEDCP(3) limits. Command: REBIND PACKAGE(MULTI\_COLL.PKG\_MULTI) VERSION('FL508') APRETAINDUP(YES) PLANMGMT(EXTENDED); query SELECT VERSION, COPYID FROM SYSIBM.SYSPACKCOPY; 2026 EOS (web:5) mandates versioning.

### Question: 908

Which log management setting controls when active log data sets are archived?

- A. ARCHIVE LOG FREQ on DSNTIPL panel
- B. COPY pending status in image copies
- C. RECOVER LOG LIMIT parameter
- D. LOGRETAIN parameter in DB2PARM

**Answer:** A

Explanation: The ARCHIVE LOG FREQ parameter configures how often active log data sets are archived to the archive log pool.

### Question: 909

In a DB2 12 for z/OS scenario, an XML column PRODUCT\_INFO in the PRODUCTS table stores complex catalogs (<product><specs><feature>Durability: High</feature></specs></product>). To query products where 'Durability: High' exists in specs, using XML data type for inline storage without auxiliary tables, which storage configuration supports efficient scalar extraction via XMLQUERY on 16K pages?

- A. Create table with PRODUCT\_INFO XML in a universal tablespace with 16K

page size

- B. Create table with PRODUCT\_INFO CLOB(2M) in a LOB tablespace
- C. Create table with PRODUCT\_INFO BLOB(2G) in a system-managed tablespace
- D. Create table with PRODUCT\_INFO VARCHAR(32704) in a segmented tablespace

**Answer:** A

Explanation: The XML data type enables native hierarchical storage in a separate XML tablespace (implicitly created), optimized for 16K pages to handle tree structures efficiently. This supports direct XMLQUERY without shredding, unlike CLOB/VARCHAR/BLOB which require parsing overhead and lack pureXML optimizations for inline queries.

**Question: 910**

In Db2 12 for z/OS HADR for SAP workloads, takeover delays 8 min due to index rebuild from LOGINDEXBUILD=NO. To minimize RTO with full logging, what parameter change and syntax during initial setup prevents rebuild?

- A. Async mode; START HADR INDEX
- B. Enable compression; UPDATE CFG INDEXCOMPRESSION
- C. Set LOGINDEXBUILD=YES; CREATE DB SAMPLE LOGINDEXBUILD YES
- D. Set LOGINDEXBUILD=NO; ALTER DB LOGINDEXBUILD NO

**Answer:** C

Explanation: Db2 12 for z/OS HADR for high-log SAP requires LOGINDEXBUILD=YES to log index changes, avoiding rebuild during takeover (8 min → <1 min). Syntax: CREATE DATABASE SAMPLE LOGINDEXBUILD YES (or ALTER for existing). Increases log volume 20-30%, but ensures RTO compliance; monitor via db2pd -indexes. NO mode

suits low-update, but SAP OLTP demands YES for integrity.

### Question: 911

In Db2 12 for z/OS with LDAP authentication (ZPARM LDAPAUTH=YES, server ldap://ldap.corp.com:636), a JDBC client connects with user JSMITH but fails SQL30082N post-bind success. The DSNLDPX exit checks LDAP group 'DB2USERS', but JSMITH belongs to 'FINANCE'. The exit logs "Group mismatch". What exit modification allows access via group equivalence without LDAP schema changes?

- A. In DSNLDPX, add `if(group=='FINANCE') return 'DB2USERS';` before mismatch log
- B. Update LDAP exit DSNLDPX to query attribute `memberOf` for 'DB2USERS' equivalence
- C. `ALTER ZPARM EXTSVR=LDAP GROUPMAP('FINANCE','DB2USERS')`
- D. `GRANT CONNECT TO JSMITH; CREATE ROLE FINROLE ADD MEMBER JSMITH`

**Answer:** A

Explanation: Db2 12 LDAP delegates bind to server, then site exit DSNLDPX for post-auth checks like group membership. Failure post-bind indicates exit rejection; coding equivalence (`if group=='FINANCE' map to 'DB2USERS'`) in DSNLDPX allows without LDAP mods. ZPARM EXTSVR sets type, no groupmap; role internal; `memberOf` query complex if not configured.

### Question: 912

A coupling facility SCA (Shared Communication Area) structure is lost on one LPAR. What is the immediate effect on your data sharing group?

- A. Statements dropped from EDM cache
- B. zIIP dispatching is halted
- C. Increased log write I/O
- D. Failure of group member communication and XCF signaling

**Answer:** D

Explanation: Loss of SCA structure disrupts communication among group members, potentially causing a split-brain or failover scenario.

### Question: 913

Deferred alters on Db2 12 UTS TS\_DEF (PAGENUM RELATIVE pending, ROTATE PART 3 pending) accumulate in AREOR, preventing ADD COLUMN. To materialize both pendings online via partition-level REORG for part 1 only, concurrent:

- A. ALTER TABLE TS\_DEF ADD COLUMN NEW\_COL INT;
- B. REORG TABLESPACE TS\_DEF SHRLEVEL STOP;
- C. ALTER TABLESPACE TS\_DEF PAGENUM ABSOLUTE IMMEDIATE;
- D. REORG TABLESPACE TS\_DEF PART 1 SHRLEVEL CHANGE;

**Answer:** D

Explanation: Partition-level REORG SHRLEVEL CHANGE materializes global pendings (PAGENUM/ROTATE) for part 1 in UTS, online concurrent, clearing advisory for that part. Full for all.

### Question: 914

You need to grant a user authority to grant and revoke object-level privileges in

Db2 12 for z/OS. Which authority is most appropriate?

- A. DBADM
- B. DATAACCESS
- C. SECADM
- D. ACCESSCTRL

**Answer:** D

Explanation: The ACCESSCTRL authority allows a user to grant and revoke object privileges without giving broad system-level or security administration rights.

**Question: 915**

Which statement correctly merges new sales data from a TEMP\_SALES staging table into the main SALES table on matching ORDER\_ID?

- A. MERGE INTO SALES USING TEMP\_SALES ON SALES.ORDER\_ID=TEMP\_SALES.ORDER\_ID
- B. SELECT \* FROM SALES JOIN TEMP\_SALES ON SALES.ORDER\_ID=TEMP\_SALES.ORDER\_ID
- C. UPDATE SALES SET VALUES FROM TEMP\_SALES WHERE ORDER\_ID=ORDER\_ID
- D. INSERT INTO SALES SELECT \* FROM TEMP\_SALES WHERE ORDER\_ID NOT IN SALES

**Answer:** A

Explanation: The MERGE statement is designed for upsert operations, specifying a source, a target, and a match condition.

## Question: 916

A Db2 12 for z/OS sales dashboard aggregates quarterly revenue using SUM(REVENUE DECIMAL(12,2)) and COUNT(DISTINCT CUST\_ID), but filters groups with AVG(ORDER\_VALUE) < 1000.00 via HAVING, and uses scalar MIN(ORDER\_DATE) and MAX(ORDER\_DATE) for range. The query groups by YEAR(ORDER\_DATE) and QUARTER(derived via CASE), excluding null REVENUE. Which GROUP BY with HAVING correctly computes without subaggregates?

- A.** SELECT YEAR(ORDER\_DATE), QUARTER(ORDER\_DATE), SUM(COALESCE(REVENUE, 0)), COUNT(\*) OVER () AS TOTAL\_CUSTS, AVG(ORDER\_VALUE) FROM ORDERS GROUP BY YEAR(ORDER\_DATE), QUARTER(ORDER\_DATE) HAVING COUNT(DISTINCT CUST\_ID) > 0;
- B.** SELECT YEAR(ORDER\_DATE) AS YR, CASE MONTH(ORDER\_DATE) WHEN 1,2,3 THEN 1 WHEN 4,5,6 THEN 2 WHEN 7,8,9 THEN 3 WHEN 10,11,12 THEN 4 END AS QTR, SUM(REVENUE), COUNT(DISTINCT CUST\_ID), MIN(ORDER\_DATE), MAX(ORDER\_DATE) FROM ORDERS WHERE REVENUE IS NOT NULL GROUP BY YEAR(ORDER\_DATE), CASE MONTH(ORDER\_DATE) WHEN 1,2,3 THEN 1 WHEN 4,5,6 THEN 2 WHEN 7,8,9 THEN 3 WHEN 10,11,12 THEN 4 END HAVING AVG(ORDER\_VALUE) >= 1000.00;
- C.** SELECT EXTRACT(YEAR FROM ORDER\_DATE) AS YR, EXTRACT(QUARTER FROM ORDER\_DATE) AS QTR, SUM(REVENUE), COUNT(DISTINCT CUST\_ID) FROM ORDERS GROUP BY EXTRACT(YEAR FROM ORDER\_DATE), EXTRACT(QUARTER FROM ORDER\_DATE) HAVING AVG(ORDER\_VALUE) < 1000.00 AND MIN(ORDER\_DATE) IS NOT NULL;
- D.** SELECT YEAR(ORDER\_DATE) AS YR, (MONTH(ORDER\_DATE) - 1)/3 + 1 AS QTR, SUM(REVENUE), COUNT(DISTINCT CUST\_ID), MIN(ORDER\_DATE), MAX(ORDER\_DATE) FROM ORDERS WHERE REVENUE > 0 GROUP BY YEAR(ORDER\_DATE), (MONTH(ORDER\_DATE) - 1)/3 + 1 HAVING AVG(ORDER\_VALUE) >=

1000.00 ORDER BY YR, QTR;

**Answer: D**

Explanation: Integer division  $(\text{MONTH}(\text{ORDER\_DATE}) - 1) / 3 + 1$  derives quarter (1-4) efficiently for GROUP BY, matching SELECT expression to avoid SQLCODE -98. SUM(REVENUE) aggregates non-nulls via WHERE >0 (excludes nulls implicitly), COUNT(DISTINCT CUST\_ID) for unique customers, MIN/MAX for range. HAVING AVG(ORDER\_VALUE) >=1000.00 filters post-group, DECIMAL precision preserved.

**Question: 917**

A Db2 12 for z/OS sequence SEQ\_GAPLESS NO CACHE for audit ID generates per-request (no gaps), but 30K/sec volume +25% CPU. What CACHE 20 balances gaps<18 with perf, and what RESTART WITH (SELECT MAX+1 -20) post-failover formula adjusts?

- A. ALTER SEQUENCE SEQ\_GAPLESS CACHE 20 RESTART WITH (SELECT MAX(ID)+1 -20 FROM AUDIT); CACHE = volume / 1500
- B. ALTER SEQUENCE CACHE 1 CYCLE; CACHE = INCREMENT \* sessions
- C. ALTER TABLE AUDIT RESTART IDENTITY; CACHE = CACHE\_HIT \* 0.1
- D. DROP SEQUENCE RESTART MINVALUE; CACHE = LASTUSED / 10

**Answer: A**

Explanation: NO CACHE CPU high; CACHE 20 = volume (30K) / 1500 reqs/gap tolerance. RESTART adjusts for lost (MAX+1 - CACHE). Alter: ALTER SEQUENCE SEQ\_GAPLESS CACHE 20 RESTART WITH (SELECT MAX(ID)+1 -20 FROM AUDIT\_TAB); monitor SYSIBM.SYSSEQUENCES.CACHE\_HIT>98%, GAP\_AVG<18. FL507 dynamic.

### Question: 918

A Db2 12 for z/OS sysprog sets IRLM parameters DEADLOCK(5,2) TIMEOUT(120) in the startup JCL. During a hot-swap to new IRLM, a transaction T1 (under CS) waits 130 seconds for a row U lock held by T2 (RR, no commit). T1 times out with -911 reason 1. What determined T1's timeout value, and how does DEADLOCK influence victim selection if a cycle formed?

- A. Base timeout 30s + IRLMRWT add-on; DEADLOCK(5,2) rolls back T2 as higher priority
- B. TIMEOUT(120)s absolute; DEADLOCK ignores, uses LOSTOCKS for victim
- C. Effective timeout = TIMEOUT \* DEADLOCK cycle; victim via random selection
- D. TIMEOUT(120) + 10% margin; DEADLOCK cycle 5s detects, victim 2 prioritizes low CPU

### Answer: D

Explanation: IRLM TIMEOUT(n) sets max lock wait in seconds (default 30); here 120s, but Db2 adds ~10% for overhead, yielding ~132s—exceeded at 130s triggers -911 reason 1 (timeout). DEADLOCK(c,v) sets detection cycle (5s) and timeout multiplier (2, so detects every 10s during waits); if cycle detected, selects victim: v=2 chooses lowest CPU-consuming process (T1 low CPU, survives; T2 rolls back). LOSTOCKS unrelated (lock loss threshold). Tune TIMEOUT lower for faster fails; monitor via IFCID 0090/0091 or -DISPLAY IRLM,DEADLOCK. Hot-swap preserves params via IRLM PROC.

### Question: 919

In a DB2 12 for z/OS financial XML (<txn><amt currency="USD">100.50</amt><fee>5.00</fee></txn>), XMLEXISTS validates multi-currency fees >1% with fn:sum(\$t/txn/amt[@currency="USD"]/@amt \*

0.01) < @fee. Which rewrite with fn:avg and conditional existence optimizes for index on /txn/amt/@currency, avoiding sum on empty sets?

- A. XMLEXISTS('\$t/txn[fee > fn:sum(amt[@currency="USD"]/@amt \* 0.01)]' PASSING TXN\_XML AS "t")
- B. XMLEXISTS('if (fn:exists(\$t/txn/amt[@currency="USD"])) then fn:avg(\$t/txn/amt[@currency="USD"]/@amt) \* 0.01 < fee else false()' PASSING TXN\_XML AS "t")
- C. XMLEXISTS('fn:sum(\$t/txn/amt[@currency="USD"]/@amt) \* 0.01 < fee' PASSING TXN\_XML AS "t")
- D. XMLEXISTS('fn:every \$a in \$t/txn/amt[@currency="USD"] satisfies @amt \* 0.01 < ../fee' PASSING TXN\_XML AS "t")

**Answer:** B

Explanation: if fn:exists guards against empty USD amts (avg undefined), then avg \*0.01 < fee uses index for existence and avg computation, sargable per txn. DB2 12 short-circuits if. sum on empty 0 false positive, every quantifier overkill, sum without guard errors.

**Question: 920**

During Db2 12 for z/OS group restart after CF failure in a 6-member data sharing setup with GDPS/HyperSwap, GBP4K1 recovery takes 20 min due to LPL (logical page list) rebuild. To parallelize recovery phases without merging logs, what automated procedure step and syntax reduces RTO to <5 min?

- A. Parallel member restart via batch job; -START DB2 MEMBER(\*)
- B. Sequential log merge; RECOVER DATABASE POSTGRES
- C. CF rebuild manual; ALTER GROUPBUFFERPOOL GBP4K1
- D. Single member restart; -STOP GROUP

**Answer:** A

Explanation: Db2 12 for z/OS data sharing group restart post-CF failure rebuilds GBP from member logs without merging, enabling parallel phases. Automate via batch: script -START DB2 MEMBER(ALL) after CF rebuild, leveraging independent log RBA checks. Syntax: //STEP1 EXEC PGM=DSNCMD, PARM='-START DB2(M1 M2 M3 M4 M5 M6)'. This cuts RTO from 20 min to <5 min by concurrent redo; GDPS/HyperSwap triggers script. Monitor via -DISPLAY GROUP STATUS; sequential risks cascade failure. No log merge needed, as each member processes own BSDS.

**Question: 921**

For object privileges in Db2 12 for z/OS, user DEV1 holds ALTER on table PROD.ORDERS but fails CREATE TRIGGER TRG1 BEFORE UPDATE ON PROD.ORDERS FOR EACH ROW... due to missing REFERENCES on ORDERS columns. To grant TRIGGER and implicit REFERENCES without full ALTER, what minimal GRANT enables trigger creation?

- A. GRANT ALTER ON TABLE PROD.ORDERS TO DEV1 WITH GRANT OPTION
- B. GRANT CREATEIN ON SCHEMA PROD TO DEV1
- C. GRANT TRIGGER ON TABLE PROD.ORDERS TO DEV1
- D. GRANT REFERENCES ON PROD.ORDERS (ORDERID, AMT) TO DEV1; GRANT TRIGGER ON PROD.ORDERS TO DEV1

**Answer: D**

Explanation: Db2 12 requires explicit TRIGGER privilege for CREATE TRIGGER, plus REFERENCES on referenced columns for integrity. GRANT TRIGGER alone insufficient; must GRANT REFERENCES on specific columns (ORDERID, AMT) and TRIGGER. ALTER includes but exceeds need; CREATEIN schema-level, not table.

**Question: 922**

A bank must tailor DSNZPARM settings to support increased parallel utility processing. Which ZPARM parameter in DSN6SPRM controls the max degree of parallelism for utilities?

- A. INDEX\_IO\_PARALLELISM
- B. PCLOSEN
- C. PARAMDEG\_UTIL
- D. INLISTP

**Answer: C**

Explanation: PARAMDEG\_UTIL controls the degree of parallelism for utilities. Setting it high allows utilities like LOAD and REORG to process in parallel, improving batch performance at the cost of increased resource consumption.

**Question: 923**

In Db2 12 for z/OS, DPSI PART\_DPSI\_SKU (non-unique on sku\_desc, including partition key cat\_id, 8K pages) on PBR TS\_CATALOG skews access. To unique DPSI with expression (UPPER(sku\_desc)), 32K compress (60% DSN1COMP), online:

- A. ALTER INDEX PART\_DPSI\_SKU KEYSEQ (SKU\_DESC); then REBUILD INDEX;
- B. ALTER INDEX PART\_DPSI\_SKU UNIQUE EXPRESSION (UPPER(SKU\_DESC)) COMPRESS YES BUFFERPOOL BP32K; then REORG INDEX PART\_DPSI\_SKU SHRLEVEL REFERENCE;
- C. ALTER TABLESPACE TS\_CATALOG SEGSIZE 64;
- D. DROP INDEX; then CREATE INDEX PART\_DPSI\_SKU ON CATALOG (UPPER(SKU\_DESC)) PARTITIONED UNIQUE INCLUDE (CAT\_ID)

COMPRESS YES BP32K;

**Answer: B**

Explanation: DPSI supports pending ALTER UNIQUE (with superset key)/EXPRESSION/COMPRESS/BUFFERPOOL, materialized by online REORG INDEX at SHRLEVEL REFERENCE.

**Question: 924**

A DBA wants to analyze the effect of buffer pool size on performance. Which command should be issued?

- A. -DISPLAY DDF STATUS
- B. -DISPLAY DATABASE()
- C. -DISPLAY BUFFERPOOL(BP1)
- D. -DISPLAY THREAD()

**Answer: C**

Explanation: DISPLAY BUFFERPOOL shows activity and hit ratios for the buffer pool, enabling analysis of buffer pool sizing and performance characteristics.

**Question: 925**

In a DB2 12 for z/OS scenario for supply chain analytics, the ORDERS table's ORDER\_DETAILS XML (<order><line item="1" qty="10">Laptop</line></order>) needs flattening to compute total quantity per supplier using XMLTABLE with nested paths. Which query joins to SUPPLIERS for sum(qty) where item contains 'Lap', leveraging column defaults?

**A.** SELECT S.SUPPLIER\_ID, SUM(T.QTY) FROM SUPPLIERS S,  
XMLTABLE('\$d/order/line' PASSING O.ORDER\_DETAILS AS "d" COLUMNS  
ITEM VARCHAR(50) PATH '@item', QTY INTEGER PATH '@qty' DEFAULT  
0) T WHERE O.SUPPLIER\_ID = S.ID AND T.ITEM LIKE '%Lap%' GROUP  
BY S.SUPPLIER\_ID

**B.** SELECT S.SUPPLIER\_ID, AVG(T.SCORE) FROM SUPPLIERS S JOIN  
XMLTABLE('\$d/order' PASSING O.ORDER\_DETAILS COLUMNS SCORE  
DECIMAL PATH 'line/@qty') T ON ...

**C.** SELECT S.SUPPLIER\_ID, XMLQUERY('sum(\$d/order/line/@qty)'  
PASSING O.ORDER\_DETAILS AS "d") FROM SUPPLIERS S, ORDERS O  
WHERE O.SUPPLIER\_ID = S.ID

**D.** SELECT COUNT(\*) FROM ORDERS WHERE XMLEXISTS('\$d/order/  
line[@item="Laptop"]' PASSING ORDER\_DETAILS AS "d")

**Answer:** A

Explanation: XMLTABLE shreds the XML into relational rows with PATH for attributes, DEFAULT 0 for missing qty, enabling efficient SUM with LIKE on ITEM and join to SUPPLIERS. XMLQUERY aggregates without flattening, XMLEXISTS counts existence only, AVG misapplies to qty.

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including Exam Questions, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.



## Practice Exam Questions Based on Current Exam Objectives

Killexams.com provides practice exam questions aligned with the latest official exam objectives and latest syllabus. Our content is reviewed and updated regularly to reflect recent changes announced by certification vendors. By studying these practice questions, candidates will cover the structure, difficulty level, and topics of the actual exam, helping them prepare more effectively and efficiently.

## Comprehensive Practice Exams (PDF Format)

Killexams.com offers multiple-choice questions (MCQs) in easy-to-read PDF format, covering all major domains of the exam. Each PDF contains a structured collection of practice questions and verified answers designed to support focused study. These MCQs help candidates reinforce key concepts, identify knowledge gaps, and improve exam readiness through consistent practice.

## Realistic Practice Tests (Online Test Engine & Desktop Test Engine)

To support hands-on preparation, Killexams.com provides practice tests through both an Online Test Engine and a Desktop Test Engine. These tools are designed to simulate a real exam environment, allowing candidates to practice under exam-like conditions, with latest syllabus and topics of the exam. Performance tracking, test history, and result analysis help users evaluate their progress and focus on areas that need improvement.

## Risk-Free Purchase Policy

Killexams.com follows a transparent and customer-friendly purchase policy. If users are not satisfied with the study materials, they may request assistance or a refund in accordance with our published terms and conditions. This policy reflects our commitment to customer satisfaction, fairness, and confidence in our preparation resources.

## Regularly Updated Content

Our practice question bank is reviewed and updated on an ongoing basis to stay aligned with the latest exam outlines and vendor updates. This ensures candidates are studying up-to-date, relevant material, and preparing with content that reflects current exam expectations, helping them stay confident and well-prepared.