Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.

KILL EXAMS

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

PASS

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 REVENU**E.** 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 (MONTH(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

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including MCQs, 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.

Exam Questions:

Killexams.com provides exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these questions, candidates can familiarize themselves with the content and format of the real exam.

Exam MCQs:

Killexams.com offers exam MCQs in PDF format. These questions contain a comprehensive collection of questions and answers that cover the exam topics. By using these MCQs, candidate can enhance their knowledge and improve their chances of success in the certification exam.

Practice Test:

Killexams.com provides practice test through their desktop test engine and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice test cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success:

Killexams.com offers a success guarantee with the exam MCQs. Killexams claim that by using this materials, candidates will pass their exams on the first attempt or they will get refund for the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exam.

Updated Contents:

Killexams.com regularly updates its question bank of MCQs to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.