


## [Q31-Q45 Master 2022 Latest The Questions Oracle Cloud Infrastructure and Pass 1Z0-1089-21 Real Exam!]



**Master 2022 Latest** The Questions Oracle Cloud Infrastructure and Pass 1Z0-1089-21 Real Exam!

**Penetration testers simulate 1Z0-1089-21 exam PDF Q31.** You are building a filesystem to maximize throughput of large files using high performance block volume and BM.Standard2.52 as file server.

Which block volume configuration should you choose to ensure aggregate Block volume throughput is higher than network bandwidth of a file server?

- \* 32 Block Volumes of 800GB
- \* 5Block Volume of 32 TB
- \* 6 Block Volumes of 32 TB
- \* 7 Block Volumes of 800GB

**Q32.** You are building a file system that needs to handle large files with a lot of nodes reading at the same time.

What should be your main goal?

- \* Maximize Latency while keeping throughput above a certain threshold.

- \* Maximize throughput.
- \* Minimize latency.
- \* Minimize throughput while keeping the latency low.

**Q33.** A two node clustered/parallel filesystem built using Baremetal Compute shape (BM.Standard2.52) and Block Volume Storage delivers maximum sustained aggregate IO throughput of 5GB/s for read and write, while storing a single copy of the data (replica=1) in filesystem.

What will be the maximum sustained aggregate IO throughput for read and write, when two copies of the data (replica=2) are stored in filesystem for High Availability?

- \* 5GB/s for both (read and write)
- \* 2.5GB/s for read and 5GB/s for write
- \* 2.5GB/S for write and 5GB/s for read
- \* 2.5GB/s for both (read and write)

**Q34.** Which is NOT an advantage associated with moving a Big Data workload to Oracle Cloud Infrastructure (OCI)?

- \* managed service offerings to support common Big Data workloads
- \* best price / performance for Big Data workloads In the Cloud
- \* availability of Oracle Airflow
- \* Object Storage as Data Lake
- \* dynamically scale capacity against workload

**Q35.** Compute nodes are on subnet2(172.16.1.0/24) and there is a BM.Dense.I02.8 on subnet1.

(1 / 2.16.0.0/24) that will serve as a bastion, file server and license server. The bastion will be accessed through a public IP with SSH.

Which security lists should you choose for an HPC cluster? If multiple options are possible, choose the one limiting access the most.

- \* Subnet1: 0.0.0.0 for all TCP/UDP ports, subnet2: 172.16.0.0/16 for all TCP/UDP ports
- \* Subnet1: 0-0.0.0 for TCP port 22, subnet 1 and 2: 172.16.0.0/23 for all TCP/UDP ports
- \* Subnet1: 172.16-1.0/24 for port 22, subnet2: 172.16.1.0/24 for all TCP/UDP ports
- \* Subnet1: 0-0.0.0 for TCP port 22, subnet 2: 172.16.0.0/16 for TCP port 22

**Q36.** A file system is built using BM.Standard2.52 Compute shape for File Servers. One 25 Gbps NIC/network card is used to connect to 10 Block Volumes of 1TB each (max. 4#0MB/s per volume). The other 25 Gbps NIC is used for sending/receiving data to/from client nodes.

File system client instances which mount the file system are provisioned using VM.Standard2.16 Compute shapes. (Network bandwidth: 16.4Gbps(2050 MB/s)) What is the max IO theoretical throughput a client node can get?

- \* 2050 MB/s
- \* 4800 MB/s
- \* 3125 MB/s
- \* 6250 MB/s

**Q37.** You have an HPC cluster running tightly coupled MPI jobs without VPN or Fast Connect.

The compute nodes are on subnet2(172.16.1.0/24), and there is a BM.Dense.I02.8 on subnet1 (172-16.0.0/24) that will serve as a bastion and file server.

What security lists should you choose to be able to run the workloads while limiting access as possible?

- \* Subnet1: 0.0.0.0 for TCP port 22, subnet 1 and 2: 172.16.0.0/23 for all protocols

- \* Subnet1: 0.0.0.0 for all protocols, subnet2: 172.16.0-0/16 for all protocols
- \* Subnet1: 172.16.1.0/24 for port 22, subnet2: 172.16.1.0/24 for all protocols
- \* Subnet1: 0.0.0.0 for TCP port 27, subnet 7: 172.16.0.0/16 for TCP port 22

**Q38.** What are two different types of Storage offerings available to build a filesystem on Oracle Cloud infrastructure (OCI)?

- \* Block Volume Service
- \* Local NVMe SSDs on DenseIO and HPC compute shapes
- \* File Storage
- \* Object Storage

**Q39.** A customer has a very busy workload. The model is very large (1 PB range) and only some small files are updated for new jobs. Throughput needed during the run is roughly 25GB/s.

What Is a fast and cost-conscious way to handle the file system?

- \* Put the data in object storage, and mount It using s3fs-fuse project.
- \* Build a file-system using NVMe on Dense shapes. Then move the data to object storage when not needed.
- \* Build a file system using Block volumes and Standard BMs, take advantage of the different block volume performances levels.
- \* Use NVMe on HPC shapes to build a File Systemwith the RDMA connection.

**Q40.** Which Is a common business problem for customers running Big Data workloads?

- \* Costs associated to process a large scale data set
- \* Cost associated with Disaster Recovery for large deployments
- \* Ability to process a small data set at the minimum cost as possible
- \* Ability to process a large data set at the maximum cost as possible

**Q41.** What three data sources can Oracle Data Science draw from?

- \* Google Cloud Object Storage
- \* Autonomous Data Warehouse
- \* On-premise Oracle database
- \* Azure ObjectStorage Blob
- \* AWS S3 buckets

**Q42.** What does the Terasort phase of the Terasort benchmark do?

- \* Itmaps and reduces source data in a parallel manner, leveraging data locality to minimize network transfer.
- \* It randomly maps source data and reduces output to a smaller data set,
- \* It maps and reduces one terabyte of data into a smaller data set.
- \* It randomly maps source data and increases output to a bigger data set.

**Q43.** You are building a RDMA cluster on a private network with 10000 cores. What CIDR will work in theVCN:

- 172.16.0.0/16?
- \* 172.16.32.0/23
  - \* 10.0.1.0/23
  - \* 172.16.32,0/16
  - \* 10.0.1.0/24

**Q44.** A customer has decided to use Oracle Cloud Infrastructure (OCI) for theirDisaster Recovery (DR) site for an on-premises Hadoop cluster. They have an RPO of 4 hours, and an RTO of 1 hour.

Which two options should they consider on OCI to meet their DR requirements?

- \* Replicate data to OCI Object Storage every hour.
- \* Replicate data to a Hadoop cluster running on OCI in near-realtime.
- \* Use Terraform to rapidly provision Hadoop Cluster on OCI.
- \* Replicate data to OCI Object Storage in near-realtime.

**Q45.** Which Message Passing Interface (MPI) distribution is NOT available today?

- \* OCI MPI
- \* Intel MPI
- \* Platform MPI
- \* Open MPI

#### Oracle 1Z0-1089-21 Exam Syllabus Topics:

TopicDetailsTopic 1- Explain how to deploy and use Data Science platform- Differentiate the various big data solutions and data services on OCITopic 2- Identify Advantages of running big data on OCI- Deploy monitoring mechanisms that ensure cost and resource optimization based on business requirementsTopic 3- Explain How to deploy and use GPU instances for AI-ML workloads- Align with business and technical requirement for scale-out workloadsTopic 4- Deploy a dynamically scalable, highly available and resilient compute solution based on technical requirements- Implement and Operate HPC solutions on OCITopic 5- Design a solution infrastructure that meets defined success measurements- Describe HPC Architecture OverviewTopic 6- Use Guidance and best practices for deploying big data products on OCI- Design, Deploy and Operate High Speed Shared File Systems

**Penetration testers simulate 1Z0-1089-21 exam:** <https://www.topexamcollection.com/1Z0-1089-21-vce-collection.html>