

KillTest

质量更高 服务更好



学习资料

<http://www.killtest.cn>

半年免费更新服务

Exam : **E20-860**

Title : EMC NAS Implementation
Engineer Expert Exam

Version : DEMO

1. Rob's Distribution Center has a Windows environment consisting of about 400 users. They are primarily using CIFS on an EMC Celerra. Users are complaining about poor performance when opening their files.

Which item should Joe's system administrator look for when running a "server_cifsstat server_2" command?

- A. An excessive network delay value
- B. High Concurrent connections values
- C. High ConnectTree values
- D. High ReadX average Service Response Times

Answer: D

2. A storage administrator is disappointed with the performance of an EMC solution. The administrator points out that Exchange server response time from the server is too slow.

How should performance have been positioned with the customer before the implementation?

- A. The customer and seller should have made a prearranged agreement to improve performance.
- B. The customer and seller should have made a prearranged agreement on bandwidth requirements.
- C. The customer and seller should have made a prearranged agreement on measurement and level of expected performance.
- D. The customer and seller should have made a prearranged agreement for a notable improvement in server and application responses measured in IOPS, application response time, or bandwidth improvement.

Answer: C

3. An Exchange administrator has consolidated the company's four Exchange stores onto a single Data Mover. The performance of the Exchange stores appears to be slow.

What should have been done to properly size the solution and set expectations for performance before implementation?

- A. A benchmark for each server needs to be set, I/O analysis needs to be completed for each server, and performance expectations need to be agreed on between the seller and company.
- B. Define how much storage is required, what type of disks are required, what RAID type will be used, and

how much money is to be spent on the project.

C. General expectations for a notable improvement in server and application responses measured in IOPS, CPU utilization, application response time, or bandwidth improvement.

D. Show the ROI for the project, what financing options are available, in which quarter will the deal be completed, and what kind of performance is expected.

Answer: A

4. Which EMC recommended tool is available for download to analyze network traffic?

A. Celerra Monitor

B. EMC ControlCenter

C. Ethereal/Wireshark

D. Iometer

Answer: C

5. A storage administrator noticed that the access to network drives on the Celerra slows down at about half past the hour every day. The local EMC service person told them that was normal, since a checkpoint is created 25 minutes past the hour every day.

Which tool could be used to show the performance degradation?

A. Celerra Monitor

B. ControlCenter Performance Monitor

C. Ethereal/Wireshark

D. Performance Monitor

Answer: A

6. Mowgli Inc. is a small motor repair company that has expanded operations to another location 100 miles away. When the new office network was connected to the main office, the iSCSI Exchange servers in the new branch were experiencing excessive response times by Microsoft performance standards.

What could cause this behavior?

A. iSCSI only works on the local subnet and not on a remote one.

B. The Exchange servers are also running backups causing the additional latency.

- C. The Exchange servers need TOE cards to decrease network latency.
- D. The network delay is too high to support a working iSCSI environment.

Answer: D

7. In a Symmetrix DMX-4 the GigE IPv4/v6 feature is a built-in data encryption processor that enables configuration on a per-port basis.

What are the two places this feature can be defined/enabled?

- A. Bin and SYMCLI
- B. Service processor and Bin
- C. SMC and service processor
- D. SYMCLI and SMC

Answer: C

8. Tracy is an administrator for a NAS gateway with eight Symmetrix 3+1 RAID 5 groups with dvol's equally distributed. Each dvol is 8.6 GB in size. Tracy needs to create a 32 GB file system.

How should she create the file system to obtain the best performance?

- A. Use d10, d11, d12, d13
- B. Use d10, d11, d12, d13, d14 with the slice option
- C. Use d10, d11, d12, d13, d14, d15
- D. Use d10, d11, d12, d13, d14, d15, d16, d17 with the slice option

Answer: D

9. A storage administrator recently purchased 32 FC drives for a DMX-4. The drives will be set up in a RAID 5 3+1 configuration.

How would you suggest the customer configure new file systems to obtain best performance?

- A. RAID 5 3+1 is not supported for the Celerra volumes.
- B. Use AVM with custom pools to configure the file systems.
- C. Use manual volume creation to maximize spindle count and minimize hotspots.
- D. Use Symmetrix metavolumes with manual creation of Celerra slice volumes to maximize spindle count and minimize hotspots.

Answer: C

10. Two Celerras are set up 10 km apart in an active/passive SRDF configuration. After a performance analysis, the recommendation is to set the copy process on one of the file systems to a higher priority.

What will be affected on the back end when setting QOS on the DMX?

- A. It is set on the cache extents that are used by that file system's volumes.
- B. It is set on the Device Group and affects all filesystems on those volumes.
- C. It is set on the front-end port and affects all addressed volume on the path.
- D. It is set on the SRDF director and thus affects all volumes on the director.

Answer: B

11. What is the name of the directory of data that is used to maintain the status of cache allocation for each logical volume in a Symmetrix?

- A. Bit Table
- B. Block Map
- C. Cache Map
- D. Track Table

Answer: D

12. Disk Director prefetching is initiated by two sequential accesses to a volume. Once prefetch is initiated, the Symmetrix will begin pulling the next two successive tracks into cache.

At what point will the Symmetrix begin to pull the next five successive tracks into cache?

- A. After 10 sequential accesses to the volume
- B. After 100 sequential accesses to the volume
- C. After 150 sequential accesses to the volume
- D. After 50 sequential accesses to the volume

Answer: B

13. A Symmetrix is initially configured (after power-up) with 200 logical volumes of the same size and emulation.

How much of the total available cache resource is distributed per volume?

- A. 0%
- B. 0.5%
- C. 1.0%
- D. 100%

Answer: B

14. How does PermaCache keep data in cache?

- A. Based on I/O
- B. Based on read rate
- C. Dynamically
- D. Statically

Answer: D

15. EMC has observed 20% to 30% improvement in back-end disk performance when Dynamic Mirror Service Policy (DMSP) is used, compared to when each of the mirrored pair serves 50% of the read I/Os. Which configuration would show this improvement?

- A. BCV creation
- B. RAID 1/0
- C. RAID 10
- D. RAID 5

Answer: C

16. One of the primary service policies on a Symmetrix is Dynamic Mirror Service Policy (DMSP). How is it set?

- A. By a user command issued via SMC
- B. A configuration change in the bin file
- C. Dynamically based on I/O characteristics
- D. Via QOS values set in the Host Component

Answer: B

17. Which Symmetrix policy is designed to minimize head movement for read operations when data is not available in cache memory?

- A. Dynamic Split Mirror Policy
- B. Enginuity Performance Algorithm Policy
- C. Interleave Service Policy
- D. Split Service Policy

Answer: D

18. An upgrade of twelve 300 GB Fibre Channel disks for an existing DMX-3 will be configured using RAID 5 3+1.

How should the resulting space be configured for optimal file system performance with the Celerra?

- A. Use AVM with a user-defined pool using the slice option to maximize the number of disks per volume
- B. Use AVM with the `symm_std` pool using the slice option to maximize the number of disks per volume
- C. Use Manual Volume Management for maximum spindle count and minimize hotspots
- D. Use Symmetrix metavolumes to minimize hotspots and do not slice the volume

Answer: C

19. In the Symmetrix DMX-4, global memory director operations are redundant by way of two global memory directors working as a pair.

Which protection methodology is this?

- A. Global memory access path protection
- B. Global memory chip-level redundancy
- C. Longitude redundancy code (LRC)
- D. Redundant global memory director

Answer: D

20. The Symmetrix supports multiple types of protection including RAID, N+2 hardware, and memory error correction.

Which method is best described as "striped pair of mirrored disks" in an open system configuration?

- A. RAID 1/0
- B. RAID 10
- C. RAID 5
- D. RAID 6

Answer: A