Liberté • Égalité • Fraternité
RÉPUBLIQUE FRANÇAISE
MINISTÈRE DE L'INTÉRIEUR

CONCOURS INTERNE ET EXTERNE DE TECHNICIEN DE CLASSE NORMALE DES SYSTEMES D'INFORMATION ET DE COMMUNICATION

- SESSION 2018 -

COLLE

Ne pas cacher le cadre d'identité. Cette opération sera réalisée par l'administration

Nom

Prénom

N° d'inscription

COLLE

Mardi 3 avril 2018

Epreuve écrite sur un document en anglais

Questions appelant des réponses écrites courtes en français, sur un document rédigé en anglais relatif aux systèmes d'information, destinées à évaluer le niveau de compréhension du candidat.

Cette épreuve obligatoire se déroule par anticipation à l'occasion des épreuves écrites d'admissibilité, mais les points sont pris en compte pour l'admission.

(Durée : 1 heure – coefficient 1)

Le dossier documentaire comporte 3 pages.

IMPORTANT

IL EST RAPPELE AUX CANDIDATS QU'AUCUN SIGNE DISTINCTIF NE DOIT APPARAITRE NI SUR LA COPIE NI SUR LES INTERCALAIRES.

ECRIRE EN NOIR OU EN BLEU - PAS D'AUTRE COULEUR

Extrait du blog iWEB

publié le 10 mai 2010 par M. Patrice GUAY : "An overview of RAID technology"

RAID (Redundant Array of Independent Disks) is a technology allowing a higher level of storage reliability and performance from disk-drive components via the technique of arranging them into arrays. A RAID array is a configuration with multiple physical disks set up to use a RAID architecture like RAID 0, RAID 1, RAID 5, etc. While the RAID array distributes data across multiple disks, it is considered as a single disk by the server operating system.

A RAID array is composed of two or more physical hard disks combined into a single logical storage unit. To give RAID array additional features compared to JBOD (Just a Bunch Of Disk), three main concepts are used:

- Mirroring

- Striping

- Error correction

Mirroring is the writing of identical data to more than one disk. The basic example of mirroring is a RAID 1 array formed by two disks. Both disks have the same content at any time. If the first drive fails, read and write operation can be done directly on the second disk. Read operations on mirrored arrays is faster compared to a single disk since the system can fetch data from multiple disks at the same time. However, write operations are slower since the data must be written to all disks instead of only one. The reconstruction of a failed mirror array is quite simple: data must be copied from the healthy disk to the new one. During reconstruction, the read performance boost of the mirror array is reduced since only the healthy disk is fully usable.

Striping is the splitting of data across multiple disks. For example, a RAID 0 array formed by two disks strips data to both disks. Striping does not provide fault tolerance, only a performance boost. Read and write operations on a striped array are faster compared to a single disks as both operation are split between the available disks.

Error correction stores parity data on disk to allow the detection and possibly the correction of problems. RAID 5 is a good example of the error correction mechanism. For example, a RAID 5 array composed of three drive strips data on the first two disks and stores parity on the third disk to provide fault tolerance. The error correction mechanism will slow down performance especially for write operation since both data and parity information needs to be written instead of data only. Moreover, the reconstruction of a failed array using parity information incurs severe performance degradation as data needs to be fetched from all drives in the array to rebuild the information for the new disk.

The design of any RAID scheme is a compromise between data protection and performance. The comprehension of your server requirements in terms of storage is crucial to select the appropriate RAID configuration.

RAID 0 is a pure implementation of striping. A minimum of two (2) disks is required for RAID 0. No parity information is stored for redundancy. It is important to note that RAID 0 was not one of the original RAID levels and provides no data redundancy. RAID 0 is normally used to increase performance. RAID 0 is useful for setups where redundancy is irrelevant.

A RAID 0 array can be created with disks of differing sizes, but the total available storage space in the array is limited by the size of the smallest disk. For example, if a 450GB disk is striped together with a 300GB disk, the usable size of the array will be 2 x min(450GB, 300GB) = 600GB.

For reads and writes operations dealing with small data blocks such as database access, the data will be fetched independently on each disk of the RAID 1 array. If the data sectors accessed are spread evenly between the two disks, the apparent seek time of the array will be half that of a single disk. The transfer speed of the array will be the transfer speed of all the disks added together, limited only by the speed of the RAID controller. For reads and writes operations dealing with large data blocks such as copying files or video playback, the data will most likely be fetch on a single disk reducing the performance gain of the RAID 0 array.

RAID 1 is a pure implementation of mirroring. A minimum of two (2) disks is required for RAID 1. This is useful when read performance or reliability are more important than data storage capacity. A classic RAID 1 mirrored pair contains two disks, which increases reliability over a single disk. Since each member contains a complete copy of the data, and can be addressed independently, ordinary wear-and-tear reliability is raised.

A RAID 1 array can be created with disks of differing sizes, but the total available storage space in the array is equal to the size of the smallest disk. For example, if a 450GB disk is mirrored with a 300GB disk, the usable size of the array will be min(450GB, 300GB) = 300GB.

The read performance of a RAID 1 array can go up roughly as a linear multiple of the number of copies. That is, a RAID 1 array of two disks can query two different places at the same time so the read performance should be two times higher than the performance of a single disk. RAID 1 is a good starting point for applications such as email and web servers as well as for any other use requiring above average read I/O performance and hardware failure protection.

RAID 5 array uses block-level striping with distributed parity blocks across all member disks. The disk used for the parity block is staggered from one stripe to the next, hence the term distributed parity blocks. A minimum of three (3) disks is required for RAID 5. This RAID configuration is mainly used to maximize disk space while providing a protection for your data in case of a disk failure.

A RAID 5 array can be created with disks of differing sizes, but the total available storage space in the array is limited by the size of the smallest disk. The parity data consumes a complete disk, leaving N-1 disks for usable storage space in an array composed of N disks. For example, on an array formed of three 450GB disks and one 300GB disk, the usable size of the array will be (4-1) x min(450GB, 300GB) = 900GB.

RAID 5 writes are expensive in terms of disk operations and traffic between the disks and the RAID controller since both data and parity information need to be written to disk. The parity blocks are not read on data reads, since this would add unnecessary overhead and would diminish performance. However, the parity blocks are read when a defective disk sector is present in the required data blocks. Likewise, should a disk fail in the array, the parity blocks and the data blocks from the surviving disks are combined mathematically to reconstruct data from the failed drive in real-time. This situation leads to severe performance degradation on the array for read and write operations.

RAID 6 extends RAID 5 by adding an additional parity block. Block-level striping is combined with two parity blocks distributed across all member disks. A minimum of four (4) disks is required for RAID 6. This RAID configuration is mainly used to maximize disk space while providing a protection for up to two disk failures.

A software RAID 6 implementation may have a significant effect on system performance especially during the reconstruction of a failed disk.

A RAID 6 array can be created with disks of differing sizes, but the total available storage space in the array is limited by the size of the smallest disk. The parity data consumes two complete disks, leaving N-2 disks for usable storage space in an array composed of N disks.

RAID 6 writes are even more expensive than RAID 5 writes in terms of disk operations and traffic between the disks and the RAID controller since both data and parity information need to be written to disk. The parity blocks are not read on data reads, since this would add unnecessary overhead and would diminish performance. However, the parity blocks are read when a defective disk sector is present in the required data blocks. Likewise, should a disk fail in the array, the parity blocks and the data blocks from the surviving disks are combined mathematically to reconstruct data from the failed drive in real-time. This situation leads to severe performance degradation on the array for read and write operations.

RAID 10 is a combination of RAID 1 (mirroring) and RAID 0 (striping) where 4N mirrored disks are striped together. A minimum of four (4) disks are required for RAID 10. One disk in each RAID 1 mirror can fail without damaging the data contained in the entire array.

A RAID 10 array can be created with disks of differing sizes, but the total available storage space in the array is limited by the size of the smallest disk. The mirroring consumes half of disk space, leaving 2N disks for usable storage space in an array composed of 4N disks. For example, on an array formed of seven 450GB disks and one 300GB disk, the usable size of the array will be (7+1)/2 x min(450GB, 300GB) = 1200GB.

RAID 10 provides better performance than all other redundant RAID levels. It is the preferable RAID level for I/O (Input/Output) intensive applications such as database servers as well as for any other use requiring high disk performance.

SUJET

QUESTIONS

Attention : Les réponses doivent uniquement se référer à ce qui est écrit dans le document.

<u>Question 1</u> *(2 points)*

Quels sont les avantages de la technologie RAID ?

Question 2 *(2 points)*

Quelle caractéristique du RAID 5 et du RAID 6 ne se retrouve pas dans les autres RAID ?

Question 3 *(2 points)*

Quel RAID offre les meilleures performances pour des applications nécessitant un haut débit d'entrée et de sortie ?

<u>Question 4</u> *(2 points)*

Pour quelle raison les temps d'écriture de RAID 5 et 6 sont-ils plus longs que ceux des autres RAID ?

Question 5 *(2 points)*

Quel est le principal avantage du volume agrégé par bandes ?

Question 6 *(2 points)*

Quel est le prérequis pour s’orienter vers la solution RAID la plus adaptée ?

Question 7 *(2 points)*

Dans quelle configuration RAID 0 est à préconiser ?

Question 8 *(2 points)*

De quoi se compose un groupement RAID ?

Question 9 *(2 points)*

Comparativement à un disque unique, quel est le niveau de performance de lecture de RAID 1 ?

Question 10 *(2 points)*

Quelle que soit la taille des disques, et quel que soit le RAID choisi, quel est le critère déterminant pour définir la capacité maximale de stockage disponible ?