

IT Infrastructure Architecture

Infrastructure Building Blocks
and Concepts

Storage

Backup schemes

- A backup scheme describes what data is backed-up, when, and how
- Backup schemes can become very complex in large environments with many applications
- Four basic backup schemes

Backup schemes

- Full backup
 - A complete copy of all data
 - Full backups are only created at relatively large intervals (like a week or a month)
 - Creating them takes much time, disk or tape space, and bandwidth
 - Restoring a full backup takes the least amount of time

Backup schemes

- Incremental backup
 - Save only newly created or changed data since the last backup, regardless of whether it is a previous incremental backup or a full backup
 - Restoring an incremental backup can take a long time
 - Especially when the last full backup is many incremental backups ago

Backup schemes

- Differential backup
 - Save only newly created or changed data since the last full backup
 - Restoring a differential backup is quite efficient, as it implies restoring a full backup and only the most recent differential backup

Backup schemes

- Continuous Data Protection (CDP)
 - Guarantees that every change in the data is also simultaneously made in the backup system
 - The RPO (Recovery Point Objective) is set to zero, because each change immediately triggers a backup process
 - Expensive technology, and therefore only used in specific situations

Backup data retention time

- Backup data retention time is the amount of time in which a given set of data will remain available for restore
- Defines how long backups are kept and at which interval
- In practice, a Grandfather-Father-Son (GFS) based schedule is often used:
 - Each day a backup is made
 - After a week, there are seven backups, of which the oldest backup is renamed to a weekly backup
 - After the second week, the same is done and the daily backups of the week before are deleted
 - Now there are eight backups: seven daily, two weekly
 - Every four weeks, the weekly backup is renamed as a monthly backup and the weekly backups are reused
 - The daily backups are the son, the weekly backups are the father, and the monthly backups are the grandfather

Archiving

- Archiving is mostly done for compliancy and regulation reasons
- Example:
 - US regulations require all medical records to be retained for 30 years after a person's death
 - This means that X-rays taken when a child was born must be kept for as much as 130 years!
- Noncompliance to law and regulation can lead to serious business disruption, fines, and even jail time

Archiving

- Archived data is read-only to protect it from being altered
 - Very important for regulatory compliance and non-repudiation
 - Some archiving systems store data in an encrypted form and use digital signatures to prove data is not tampered with
 - Some systems allow data to be written to it for archiving, but disallow changing or deleting data
 - CD / DVD/ Blu-ray
 - WORM tapes


Archiving

- Data must be kept in such a way that it is guaranteed the data can be read after a long time
 - Digital format (like a Microsoft Word file or a JPG file)
 - Physical format (like a DVD or a magnetic tape)
 - Storage environment (temperature, humidity)
- Use open standards for storing archived data
 - Open standards are well documented
 - Reading data will always be feasible, using emulation software if needed
 - Storing all documents in structured human-readable XML text files is one way to ensure data can be read for many decades
- Transfer data that is to be kept for a long time to the latest storage media standard every 10 years

Storage performance

Disk performance

- Disk performance is dependent on:
 - Disk rotation speed
 - Seek times
 - Interface protocol
- Some common examples of rotation delay:



Disk RPM	Average rotational delay (ms)
5,400	5.6
7,200	4.2
10,000	3
15,000	2

Disk performance

- Disks cannot spin much faster than 15,000 RPM
 - At this speed the velocity at the edge of a 3.5" disk is 250 km/h!
 - Increasing this velocity would physically destroy the disk
- Seek time is the time it takes for the head to get to the right track
 - Average seek times:
 - 3 ms for high-end disks
 - 9 ms for low-end disks

IOPS

- Input/output Operations Per Second (IOPS) is a measure of how many read and write operations a disk can complete in one second


$$\frac{1000}{\text{Rotational delay (ms)} + \text{Seek time(ms)}}$$

- Writing is typically a bit slower than reading
- Typical IOPS:

Disk RPM	IOPS
7,200	50
10,000	120
15,000	160
SSD	2,500 to 10,000

RAID penalty

- In RAID sets multiple disks are used to form one virtual disk (LUN)
- Writing data on multiple disks introduces some delay, known as the RAID penalty



Interface throughput

- Storage performance is also dependent on how fast the interface can move data from the disks to the systems consuming the data and vice versa
- An overview of the various interface speeds:

Interface	Speed
IDE (Parallel ATA)	100 MB/s 133 MB/s
SATA	1.5 Gbit/s (192 MB/s) 3 Gbit/s (384 MB/s) 6 Gbit/s (768 MB/s)
SCSI	160 MB/s (Ultra-160) 320 MB/s (Ultra-320)
SAS	1.5 Gbit/s (192 MB/s) 3 Gbit/s (384MB/s) 6 Gbit/s (768 MB/s)
FC	1 Gbit/s (128 MB/s) 2 Gbit/s (256 MB/s) 4 Gbit/s (512 MB/s) 8 Gbit/s (1024 MB/s) 16 Gbit/s (2048 MB/s)

Caching

- A caching system in disk controllers can improve performance by several orders of magnitude
 - Read-cache acts as a buffer for reads. When the same data is read multiple times, it is served from cache
 - Write-through cache: data is written to cache and then to disk, and only acknowledged as written when the data is physically written on the disk
 - Write-through cache: allows the disk controller to acknowledge the data as written as soon as it is held in cache. This allows the cache to buffer writes quickly and then write the data to the slower disk when the disk is ready to accept new I/O operations
- The type and amount of cache needed depends on what applications need
 - A web server, for instance, will mostly benefit from read-cache, whereas most databases are better off with write cache

Storage tiering

- Tiered storage creates a hierarchy of storage media, based on cost, performance requirements, and availability requirements
- Example:
 - Tier 1: Production data (SSD and SAS disks)
 - Tier 2: Seldom used data, like email archives (NL-SAS disks)
 - Tier 3: Backups (Virtual Tape Libraries on NL-SAS disks)
 - Tier 4: Archived data (Tape or NL-SAS disks)
- The more tiers are used, the more effort it takes to manage the tiers
- Automated tiering usually checks for file access times, file creation date, and file ownership, and automatically moves data to the storage medium that fits best

Load optimization

- Storage performance is highly dependent on the type of load
- Most vendors recommend a specific storage configuration for their systems or applications
 - For example, Oracle recommends a combination of RAID 1 and 5 for its database in order to optimize performance

Storage security

Protecting data at rest

- Data can be:
 - In transit (transported over a network)
 - In use (by an application or a cache)
 - At rest (on a disk or a tape)
- Data at rest can be secured using encryption techniques
 - Prevent reading or writing data to disk or tape without the correct encryption/decryption key
- Disk encryption in the datacenter has limited benefits:
 - Databases and applications need to work with unencrypted data to perform useful work
 - Disk encryption is only useful when the disks are physically lost or stolen (laptops, desktops, or removable media)
 - Disks in the datacentre are in a physically secure area

Protecting data at rest

- Disk encryption in the datacenter is useful:
 - A disk drive might get in the wrong hands – for instance because it was removed after it was marked "faulty" and was never destroyed
 - In case of disk failure, having the data encrypted solves the issue of having potentially sensitive data on a disk that can't be accessed anymore, as it is defective
 - Maintenance contracts often require that a failed disk must be sent back to the vendor after replacing it with a new one. Without disk encryption, returning disks may not be possible since a failed disk cannot be erased anymore.
 - Full disk encryption makes it harder for an attacker to retrieve data from the "empty" space on the disks, which often contains traces of previously stored data.

Protecting data at rest

- Self-Encrypting Drives (SEDs):
 - Use in laptops and desktops
 - When an SED is powered up, authentication is required to access data – the user must type in a password to start the boot sequence of the computer
 - Encryption is built into the disk drive's hardware
 - Encryption keys are stored on the disk
- Cryptographic Disk Erasure (CDE):
 - Deletes the encryption key on the disk
 - This has the same effect as erasing all disk contents
 - Without the key, unencrypted data can no longer be read from the disk
 - One of the best ways to fully wipe a disk's contents

SAN zoning

- SAN zoning is a method of arranging Fibre Channel devices into logical groups on a SAN fabric for security purposes
 - SAN zoning is implemented in the SAN switches
 - SAN zones are comparable with VLANs in Ethernet networks
 - Fibre Channel devices can only communicate with each other if they are members of the same zone

SAN LUN masking

- In a SAN, LUN masking makes a LUN available to some hosts and unavailable to other hosts
- LUN masking is implemented primarily at the HBA level, not in the SAN switches
- It is good practice to use a combination of SAN zoning and LUN masking