

Storage Systems

Vidhya Sagar
SQL Server MVP



Professional Association for SQL Server

What's in this Session?

- Intro to Storage system
- Hierarchy of storage
- Storage Sharing
- RAID
- Storage Best Practices for SQL Server
- Q & A

What's in this Session?

- Intro to Storage system

Intro to Storage system

What are we going to store?

Data

Where Data is stored?

Tables

Where the table resides?

File Groups

Where FG resides?

Files (MDF, LDF
& NDF)

How the files are stored?

Binaries (0's & 1's)

Where the binaries are stored?

Storage
Medium

Intro Storage system contd..

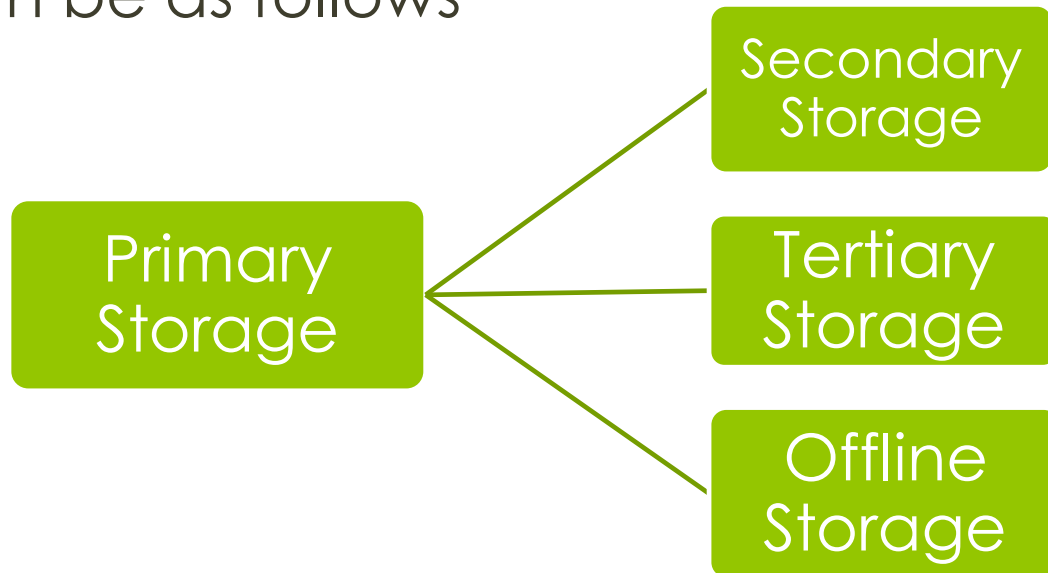
- A computer system can store and retrieve the data only in binary format.
- All the data (audio, video, documents, etc) stored will be converted into bits and stored in the storage medium.
- These binary units are stored as blocks into the storage medium and the most common unit is a byte (nothing but 8 bit as 10101010)
- All these units of bytes decides the size of the storage medium. Ex 8 million bits is equal to one MB
- To satisfy large storage & usage requirement, we have some different types of storage mediums

What's in this Session?

- Intro to Storage system
- Hierarchy of storage

Hierarchy of storage

- Based on the usage storage hierarchy can be as follows



Hierarchy of storage contd..

- Primary Storage

- Simply referred to as Memory
- Only One directly accessible to CPU




- Offline Storage

- It's a medium or device which is not under the control of CPU



Hierarchy of storage contd..

- Tertiary Storage
 - 3rd level of storage.
 - Involves robotic mechanism to mount or dismount storage medium based on requirement
 - Usage is mainly for archival process



Hierarchy of storage contd..

- Secondary Storage
 - 3rd level of storage.
 - Involves robotic mechanism to mount or dismount storage medium based on requirement
 - Usage is mainly for archival process



What's in this Session?

- Intro to Storage system
- Hierarchy of storage
- Storage Sharing

Storage Sharing

- Refers to how you are accessing the storage from CPU.
- Choosing a sharing method is based on personal and individual decision. However some of the most common factors to consider are
 - Capacity
 - Performance
 - Scalability
 - Availability and Reliability
 - Data Protection &
 - Budget Concerns
- Storage sharing is segregated mainly into 3 categories
 - DAS – Dedicated Attached Storage
 - NAS – Network Attached Storage
 - SAN – Storage Area Network

DAS – Dedicated Attached Storage

- Most basic level of storage
- Storage is directly connected to the server or host machine
- Reads data at block level
- Less Expensive than NAS or SAN
- Sharing is not possible
- Doesn't have advanced features as replication or snapshot
- It's not scalable
- Performance is less compared to SAN or NAS

NAS – Network Storage Access

- Can be connected virtually through SMB or NFS protocols
- Reads at file level
- It's Shareable
- NAS is scalable in terms of capacity and performance
- It's costly compared to DAS
- Provides advanced features

SAN – Storage Area Network

- Can be connected virtually through AoE, FC, iSCSI etc
- Reads at block level
- It's Shareable
- SAN is scalable in terms of capacity and performance
- It's costly compared to DAS
- Provides advanced features with synchronous replication

What's in this Session?

- Intro to Storage system
- Hierarchy of storage
- Storage Sharing
- RAID

RAID

- Redundant Array of Inexpensive Disks or Redundant Array of Independent Disks
- High level storage reliability with low cost
- Distributes data across multiple disks but array is seen as single disk in OS
- Combines two or more physical hard disk into one single logical disk, this can be achieved using hardware or software
- There are three main key concepts for RAID
 - Mirroring
 - Striping
 - Error Correction
- Based on these 3 concepts there are many types of RAID schemes available.
 - RAID 0,1,2,3,4,5,6,7,0+1,1+0,0+3,3+0,0+5,5+0,1+5
- We are going to discuss only most commonly used RAID schemas, which are RAID 0,1,5 and 10

RAID 0

- Strips the data across the disks
- Requires minimum of two disks
- This scheme is the fastest of all the RAID's since no burden on read or write, so more IOPS
- Easiest configuration
- Data cannot be retrieved if any of the disks get failed
- If N number of disks are there then the usable disks will be N


Reads	Writes
Good	Good

RAID 0

contd..

Striping	Mirroring	Error Correction
✓	✗	✗

RAID 0
striping



RAID 1

- Mirroring the data across the disks
- Data can be available even if a disk gets failed
- Requires minimum of two disks
- If N number of disks are there then the usable disks will be $N/2$


Reads	Writes
Good	Fairly Ok

RAID 1

contd..

Striping	Mirroring	Error Correction
x	✓	x

RAID 1 mirroring



RAID 5

- Strips the data across disk and store with parity for error correction
- Commonly used RAID method since it achieves good balance between performance and availability
- Requires minimum of three disks
- If N number of disks are there then the usable disks will be N-1


Reads	Writes
Good	Poor

RAID 5

contd..

Striping	Mirroring	Error Correction
✓	✗	✓

RAID 5 parity across disks



RAID 10


- It's a hybrid or Nested RAID system with RAID 1 mirrors with a RAID 0 stripe
- Provides high performance fault tolerance system
- Requires minimum four disks
- If N number of disks are there then the usable disks will be $N/2$

Reads	Writes
Good	Good

RAID 10

contd..

Striping	Mirroring	Error Correction
✓	✓	✓



What's in this Session?

- Intro to Storage system
- Hierarchy of storage
- Storage Sharing
- RAID
- Storage Best Practices for SQL Server

Storage Best Practices for SQL Server

- Study the IO characteristics of SQL Server & specific IO requirements
 - Average read bytes/sec, average write bytes/sec
 - Reads/sec, writes/sec
 - Disk read bytes/sec, disk write bytes/sec
 - Average disk sec/read, average disk sec/write
 - Average disk queue length
- Make sure you have more disk drives for large size LUN, however some cases performance goes down after 8 drives.
- Faster spindles for better performance.

Storage Best Practices for SQL Server contd..

- Isolate log, data and tempdb files at physical disks
- Place log files in any of the RAID shown below
 - RAID 10 - Preferred RAID
 - RAID 1 - If not RAID 10, then choose this
 - RAID 5 - If not RAID 1, then choose this
- Place data files in any of the RAID shown below
 - RAID 10 - Preferred RAID
 - RAID 5 - If not RAID 10, then choose this
 - RAID 1 - If not RAID 5, then choose this

Storage Best Practices for SQL Server

contd..

- Place tempdb files in any of the RAID shown below
 - RAID 1 - Preferred RAID
 - RAID 10 - If not RAID 1, then choose this
 - RAID 0 - Use with caution*
- Pre-size your data and log files so that it won't screw up disk IO
- If possible turn on Instant initialization
- Always run storage array with latest update.

What's in this Session?

- Intro to Storage system
- Hierarchy of storage
- Storage Sharing
- RAID
- Storage Best Practices for SQL Server
- Q & A

Q & A

Contact me @

kvs1983@gmail.com

www.sql-articles.com/blogs/

www.sql-articles.com/forums/

www.twitter.com/sqlarticles/

<http://groups.google.com/group/chnsqlug>

Feedback @ <http://speakerrate.com/kvs1983>