

SQL Server, Storage And You

Part 2: SAN, NAS and IP Storage

What we are going to learn

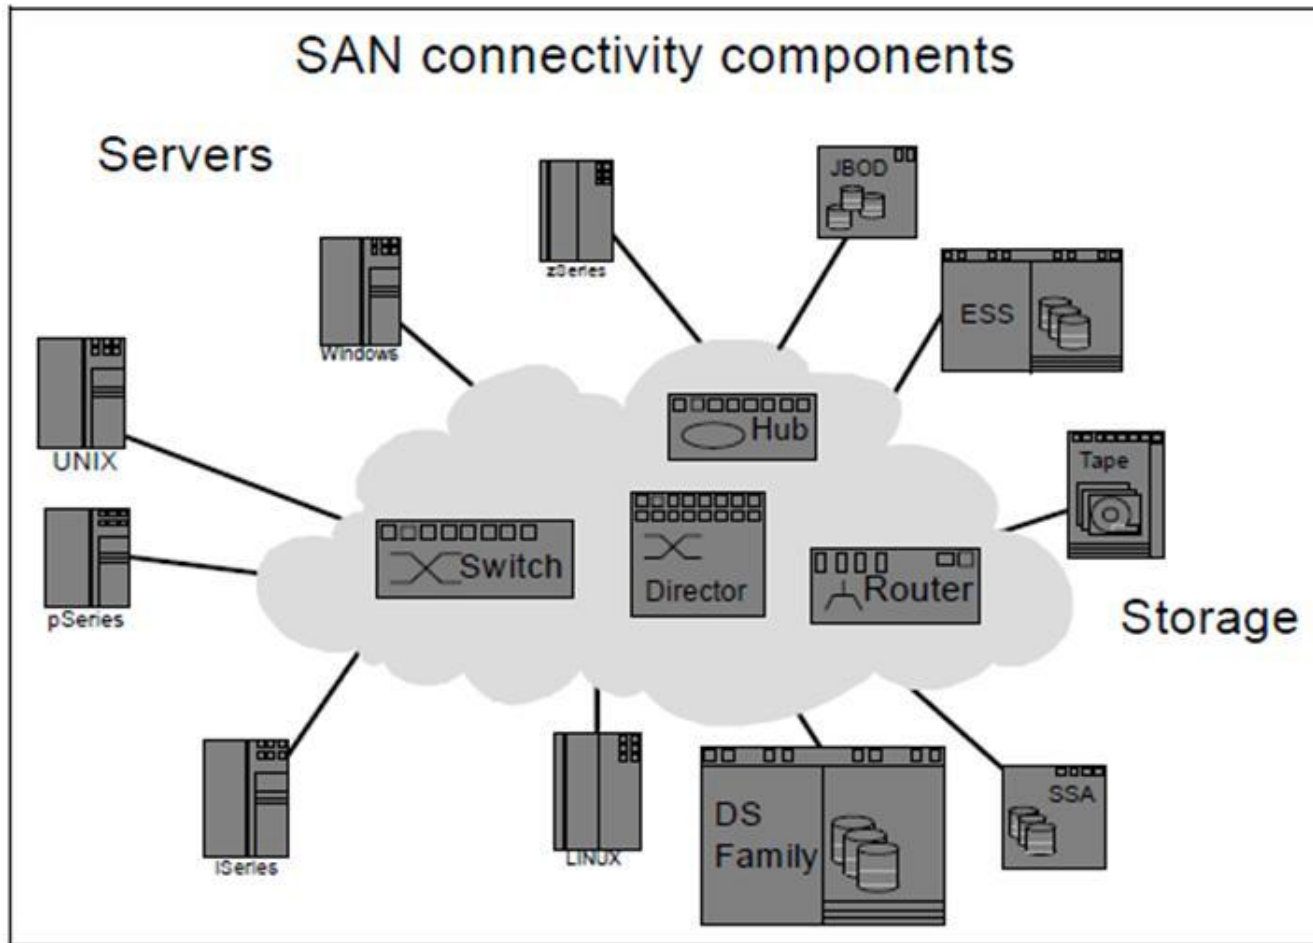
- What makes up a Storage Area Network
- Network Attached Storage, sharing files
- IP Storage, blurring the lines

In The Beginning...

- Islands of disks
 - Disjointed
 - Over utilization
 - Under utilization
- Unreliable
 - Difficult to make highly available
 - Limited disaster recovery options
 - Hard to manage



A Network By Any Other Name



Network Behind Your Servers

- SAN makeup
 - Dedicated Network
 - Switches, routers, directors oh my!
 - Consolidation
 - Utilization
 - Server-less backups
 - Speed *
 - SAN != Fibre Channel
 - FC copper
 - Infiniband

More Than Disks

- Tape Systems
 - Drives
 - Autoloaders
 - Libraries
- Chameleon storage
 - Expose as NAS
 - Expose as iSCSI

More Than Hardware

- Storage Processors
 - Dedicated
 - Redundant
- Software
 - Disk management
 - JOBD
 - RAID
 - Replication
 - Block level
 - LAN/WAN

Serving Files Network Style

- Commodity Servers
 - X86/X64
 - Dense
- Standard Networking
 - TCP/IP Stack
- Ease Of Implementation
 - Lower Management Overhead

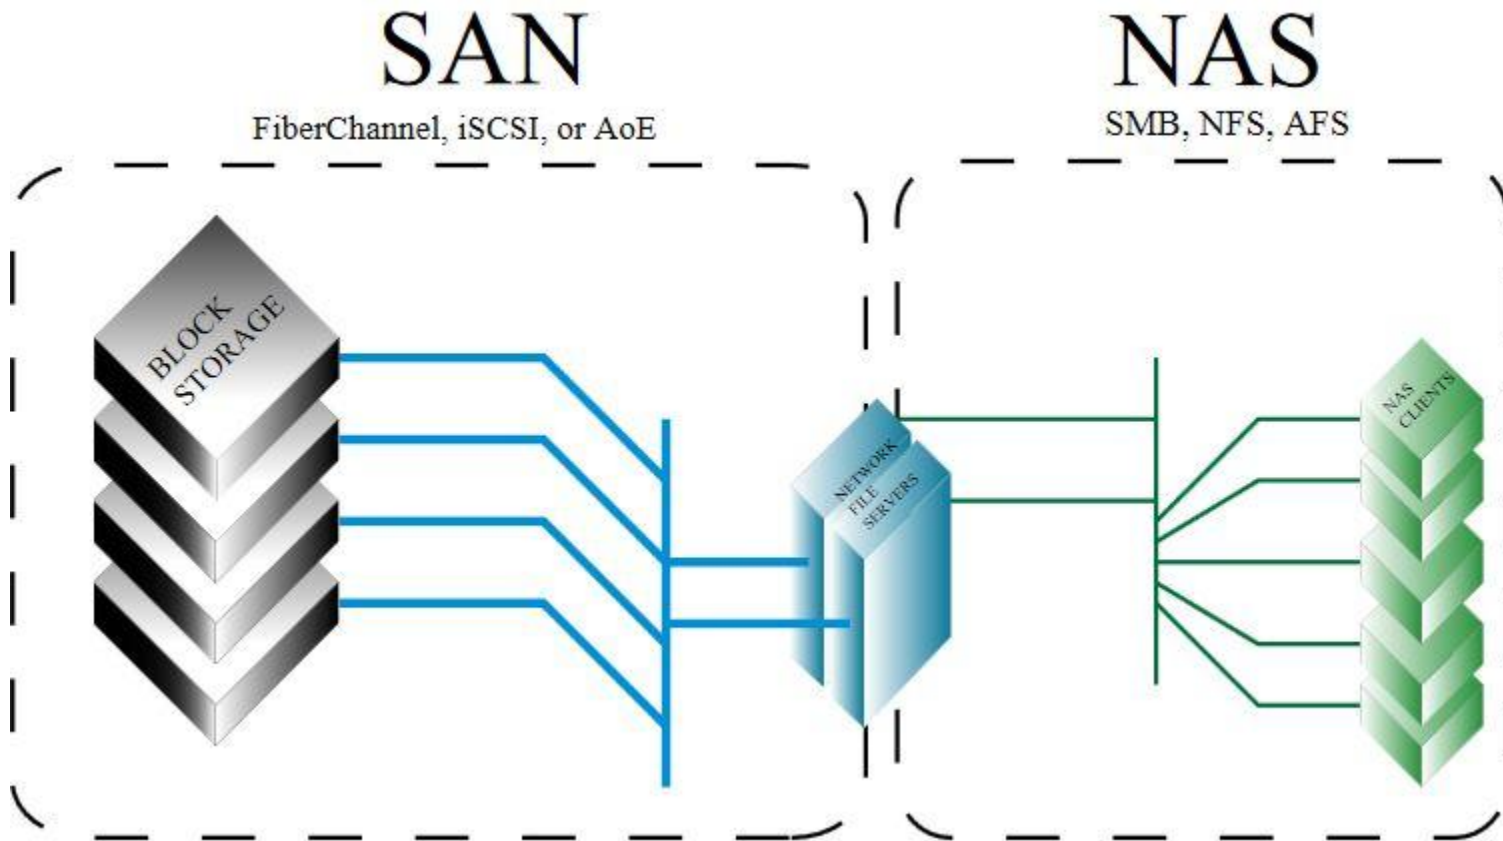
Server Message Block

- Meant to live on top of TCP/IP
 - Build initially by IBM
- Offer abstraction
 - File shares
 - Printers
 - IPC
- Not supported for SQL Server prior to 2008 R2

The Similar But Different

- Redundant
 - Both use RAID
 - May be twists on a theme
 - Both have replication
 - Block Vs. File
 - Both use a “network”
 - One built specifically
 - One uses protocol layers on IP
- !Danger!
 - Not completely safe for SQL Server *

SAN Behind, NAS In Front



iSCSI, Building Bridges

- Built on TCP/IP
 - Bundles SCSI commands
 - Common network infrastructure
 - != SAN
- Block level storage abstraction
 - Direct attached storage
 - stand alone server acts as storage head
 - Network attached storage
 - NAS can still serve SMB/NFS
 - Storage Area Network
 - iSCSI being built into bridge heads or directly into storage head
 - Support for MPIO and other SAN features

SAN Pitfalls

- High utilization
 - Servers outside your control effect your performance
- Poor Configuration
 - Striping data “wide and thin”
 - Over subscribing single disks “hot spots”
 - Fabric misconfiguration
 - SAN replication over large distance
- Expensive to scale
 - SAN disk can be 10x DAS/NAS
 - Multiple FC HBA’s needed for high throughput
 - Dedicated network infrastructure

NAS Pitfalls

- Meant to serve files
 - Generally, block level storage
 - Higher latency than SAN/DAS
- May not honor “no cache” flags
 - Puts data at risk
- Limited disk configurations
 - Almost always “wide and thin”
- VLAN != separate network
 - Often share traffic with other apps
 - QoS doesn't fix this

iSCSI Pitfalls

- Hides back end storage
 - Are you on a fibre SAN or NAS?
 - NAS vendors offering iSCSI calling it SAN
- Easy to setup, hard to get it right
 - Should have dedicated network
 - Should have ToE or initiator HBA's
- Back end may not honor “no cache” flags
 - Puts data at risk

SAN General Configuration

- Always ask for IO's not gigabytes
 - Space is cheap
 - SQL Server eats IO
- Dedicated drives
 - Request LUN's dedicated to data files
 - Request LUN's dedicated to log files
 - Request LUN's dedicated to tempdb
- Multiple FC ports
 - Separate IO traffic
 - Provide redundancy

Network Attached Storage

- Verify back end storage
 - Is it a DAS or SAN?
- Request drive pool separation
 - Pool data together
 - Separate logs from data
- Request multiple NIC ports
 - Provides load balancing
 - Provides redundancy
 - Separate IO workloads
 - Configure for jumbo frames
 - Separate client requests

iSCSI General Configuration

- Separate network
 - VLAN isn't separate
 - Reduce routing
 - Configure for jumbo frames
- Require ToE/Initiator HBA's
 - Reduces server load
 - Can speed up IO
- Request 10 gigabit
 - MUCH larger pipe 125MB/sec Vs. 1250MB/sec
- Request multiple ports
 - MPIO
 - Separate IO workloads

Monitoring IO Health

- Very few vendor neutral tools
- Response Time = Service Time + Wait Time
- Disk Queue Length
 - Caches mask DQ
 - Focus on latency and waits
 - `sys.dm_io_virtual_file_stats`
 - Gives you time to read and write IO's
 - Gives you amount of data written and read at the file level
 - Great for finding SAN hot spots
 - <http://sqlserverio.com/2011/02/08/gather-virtual-file-statistics-using-t-sql-tsql2sday-15/>
 - `sys.dm_os_wait_stats`
 - Gives you what SQL Server is doing besides IO
 - Only at a instance level

Questions?

Email: wes@planetarydb.com

Twitter: @WesBrownSQL

Blog: <http://www.sqlserverio.com>