

Storage Triage: Troubleshooting Slow SQL Servers

Brent Ozar

Brent Ozar PLF

BrentO@BrentOzar.com

SQL Server
CONNECTIONS

BrentOzar.com/go/san

Brent Ozar PLF

Your technology pain-relief experts.



[HOME](#) [WHO WE ARE](#) [WHAT WE DO](#) [COMMUNITY](#) [CONTACT US](#)

You Are Here: [Home](#) » [SQL Server Articles](#) » SAN Storage Best Practices for SQL Server

SAN Storage Best Practices for SQL Server

I bet you're here because you don't trust your SAN administrator. The SAN admin's been telling you everything's fine, and that it must be a SQL Server problem, right?

Odds are, it's not. Time to learn what SQL Server needs from its storage, and what SAN admins need from us database administrators.

SQL Server Storage Performance Tuning

Whew! We've got the SAN controllers, pathing, and Windows set up correctly, and now it's time to install SQL Server and put our databases on it. Here's my thoughts on how to do it:

- [SQL University Part 1: The Basics of Storage](#) – the basics of data files, log files, and cache.
- [SQL University Part 2: Where We Store Stuff](#) – RAID levels, magnetic drives versus solid state, and where to put your data and log files.
- [SQL University Part 3: Scary Storage Scenarios](#) – why DBCC is so important in SANs and virtualization.

About Brent Ozar PLF



Brent Ozar PLF is a boutique consulting firm focused on understanding your environment and strategy.

...

About @BrentO

Microsoft
CERTIFIED
Master



Today's Big 5 Questions

1. Is storage really a big bottleneck?
2. Can we fix it easily?
3. If not, how do we prove the problem?
4. Can solid state help?
5. How do we do a storage renovation?

How SQL Server Schedules

What's Running Now

What's Waiting (Queue)

How SQL Server Schedules

What's Running Now

- `SELECT * FROM
dbo.Restaurants`
(By Brent)

What's Waiting (Queue)

How SQL Server Schedules

What's Running Now

- `SELECT * FROM
dbo.Restaurants`
(By Brent)


What's Waiting (Queue)

- `SELECT * FROM
dbo.LoveSongs`
(By Celine)
- `SELECT * FROM
dbo.MagicTricks`
(By Penn)

How SQL Server Schedules

What's Running Now

- `SELECT * FROM dbo.Restaurants`
(By Brent)



What's Waiting (Queue)

- `SELECT * FROM dbo.LoveSongs`
(By Celine)
- `SELECT * FROM dbo.MagicTricks`
(By Penn)

How SQL Server Schedules

What's Running Now



What's Waiting (Queue)

- `SELECT * FROM dbo.LoveSongs`
(By Celine)
- `SELECT * FROM dbo.MagicTricks`
(By Penn)
- `SELECT * FROM dbo.Restaurants`
(By Brent)



SQL Server
CONNECTIONS

That's the Simple Version

What's Running Now

- `SELECT * FROM`
`dbo.LoveSongs`
(By Celine)

What's Waiting (Queue)

- `SELECT * FROM`
`dbo.MagicTricks`
(By Penn)
- `SELECT * FROM`
`dbo.Restaurants`
(By Brent)



SQL Server
CONNECTIONS

Slightly More Complicated

Running

- SELECT *
FROM
dbo.LoveSongs
(By Celine)

Runnable

- SELECT *
FROM
dbo.MagicTricks
(By Penn)

Waiting (Queued)

- SELECT * FROM
dbo.Restaurants
(By Brent)
- SELECT * FROM
dbo.Deodorants
(By Mennen)

What We Wait For

- **Resources**
 - CPU, memory, network, locks, storage (disk)
- **System Tasks**
 - Lazywriter, trace, full text search
- **Stuff Outside SQL Server (Preemptive)**
 - COM, OLEDB, CLR

How to Track Waits Over Time

What's happening right now:

- DMV: `sys.dm_os_waiting_tasks`
- Tool: `sp_WhoIsActive` – BrentOzar.com/go/active

What's been happening over time:

- DMV: `sys.dm_os_wait_stats`
- Tool: Glenn Berry's DMV queries – BrentOzar.com/go/glenn
- Drawback: doesn't tie back to individual queries

Storage Waits To Look For

- PAGEIOLATCH*
- WRITELOG
- IO_COMPLETION
- ASYNC_IO_COMPLETION

Question #1:

Is storage really a big bottleneck?

Answer:

If those waits show up a lot, yes.

Question #1:

Is storage really a big bottleneck?

Answer:

If those waits show up a lot, yes.

Next Question:

Can we fix it easily?

When I say “Can we fix it easily?”,

“we” doesn’t mean the SAN admin.

“We” means us – the database people.

<http://flickr.com/photos/gowestphoto/3921760653/>



Why Storage Struggles

- **Doing too many reads**
 - Insufficient memory to cache data
 - Missing indexes (incl. clustered indexes)
 - Non-sargable queries (LIKE '%smith%')
 - One misbehaving database dominating RAM
- **Doing too many writes**
 - More indexes than we need
 - Page splits

DDR: Doing Database Research

1. Run `sp_Blitz`
2. Run Glenn Berry's DMV queries
3. Run my index tuning queries

Get 'em all at BrentOzar.com/go/san

**Your mission:
spend just one day
looking for low-hanging fruit.**

**Question #2:
Can we fix it easily?**

**Answer:
Spend 10 hours tuning and/or
\$2,000 on memory to find out.**

**Question #2:
Can we fix it easily?**

**Answer:
Spend 10 hours tuning and/or
\$2,000 on memory to find out.**

**Next Question:
If not, how do we prove there's a
storage problem?**

Providing Proof

- Wait stats prove we're waiting on storage
- Perfmon proves whole volumes are slow
- DMVs prove which SQL files are suffering

Perfmon Physical Disk Counters

- Avg Disk Sec/Read
- Avg Disk Sec/Write
 - Measured in whole seconds,
so 3 decimal places is milliseconds
- Avg Disk Reads/Sec
- Avg Disk Writes/Sec

Sys.dm_io_virtual_file_stats

- One row per database file
- Resets on instance restart, restore, attach
- io_stall_read_ms, io_stall_write_ms:
Total latency over time
- num_of_reads, num_of_writes
- Divide time by quantity to get average

General Internet Guidelines

OLTP databases (transactional apps)

- Reads should be under 20 milliseconds
- Log writes should be 1-5 milliseconds

OLAP databases (data warehouses)

- Reads should be under 50 milliseconds
- Writes should still be ≤ 5 milliseconds

Brent's Real-World Guidelines™

- Reads should be under 100 milliseconds
- Writes should be under 20 milliseconds

Question #3:

How do we prove there's a storage problem?

Answer:

Use wait stats, Perfmon, and DMVs compared to my guidelines.

Question #3:

How do we prove there's a storage problem?

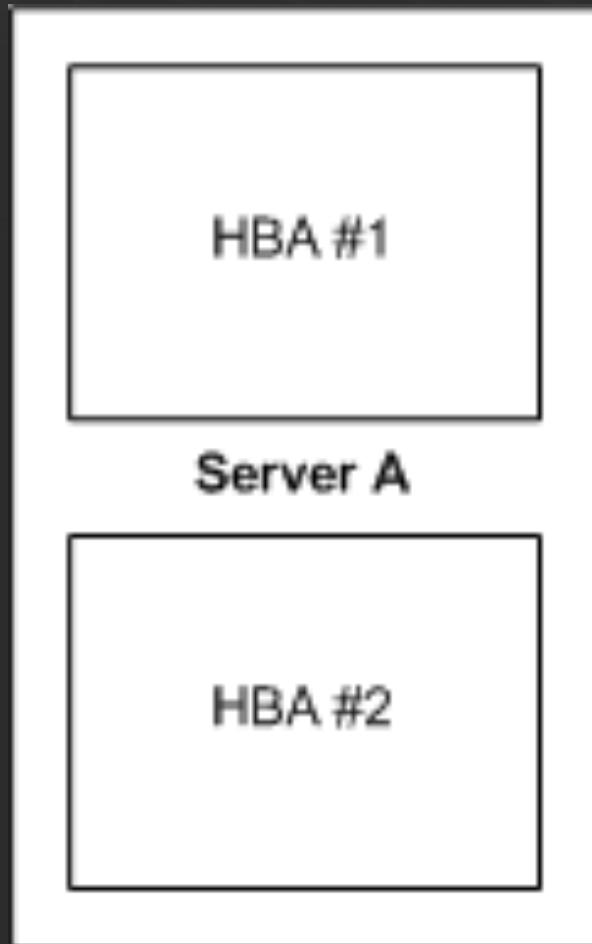
Answer:

Use wait stats, Perfmon, and DMVs compared to my guidelines.

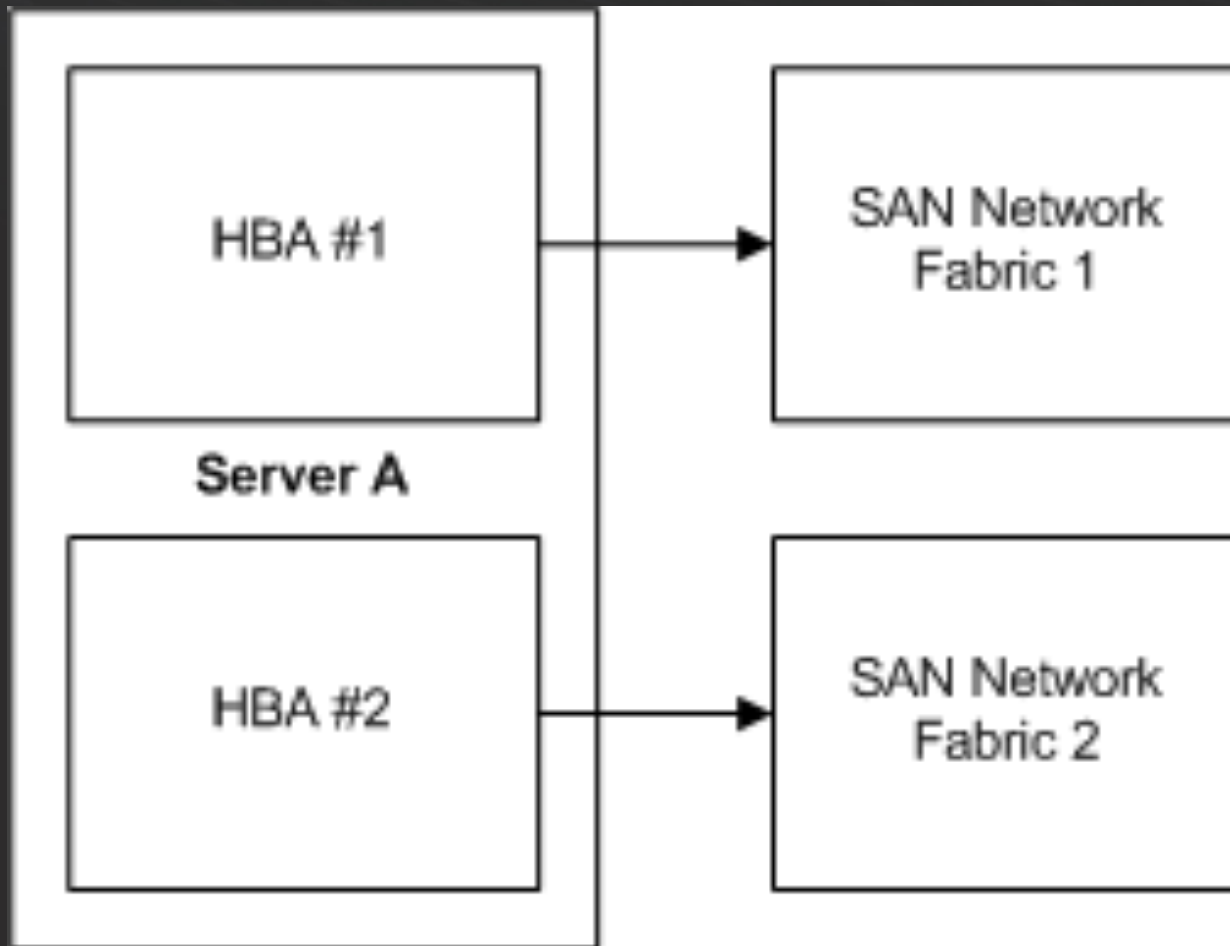
Next Question:

Can solid state help?

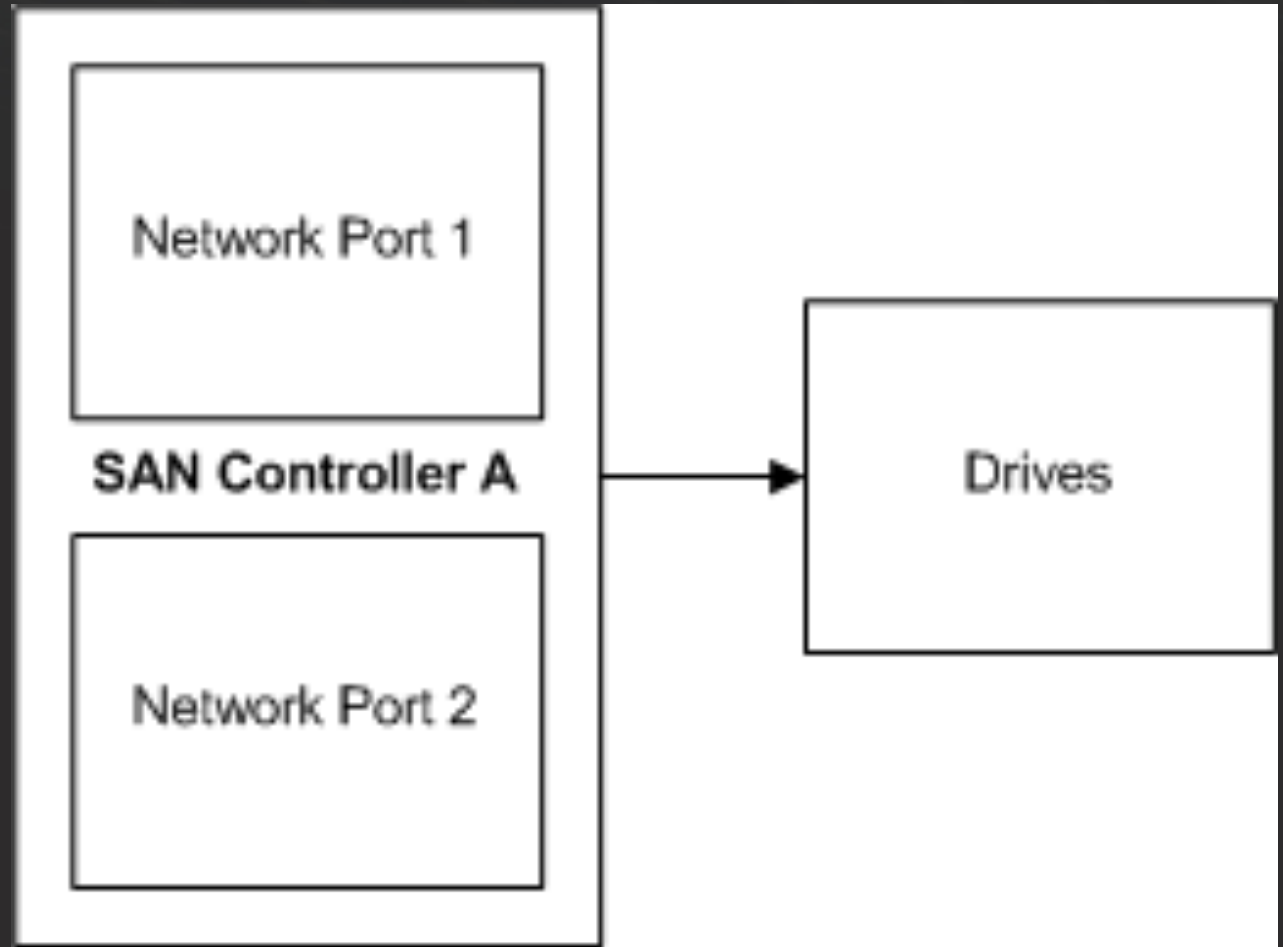
Simple SAN Paths



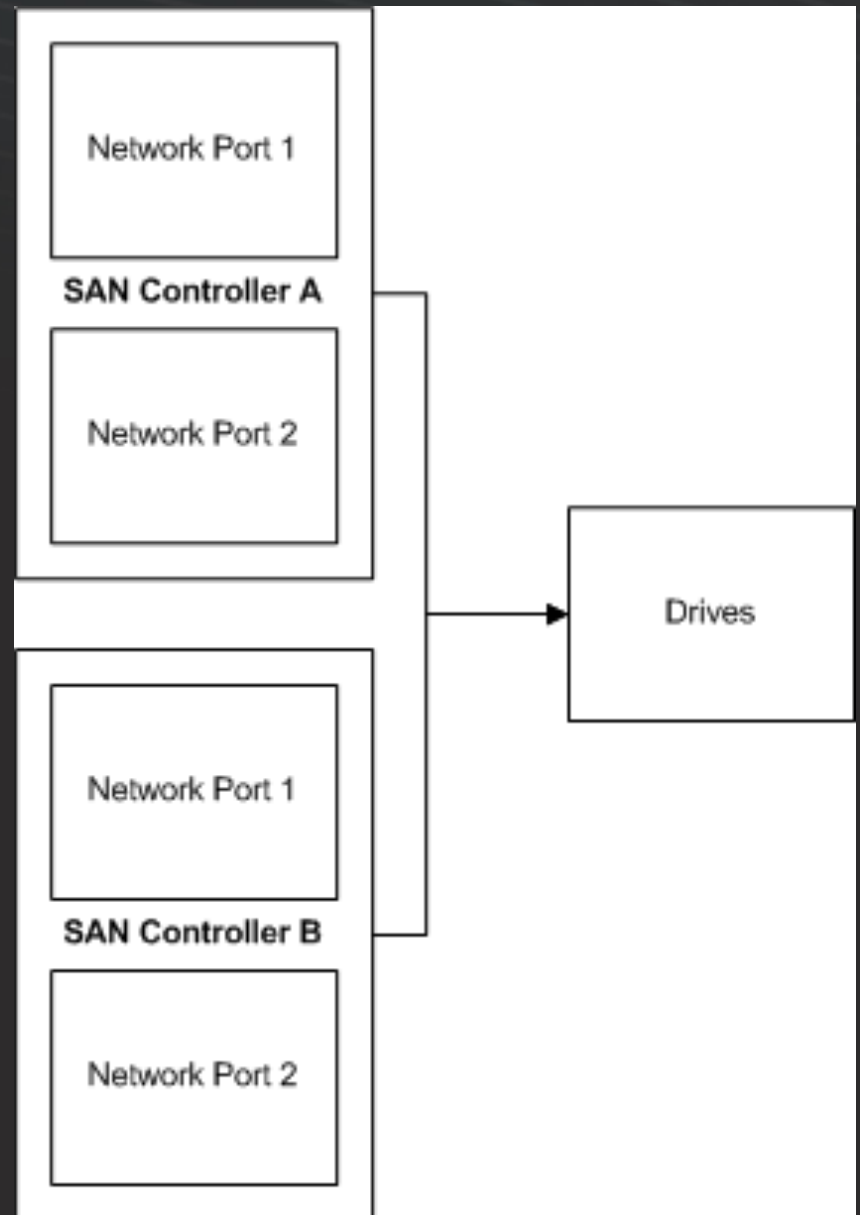
Simple SAN Paths



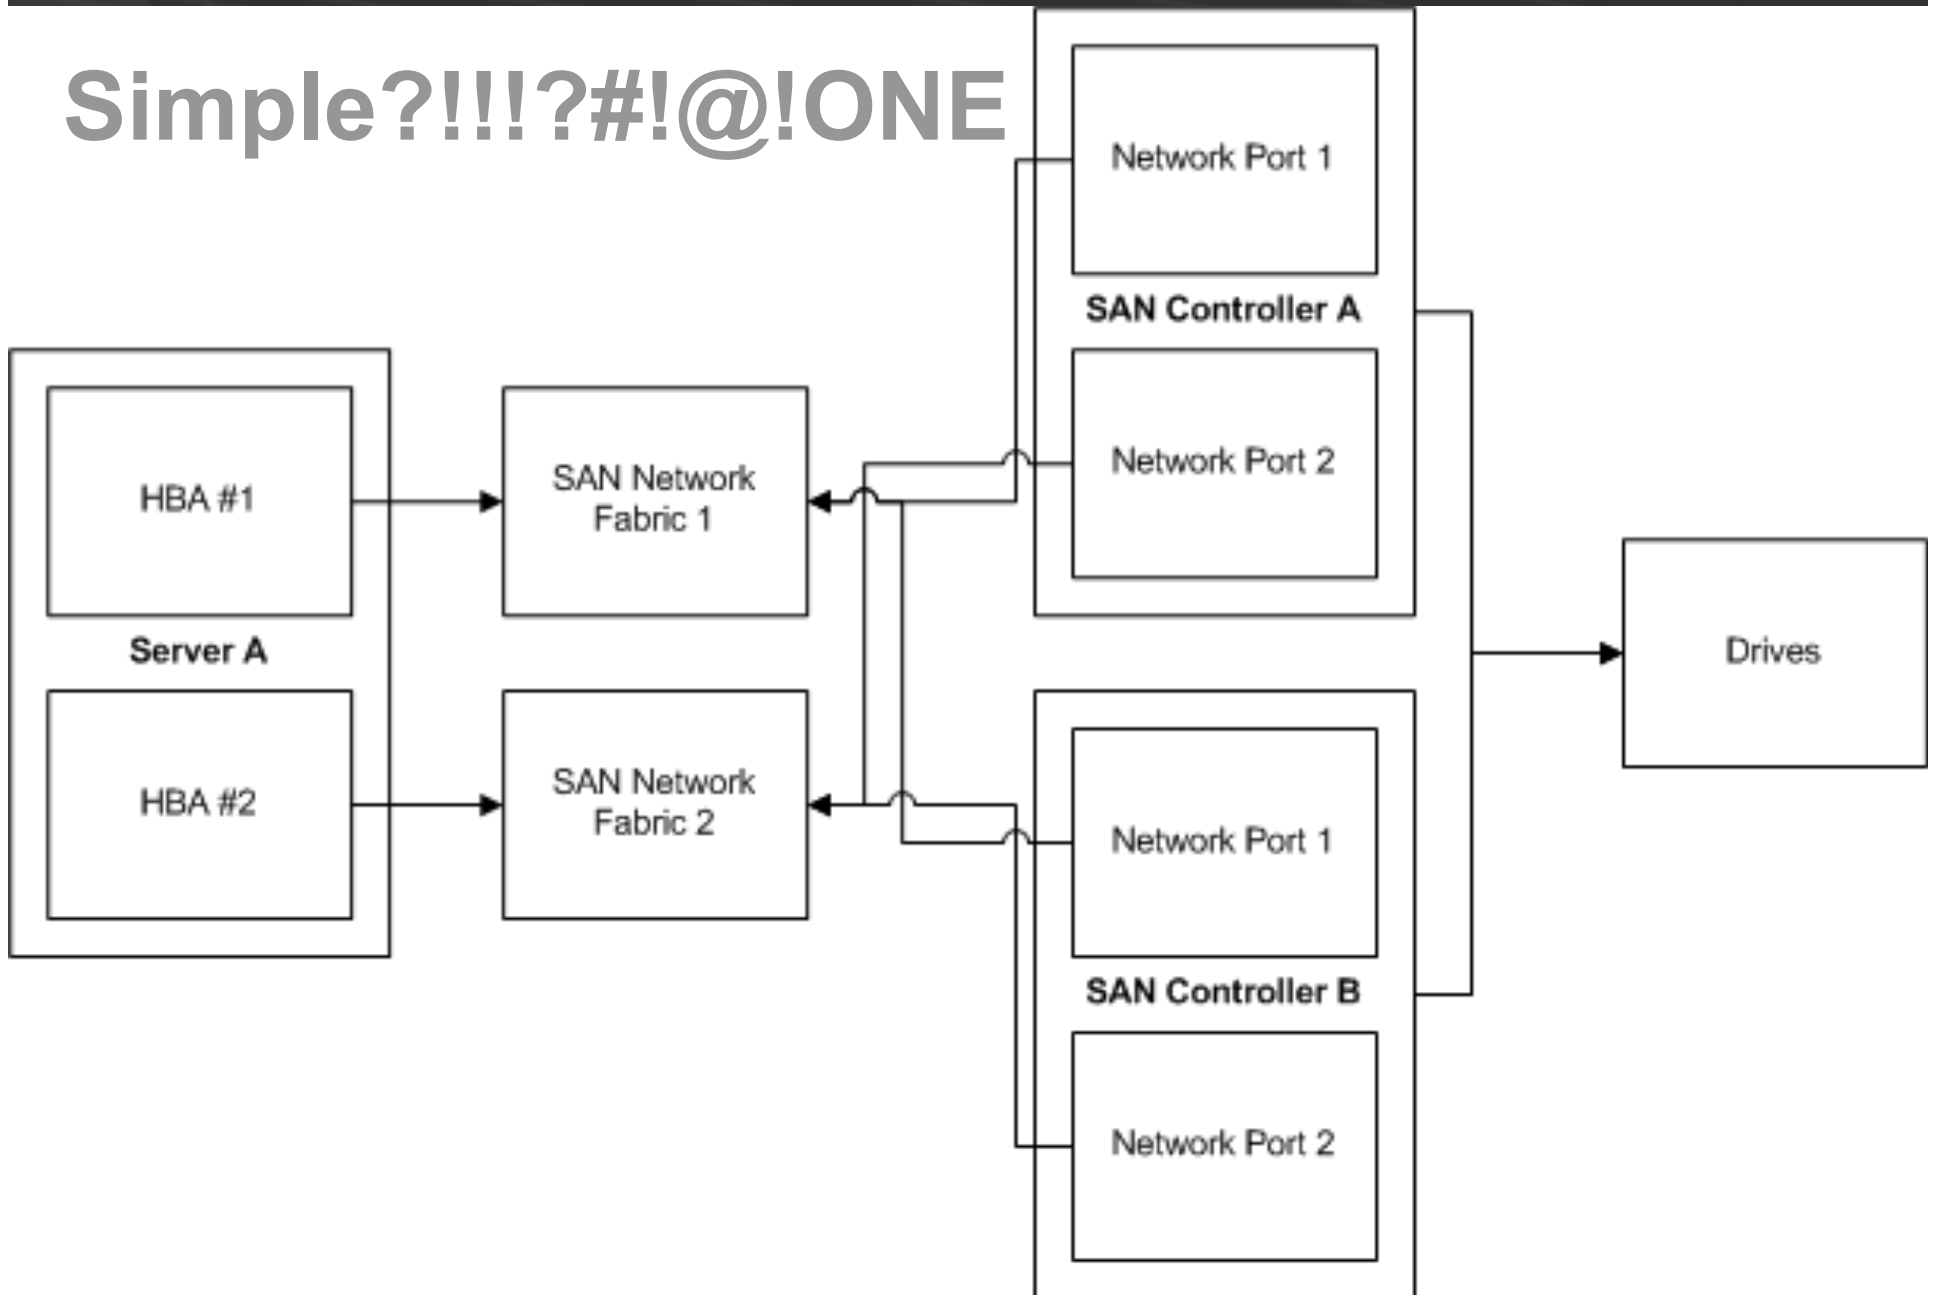
Simple SAN Paths




Simple SAN Paths


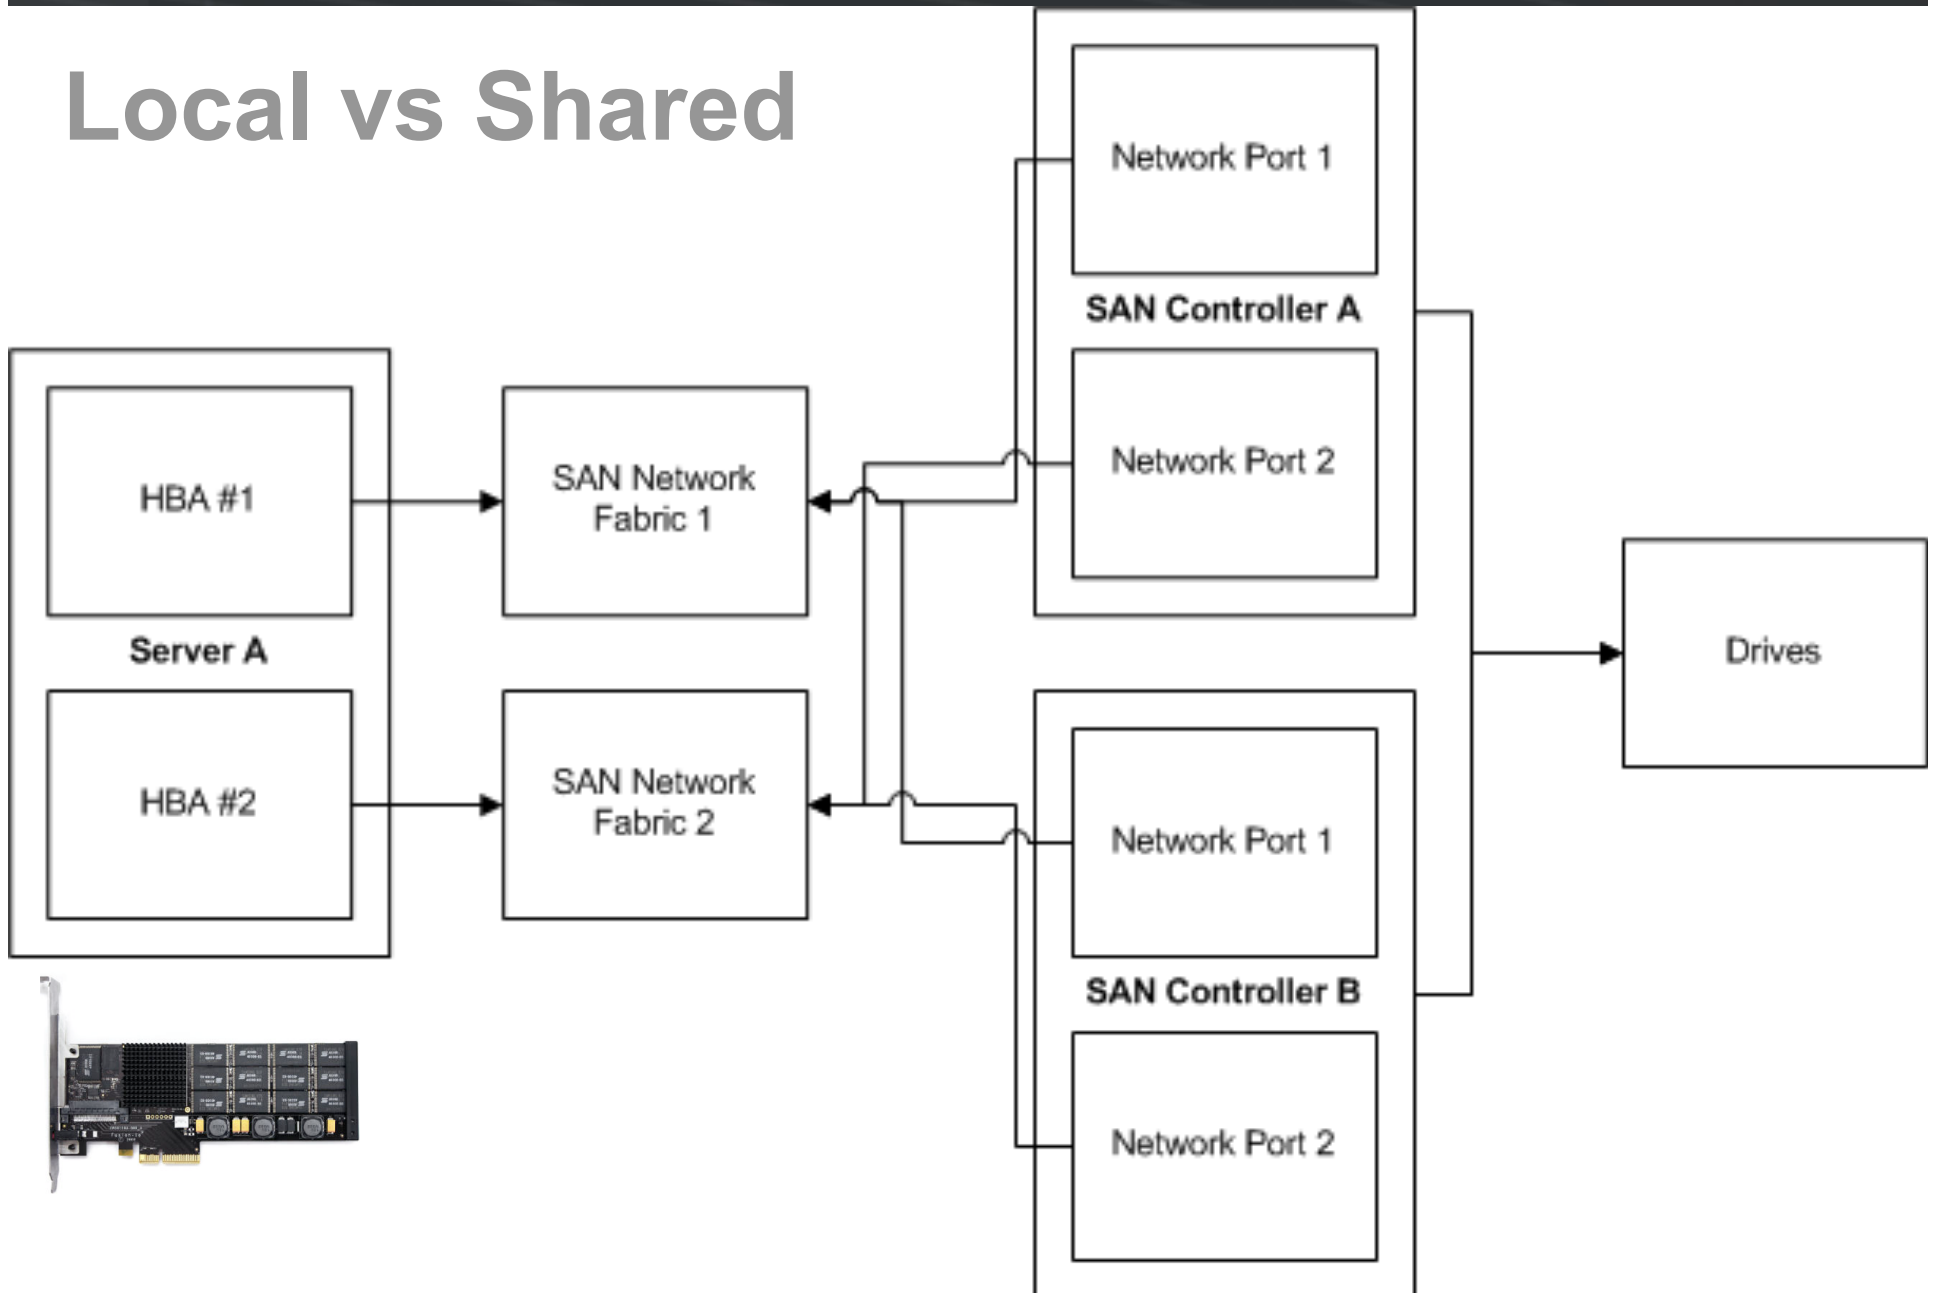


Simple?!?!?#!@!ONE





Local vs Shared



**Question #4:
Can solid state help?**

**Answer:
Usually yes, but know where to put
them. PCI Express drives rock.**

**Question #4:
Can solid state help?**

**Answer:
Usually yes, but know where to put
them. PCI Express drives rock.**

**Last Question:
How do we do a storage renovation?**

Very Typical Storage Problems

- Partitions aren't aligned
- 4kb NTFS allocation unit sizes
- Fragmented data files
- One ginormous data file
- Lots of virtual log files
- Small SAN stripe sizes
- Concatenated LUNs

Extreme Storage Makeover

- Put the databases in read-only mode
- Back up user databases (and test 'em!)
- Blow away the drives
- Rework them according to best practices
- Restore the databases
- Build 4 data files per filegroup
- Rebuild all indexes to balance data files
- Shrink/regrow the log files to fix VLFs

Question #5:

How do we do a storage makeover?

Answer:

Very carefully.

Question #5:

How do we do a storage makeover?

Answer:

~~Very carefully.~~ Take a weekend maintenance window and get a 30%, 50%, 100%, or higher improvement.

Top 5 Things To Remember

1. >1TB databases and TempDB need at least 4 data files on 4 LUNs
2. Sys.dm_os_wait_stats shows bottlenecks
3. Sys.dm_io_virtual_file_stats shows stalls
4. Track Perfmon Physical Disk Averages:
Sec/Read, Sec/Write, Reads/Sec, Writes/Sec
5. Learn more at BrentOzar.com/go/san

Your Feedback is Important

Please fill out a session evaluation and drop it off at the conference registration desk.

Thank you!