

Intro to SharePoint for SQL Server DBA's

Steve Knutson | Microsoft MVP Office Servers and Services

Stratos Technology Partners Managing Director

#614 SQL Saturday South Island




Who am I?



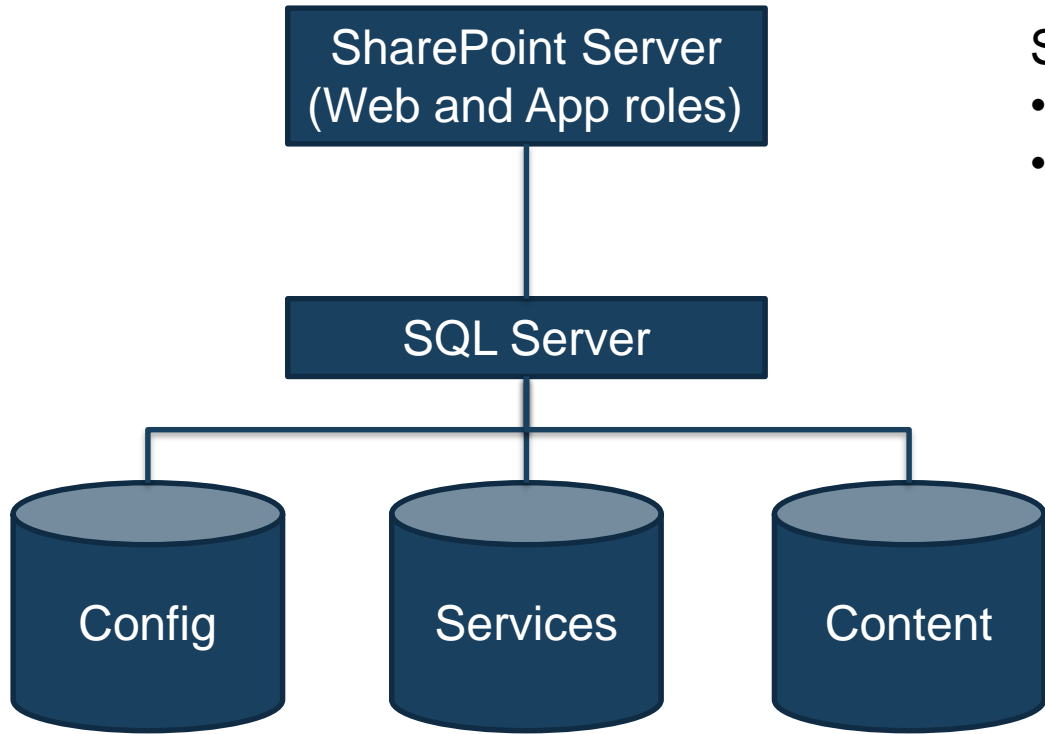
steve@stp.co.nz | @nztechtweet

Agenda

- A bit about SharePoint Architecture
- SharePoint Databases
- How big, how many and how fast?
- ~~Backup Recovery~~
- Maintenance do's and don'ts
- If you only remember a few things



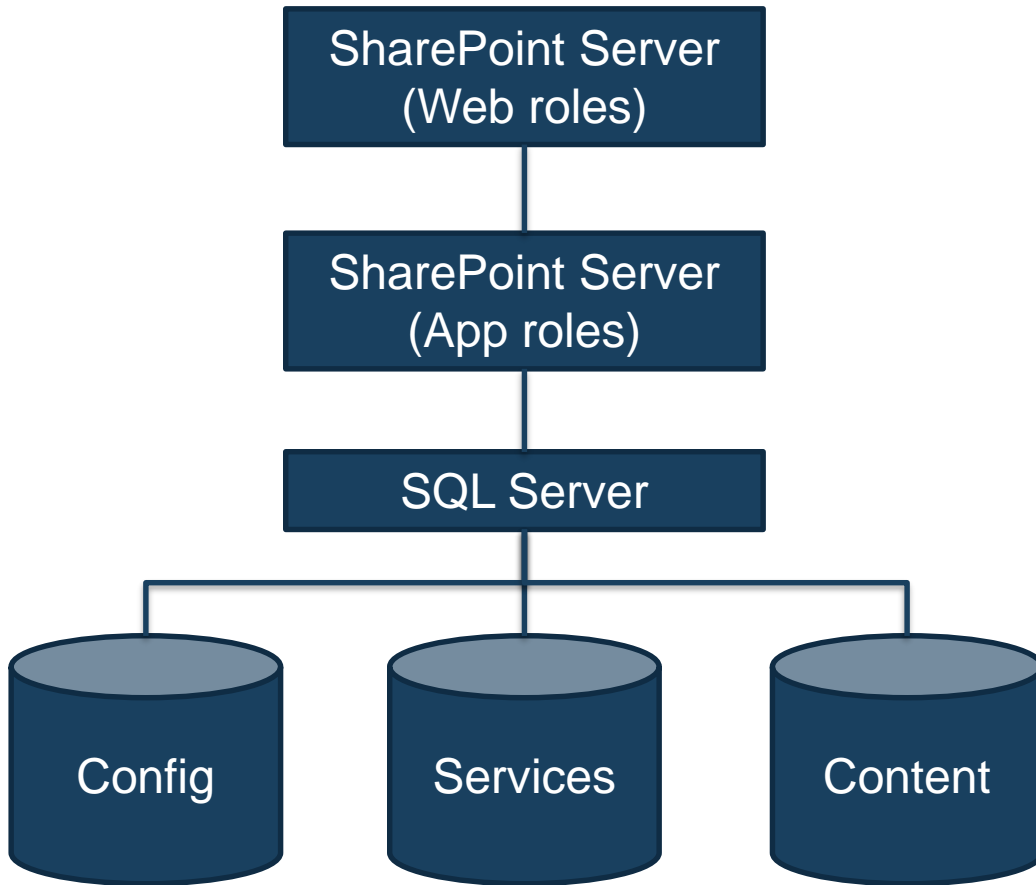
SharePoint Architecture 101



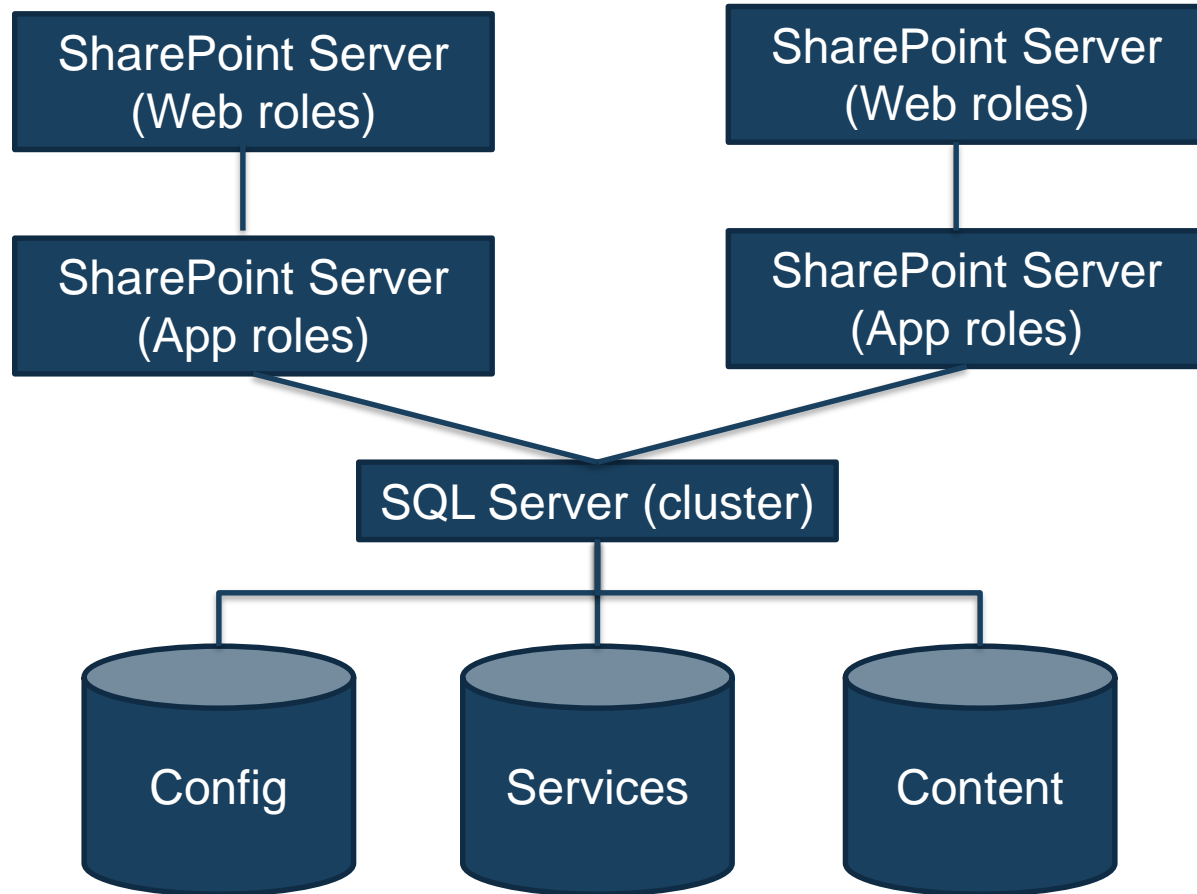
Suitable for:

- Dev and UAT
- Small Farms

Or this...




Or this!



What are all these databases for?

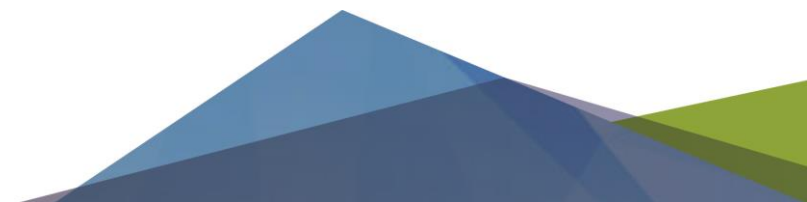
SharePoint has databases for three purposes

- Farm configuration – contains the configuration of the entire SharePoint farm
- Service Applications – configuration and data for SharePoint application services e.g. Search, User Profile, Managed Metadata etc
- Content databases – contains user data e.g. site configuration and data




A transaction is a document

SharePoint workloads are different in several ways

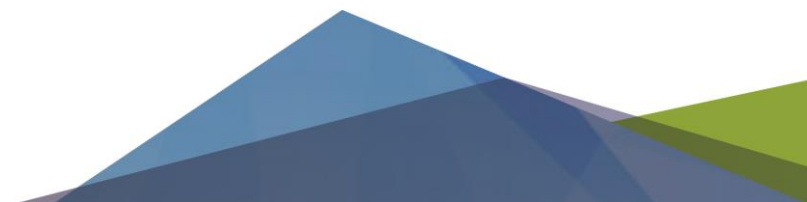
- Content databases can contain documents. An “HTTP Get” from SharePoint can return a document (they can be big!)
 - Read to Write ratios can vary – rule of thumb 80% read, 20% write
 - A ‘Full’ Search Crawl will read every item and document in the Content Databases, generating high I/O in SQL.
- 

Temp DB

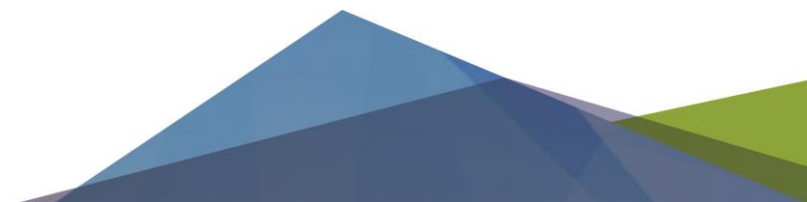
- Put on the fastest storage
- Use a dedicated LUN if possible
- Use “MB” rather than “%” auto-grow
- **TempDB size** = [MAX DB SIZE (KB)] X [0.25] / [No. CORES] = DATA FILE SIZE (KB)



Capacity Planning


- Soft limits (there are no hard limits)
 - 50,000 site collections per content db
 - 1TB per content db (but smaller is recommended)
 - 8 Web Front Ends per SQL server
 - Understand the size and growth profile of content
 - Versioning requirements
 - High growth vs archival
 - Access frequency
 - 80% Read, 20% write
- 

SQL Server Memory, Disk and Network

- RAM – Minimum 4GB (Dev), Good 8-16GB, Better >16GB
 - CPU cache - >2MB L2 Cache
 - CPU cores – Good 4 cores, Better >4 cores
 - Disk latency – Good < 10msec, Better <10msec (Data) <5msec (Tlog)
 - Network latency (SQL to SP server) – Less than 1ms
- 


SQL and SharePoint Server Config

- Keep up to date with patches but confirm SP supported version
- Use a dedicated SQL server if possible
- Plan your content databases and “scale out” if necessary
- Use SQL Aliases




Content Databases

- Using dedicated content databases per site collection allows you to recover individual site collections and other management flexibility
- This is important for recovery and DR



Remote BLOB Storage

- Has some limitations
 - No encryption or compression
- Best for...
 - Few large BLOBs (mostly read access)
- Worst for...
 - Lots of smaller BLOBs (or intensive write access)
- Requires Enterprise licensing
- Third-party add-ons make this much better
- Adds complexity




Business Data Connectivity

- Can present SQL views / tables in SharePoint
- Can write data back into SQL

SharePoint also supports oData and REST


- Use SSIS to pull data from lists (on-premises & cloud)



Access Services

- Access databases published to SharePoint
- Uses SQL to host tables NOT Access

- NB: Access Services is end of life for SharePoint Online




Business Intelligence

Services (on-premises)

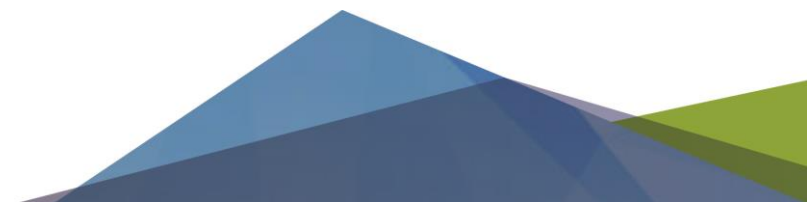
- PowerPivot
- PerformancePoint
- SQL Server Reporting Services

Cloud

- PowerBI (Pro license)



Secure Store

- Manages credentials for BI Services
 - Works with Microsoft and non-Microsoft data sources
 - Authentication:
 - Pass-through (logged in user)
 - Data access account
 - Non-windows accounts
 - Reduces need for Kerberos (but isn't a complete replacement)
- 

Which Recovery model?


Microsoft recommend using both the Simple and Full recovery models.

Full Model for these databases

- Configuration, Admin Content, Content, Secure Store, BDC and PerformancePoint

Simple Model for these databases:

- Search and User Profile Service



Backup and Restore


- Understand what you might need to recover
 - Library, List, File or list item
 - Site
 - Site Collection
 - Web application
 - Farm
- Use the recycle bins first
- Unattached database recovery
- Get sign-off on your RTO and RPO's, don't assume they are the same as other systems



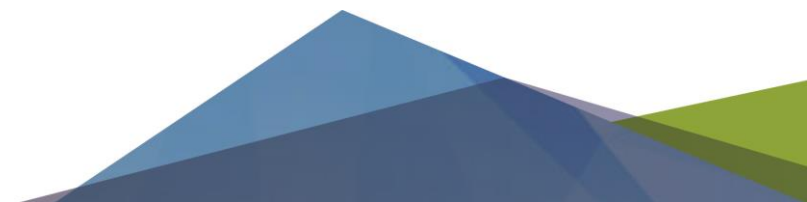
Maintenance

Indexes - Fragments happen! (who said by design?)

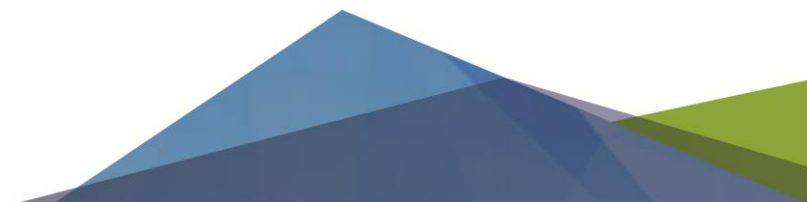
- Performance degrades as fragmentation increases – more I/O
- Content and Search DB's are most susceptible
- Use `sys.dm_db_index_physical_stats`, it is more accurate than `DBCC Showconfig`



Maintenance Do's

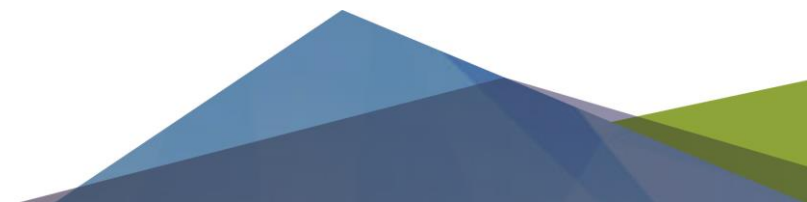
- Auto-defrag (Content Databases only)
 - Shrink content db's (but not other database types)
 - Update statistics
 - Change the server wide fill factor to 70%
 - Use DBCC CheckDB to check for consistency errors
 - Check that your backups are working
 - Do maintenance during off-peak hours
- 

Maintenance Don'ts

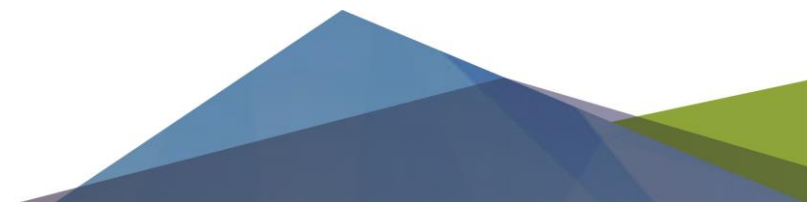
- Use maintenance plans to shrink databases
 - Drop and recreate indexes! (definitely don't do this...)
 - Shrink any databases other than Content DB's
 - Auto-shrink databases
 - Run maintenance tasks like DBCC or index rebuilds during business hours
 - Directly update SharePoint databases!!!
- 

Know your environment


SharePoint is a platform with a great deal of flexibility and wide range of uses. Your environment will be unique, so take time to understand normal:

- Setup performance counters
 - Monitor your environment doing Search Crawls
 - Be proactive about database growth
 - Prepare for bulk data migrations
 - Practice recovery scenarios
- 

If you only remember a few things

- Always use a SQL Alias for SharePoint
 - Use multiple content databases if you have lots of data (plan for scale)
 - You can scale out or scale up
 - Capacity Planning is important
 - Measure “YOUR” environment
 - Understand “YOUR” recovery options
- 

Q & A time



Please fill out your evaluation forms.

You have them in your A4 pack from registration.
Please put them in the box at the front of the room.

There are spot prizes for completed evaluation forms.

Steve Knutson | Microsoft MVP

SQL SATURDAY | #614 | South Island 2017



Please fill out your evaluation forms.

You have them in your A4 pack from registration.
Please put them in the box at the front of the room.

There are spot prizes for completed evaluation forms.

Steve Knutson | Microsoft MVP

SQL SATURDAY | #614 | South Island 2017



Thank you to our sponsors:

Gold Sponsors



Silver Sponsors



Bronze Sponsors

