


Platinum Sponsor

**Microsoft®**

Gold Sponsors

 **AvePoint®**  
Unleashing the Power of SharePoint™



**Kodak**

 **QUEST  
SOFTWARE®**

**SOUTHEAST ASIA  
SHAREPOINT CONFERENCE**

SINGAPORE 26 + 27 OCTOBER 2010

**Pinal Dave**

**MAINTAINING SQL SERVER AT  
OPTIMAL PERFORMANCE FOR  
BLAZING FAST SHAREPOINT SITE!**



**SOUTHEAST ASIA  
SHAREPOINT CONFERENCE**

SINGAPORE 26 + 27 OCTOBER 2010





# AGENDA

- Addressing Variable Factor of Query
- Database Layout and Configuration
- Optimal Database Configuration
- TempDB and Performance
- Index, Statistics & Optimizations

# KEY TAKE AWAY

- Understanding the Difference of MDF and LDF and performance
- Script to Change Autogrowth to Optimal Settings
- Script to Change Your Auto Stats
- First Step to Check your Database Health

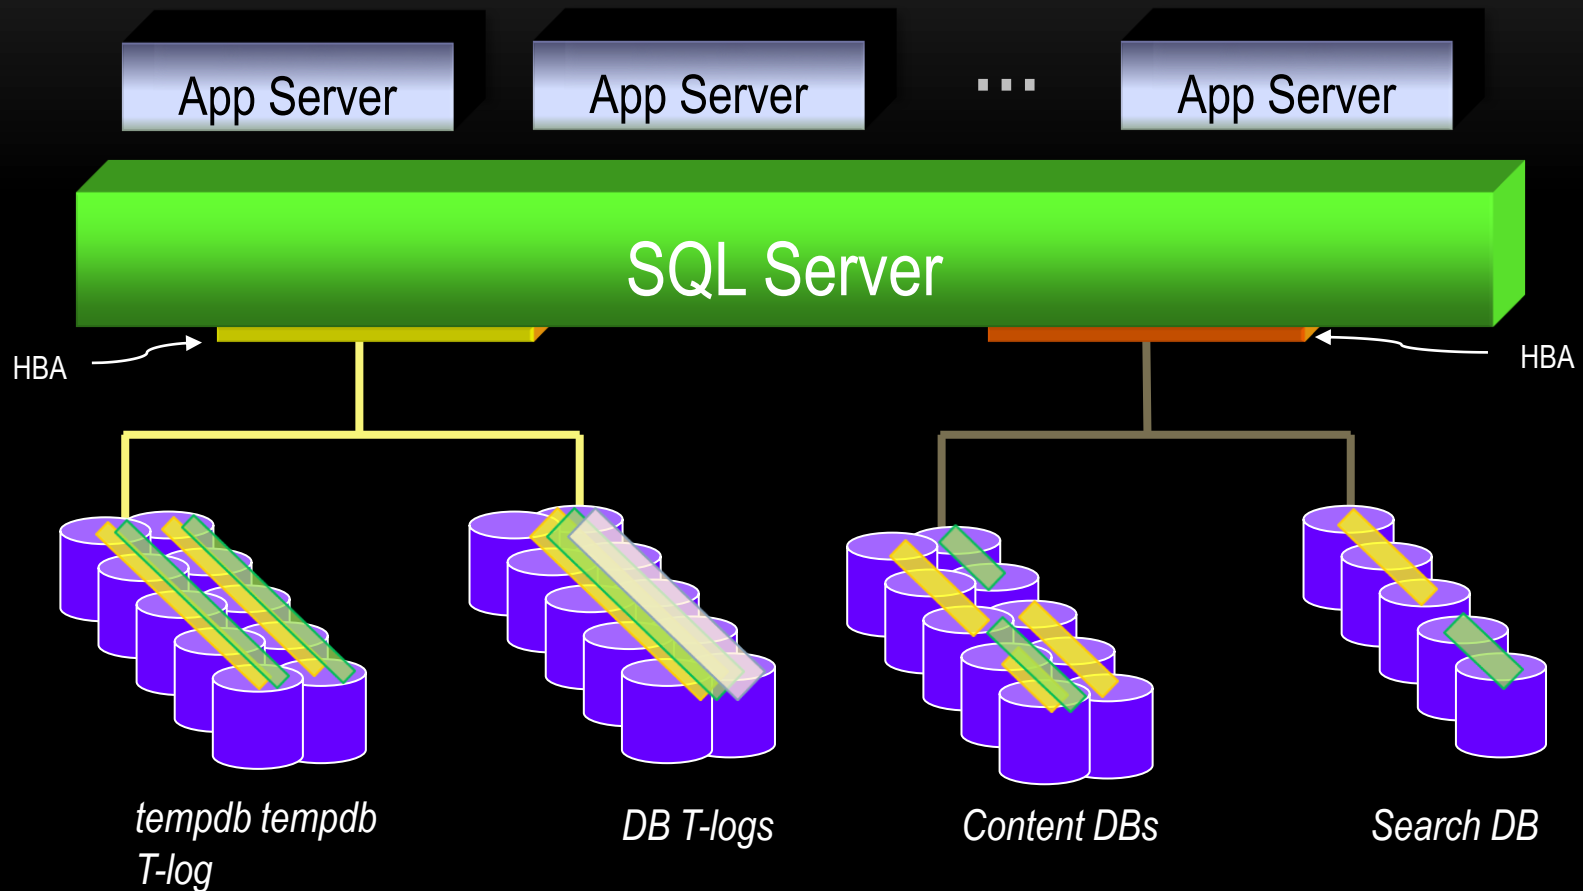
# ABOUT PINAL DAVE

- SQL Server MVP 
- Founder – SQLAuthority.com 
- Mentor – Solid Quality India 
- SQL Server Performance Tuning Consultant
- Regional Mentor – Professional Association for SQL Server (PASS) Asia Pacific Middle East 
- MCT, MCP, MCTS, MCDBA, MCAD
- [pdave@solidq.com](mailto:pdave@solidq.com) , [pinal@sqlauthority.com](mailto:pinal@sqlauthority.com)

**SOUTHEAST ASIA  
SHAREPOINT CONFERENCE**

SINGAPORE 26 + 27 OCTOBER 2010

# GENERIC STORAGE ARCHITECTURE






# THREE IMPORTANT THINGS

- CPU
  - Memory
  - Storage Drive
- 
- My Suggestion – Get the Best Ones!

# THE VARIABLE FACTOR

- Input and Output Query



**SOUTHEAST ASIA  
SHAREPOINT CONFERENCE**

SINGAPORE 26 + 27 OCTOBER 2010


# THE BIG NO!

- Do not do any of the following:
  - Manually or programmatically
  - Adding new information
  - Adding or deleting
  - Changing
  - Calling
  - Adding
  - Adding
  - Adding
  - Making
  - Adding
  - Changing



STOP

WHY NOT?



**Balance is the Key to Life**

# PRIORITIZING DATABASE VOLUMES

- Separate database volumes into unique logical unit numbers (LUNs) consisting of unique physical disk spindles
- Prioritize data among faster disks with ranking:
  - SQL TempDB data files
  - Database transaction log files
  - Search database
  - Content databases
- In a heavily read-oriented portal site, prioritize data over logs
- Separate out Search database transaction log from content database transaction logs

# DEMO

## LOG FILE VS DATA FILE PLACEMENT




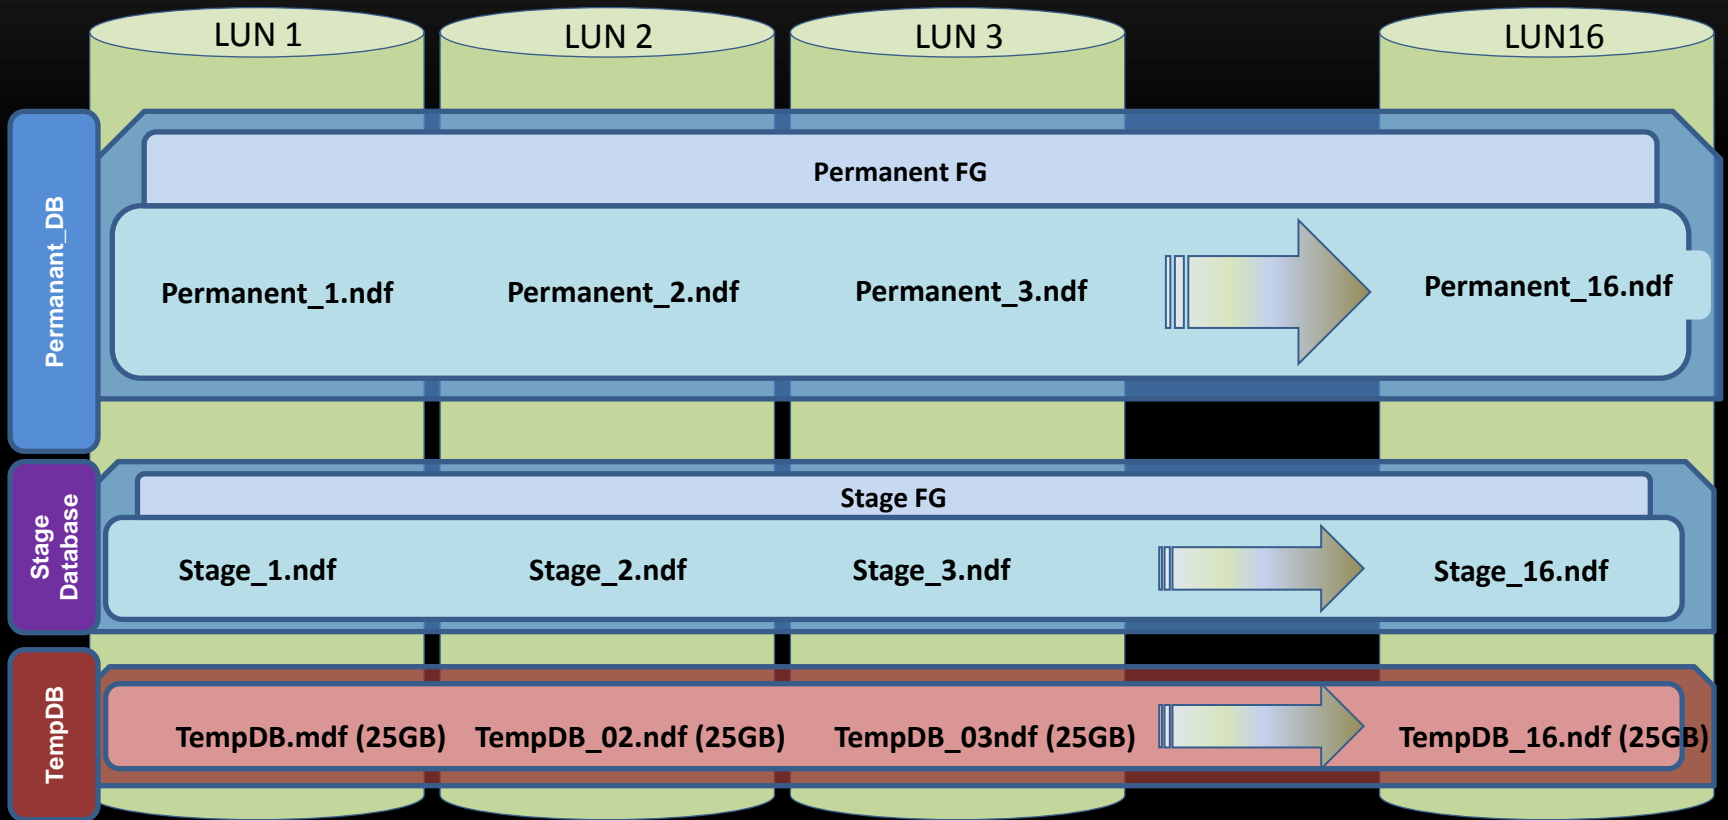
**SOUTHEAST ASIA  
SHAREPOINT CONFERENCE**

SINGAPORE 26 + 27 OCTOBER 2010

# SQL SERVER TEMPDB DATA FILES

- Data files allocated = number core CPUs in SQL Server
- Data file sizes consistent across all data files
- Data files spread across unique LUNs
  - Separated from Content DB, Search DB, etc.
- Log file separated to unique LUN

# SQL SERVER FILE LAYOUT



SOUTHEAST ASIA

SHAREPOINT CONFERENCE

SINGAPORE 26 + 27 OCTOBER 2010

# SQL SERVER TEMPDB DATA FILES

- Optimal TempDB data file size formula:

$$[\text{MAX DB SIZE (KB)}] \times [.25] / [\text{\# CORES}] = \text{DATA FILE SIZE (KB)}$$

- Result (starting size) should be roughly equal to 25% largest content or search DB or 10% of total content DB size
- Use RAID 10; separate LUN from other database objects (content, search, etc.)
- “Autogrow” feature set to a fixed amount; if auto grow occurs, permanently increase size

# MANAGE CONTENT DB - GENERIC

- Pre-Construct and Pre-Size
- Script generation of empty database objects
- “Autogrow” feature on for safety
- Use RAID 5 or RAID 10 logical units
- Multi-core computer running SQL Server
  - Primary file group should consist one File per CPU Core
  - Best Practices is to use an additional file group with one File per CPU Core, but it’s easier to manage a database with one filegroup

# SEARCH DATABASE

- Pre-Construct and Pre-Size
- Script generation of empty database objects
- “Autogrow” feature on for safety
- Use RAID 10 logical units
  - Should be a requirement for large-scale systems
- Multi-core computer running SQL Server
  - Primary file group with x Datafiles
  - Secondary file group with x Datafiles (x=CPU Core)

# DEMO

PRE-SIZE DATABASE &  
AUTOGROW SETTINGS



**SOUTHEAST ASIA  
SHAREPOINT CONFERENCE**

SINGAPORE 26 + 27 OCTOBER 2010

# SEARCH DATABASE - CONTINUED

- Search database is **VERY** read/write intensive!
- Do not place any other database data files on any logical unit where search database files reside
- Place the search database log files on an independent logical unit
- Place SharePoint Search crawl and query processing tables on separate spindles



# INDEX DEFRAG

- Fragmentation occurs by design on SharePoint
  - Clustered Indexes (PK) on GUID Columns
- Increase space utilization & I/O → degrades performance
- Content and Search dbs most susceptible
- Rebuild / Reorganize indexes to eliminate fragmentation
  - Reorganize index when fragmentation between 10-30%
  - Rebuild index when fragmentation > 30%
    - Existing index options used before the rebuild operation are not restored. Issue resolved in SQL Server 2005 SP2 and higher
- Use `sys.dm_db_index_physical_stats` to measure
  - More accurate than DBCC SHOWCONTIG, often reports higher fragmentation numbers

# DATABASES MAINTENANCE

- Do's
  - Have reliable backups for all databases before implementing maintenance operations
  - Check for and repair consistency errors by using DBCC CHECKDB
  - Use *sys.dm\_db\_index\_physical\_stats* dynamic view to measure fragmentation
  - Defragment indexes by either reorganizing them or rebuilding them, or use the *dbo.proc\_DefragIndexes* procedure
  - Change the server-wide fill factor setting to 70
  - Update statistics – don't rely on the timer service

# DEMO

## EFFECT OF STATISTICS UPDATE ON PERFORMANCE



**SOUTHEAST ASIA  
SHAREPOINT CONFERENCE**

SINGAPORE 26 + 27 OCTOBER 2010

# DATABASES MAINTENANCE

- Don'ts
  - Drop and re-create indexes
  - Rebuild indexes or run consistency checks during business hours
  - Set fill factor for individual tables or indexes
  - Shrink any databases other than content databases
  - Auto-shrink databases
  - Shrink databases at all unless you really need to
  - DBCC Checkdb REPAIR\_ALLOW\_DATA\_LOSS not supported (REPAIR\_REBUILD supported, but not always possible)
  - Time consuming operation, run during non-peak hours

# REFERENCE

- Roger Breu - Technology Solution Professional Data Platform, Microsoft
- Antonio Soto – Solid Quality Mentors, Spain
- Gianluca Hotz – Solid Quality Mentors, Italy
- Gilberto Zampatti – Solid Quality mentors, Italy
- Pinal Dave – Solid Quality mentors, India – [SQLAuthority.com](http://SQLAuthority.com)

# THANK YOU!

- [pinal@sqlauthority.com](mailto:pinal@sqlauthority.com)
- [pdave@solidq.com](mailto:pdave@solidq.com)
- <http://blog.sqlauthority.com>



**SOUTHEAST ASIA  
SHAREPOINT CONFERENCE**  
SINGAPORE 26 + 27 OCTOBER 2010


# Microsoft®

Thank you to our Sponsors

## Gold Sponsors



## Silver Sponsors



SYNERGY



## Bronze Sponsors



# SOUTHEAST ASIA SHAREPOINT CONFERENCE

SINGAPORE 26 + 27 OCTOBER 2010