DBADuck
ben miller

Big Mountain Data - Utah Geek Events

SQL Server Configuration Best Practices

Ben Miller, SQL Architect
DBADuck Consulting

Microsoft CERTIFIED
Master

Thanks to our Sponsors!

Premier

Marquee

Logo

Wifi:
Network:
lhm_conf

![DBADuck ben miller logo]

![Microsoft MVP Most Valuable Professional logo]

ben@benmiller.net
https://www.dbaduck.com

f www.linkedin.com/in/dbaduck

🐦 www.twitter.com/DBAduck

in www.facebook.com/DBAduck

SQL Server Certified Master
The first MCM in Utah. Has been working with SQL Server since 4.2, even on O/S2.

Microsoft Data Platform MVP
Awardee since 2009 for SQL Server. Contributes to Experts-Exchange and SQLServerCentral forums. Speaker at many SQL Saturdays around the country. I have led local PASS Chapters as well as founded the Virtual PowerShell Group.

PowerShell DBA
I have always loved automation and have used PowerShell since v1.0 I have spent many years using SMO and love to automate anything that I can to get the computer to do my work. I even wrote a chapter for the PowerShell Deep Dives book. On the Quest to become a PowerShell DBA.

![DBADuck logo]

Agenda

- Operating System
- SQL Server Engine
- Database Configuration
- Management Best Practices
- Agent Best Practices

DBADuck

Operating System

- Windows Server 2016+
- Drives at 64KB block size
- Turn off Indexing on the drive
- Data, Log, TempDb drives
- NO databases on C:\
- Page file fixed at 4GB on C:\
- Time to wait on OS = 0 (except for clusters and AG)
- Memory Dump to mini-memory dump
- Service account as domain account
- Lock Pages in Memory
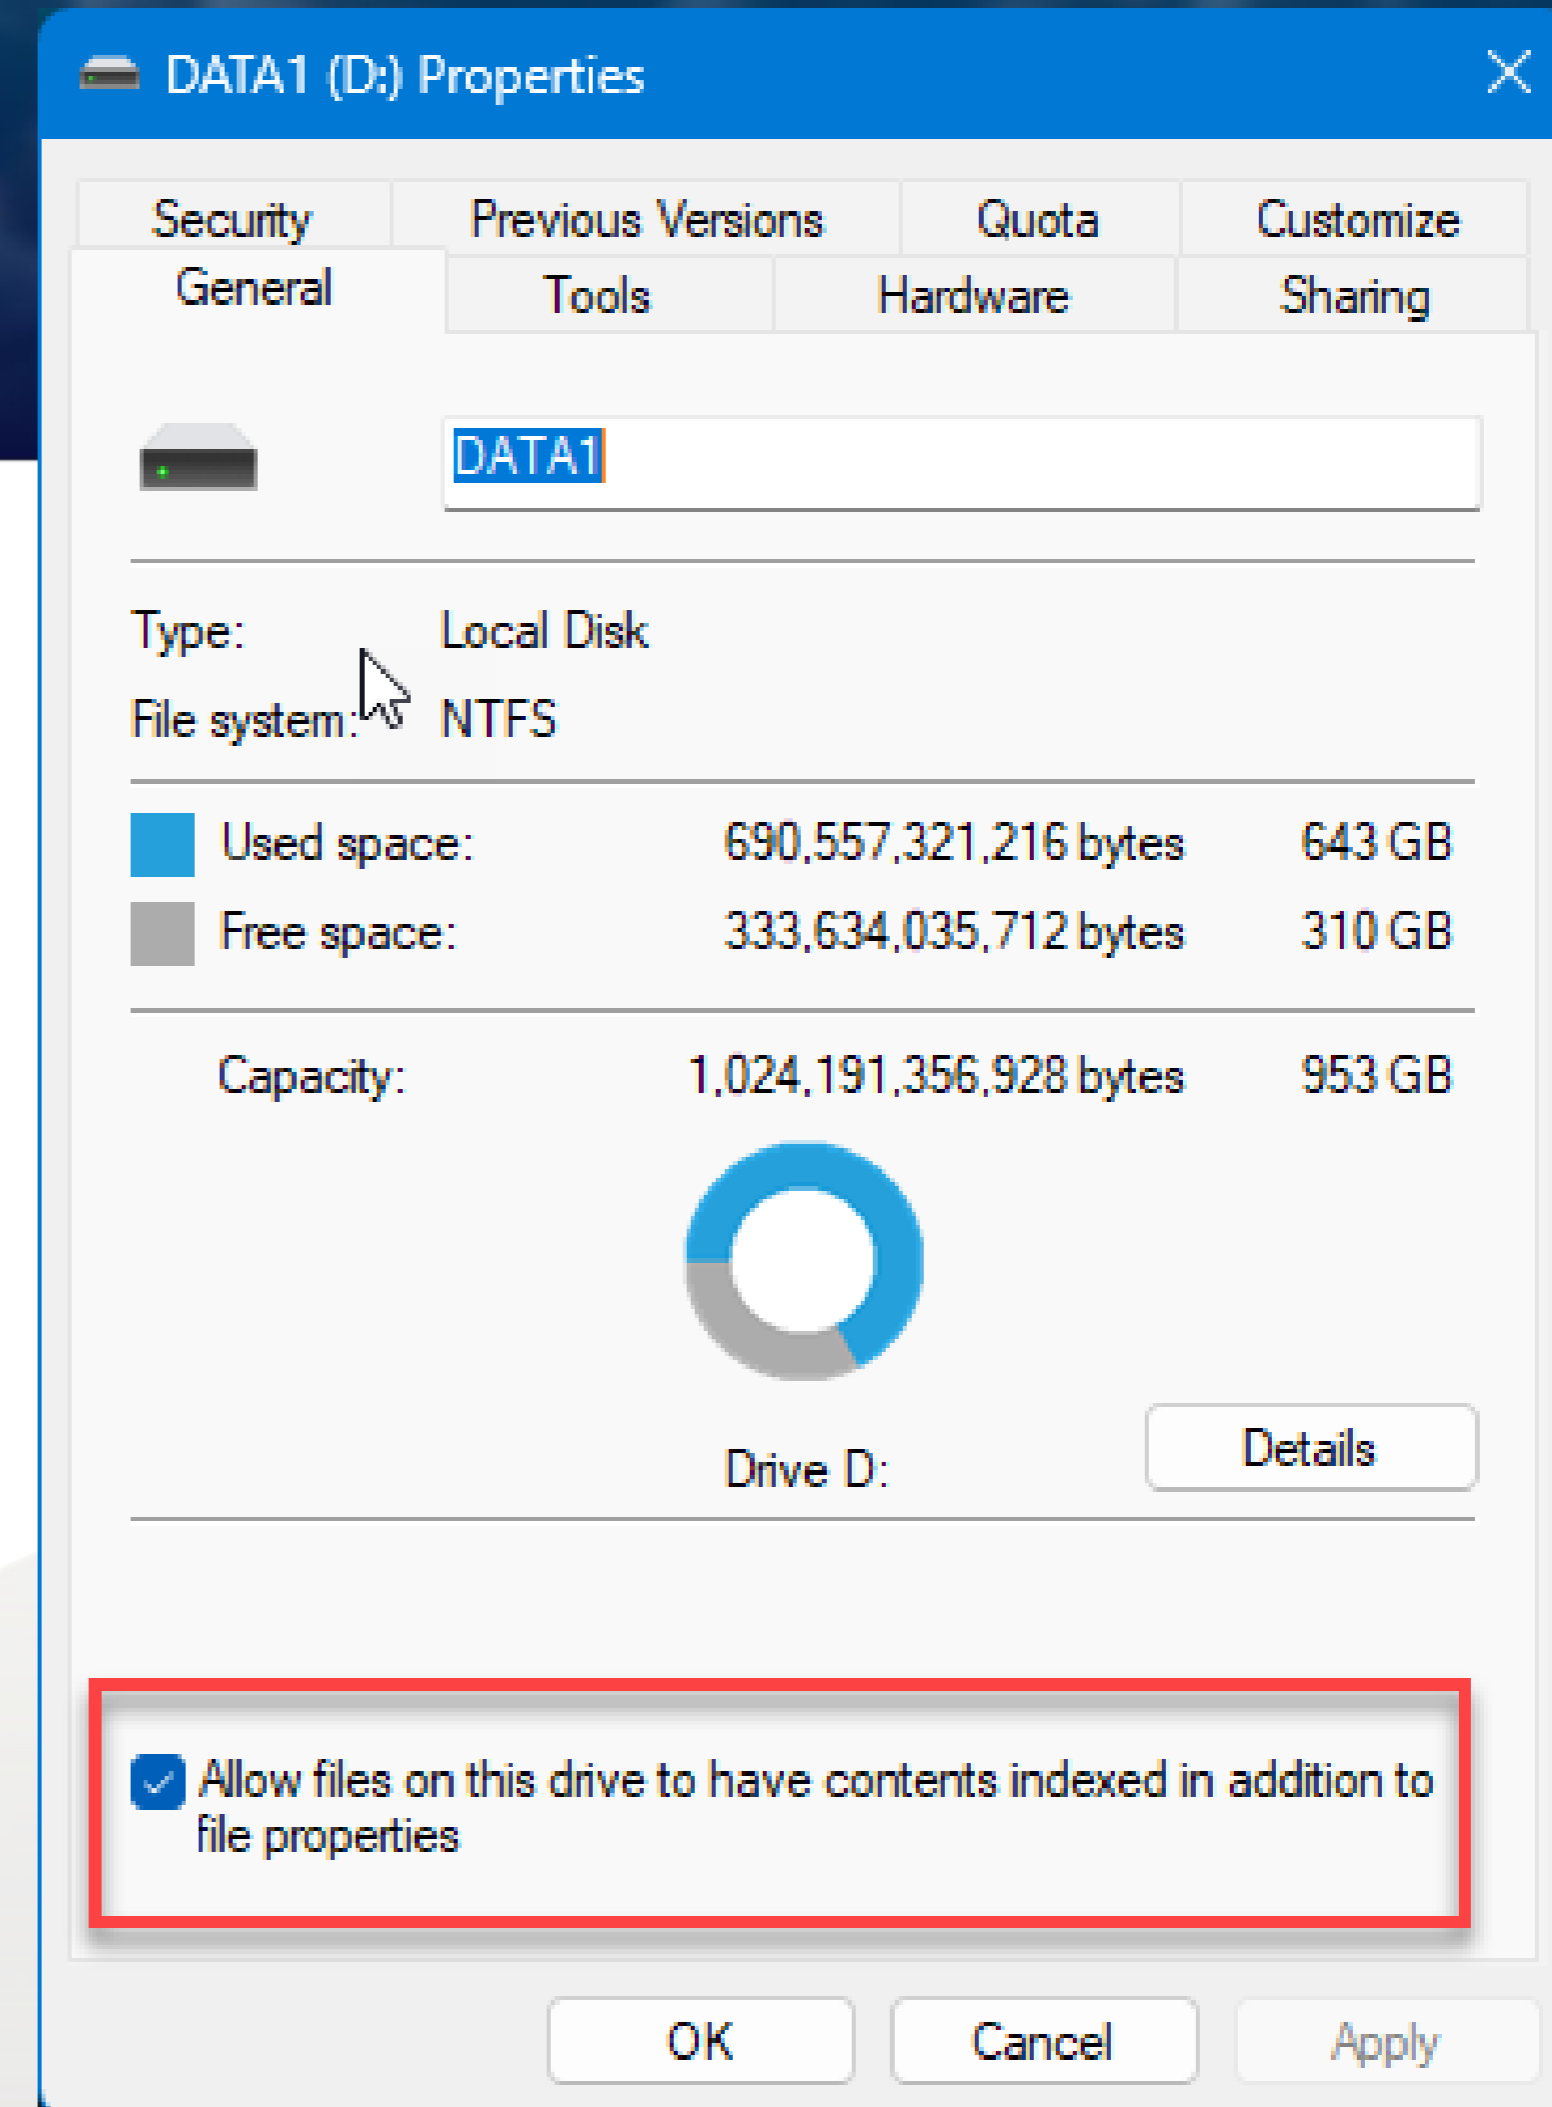- Perform Volume Maintenance – Instant File Initialization

DBADuck

OS Detail

- Windows Server 2016+
  - SQL and Windows co-developed and has benefits together
- Windows Server 2019
- Windows Server 2022

- Drives at 64KB – NTFS
  - SQL works with data in 8x8KB pages (64KB)
  - If at default 4K then it is 16 reads of that many blocks to get to 64KB

DBADuck

OS Detail

- Turn off Indexing on the drive

- NO databases on C:\

OS Detail
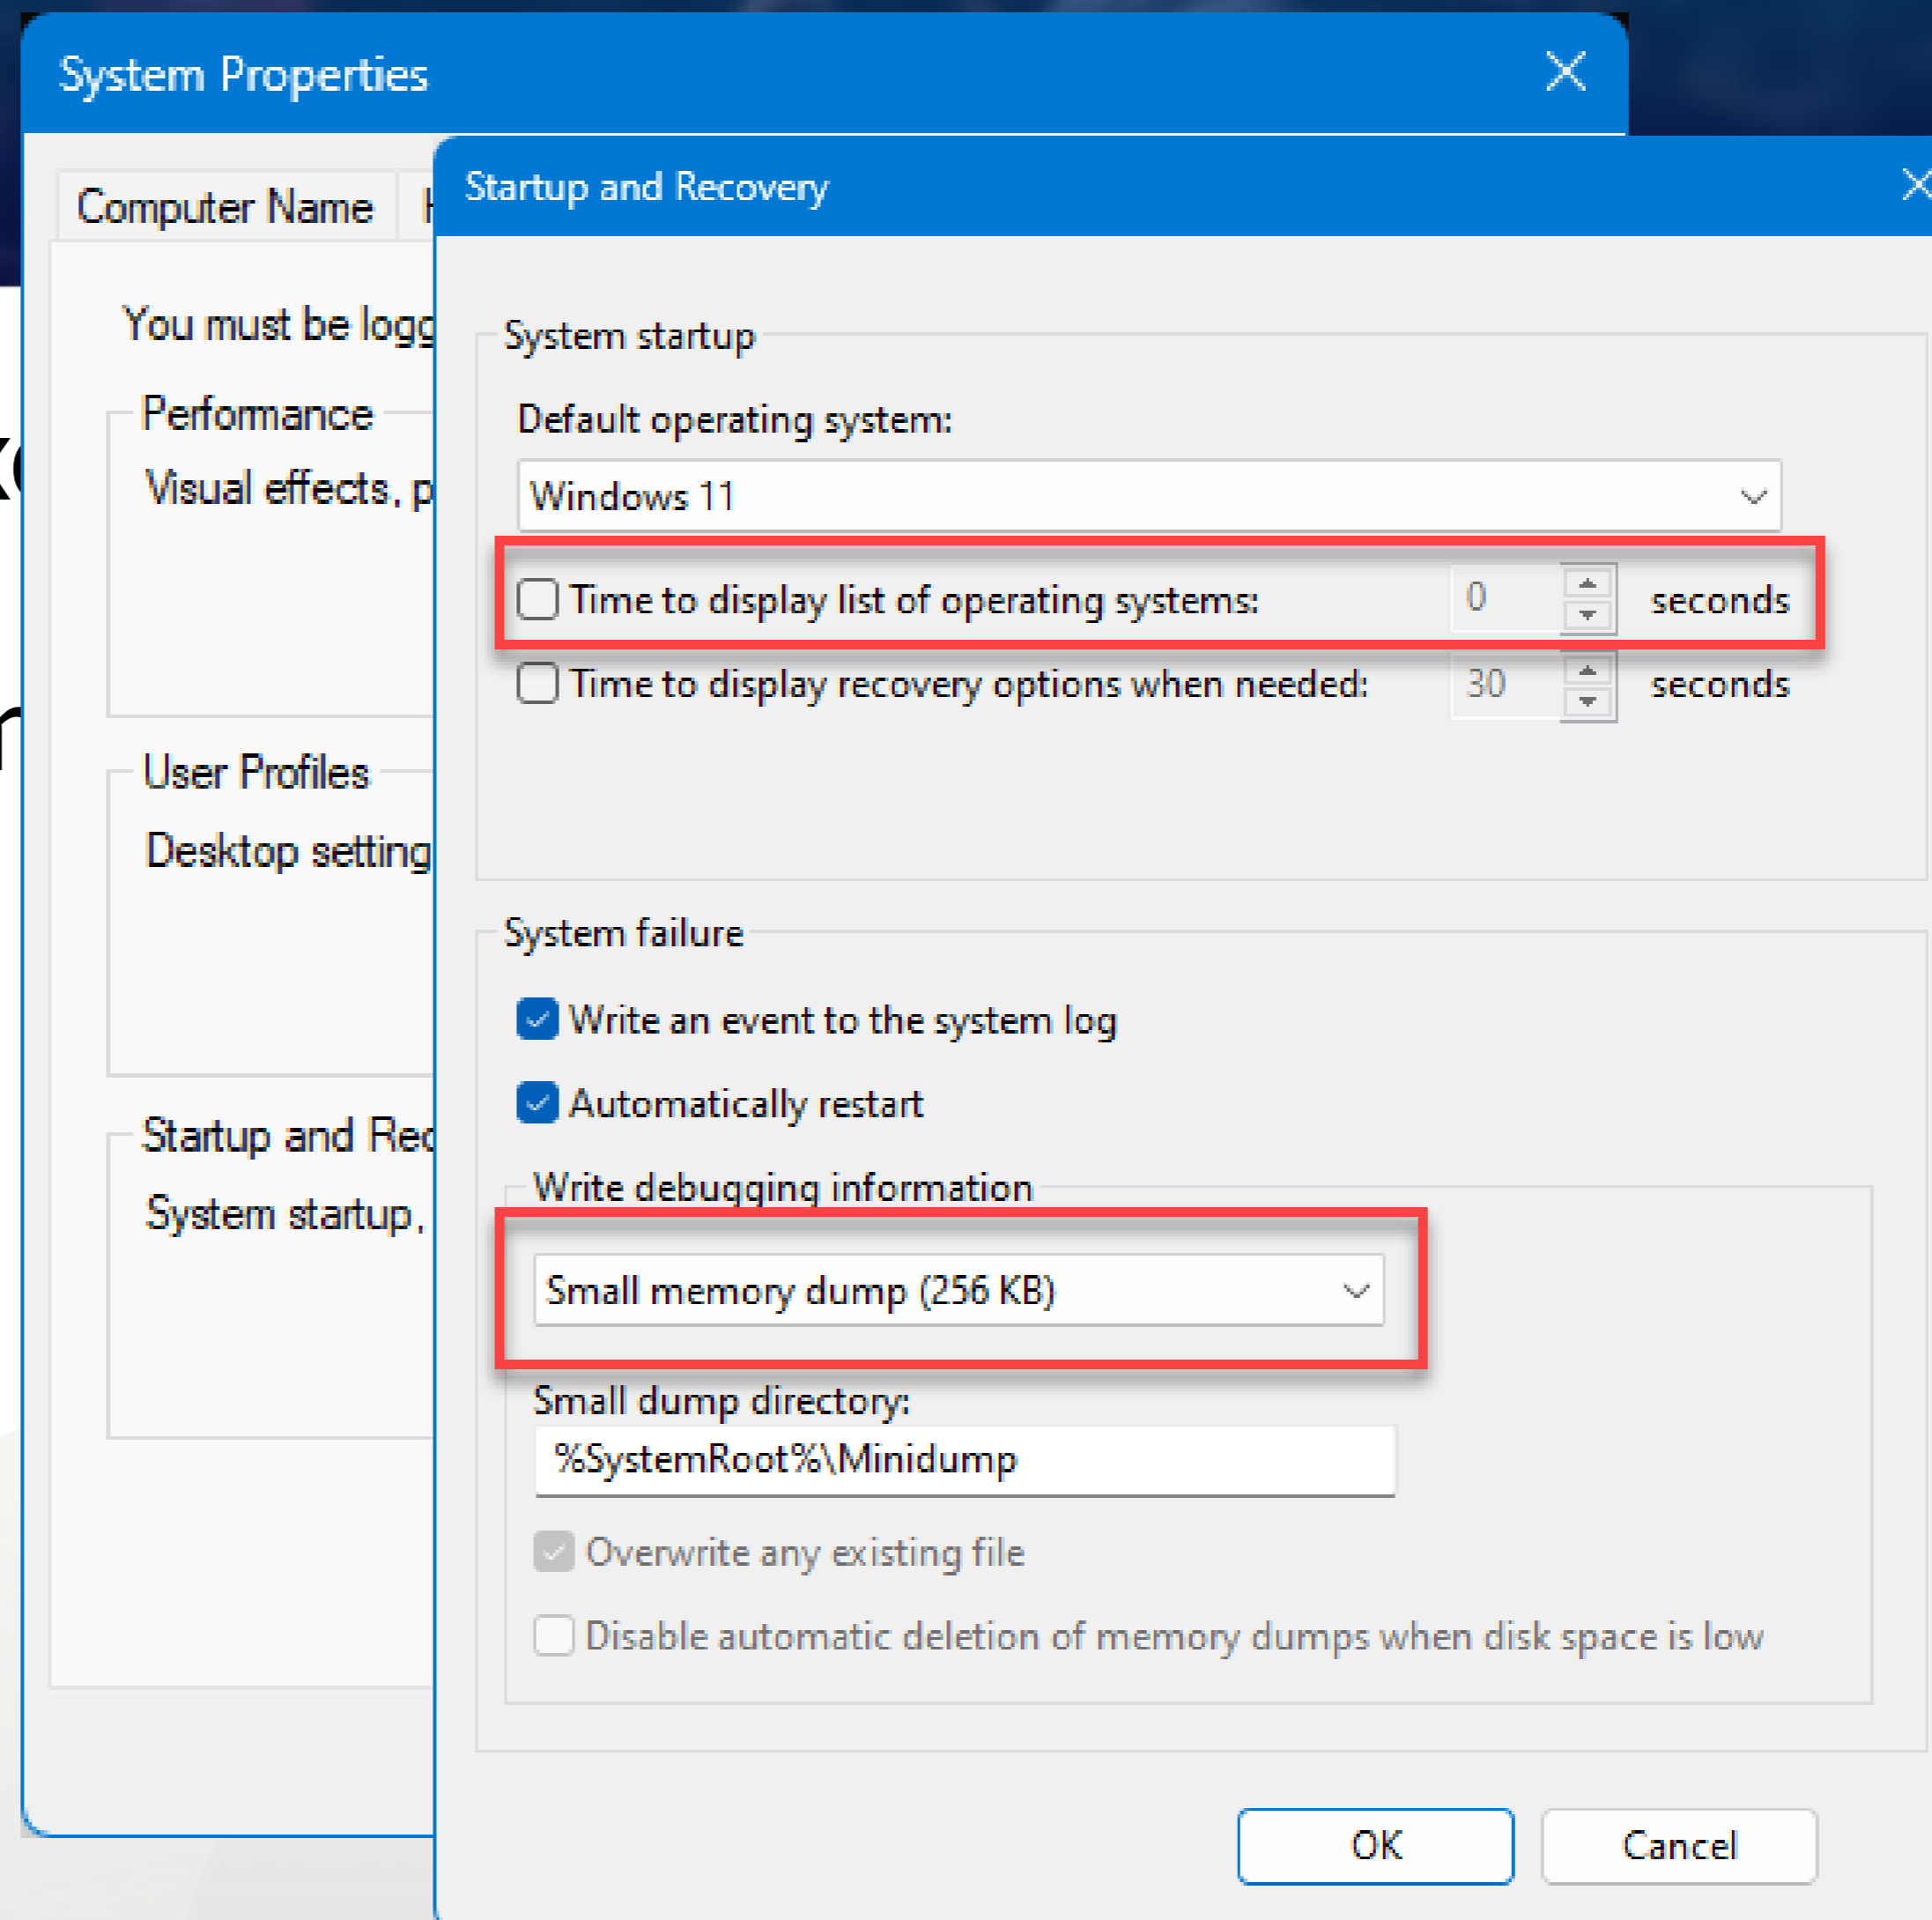
- Data, Log, TempDb drives
- Page file fixed at 4GB on C:\

OS Detail

- Time to wait on OS = 0 (exc

- Memory Dump to mini-mem

System Properties

Computer Name

You must be logg

Performance
Visual effects, p

User Profiles
Desktop setting

Startup and Rec
System startup,

Startup and Recovery

System startup

Default operating system:

Windows 11

☐ Time to display list of operating systems: 0 seconds

☐ Time to display recovery options when needed: 30 seconds

System failure

☑ Write an event to the system log

☑ Automatically restart

Write debugging information

Small memory dump (256 KB)

Small dump directory:

%SystemRoot%\Minidump

☑ Overwrite any existing file

☐ Disable automatic deletion of memory dumps when disk space is low

OK   Cancel

DBADuck

OS Detail

- Service account as domain account
  - Especially the SQL Server Service Account
- Lock Pages in Memory
  - Typically set on VM computers, but I do it on physical too
  - Locks pages used in SQL Server in memory to protect against balloon driver reapers
- Perform Volume Maintenance – Instant File Initialization
  - This is an option in Setup for 2016+ but always set it for quick expansion of database files (not applicable completely to Log files)

DBADuck

SQL Server Engine

- Sp_Configure options
  - Max server memory
  - Max DOP
  - Cost Threshold
  - Optimize for Adhoc Workloads
  - Backup Compression
  - Backup Checksum (later versions)
  - Remote Admin Connections

DBADuck

sp_configure

- Max server memory
  - Do not leave it default (2PB) as you do not want tug of war between OS and SQL for memory
- Max DOP
  - Discussion and debatable. I typically set it to 4 for anything above 6 cores. And 2 for below 6 cores.
- Cost Threshold
  - NEVER leave it at default of 5 (story of how that came to be the number)
- Optimize for Adhoc Workloads
  - Always set it on, unless you prove that it is a bad thing. There is more to it but for the most part, this is my rule. (Query for adhoc)

DBADuck

sp_configure

- Backup Compression
  - For when someone uses BACKUP DATABASE and does not specify to compress.
- Backup Checksum (later versions)
  - Introduced in SQL 2017 to allow you to have BACKUP DATABASE automatically put checksums in the backup if you don't specify it
- Remote Admin Connections
  - VERY important to ensure that you can use the Remote Admin connection remotely instead of having to RDP into the instance to fix stuff.

DBADuck

Trace Flags

- Trace Flags are a great topic
  - 1117 – Auto grow all files in a file group at the same time
  - 1118 – Eliminate Shared GAM allocations and only get Uniform Extents
  - 2371 – Auto update stats algorithm changes from 20% + 500
- The above are auto on in 2016+

DBADuck

TempDB

- Get Ready for some fun It Depends action….

- How many files?
  - Anywhere between 4 and 12. With servers that have 4-6 processors the count can be anywhere from 4 to 6. When you get to 8 processors, you can go up to 12. Beyond 8, I would start with 8 and see how the contention gets.
- How big?
  - All files should be the same size with the same auto grow. In SQL 2014 and below, ensure that you use Trace Flag 1117.

Database Configuration

- Owner should be other than Windows Account
  - Owners not existing can cause the database to not start up
  - (especially when there is CLR involved)
- Auto Close Off
  - This is a setting typically set on Express Instances. When migrating to SQL Standard and above, this should be changed
- Auto Shrink Off
  - This is also a setting typically on Express Instances. This should be turned off.

DBADuck

Database Configuration

- Auto Create Statistics
  - Should be turned on. It is default.
- Auto Update Statistics
  - Should be turned on, typically.
- Allow Snapshot
  - Alternative to NOLOCK
- RCSI On
  - Default in Azure. Allows 3 of 4 contention points to resolve.

Database Configuration

- Page Verify CHECKSUM
  - Corruption Detector – make sure it is on.
- Query Store ReadWrite
  - Query store options – 1000+ in size, 30 days of Stale Query
- Target Recovery Time (except master database)
  - All databases should be something other than 0 and master at 0. This is about log hammer.

DBADuck

Management Configuration

- SQL Server Errorlogs – 31
  - Default is 6 and that can pass fast on a flip flop start stop storm
- Daily sp_cycle_errorlog at midnight
  - When reading logs with the Log Viewer having multiple days in the logs can cause slowness.
- Do NOT use Maintenance Plans
  - They are old, deprecated and have been known to generate invalid scripts. They also lack many options that you can do in other maintenance solutions.

DBADuck

Management Configuration

- Ola Hallengren Maintenance Solution
  - Backups
  - Index Optimize
  - Output files for each steps
- Configure DBMail even with multiple profiles/accounts
  - Ensure that you have this enabled for Alerts, Job Failures, etc.

DBADuck

SQL Agent

- Agent Properties

- Alerting
  - First configure settings in Agent P
- Operator
  - Create one which allows emails to
- Fail Safe Operator
  - Make one either the one above or

SQL Agent

- Job Notifications
  - Should be set on jobs that are critical enough to insure you are notified
- Job History specifications
  - By default, it is 1000 total lines of completion. This should be changed to something a little more relevant to today.
- Backup History management
  - Keep in mind sometimes ISO and SOC may need proof of backups happening, and you may want to keep it longer. If that is true and you need to remove history, you may want to archive to another database with the same schema to ensure that you have auditability.

DBADuck

Questions?

ben@benmiller.net
https://www.dbaduck.com

THANK YOU!

UTAH GEEK EVENTS

training for geeks

DBADuck