boldonjames
by HelpSystems

Database Performance Guide
Classifier Reporting

May 2021




Copyright Terms and Conditions

202105280831

Table of Contents

About this guide

This guide examines the performance characteristics of a SQL Database containing approximately 150 million Classifier Reporting Events where the virtual hardware specification (RAM, CPU and Disk IOPS) is varied.

This document is for informational purposes only, and Boldon James cannot guarantee the precision of any information supplied.

Terms used in this guide

Term	Definition
IOPS	Input/output Operations Per Second, a disk performance measurement.
LDF	Log Database File (SQL)
MDF	SQL Master Database File
SP	Stored Procedure
VM	Virtual Machine

Performance Characteristics

For a SQL database with some 150 million Classifier Events. the greatest limitation on the speed of executing the test Stored Procedure is the disk I/O where any rating below approximately 3000 IOPS significantly limits throughput.

There seems to be a 'sweet spot' for the amount of system RAM, this being 32GB, and the time taken to execute the first query, although greater amounts of RAM has little effect on the time taken for subsequent queries.

Generally, the more CPU cores that are available, the quicker the query times and the less system memory is used. System utilisation is also reduced, so allowing the system to be available for other work.

About the testing

The following tests were completed on a SQL Database containing approximately 150 million Classifier Events, using the Stored Procedure (SP) "usp_ReportEmailsByDomain" to return a consistent 14,374,260 rows.

To allow hardware characteristics to be easily changed, the tests are run on a VM hosted by a server that has no other software running (apart from Windows) or other user access.

The VM Operating System disk drive image resides on the host server Windows drive (C:), SSD1. To maximise the SQL Database performance on the VM, the SQL Database MDF and LDF files accessed by the VM reside on the host server's SSD2 and SSD3 respectively, see *Virtual Machine host specifications* on page 10.

The VM has HyperV hardware settings changed for each test: memory allocation, number of virtual processors, and the disk Input/output Operations Per Second (IOPS).

On the VM, to minimise resource usage interfering with the performance tests, several Windows services are stopped, see *Windows Services stopped* on page 10.

During a test 'cycle' the following are recorded:

- The time taken for the first run of the query
- The mean time of 5 subsequent queries

The test cycle is repeated for both a static database, and a database where Classifier Events are added at a rate of approximately 82.5 Classifier Events/sec.

Memory

The amount of RAM allocated to a VM can be configured in HyperV under the Memory options. Tests were run with the following RAM allocations: 8, 16, 24, 32, 40 and 64GB.

Looking at the figure below, lines 3 and 4), there is a sudden reduction in the query time between 24 and 32GB RAM allocation. It was found that the SQLservr.exe process grows to ~28GB when executing the query for the first time. Subsequent queries are completed much quicker.

NOTE: This behaviour changes as the CPU count is increased.

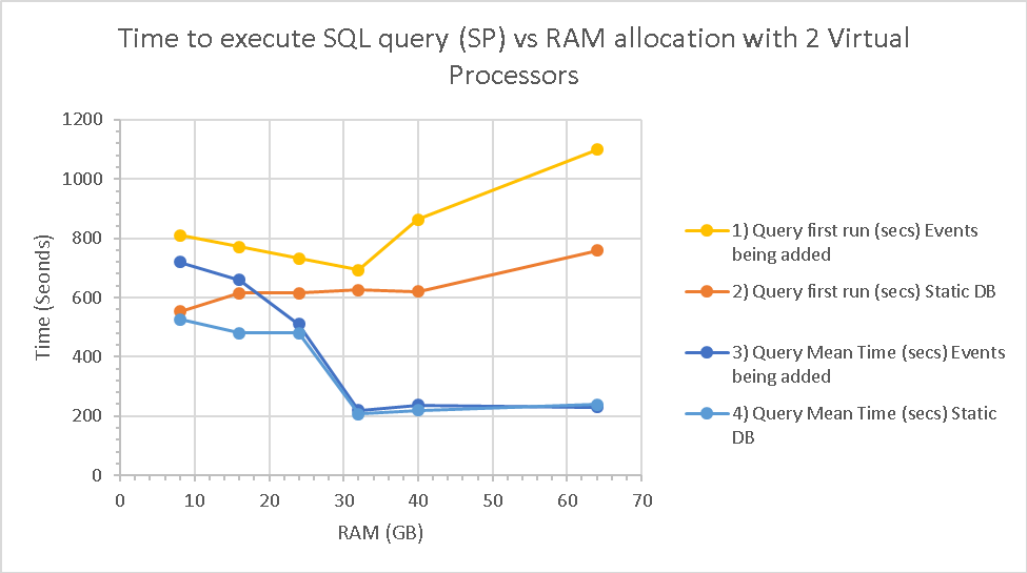

Processor (CPU) Count

The number of virtual processors of a VM can be defined by HyperV under the Number of Virtual Processor options. Tests were run against the following number of virtual processors: 2, 3, 4, 5, 6, 12 and 24.

The RAM allocated was set to a fixed 32GB – from previous tests. The SQLservr.exe process is able to hold in memory all information required to complete the query (the SQLservr.exe process grows to approximately 28GB the first time the query is run).

Looking at the figure below, lines 1-4), the processor count increases, the difference in query response time is reduced between a static database and one where events are being added (a more realistic, real world situation) for both the first run and subsequent queries.

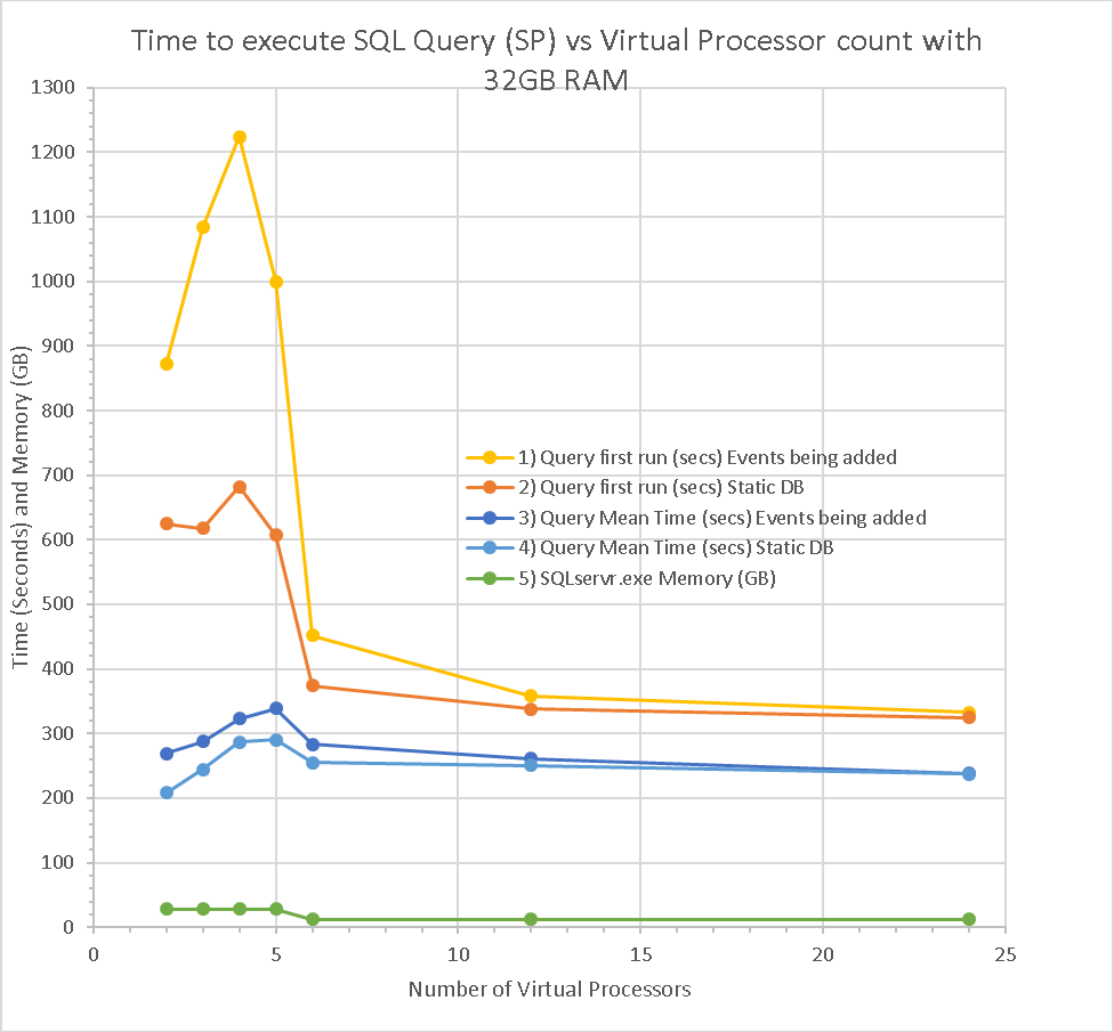

Time to execute query vs number of Virtual Processors

In the figure above, with a virtual processor count up to 5, the process used ~28GB, however with 6 or more virtual processors, the process memory usage drops to ~12GB. This process memory is approximate and was similar with both a static and events being added to the database. With a static database there was no MDF access activity even when the process memory fell to 12GB, implying that SQL is changing its behaviour depending upon CPU resources.

As the processor count is increased, the approximate CPU % Utilisation falls, per the following figure. This directly affects how 'usable' / responsive the server is, and the ability to do more work.

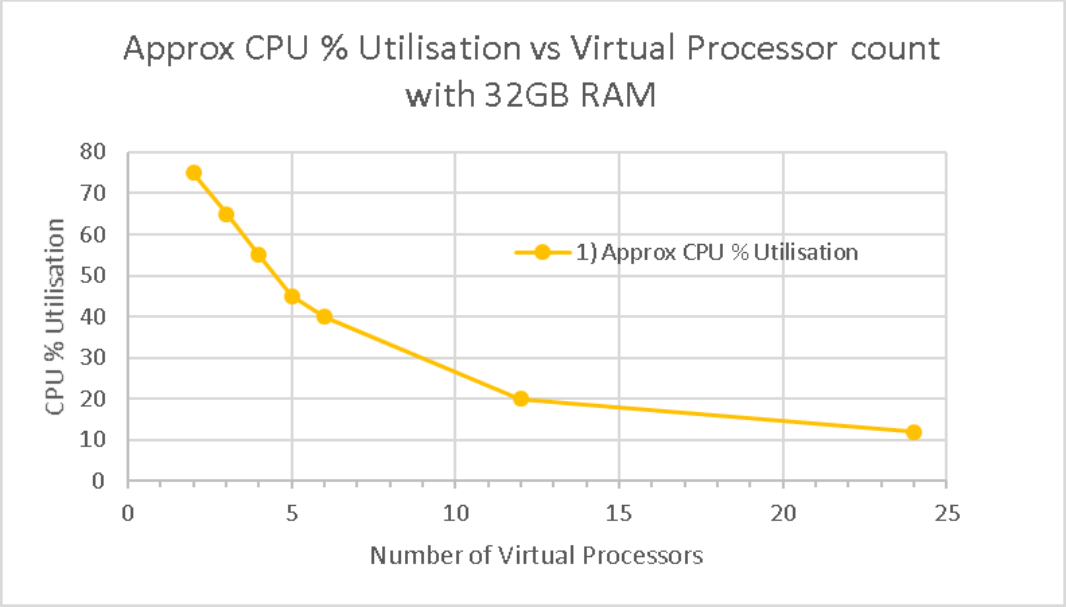

Disk IOPS

The disk IOPS of a running VM can be limited by HyperV under the SCSI Controller, disk 'Quality of Service' options.

> **NOTE:** The measured IOPS value, see *Disk Performance Tool* on page 11, appears to be 1/8 of the setting value in HyperV.

The manufacturers specification of SSD2 hosting the SQL MDF file list the following:

- Random Read: 97K IOPS
- Random Write: 24K IOPS

Using DiskSpd, the VM random read access of the SSD was measured at ~93K IOPS, reasonably close to the manufacturer's specification.

Tests were run with the disk IOPS limited to the following values: 50K, 25K, 12.5K, 6250, 3125, 1563 and 781 (x8 for the actual HyperV setting).

In the figure below, the query execution time suddenly increases with a disk IOPS below approximately ~3000 IOPS for both a database with data being written, line 1 and for a static database, line 2.

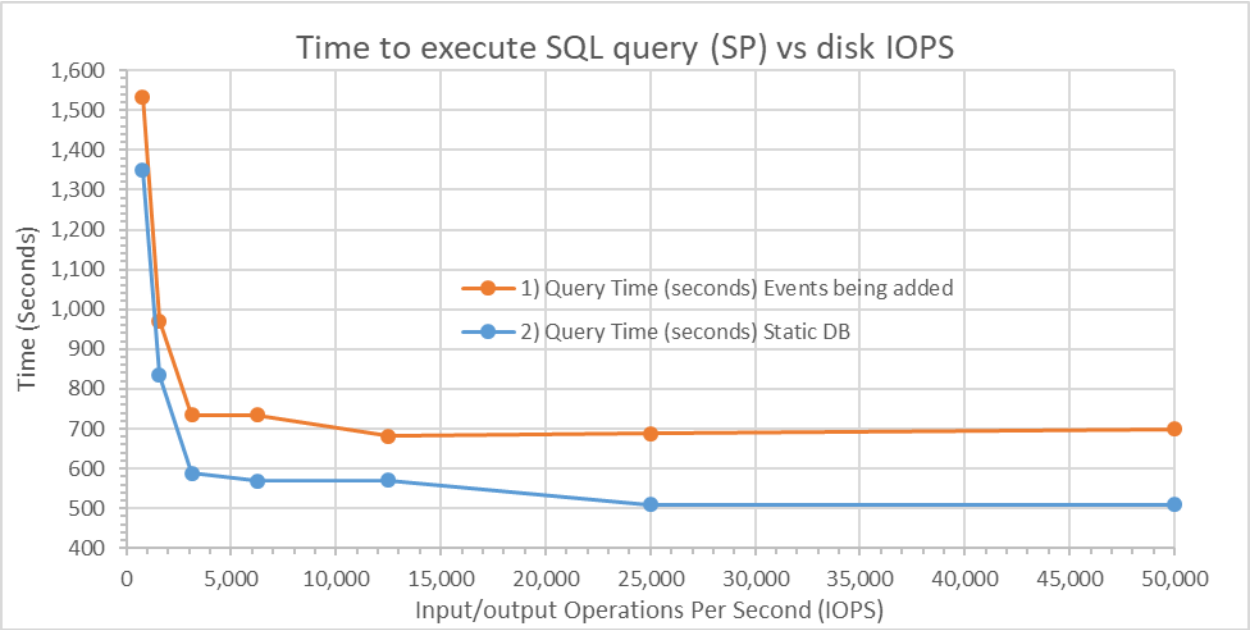
Time to execute SQL query (SP) vs disk IOPS
Time (Seconds)
1,600
1,500
1,400
1,300
1,200
1,100
1,000
900
800
700
600
500
400
0
5,000
10,000
15,000
20,000
25,000
30,000
35,000
40,000
45,000
50,000
Input/output Operations Per Second (IOPS)
1) Query Time (seconds) Events being added
2) Query Time (seconds) Static DB

Appendix

Virtual Machine host specifications

Attribute	Value
CPU	Intel Xeon Silver 4114 @ 2.20GHz
- sockets	2
- cores	20
- logical processors	40
Memory	128GB
Mass storage	
- SSD1 (host server and VM OS drive)	SAMSUNG MZ7LM240HMHQ (~240GB)
- SSD2 (VM MDF file)	SAMSUNG MZ7LM960HMJP (~1TB)
- SSD3 (VM LDF file)	SAMSUNG 960 CXE7 (~250GB)
Operating System	Windows Server 2016 Standard
- WinVer	Version 1607 (OS Build 14393.2999)

Virtual Machine specifications

Attribute	Value
Operating System	Windows Server 2016 Standard
- WinVer	Version 1607 (OS Build 14393.2999)
SQL Server	MS SQL Server 2016 Standard Edition

Windows Services stopped

The following Windows Services are stopped on the VM to reduce interference of the performance tests:

- MS Defender (via Group Policy)
- Windows Update (wuauserv) service
- Windows Modules Installer (TrustedInstaller) service

SQL Query

Query used in tests:

```
USE [ClassifierEventsDB] GO

exec "ClassifierReporting3"."usp_ReportEmailsByDomain" @FromDate='2000-01-01
00:00:00',@ToDate='2020-01-01 00:00:00',@Department='%',@Classification='%'
```

Disk Performance Tool

Disk IOPS is measured using the “DiskSpd” software tool:

https://github.com/Microsoft/diskspd