


Microsoft® SQL Server® Tuning Guide for Online Transaction Processing workload on 3rd Generation Intel® Xeon® Scalable Processors Based Platform



Revision Record	2
1. Introduction	3
2. Hardware Tuning.....	4
2.1. BIOS Setting	4
2.2. Memory Configuration/Settings	5
2.3. Storage/Disk Configuration/Settings	5
2.4. Network Configuration/Setting.....	5
3. Software Tuning	5
3.1. Microsoft Windows 2019 settings.....	6
3.2. OLTP Architecture	6
3.3. Tuning SQL Server for OLTP Workload	7
4. Related Tools.....	9
5. Best Practices for Testing and Verification	9
6. Conclusion.....	9
7. References	9
8. Feedback.....	9

Revision Record

Date	Rev.	Description
08/15/2021	1.0	Initial public release

1. Introduction

This guide is targeted towards users who are already familiar with Microsoft* SQL Server* and provides pointers and system settings for hardware and software that will provide the best performance for most situations. However, please note that we rely on the users to carefully consider these settings for their specific scenarios, since Microsoft SQL Server can be deployed in multiple ways and this is a reference to one such use-case.

Microsoft SQL Server is a relational database management system developed by Microsoft. SQL Server is offered in many different editions, but this guide will focus on the SQL Server Enterprise edition. SQL Server's database architecture is based on a Client-Server architecture model. This tuning guide will focus on Online Transaction Processing (OLTP) workloads (TPCE like) while using SQL Server for Microsoft Windows*.

3rd Gen Intel® Xeon® Scalable processors deliver industry-leading, workload-optimized platforms with built-in AI acceleration, providing a seamless performance foundation to help speed data's transformative impact, from the multi-cloud to the intelligent edge and back. Improvements of particular interest to this workload applications are:

- Enhanced Performance
- More Intel® Ultra Path Interconnect
- Increased DDR4 Memory Speed & Capacity

Tested hardware and software environment for this tuning guide:

Server Configuration	Hardware	Server Platform Name/Brand/Model	Intel Wilson City Reference Platform
		CPU	Intel® Xeon® PLATINUM 8358 CPU @ 2.60GHz
		BIOS	Intel Corporation WLYDCRB1.SYS.0020.P95.2104010233
		Memory	32*64 GB DDR4, 3200 MT/s
		IO Controllers	4x RS3SC008 Intel® Raid
		Disks	1x Intel 1.92TB SSD OS Drive 68x Intel SSD DC-S3700 (800GB) DB 5x Intel SSD DC-S3700(800GB) Temp DB 7x Intel SSD DC S4600 (1.92TB) Backups 4x Intel SSD DC-S3700 (800GB) Log
	Software	NIC	2x Intel X520-2 10GBASE-T
		Operating System	Microsoft Windows Server 2019 Datacenter
		Kernel	10.0.17763 Build 17763.1637
		Workload	Microsoft SQL Server 2019 RTM Cumulative Update 11
		Workload Kit	TPCE Kit 1041
		Benchcraft	2.5.6

Note: The configuration described in this article is based on 3rd Generation Intel Xeon processor hardware. Server platform, memory, hard drives, network interface cards can be determined according to customer usage requirements.

2. Hardware Tuning

2.1. BIOS Setting

Please note, that all BIOS settings outlined below are based on Intel's Software Development Platforms designed for the 3rd Gen Intel Xeon Scalable processors. Begin by resetting your BIOS to default setting, then follow the suggestion below for changes:

Configuration Item	Recommended Value
EDKII Menu ->Socket Configuration-> Processor Configuration-> Hardware Prefetcher	Disabled
EDKII Menu ->Socket Configuration-> Processor Configuration->Adjacent Cache Prefetch	Disabled
EDKII Menu ->Socket Configuration-> Processor Configuration->DCU Streamer Prefetcher	Enabled
EDKII Menu ->Socket Configuration-> Processor Configuration-> DCU IP Prefetcher	Enable
EDKII Menu ->Socket Configuration-> Processor Configuration-> LLC Prefetch	Disabled
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU P State Control->Energy Efficient Turbo	Disabled
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU P State Control->Turbo Mode	Enabled
EDKII Menu->Socket Configuration->UPI Configuration->UPI General Configuration->Boot performance Mode	Max perf
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->Hardware PM State Control->Hardware P State	Disabled
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU – Advanced PM Tuning->Energy Perf BIAS->Power Performance Tuning	BIOS control EPB
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU – Advanced PM Tuning->Energy Perf BIAS->ENERGY_PERF_BIAS_CFG_mode	Performance
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU – Advanced PM Tuning->Energy Perf BIAS->Workload Configuration	I/O sensitive
EDKII Menu->Socket Configuration->Memory Configuration->Memory RAS Configuration->Patrol Scrub	Disabled
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU C State Control->Package C State	CO\C1 state
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU C State Control->Enhanced Halt State (C1E)	Disable
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU C State Control->CPU C6 Report	Disable
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU C State Control->CPU C1 auto demotion	Disable
EDKII Menu->Socket Configuration->Advanced Power Management Configuration->CPU C State Control->CPU C1 auto undemotion	Disable
EDKII Menu->Socket Configuration->Common RefCode Configuration->UMA-Based Clustering	Disable (All2All)
EDKII Menu->Socket Configuration->Uncore Configuration->Uncore General Configuration->SNC (Sub NUMA)	Enable SNC2 (2-clusters)

2.2. Memory Configuration/Settings

This workload runs best with 2 DIMMS per channel populated with 3200 MTS DDR4 registered memory. For specifics consult your platform’s manual.

2.3. Storage/Disk Configuration/Settings

The table below details the type and number of drives used. Each type of drive sets (DB, Temp DB, etc....) are organized into RAID volumes, with the DB drives being evenly split into RAID0 volumes per IO controller. The only exception to this is the LOG volume which is configured as RAID10. The following are the recommended settings for the RAID volumes and can be set via the Intel® RAID Web Console 3 for Windows.

Drive Purpose	RAID Level	Stripe Size (KB)	Read Ahead	Write Back Cache Policy	Disk Cache Policy
DB	0	64	No	Write Through	Enabled
Temp DB	0	64	No	Write Through	Enabled
Backups	0	64	Yes	Write Through	Enabled
Log	10	64	No	Write Back	Enabled

2.4. Network Configuration/Setting

This workload runs best with the following settings for all adapters connecting the Server to the Client:

The following PowerShell script is used as an example with 'Ethernet 2':

```
Set-NetAdapterAdvancedProperty -Name "Ethernet 2" -RegistryKeyword '*JumboPacket' -
RegistryValue '9014'
Set-NetAdapterRss
-Name "Ethernet 2" -Enable 1 -BaseProcessorGroup 0 -BaseProcessorNumber 0 -MaxProcessors 16 -NumaNode
65535 -MaxProcessorGroup 0 -MaxProcessorNumber 16 -NumberOfReceiveQueues 8 -Profile Closest
```

In addition, the following must be set in SQL Server using the SQL Server sp_configure command.

```
sp_configure network_packet_size,8192
go
RECONFIGURE WITH OVERRIDE
go
```

3. Software Tuning

Software configuration tuning is essential. From the Operating System to SQL Server configuration settings, they are all designed for general purpose applications and default settings are almost never tuned for best performance.

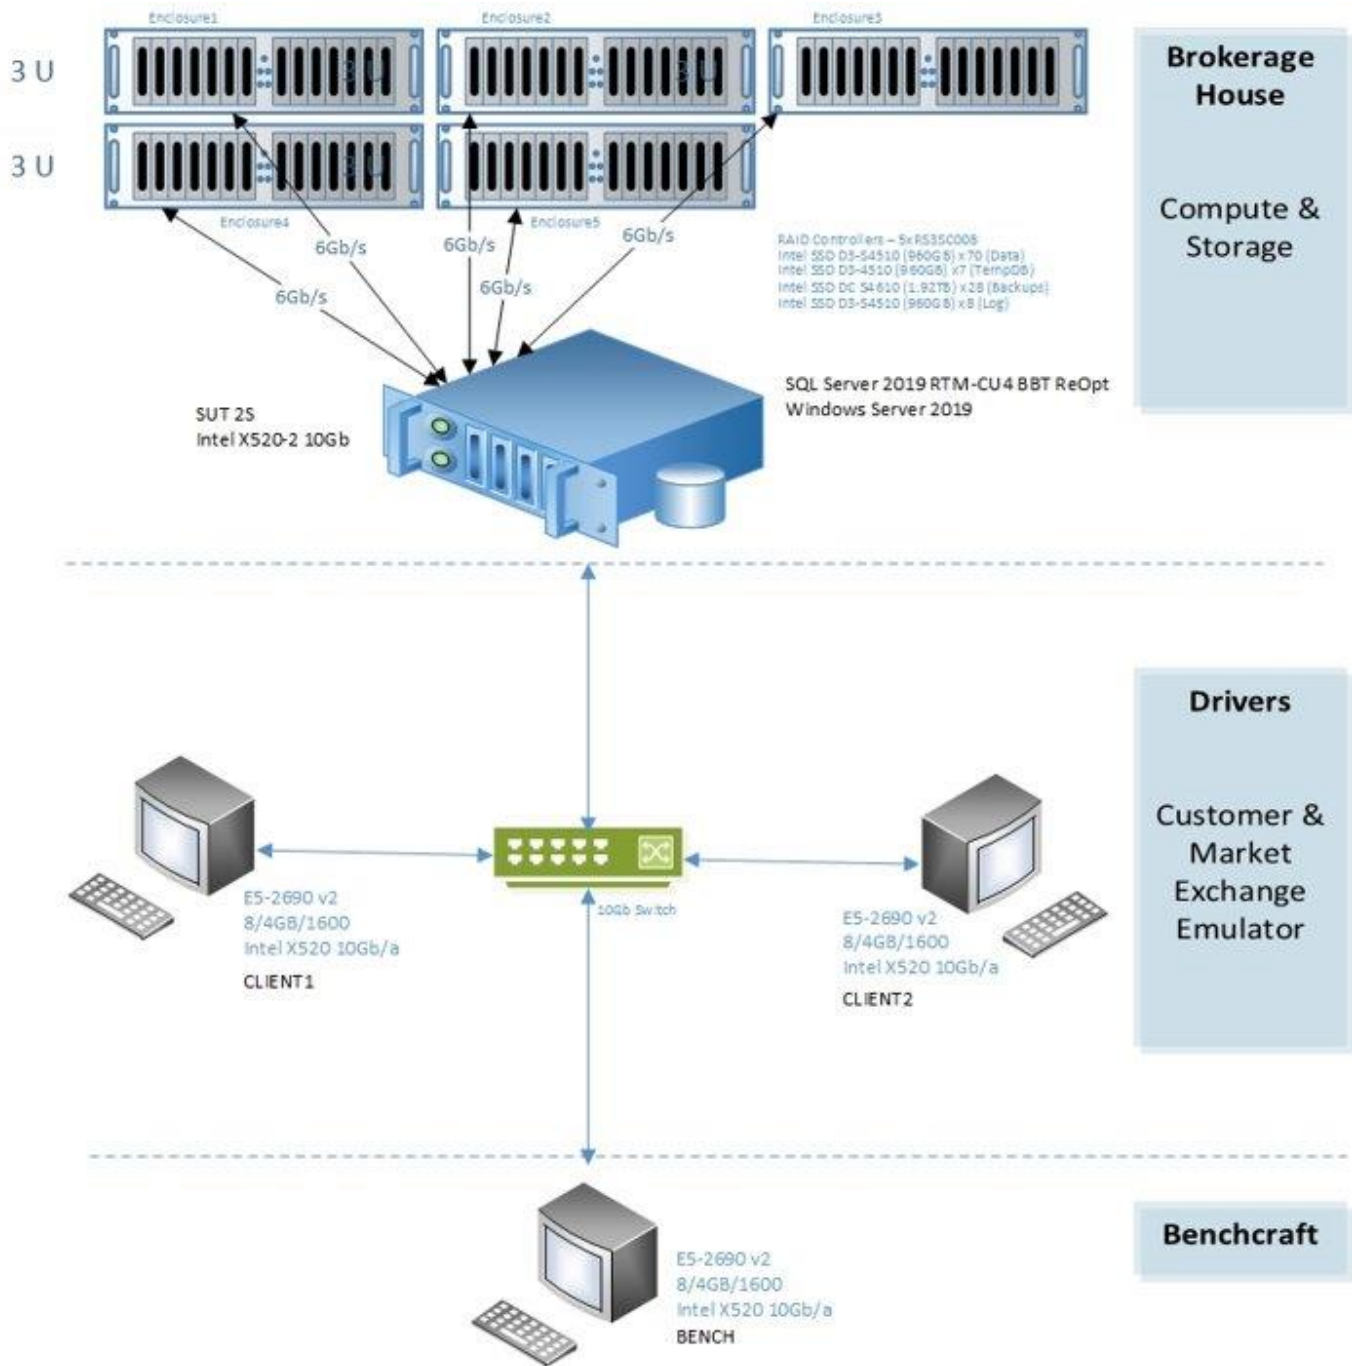
3.1. Microsoft Windows Server* 2019 settings

Configuration Item	Recommended Value
Windows Firewall	Control Panel -> Windows Firewall -> "Turn Windows Firewall on or off" (left side pane) -> select "Turn off Windows Firewall" for both private and public networks.
Visual Performance Setting	Computer -> properties -> advanced system settings-> advanced ->performance -> settings -> visual: adjust best performance
Virtual Memory Setting	Computer -> properties -> Advanced tab, virtual memory change: custom size: 4096 initial and maximum (not system managed)
Power Profile	Control Panel->Hardware and Sound->Power Options → High Performance Option
Large Page Enable	Large Page Enable (Reboot to take effect): Control Panel >>Administrative tools → local security policy → local policy → user rights assignment → Lock pages in memory → properties → add user or group → Administrators (or the user to be used with SQL Server)
Remove Windows Defender	PowerShell command: Remove-WindowsFeature -Name Windows-Defender-Features

3.2. OLTP Architecture

The TPC-E like workload consists of transactions that simulate the interchanges commonly associated with brokers, customers, and a real-time stock exchange. These transactions are intended to represent a balanced mixture of disk input/output and CPU usage. Performance is measured in transactions per second and reported as tpsE.

Example of an OLTP benchmarking hardware configuration:



3.3. Tuning SQL Server for OLTP Workload

The following `sp_configure` commands should be used to configure SQL Server:

```
sp_configure show_advanced_options,1
go
RECONFIGURE WITH OVERRIDE
go
```

```
sp_configure backup_compression,1
go
RECONFIGURE WITH OVERRIDE
go
sp_configure "default trace enabled",0
go
RECONFIGURE WITH OVERRIDE
go
sp_configure lightweight_pooling,1
go
RECONFIGURE WITH OVERRIDE
go
sp_configure max_degree_of_parallelism,1
go
RECONFIGURE WITH OVERRIDE
go
sp_configure max_serv,<This number should be equal to 90% of system memory>
go
RECONFIGURE WITH OVERRIDE
go
sp_configure max_worker_threads,3000
go
RECONFIGURE WITH OVERRIDE
go
sp_configure priority_boost,1
go
RECONFIGURE WITH OVERRIDE
go
sp_configure recovery_interval,32767
go
RECONFIGURE WITH OVERRIDE
go
sp_configure remote_query_timeout,0
go
RECONFIGURE WITH OVERRIDE
go
sp_configure set_working_set_size,1
go
RECONFIGURE WITH OVERRIDE
go
sp_configure network_packet_size,8192
go
```



```
RECONFIGURE WITH OVERRIDE  
go
```

The following flags should be used on the command line when starting up SQL Server prior to running the OLTP workload.

```
Sqlservr.exe -c -x -T3502 -T834 -T661 -T8744 -T652 -s SQLOLTP
```

4. Related Tools

Microsoft Windows Performance Monitor tool (Perfmon) can be used to monitor overall system and disk performance metrics.

5. Best Practices for Testing and Verification

This workload is best run just after a fresh restoration of the database from backup and restart of the system. This workload should drive CPU Utilization to 99% across all logical processors. CPU kernel utilization should be around 12-13%. This is an I/O intensive workload that is characterized by 600K IOPs, Random Access, and a 90/10 Read/Write ratio. The clients for this workload initiate transactions with 500+ simultaneous connections to the system under test with 500+ transactions in flight. This workload is more sensitive to memory speed than capacity. Typically, this workload reaches peak performance with memory size equal to about 3% of the database size. For verification of the database integrity, refer to the TPCE kit mentioned in resources for details. The TPCE kit provides verification tools suitable for an audit per the TPCE benchmark spec.

6. Conclusion

We have shared our best-known methods to optimally benchmark our 3rd Generation Intel Xeon Scalable Processors using an OLTP benchmark. We have covered both software and hardware configuration considerations to get the best performance.

7. Additional Resources

[1] Microsoft SQL Server: <https://www.microsoft.com/en-us/sql-server/sql-server-downloads>

[2] Transaction Processing Council: <http://www.tpc.org/>

[3] TPCE benchmark kit version 1.14.0:

http://tpc.org/TPC_Documents_Current_Versions/download_programs/tools-download-request5.asp?bm_type=TPC-E&bm_vers=1.14.0&mode=CURRENT-ONLY

8. Feedback

We value your feedback. If you have comments (positive or negative) on this guide or are seeking something that is not part of this guide, [please reach out](#) and let us know what you think.

Notices & Disclaimers

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.