Hewlett Packard Enterprise

Reference Architecture

HPE Reference Architecture for deploying Microsoft SQL Server 2022 on HPE ProLiant DL365 Gen11 Server

CONTENTS

EXECUTIVE SUMMARY

Focusing on your data strategy is an imperative for many businesses not only to manage the massive growth of data and devices, future-proof your business, all while keeping an eye on costs, and ensure it is secure, performant, and available. With the growing data consumption, the demand for high compute processing is also growing. This all requires a more capable data strategy than ever before. It must also support a storage system that can easily scale capacity and/or performance as CPU and memory resources are increased. This is the level of flexibility that many businesses are seeking.

HPE ProLiant DL365 Gen11 server is rack optimized 1U 2P dense solution to meet the demands of high-performance workloads. It delivers exceptional compute performance, upgraded high-speed data transfer rate and memory depth at 2P compute capability. Enhanced security features with the HPE silicon root of trust are built into the firmware, creating a digital fingerprint for the AMD Secure Processor to validate safe operation prior to boot. The HPE ProLiant DL365 Gen11 server is an excellent choice for those who require compute density with built-in security and flexibility. The HPE ProLiant DL365 Gen11 server coupled with the HPE Alletra storage array is an ideal solution to meet the demands of growing database requirements.

This solution demonstrates the HPE ProLiant DL365 Gen11 server can support an Online Transactional Processing (OLTP) workload with Microsoft SQL Server 2022, performing at 10.3 million transactions per minute with about 80 % CPU utilization. With high-performance NVMe drives, the HPE Alletra Storage arrays used in our internal test complements the HPE ProLiant DL365 Gen11 server with its ability to deliver these transaction rates with little to no measurable disk latency. The HPE Alletra 6090 delivers on this promise of increased performance, as well as new capabilities for cloud-native management.

Target audience: CIOs, IT architects, IT managers, database engineers, and administrators. Working knowledge of server architecture, networking architecture, and storage design is recommended.

Document purpose: The purpose of this document is to describe a Microsoft SQL Server 2022 on Windows Server 2022 Reference Architecture using an HPE ProLiant DL365 Gen11 server platform and HPE Alletra 6090, highlighting recognizable benefits to technical audiences.

This Reference Architecture describes solution testing performed in June 2023.

SOLUTION OVERVIEW

This solution we tested in the HPE lab included an HPE ProLiant DL365 Gen11 server combined with an HPE Alletra Storage running Microsoft SQL Server 2022. We performed OLTP test case, to provide insights into OLTP workload performance with Microsoft SQL Server 2022 database.

HPE Alletra, the Hewlett Packard Enterprise answer to a new set of requirements for high-end data storage infrastructure, is a modern data storage platform that delivers the performance, availability, and scalability needed for mission-critical primary storage workloads.

Figure 1 shows the solution hardware diagram.

FIGURE 1. Solution hardware diagram

HPE ProLiant DL365 Gen11 server

HPE ProLiant DL365 Gen11 server is a highly scalable, compute-dense solution for high performance workloads. The new HPE ProLiant DL365 Gen11 server is a rack-optimized 1U 2P dense solution that delivers exceptional compute performance, upgraded high-speed data transfer rate

and memory depth at 2P compute capability. Powered by 4th Generation AMD EPYC Processors with up to 96 cores, increased memory bandwidth (up to TB), and high-speed PCIe Gen5 I/O, the HPE ProLiant DL365 Gen11 server is a superb rack-optimized, 1U 2P, dense solution.

The silicon root of trust anchors the server firmware to an HPE-exclusive ASIC, creating an immutable fingerprint for the AMD Secure Processor that must be matched exactly before the server boots. The HPE ProLiant DL365 Gen11 server is an excellent choice for those who require increased compute density with built-in security and flexibility.

FIGURE 2. HPE ProLiant DL365 Gen11 server

What's New

- Supports the 4th Generation AMD EPYC Processors that supports up to 96 cores, 400W, and 384MB of L3 Cache.
- 12 DIMM channels per processor for up to 6TB total DDR5 memory 4800MT/s.
- Advanced data transfer rate and higher network speed from the PCIe Gen5 serial expansion bus.
- Includes HPE Integrated Lights-Out 6 (iLO 6) server management software.
- Supports hot-pluggable, high-availability RAID M.2 boot options.
- Supports up to 20 EDSFF drive bay.
- Supports up to 2x Single Width or 2x Double Width GPU cards at the front chassis.
- The HPE ProLiant DL365 Gen11 server is tied into the silicon root of trust and the AMD Secure Processor, a dedicated security processor embedded in the AMD EPYC system on a chip (SoC), to manage secure boot, memory encryption, and secure virtualization.
- HPE ProLiant DL365 Gen11 servers use the silicon root of trust to anchor the firmware of an HPE ASIC, creating an immutable fingerprint for the AMD Secure Processor that must be matched exactly before the server will boot. This helps ensure malicious code is contained and healthy servers are protected.

HPE Alletra 6090

For business-critical workloads, HPE Alletra 6000 delivers fast, consistent performance and industry leading data efficiency. It enables IT to shift from owning and maintaining data infrastructure to simply accessing and utilizing it on-demand, as-a-service. Eliminate performance and efficiency trade-offs with no knobs or configurations to adjust and always-on data services. Get resilient storage with intelligence and a no single point of failure platform that together deliver 6-nines availability guaranteed. Deliver on recovery SLAs with fast, integrated app aware backup and recovery-on-premises and in the cloud.

FIGURE 3. HPE Alletra 6090

What's new

- Eliminate complexity by unifying infrastructure management silos under a cloud managed single pane of glass - accessible from anywhere and from any device.

- Developed from the foundation of HPE Nimble Storage architecture, proven to deliver speed, efficient, and resiliency with cloud agility.
- HPE Store More Guarantee1 delivers more effective capacity per terabyte of raw flash than competitive all-flash arrays.
- Eliminate forklift upgrades, non-disruptive controller upgrades, flat support pricing, and flexible consumption options.

Absolute resiliency

- 99.9999% (6-nines) guaranteed availability.
- Triple+ Parity RAID tolerates three simultaneous drive failures plus additional protection through intra-drive parity.
- App-granular, FIPS-certified encryption provides data at rest and over-the-wire protection. Secure data shredding is built in.
- Native application-consistent snapshots and replication plus integration with leading backup software.
- Redundant, hot-swap components including controllers, power supplies, SSDs, and IO cards.

AI-driven

- Predict and prevent disruptions across storage, services, and virtual machines.
- Redefine the support experience with predictive support automation that delivers an unprecedented experience with automated L1 and L2 and direct access to the resources you need.
- Pinpoint issues between storage and VMs and underutilized virtual resources without effort.
- Take the guesswork out of managing data infrastructure with AI-driven recommendations.

Built for cloud

- Set up in minutes because systems are automatically discovered, on-boarded, and configured.
- Say goodbye to time-consuming LUN-centric provisioning with AI-driven, intent-based provisioning of application workloads on infrastructure best suited for optimizing SLAs.
- Experience faster access to innovation with no disruptions because new features and enhancements are instantly available through self.
- Service upgrades.
- Manage from anywhere with simple global management driven through a SaaS-based user experience.

As-a-service

- Consume data infrastructure as-a-service via HPE GreenLake, eliminating up-front capital costs with a pay-per-use model.
- Shift from owning and maintaining data infrastructure to simply accessing and utilizing it on-demand.
- Free up your cash flow and increase financial agility with the right mix of subscription and consumption-based services.

Microsoft SQL Server 2022

Microsoft announced the general availability of SQL Server 2022 in November 2022. Microsoft SQL 2022 is designed to be a hybrid data platform and is the release that brings cloud innovations to customers.

There are many newly available features in Microsoft SQL Server 2022, but a few to highlight include the following:

- Improvements in intelligent query processing (IQP) and Query Store continuing the tradition of release over-release performance. Improvement in the performance of buffer pool scan operations on large-memory machines by utilizing multiple CPU cores.
- Ordered clustered Columnstore index (CCI) sorts the existing data in memory before the index builder compresses the data into index segments. This has the potential of more efficient segment elimination, resulting in better performance as the number of segments to read from disk is reduced.
- Improved memory management in large memory servers to reduce out-of-memory conditions.
- Enhanced spinlock algorithms. Spinlocks are a huge part of the consistency inside the engine for multiple threads. Internal adjustments to the Database Engine make spinlocks more efficient. This change was introduced in Microsoft SQL Server 2022 (all editions) and included in Microsoft Azure SQL Database and Microsoft Azure SQL Managed Instance.
- Instant file initialization for transaction log file growth events. In general, transaction log files cannot benefit from instant file initialization (IFI). Starting with Microsoft SQL Server 2022 (16x) (all editions) and in Azure SQL Database, instant file initialization can benefit transaction log growth events up to 64 MB. The default auto growth size increment for new databases is 64 MB. Transaction log file autogrowth events larger than 64 MB cannot benefit from instant file initialization.

See the Resources and additional links section at the end of this Reference Architecture for links to find more information about the Microsoft SQL Server 2022 release.

SOLUTION COMPONENTS

The following components were used for this solution, and the configuration can be varied based on needs.

Hardware

Table 1 summarizes hardware components that were utilized in the design and construction of this Reference Architecture.

TABLE 1. Hardware components utilized in this solution

Component	Qty	Description
HPE ProLiant DL365 Gen11 8SFF CTO server • 2 x AMD EPYC 9354 (3.25GHz/32-core/280W) Processor • 24 x HPE 64 GB Dual Rank DDR5 4800 MT/s RDIMM • HPE 64GB Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit • 2 x HPE 960 GB SAS 12G MU SFF SSD • HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter	1	
HPE Alletra 6090 CTO Base Array • 2 Controller nodes • 24 x 7.68TB NVMe SSD • 2 x HPE Alletra 6000 2x10/25GbE 2-port FIO Adapter Kit	1	Storage for Microsoft SQL Server database
Aruba 8325-32C 32-port 100G QSFP+/QSFP28	2	Switch for providing connectivity between HPE Alletra and Servers
Aruba 6300M 48-port 1GbE and 4-port SFP56	2	Management Switch

Software

Table 2 shows the software components used in this solution configuration.

TABLE 2. Software used in this solution

Component	Version	Description
Microsoft Windows Server 2022	Datacenter Edition	Operating System
Microsoft SQL Server 2022	CU3	Microsoft SQL Server 2022 Enterprise Evaluation edition for database

BEST PRACTICES AND CONFIGURATION GUIDANCE FOR THE SOLUTION

For a Microsoft SQL Server solution, there are a few recommendations listed below:

HPE ProLiant DL365 Gen11 server

Workload Profiles is one of the HPE Intelligent System Tuning (IST) features and allows you to tune the resources in your HPE ProLiant server by choosing a preconfigured workload profile. The server will automatically configure the BIOS settings to match the selected workload.

Transactional Application Processing profile is intended to be used for business processing environments, such as OLTP applications that require a database back-end.

For HPE ProLiant DL365 Gen 11 server Workload Profile is set to Transactional Application Processing.

Microsoft Windows Server 2022 and Storage settings

Some of the customized settings are listed as follows.

- Windows Server Power plan set to High Performance.
- Lock Pages in memory user right was enabled for Microsoft SQL server service account to keep data in physical memory, preventing the system from paging the data to virtual memory on disk.
- Format the disks with NTFS file system with 64 K allocation unit size.
- For multipathing MPIO was enabled. MPIO and iSCSI setup was configured using HPE Connection manager.

For steps to enable d Lock Pages in memory privilege and MPIO and iSCSI setup, see Appendix B: Configuration Settings for more details.

HPE Alletra 6030 storage array configuration

This configuration assumes that the storage array is already registered and configured through HPE's Data Services Cloud Console (DSCC) and is online and accessible from the local network. See this preinstallation link.

HPE Alletra 6090 volumes were configured for SQL with below performance policy:

- Microsoft SQL Server for database data volume
- Microsoft SQL Server Logs for database log volume

This configuration used iSCSI to map volumes from the HPE Alletra 6090 array, so an Initiator Group has been created and initiators from the windows server added. After configuring the Initiator Group, array volumes mapped to it.

Microsoft SQL Server settings

Below Microsoft SQL Server settings have been configured.

- Multiple tempdb data files have been created.
- Database instant file initialization enabled during the Microsoft SQL server installation by selecting the check box Grant Perform Volume Maintenance Task privilege to Microsoft SQL Server Database Engine Service in the installation wizard.
- Set max server memory to 90% of total memory, this ensures enough memory left for OS operations.
- Soft Numa disabled using the following command:

```
ALTER SERVER CONFIGURATION SET SOFTNUMA OFF
    RECONFIGURE WITH OVERRIDE
    GO
```

  Changing the value of this setting requires a restart of the database engine to take effect.
- Below mentioned Trace Flag was enabled:
  - TraceFlag 834- Use large pages for buffer pool.
  - TraceFlag 652 - Disables page prefetching scans.
- Max degree of parallelism set to 1. This option limits the number of processors used in parallel plan execution and is often recommended in OLTP environments where we need to perform as many transactions as possible, and they rely on short and simple queries.

NOTE
Higher values for Max degree of parallelism (MaxDOP) can be used as needed but as always you must test your workload to find optimal values.

CAPACITY AND SIZING

Workload description

The OLTP workload tests were conducted using HammerDB, an open-source tool. It implements a computer system to fulfil orders from customers to supply products from a company. The company sells 100,000 items and keeps its stock in warehouses. Each warehouse has 10 sales districts, and each district serves 3000 customers. The customers call the company whose operators take the order, each order containing a number of items. Orders are usually satisfied from the local warehouse however a small number of items are not in stock at a particular point in time and are supplied by an alternative warehouse. For this Reference Architecture, HammerDB 4.7 was used to implement an OLTP workload. The testing was performed using multiple databases depending on how large the configuration was.

Each database was based on a schema consisting in 5000 warehouses for a 500 GB size.

Table 3 show the database file location. Separate volumes and disks were created for system databases, data, log, and tempdb. Tempdb database was hosted on separate drive (I:\Data and H:\Log). The environment was tuned for maximum user transactions. After the database was tuned, the transaction rates were recorded for various SQL user connections.

TABLE 3. Database file location configuration

Database Name	Database Size	Data file location	Log file location
Tpcc	500 GB	J:\Data	N:\Log
tpcc2	500 GB	K:\Data	O:\Log
tpcc3	500 GB	L:\Data	P:\Log
tpcc4	500 GB	M:\Data	Q:\Log

The workload is generated from external machines to find out an optimal number of active users driving OLTP transactions on the Microsoft SQL Server database. Transactions were driven from the HammerDB testing tool by a varying number of active users and observed the overall performance. It is highly CPU intensive and moderately I/O intensive.

Microsoft SQL Server Resource Governor was enabled and configured such that each database had its CPU/Memory affinity isolated at CPU socket level. For the most part the default resource pool settings were used to create the resource pools, except the AFFINITY NUMANODE value was uniquely set for each database. For the OLTP workload, to have consistent and reproducible results across multiple workloads runs, we set the MAXDOP value to 1. Unlike DSS/Warehouse queries, OLTP transactions do not yield a performance benefit with increased parallelism. The recovery interval (minutes) was set to a max value of 32767 to avoid checkpoints during the test.

Tables 4 show how the Resource Governor was configured to funnel workload sessions with unique logins to use the corresponding workload group, which in turn, used the corresponding resource pool. Two users were used for access database workloads. A script was created for ease of configuration. Resource Governor configuration is given in Appendix C: Configuration Settings.

TABLE 4. Database file location configuration

UserName	Workload Group	Resource Pool	Affinity Numa Node	Databases
dbuser1	Wg1	RP1	0	tpcc, tpcc2
dbuser2	Wg2	RP2	1	tpcc3, tpcc4

Microsoft SQL Server maximum memory was set at 90% of physical memory for our workload. Once the SQL databases were created and laid out optimally on the storage, Microsoft SQL Server memory was set appropriately, our OLTP SQL environment was ready for SQL OLTP workload performance testing.

For the OLTP workload, HammerDB provides a real-world type of scenario that consumes both CPU for the application logic and I/O. The HammerDB tool implements an OLTP-type workload with small I/O sizes of a random nature. Each test started with freshly restored OLTP databases and Microsoft SQL server restart in order to have a common start point. Ramp up time was set to two (2) mins and test duration was set to 10 mins in HammerDB. System performance monitoring was captured for the test time duration.

Workload test results

The Transactions Per Minutes (TPM) metric was captured with HammerDB while CPU, physical disk, network resources were monitored with window's Performance monitor and Resource Monitor. Also, Microsoft SQL Server Activity monitor was used for monitoring.

For testing, four databases of 500 GB each were used, and different user workloads were driven. For the test peak transaction per minute of 10.30 million was achieved with 110 user workloads for each database with CPU utilization at about 80%.

Always On Availability Groups

In this Reference Architecture, primary and standby databases are configured with Microsoft SQL server Always On Availability Groups. The databases can be located on the same premises or in geographically dispersed data centers and data can be replicated by Microsoft SQL server. Microsoft SQL Server Always On Availability Group was configured to validate the high availability.

Analysis and recommendations

- HPE ProLiant DL365 Gen11 server is a rack optimized dense solution, with high performance compute.
- The I/O subsystem is especially important for database workloads because data needs to be read and written at extreme speeds to achieve high performance. HPE Alletra 6090 provides the faster response time which allows system to perform more transactions.
- Optional deployment with HPE GreenLake provides flexible as-a-service consumption while maintaining on-premises control.

SUMMARY

The goal of this Reference Architecture is to provide a high-performance hardware platform combined with the most available database platform possible, offering the best experience for users of business-critical applications.

HPE ProLiant DL365 Gen11 server is a rack optimized 1U 2P dense solution that deliver high compute performance, reliable, and secure server. Running Microsoft SQL server, the HPE ProLiant DL365 Gen11 server combined with an HPE Alletra 6090 provides a significant OLTP performance. The HPE Alletra 6090 is the new all-NVMe storage array that delivers on the promise of increased performance as well as new capabilities for cloud-native management.

APPENDIX A: BILL OF MATERIALS

NOTE

Part numbers are at time of publication/testing and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult with your Hewlett Packard Enterprise Reseller or Hewlett Packard Enterprise Sales Representative for more details. hpe.com/us/en/services/consulting.html

TABLE A1. Bill of materials

Part number	Quantity	Description
		Rack and Switches
P9K10A	1	HPE 42U 600mmx1200mm G2 Kitted Advanced Shock Rack
P9K10A #001	1	HPE Factory Express Base Racking Service
JL627A	2	Aruba 8325-32C 32-port 100G QSFP+/QSFP28 Back-to-Front 6 Fans and 2 Power Supply Bundle
JL762A	2	Aruba 6300M 48-port 1GbE and 4-port SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle
		HPE ProLiant DL365 Gen11 server
P53933-B21	2	HPE ProLiant DL365 Gen11 8SFF Configure-to-order Server
P53701-B21	4	AMD EPYC 9354 3.25GHz 32-core 280W Processor for HPE
P53701-B21#0D1	4	Factory Integrated
P47789-B21	2	HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller
P47789-B21#0D1	2	Factory Integrated

Part number	Quantity	Description
P56901-B21	2	HPE ProLiant DL365 Gen11 8SFF OROC x1 SAS/SATA/NVMe Cable Kit
P56901-B21#0D1	2	Factory Integrated
P58462-B21	14	HPE ProLiant DL3XX Gen11 1U Performance Fan Kit
P58462-B21#0D1	14	Factory Integrated
P55020-B21	2	HPE ProLiant DL365 Gen11 8SFF Tri-Mode U.3 x1 BC FIO Backplane Kit
P08040-B21	2	HPE iLO Common Password FIO Setting
P40510-B21	4	HPE 960GB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD
P40510-B21#0D1	2	Factory Integrated
P31348-B21	2	HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter
P31348-B21#0D1	2	Factory Integrated
P50312-B21	48	HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit
P50312-B21#0D1	48	Factory Integrated
P57849-B21	2	HPE ProLiant DL3X5 Gen11 x16 OCP1 OCP2 2P Upgrade Cable Kit
P57849-B21#0D1	2	Factory Integrated
P38997-B21	4	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
P38997-B21#0D1	4	Factory Integrated
P58457-B21	4	HPE ProLiant DL3X5 Gen11 CPU Performance 1U Heat Sink Kit
P58457-B21#0D1	4	Factory Integrated
P52351-B21	2	HPE DL3XX Gen11 Easy Install Rail 2 Kit
P52351-B21#0D1	2	Factory Integrated
P48922-B21	2	HPE ProLiant DL3XX Gen11 Intrusion Cable Kit
P48922-B21#0D1	2	Factory Integrated
P50227-B21	2	HPE 1.6TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD
		HPE Alletra 6090 Storage
R4U32A	1	HPE Alletra 6090 CTO Base Array
R0R12A	2	HPE Alletra 6000 2x10/25GbE 2p FIO Adapter
R7S86A	1	HPE Alletra 6000 184TB 24x7.68TB FIO Bdl
R9D23A	4	HPE C13 - C14 2m WW PDU FIO Power Cord
R9X15A	1	HPE Alletra Tier 1 Storage Array Std Trk
R7G13A	1	HPE Alletra 6000/H 4x1600W FIO AC PS Kit
Q8G27B	1	HPE Tier 1 Storage OS Default FIO SW
R7N52AAE	1	HPE Alletra SW/Sup SaaS

APPENDIX B: CONFIGURATION SETTINGS

Lock pages in memory (LPIM) enablement

To enable the Lock pages in memory option by granting the privilege to the service account for the instance of Microsoft SQL server, follow these steps:

1. Click **Start** > **Run**.
2. In the Open box, type **gpedit.msc**.
3. From the Local Group Policy Editor console, expand **Computer Configuration** > **Windows Settings** > **Security Settings** > **Local Policies**.

4. Select the Users Rights Assignment folder. The policies will be displayed in the details pane.
5. In the details pane, double-click Lock Pages in Memory.
6. In the Local Security Setting dialog box, select Add User or Group. Add Microsoft SQL Server service account and select OK.
7. Restart the Microsoft SQL Server service for this change to take effect.

Multipathing and iSCSI setup

1. On the command line, check for previously installed multipathing software:

```
Get-WindowsOptionalFeature -Online -FeatureName MultiPathIO
```

2. Open the Server Manager. Click Server Manager > Dashboard > Add Roles and features. This starts the Add Roles and Features wizard. Enable the MPIO feature using the wizard. Restart the windows server.

Using HPE connection manager for MPIO and iSCSI setup:

Instead of manually configuring MPIO and each network path from within the Windows iSCSI Initiator, the HPE Storage Connection Manager tool should be used to automatically configure them.

The HPE Storage Connection Manager is designed to simplify making and maintaining the optimal number of iSCSI connections (also known iSCSI sessions) between the Windows host and the array. The HPE Storage Connection Manager feature is designed for arrays running array OS 2.x and later that are using the Automatic iSCSI Connection Method under Network Configuration Group settings.

Once HPE connection manager installed, start the HPE connection manager GUI and perform the following tasks:

- Include or exclude Host IP address for MPIO: This setting enables you to select which IP addresses the HPE Storage Connection Manager feature should use for host MPIO connections.
- Add discovery target portal: A discovery target portal refers to the Discovery IP address of a storage array group. A host uses the discovery target portal of the array to discover the iSCSI targets available on the array. You can add the target portal of a Storage array or array group even if it is not in the HPE Storage Connection Manager list.
- View Discovered Volumes: From the Volume tab you can see the volume group which was created on HPE Alletra array.
- Connect to an iSCSI Target: On the Volume tab in the Discovered Volumes list, highlight the target you want, and click Connect.
- The Connect to Target dialog box appears: It displays the iSCSI target name and the target IQN.

After configuring the HPE Storage Connection Manager tool and rebooting, the Nimble server device should have been added to MPIO. The Windows iSCSI Initiator should now have all network paths between the storage array and the server added in the Favorite Targets tab.

For detailed steps to perform these task, please refer to the HPE Storage connection Manger for Windows section from this reference link *Windows Integration Guide (hpe.com)* .

APPENDIX C: RESOURCE GOVERNOR SETTINGS

Following T-SQL scripts were used to configure the Resource Governor. It performs three tasks- create a user login for a specific workload group, create a resource pool that has affinity to NUMA Node 0, and create a workload group that will use the NUMA Node 0-affinitized resource pool. Substitute the name of login user, resource pool, and workload group to create the remaining Resource Governor profiles as mentioned in Table 4.

```
--- Create or Alter Login dbuser1
IF NOT EXISTS (SELECT loginname FROM master.dbo.syslogins WHERE name = 'dbuser1')
        CREATE LOGIN dbuser1
    WITH    PASSWORD='Password!234',
            DEFAULT_DATABASE=master,
            DEFAULT_LANGUAGE=us_english,
            CHECK_EXPIRATION=OFF,
            CHECK_POLICY=OFF
ELSE
        ALTER LOGIN dbuser1 WITH PASSWORD = 'Password!234'
GO
ALTER LOGIN dbuser1 ENABLE
GO

ALTER SERVER ROLE sysadmin ADD MEMBER dbuser1
GO

--- Create or Alter Resource Pool DBRP1
IF EXISTS ( SELECT name FROM sys.resource_governor_resource_pools WHERE name = 'DBRP1' )
    ALTER RESOURCE POOL DBRP1 WITH(min_cpu_percent=0,
                max_cpu_percent=100,
                min_memory_percent=0,
                max_memory_percent=100,
                cap_cpu_percent=100,
    AFFINITY NUMANODE = (0),
                min_iops_per_volume=0,
                max_iops_per_volume=0)
ELSE
        CREATE RESOURCE POOL [DBRP1] WITH(min_cpu_percent=0,
                        max_cpu_percent=100,
                        min_memory_percent=0,
                        max_memory_percent=100,
                        cap_cpu_percent=100,
        AFFINITY NUMANODE = (0),
                        min_iops_per_volume=0,
                        max_iops_per_volume=0)
GO
--- Create or Alter Workload Group WGDB01
IF EXISTS ( SELECT name FROM sys.resource_governor_workload_groups WHERE name = 'WGDB01' )
        ALTER WORKLOAD GROUP WGDB01 WITH(group_max_requests=0,
                importance=High,
                request_max_cpu_time_sec=0,
                request_max_memory_grant_percent=25,
                request_memory_grant_timeout_sec=0,
                max_dop=1) USING WGDB01
ELSE
        CREATE WORKLOAD GROUP [WGDB01] WITH(group_max_requests=0,
                        importance=High,
                        request_max_cpu_time_sec=0,
                        request_max_memory_grant_percent=25,
                        request_memory_grant_timeout_sec=0,
                        max_dop=1) USING [DBRP1]
GO
```

```
ALTER RESOURCE GOVERNOR RECONFIGURE;
GO
```

Below is the sample classifier function script that assigns the user to a workload group which is defined by the script above to use the corresponding resource pool.

```
USE [master]
GO
SET ANSI_NULLS ON
GO
SET QUOTED_IDENTIFIER ON
GO

CREATE FUNCTION [dbo].[OLTPClassifier]() RETURNS sysname
    WITH SCHEMABINDING
AS
BEGIN
        DECLARE @WorkloadGroup AS SYSNAME
        IF(SUSER_NAME()= 'dbuser1')
                SET @WorkloadGroup='WGDB01'
        ELSE IF (SUSER_NAME()= 'dbuser2')
                SET @WorkloadGroup='WGDB02'
RETURN @WorkloadGroup
END
GO
ALTER RESOURCE GOVERNOR WITH (CLASSIFIER_FUNCTION = [dbo].[OLTPClassifier]);
GO
ALTER RESOURCE GOVERNOR RECONFIGURE;
GOGO
```

RESOURCES AND ADDITIONAL LINKS

HPE Reference Architectures, hpe.com/info/ra

HPE Servers, hpe.com/servers

HPE Storage, hpe.com/storage

HPE Networking, hpe.com/networking

HPE GreenLake Advisory and Professional Services, https://www.hpe.com/us/en/services/consulting.html

Microsoft SQL Server 2022, aka.ms/sqlserver2022

To help us improve our documents, please provide feedback at hpe.com/contact/feedback.

LEARN MORE AT

https://www.hpe.com/us/en/servers/superdome.html

a50009050enw, version 1.0, July 2024

Hewlett Packard Enterprise