



Hewlett Packard
Enterprise

HPE Reference Configuration for Microsoft SQL Server 2016 on HPE Synergy composable infrastructure

Contents

Executive summary.....	3
Introduction.....	3
Solution overview.....	4
Solution components.....	5
Hardware.....	5
Software.....	11
Microsoft SQL Server 2016 and HPE Synergy.....	11
Microsoft SQL Server 2016 configuration guidance and overview.....	11
Capacity and sizing.....	13
Analysis and recommendations.....	14
Summary.....	15
Appendix A: NUMA Affinity bitmasks and ports.....	16
Resources and additional links.....	17

Executive summary

In today's IT environment, organizations are faced with the challenge of supporting two models of IT. The first model is the traditional IT environment – a rigid, inflexible and stove-piped IT environment that is designed application by application and can sometimes lead to extensive unused capacity and inefficiency. The second model is aligned to the idea economy built on the fundamental requirements of speed, agility, and the responsiveness needed for new mobile and cloud applications which have become foundational to growing the business. The challenge for many organizations is to maintain the structure and security of traditional IT, while allowing flexibility and speed for new applications and operating models such as virtualization, hybrid cloud, and DevOps. Organizations today are hindered by the time consuming process of manual IT operations which cause significant delays in application deployment. To solve these challenges, Hewlett Packard Enterprise has created a way for organizations to meet the growing demand for faster, more flexible IT infrastructure, also known as Composable Infrastructure.

Composable Infrastructure consists of three elements:

- Fluid resource pools – Compute, Storage, and Fabric
- Software Defined Intelligence – automatically manage and quickly provision hardware resources based on an application's requirements
- A Unified API to control the integration of these resources with a single line of code

Database management systems can be one of the more expensive and typically over-provisioned systems in the data center. Database systems that are highly transactional and business critical may also require large amounts of compute and memory, as well as a seemingly endless need for high performance storage. These systems also generally require higher levels of availability, which may also add complexity and cost to the solution. However the larger expense to the business may be with labor intensive processes and the length of time that may challenge IT to deliver.

In this Reference Configuration, we showcase a Microsoft® SQL Server Database workload running on the HPE Synergy architecture – a composable infrastructure. The test workload is highly transactional and treated as business critical in management and operation. The workload simulates both an increase in numbers of transactions as well as an increase in data volumes while maintaining availability. In addition, read only replicas are also added, to showcase changing workload requirements. Read only replicas are used to provide a mechanism to offload queries from the transactional system. As computing and memory requirements change, HPE Synergy is equipped to add compute modules which can be easily provisioned, scaling from a two socket HPE Synergy 480 Gen9, up to a four socket HPE Synergy 680 Gen9. This paper also showcases the ability to use Microsoft SQL Server - AlwaysOn technologies to create high availability, as well as the opportunity to have offloaded read-only isolated workloads. This paper features a reference configuration for HPE Synergy and its flexible, composable infrastructure for Microsoft SQL Server Enterprise workloads.

Target audience: Chief information officers (CIOs), chief technology officers (CTOs), data center managers, enterprise architects, systems engineers, database administrators and others wishing to learn more about this reference configuration from HPE.

Introduction

This white paper is part of a series of technical documents to showcase the HPE Synergy composable infrastructure, and the flexibility it offers for many business and mission critical enterprise applications. Composable infrastructure is designed around three core principles:

- Fluid resource pools – Including compute, storage, and networking that can be deployed, managed and reconfigured as business and application needs change.
- Software Defined Intelligence – The ability to use templates and profiles to automatically manage and quickly provision hardware resources based on an application's requirements.
- A Unified API – Exposing infrastructure as code and providing access to hardware resources and management of these resources over a common API.

With HPE Synergy, resources (Compute, Storage, and Networking) can be rapidly provisioned for a specific purpose or function. When those resources are no longer needed, they can return to the pool ready for the next job. This approach allows organizations to deploy new SQL Server databases and applications quickly, more efficiently, and still deliver the performance necessary for these workloads. HPE Synergy allows IT to become more of a business partner, and not just a service provider, enabling rapid delivery of the applications necessary to the business.

This reference configuration is intended to demonstrate how HPE Synergy and the composable infrastructure provide solutions for a number of common needs IT organizations must consider when deploying SQL Server (or other business critical database workloads). Figure 1 below highlights a number of these key considerations and inhibitors IT faces when deploying SQL Server.

Database infrastructure needs

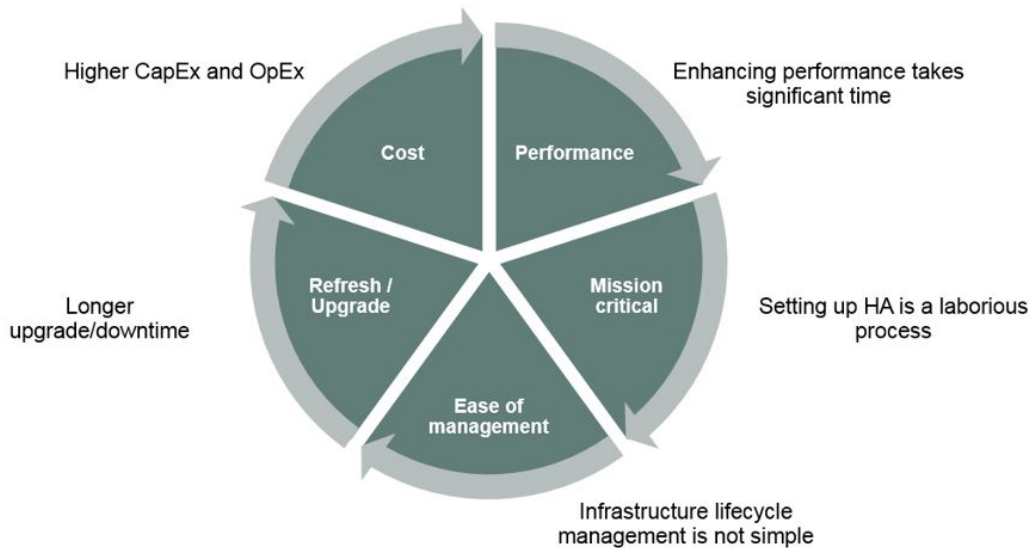


Figure 1. Database Infrastructure needs

When deploying new SQL Server database workloads there are a number of factors that need to be addressed by IT and the line of business (LOB) requestor that will drive the overall IT requirements. An understanding of the applications' performance, SLAs, and lifecycle (growth/upgrade) requirements are all key factors that drive variability and complexity into the overall solution design and cost.

The HPE Synergy composable infrastructure model provides the design fluidity (both in terms of flexibility and scalability) needed to address the wide range of SQL Server database deployment scenarios for the business. Performance, high availability, and scalability can be more easily, efficiently, and effectively managed with a composable infrastructure.

Solution overview

The HPE Synergy 12000 Frame offers the following composable infrastructure components for deploying Microsoft SQL Server.

- Two and/or four socket HPE Synergy Compute Modules including the HPE Synergy 480, HPE Synergy 660, and HPE Synergy 680
- Virtual Connect (HPE Virtual Connect SE 40Gb F8 Module for HPE Synergy)
- HPE Synergy D3940 Storage Module supporting up to 40 hot-pluggable, small form factor drives (12G/6G SAS, SATA, SSD or a combination)
- HPE Synergy Composer powered by HPE OneView

Demands for high performance and availability are two key drivers when designing SQL Server database environments. The HPE composable infrastructure provides a foundation for deploying applications and database workloads across physical or virtual environments, at small or large scale; and with the right type of storage to meet various performance and SLA requirements.

In many IT organizations there can be a complicated synchronization between teams managing and provisioning resources which are utilized by a database system. HPE Synergy's composable infrastructure offers tools and intelligence to provide a single view into the utilization of resources – and a template based deployment and management stack for ease of implementation. IT organizations can leverage template-based provisioning to simplify the initial database server deployment. Using templates to specify items, such as firmware, BIOS, boot settings, networking, storage, etc. can provide a mechanism to quickly provision consistent resources as the business demands and workload change.

This reference configuration focuses on deploying a Microsoft SQL Server 2016 database environment designed to run a transactional workload. High Availability is additionally provided by enabling Microsoft SQL Server's "AlwaysOn". This example was selected to provide a real-world, complex solution demonstrating the scale-up, and scale-out fluidity of the HPE Synergy composable infrastructure. The workload was first built on the HPE Synergy 660 Gen9 Compute Module. Pressure was then applied to the system by increasing the volume of transactions and the volume of data. The simulated transaction rates and volumes, were progressively increased until performance issues began to surface. The HPE Synergy 680 and 480 Compute Modules were composed and deployed in the HPE Synergy system, and then joined to the SQL Server availability replica. The databases were failed-over to the newly added HPE Synergy 680 Compute Module. Additionally demonstrated, was the concept of leveraging less performant resources for less demanding workloads. In this particular case, the premise was to off-load extracts or reporting functions to a read-only replica on the HPE Synergy 480 Gen9.

Note, for this white paper the focus is specifically on the transactional database workload tier. However, the flexibility and composability of the HPE Synergy infrastructure provides a platform for supporting all types of enterprise applications; including the front-end and mid-tier application servers, along with decision support/data warehouse workloads.

Solution components

The HPE Synergy components used in this solution were selected to demonstrate an example proof point to highlight the fluidity of the resource pools within the HPE Synergy environment for Microsoft SQL Server 2016 database workloads. For simplicity, the reference configuration (RC) focuses on a design that resides within a single HPE Synergy 12000 Frame. Depending on the level of HA required for the workload, multiple frames can be considered for additional resources or isolation of SQL Server databases and replicas.

Hardware

The solution design in this RC provides an example of a highly available SQL Server 2016 DB workload environment. The design was modeled to support a combined database capacity of 3 TB. The solution follows SQL Server deployment best practices as outlined in the "Microsoft SQL Server 2016 configuration guidance and overview" section below. The specific hardware components selected were intended to provide an enterprise class database system infrastructure.

HPE Synergy 12000 Frame

This RC provides a simplified tested example for deploying an HA SQL Server 2016 environment using a single frame. The HPE Synergy 12000 Frame consists of a mixture of two and four socket compute modules and internal DAS storage. The intent is to provide visibility to the range of resources available as part of the compute and storage resource pools. Figure 2 provides the specific tested solution configuration.




Figure 2. HPE Synergy frame configuration for Microsoft SQL Server 2016

The specific slot location of the compute and storage resources in the frame are specifically relevant to this RC configuration. Please review the official HPE Synergy documentation for any module placement location best practices.

HPE Synergy Compute Modules

For this solution, three different compute modules were utilized – HPE Synergy 480 Gen9 two socket, HPE Synergy 660 four socket, and HPE Synergy 680 four socket – shown in Figure 3 below. However, the intent of this RC isn't to specify that three different compute modules should be utilized in the SQL Server design. The intent is to showcase the progression of compute resources available for database deployments within the HPE Synergy solution. This variability allows you to match the workload to the right compute resource, ensuring that organizations don't have to over or under provision the hardware resources, supporting both scale-out and scale-up as well as consolidation deployment scenarios.



Figure 3. HPE Synergy Compute Modules

The HPE Synergy 480 Gen9 is a two socket compute module and the HPE Synergy 660 Gen9 is a four socket compute module, based on the Intel C610 Series chipset. The HPE Synergy 680 Compute Module is a four socket system based on the Intel C602J Series Chipset. The HPE Synergy 680 is a compute module providing higher memory and storage density with additional fabric attach capability. Each compute module has 2 (or 4, depending on model) internal drives configured as a RAID1 LUN hosting the Microsoft Windows Server® 2016 and Microsoft SQL Server 2016 binaries. For additional fluidity in the HPE Synergy environment, the OS binaries can also be run on the storage fabric, enabling a stateless compute environment, and providing administrators the freedom to move SQL Server workloads between compute modules as workload scaling requirements change.

Each compute module also has an HPE Smart Array P542D controller installed in mezzanine slot 1. This controller provides SAS connectivity to the internal DAS storage module in the Synergy Frame where the SQL data and transaction log files were installed.

HPE Synergy Interconnect Modules

As tested, the HPE Synergy 12000 Frame had one HPE VC SE 40Gb F8 Module for HPE Synergy, and two HPE Synergy 12Gb SAS Connection Modules. Different configurations of interconnects can be leveraged for specific application or user requirements. Pairs of interconnects are recommended for redundancy and high availability. Figure 4 shows an excerpt from the HPE Composer displaying the interconnect map for the F8 Module and Table 1 provides interconnect to mezzanine mapping highlighting the relationship between the Synergy Frame and compute modules.

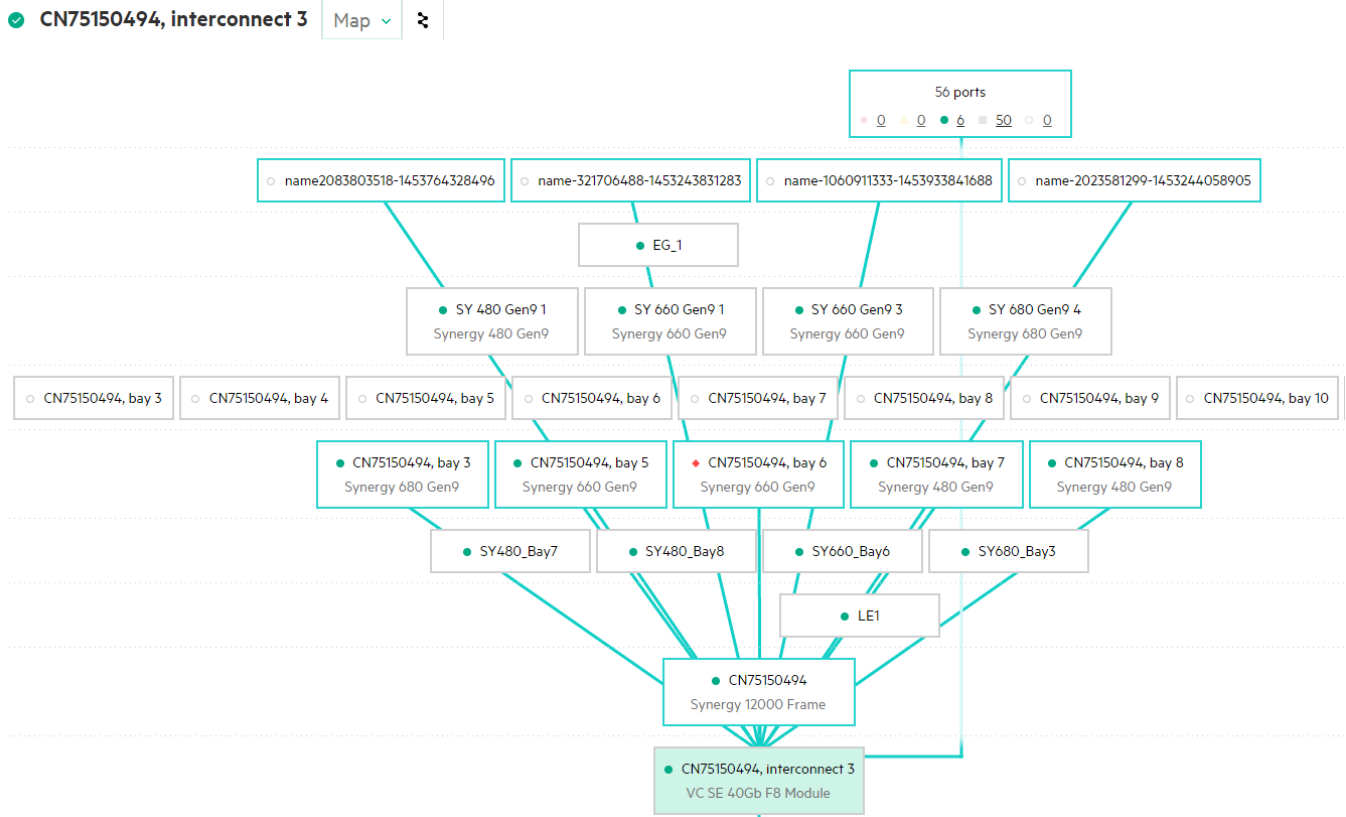


Figure 4. HPE Synergy interconnect map

Table 1. Interconnect module slot

BAY	INTERCONNECT	COMPUTE MODULE MEZZANINE SLOT	DESCRIPTION
1	CN75150494, interconnect 1	Mezz 1	Synergy 12Gb SAS Connection Module
2	empty		none
3	CN75150494, interconnect 3	Mezz 3	VC SE 40Gb F8 Module
4	CN75150494, interconnect 4	Mezz 1	Synergy 12Gb SAS Connection Module
5	empty		none
6	empty		none

HPE Synergy D3940 Storage Module

The HPE Synergy 12000 Frame used in this solution was configured with one HPE Synergy D3940 Storage Module. This provided the storage resource pool for the SQL Server database workloads. The HPE Synergy D3940 Storage Module contained 40 SFF drives, consisting of a mixture of eight, 450GB SSD drives and thirty-two, 450GB SAS 15K drives. These drives were provisioned to the compute modules as direct attached (DAS) storage.

As shown in Figure 5 below, the HPE Synergy D3940 storage module opens and pulls out in the front, providing access to the hot pluggable drives, in the event a drive needs to be replaced. As storage needs increase, additional HPE Synergy D3940 Storage Modules can be inserted into the HPE Synergy frame. A single HPE Synergy 12000 Frame can support up to five D3940 storage modules, providing a significant amount of capacity for storage requirements



Figure 5. HPE Synergy D3940 Storage Module

HPE Synergy Composer

Powered by HPE OneView, the redundant HPE Synergy Composer modules (Figure 6) manage the Synergy Frame and associated components. The HPE Synergy 12000 Frame, for redundancy, has two available slots for the HPE Synergy Composer modules.



Figure 6. HPE Synergy Composer

Simply stated, the HPE Synergy Composer is leveraged during initial deployment. Templates are utilized to provision specific compute, storage, and fabric information which can then be applied to one or more compute modules. An example profile for the HPE Synergy 660 compute module is shown in Figure 7. After the initial deployment, the HPE Synergy Composer is utilized to maintain and control infrastructure, firmware, and driver updates along with overall operational management of the HPE Synergy environment. For ease of use and automation, the HPE Composer has a unified API which can be leveraged by customizable scripts for provisioning, updates and other lifecycles operations.

The screenshot displays the HPE OneView interface for configuring a server profile. At the top, there is a search bar with the text 'status:OK'. Below it, the 'Server Profiles' section shows a list of profiles: SY480_Bay7, SY480_Bay8, SY660_Bay5, SY660_Bay6 (selected), and SY680_Bay3. The main area shows the configuration for 'SY660_Bay6' under the 'General' tab. The configuration includes:

- General:** Server hardware type (SY 660 Gen9 1), Enclosure group (EG 1), Affinity (Device bay), Server power (On), Serial number (VCGUWT8004), UUID (ba471c52-72be-4dfb-95fd-9fa55fb90cd2), and iSCSI initiator name (iqn.2015-02.com.hpe:oneview-vcguwt8004).
- Firmware:** Firmware baseline (managed manually).
- Connections:** A table showing network connections. One connection is listed with ID 1, Name VLAN210, Network V210 vlan210, Port Mezzanine 3:1-a, and Boot status Not bootable.
- Local Storage:** Integrated storage controller mode is RAID. A table shows two logical drives, both RAID 1, each with 2 drives, and no bootable status.

ID	Name	Network	Port	Boot
1	VLAN210	V210 vlan210	Mezzanine 3:1-a	Not bootable

Name	Logical Drive	RAID Level	Number of Drives	Min GB	Max GB	Drive Technology	Boot
Logical Drive 1	1	RAID 1	2	n/a	n/a	SAS HDD	No
Logical Drive 2	2	RAID 1	2	n/a	n/a	SAS HDD	No

Figure 7. HPE Synergy 660 Gen9 profile

Software

Microsoft is launching two new releases in 2016 that will provide new capabilities across both the OS and application. This white paper describes an environment using Microsoft SQL Server 2016 (CTP3.3) – 13.0.1000.281 (X64) Enterprise Evaluation Edition (64-bit) on Windows Server 2016 Technical Preview 4 6.3 (X64) Build 10586). The reason for using preview software is to highlight the intersection of technology innovation for both HPE and Microsoft in this solution. However, from an HPE Synergy perspective, this solution can support Microsoft Windows Server 2012 R2 and later, along with supported versions of SQL Server by Microsoft. In many environments there is a combination of Windows® and SQL versioning that exists for the database tier as applications are added and upgraded over different lifecycles.

Microsoft SQL Server 2016 and HPE Synergy

Database workloads require infrastructure that is capable of providing high performance, scalability and an ability to support a range of high availability requirements. The intent of this reference configuration is to highlight proof points and benefits of a composable infrastructure with HPE Synergy for running Microsoft SQL Server 2016 database workloads. With HPE Synergy, SQL Server can be provisioned on physical or virtualized servers, it can be scaled up to a 4-processor compute node with up to 6 terabytes of memory, and can leverage internal storage or external storage solutions for maximum performance. This enables organizations that don't need to size for maximum growth on day one, to save on licensing and operational costs as the infrastructure grows alongside the workload.

Microsoft SQL Server 2016 configuration guidance and overview

The tested component of this reference configuration focused on a three server configuration using AlwaysOn Availability Groups for HA and off-loading of read-only queries on the secondary replica. The intent of the white paper is not to provide overall performance or benchmark numbers, but to simulate a continuous transactional workload and provide insight on design points and configuration guidance for running SQL on HPE Synergy.

Initial profile configuration and deployment is driven by the core infrastructure management component, HPE Composer. Using HPE Composer, three server profiles were created and provisioned to the HPE Synergy 480, HPE Synergy 660 and HPE Synergy 680 Compute Modules, with each module running the latest technical preview of Windows Server 2016 (TP4) as the base OS. To optimize for the SQL DB workloads, the following BIOS settings can be configured in the HPE Synergy Composer compute profiles:

- Hyper-Threading—Enabled
- Intel Turbo Boost—Enabled
- HPE Power Profile—Maximum Performance

Additionally, the compute module's Smart Array disk controller write back cache was set to 100% write for logs and checkpoints due to the low intensity of reads given the memory available.

The database configuration was modeled around a multi-database scenario for a medium sized organization. One of the key use cases for HPE Synergy and composable infrastructure is SQL database consolidation and broader Database as a Service (DBaaS) platforms. This RC highlights a SQL instance consolidation platform where multiple OLTP DBs have been consolidated into a single SQL instance in an AlwaysOn highly available configuration. A combination of two mid-sized databases (900 GB) and 2 small-sized databases (500GB) were deployed to the primary database server. In addition, AlwaysOn Availability Group replicas were deployed on two additional servers, with one replica acting as a read/write secondary and one as read-only secondary for the small databases.

The following figure (Figure 8) shows the databases and their replication relationships among compute modules. The HPE Synergy 680 Compute Module (module 1) is the primary DB server with the HPE Synergy 660 Compute Module (module 2) serving as a secondary replica for all databases. Module 3 is the HPE Synergy 480, two socket system, serving as a read only replica for the two smaller databases.

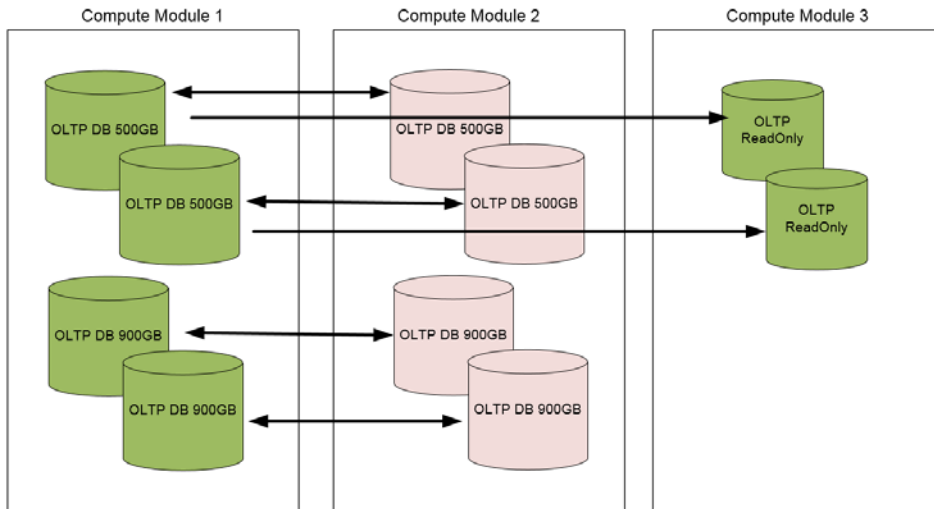


Figure 8. Logical representation of SQL databases and Compute Modules

The compute modules were joined to an external lab domain, although this is not required for AlwaysOn, as the Windows failover server can be set up in Windows Server 2016 using a workgroup cluster. The external domain controller is closer to what a SQL database environment might encounter in production environments, but hosting it within the Synergy Frame is certainly supported if needed.

The basic steps of creating availability groups, and preparing secondary databases and listeners is well documented and beyond the scope of this reference configuration. A link to deployment steps in Windows PowerShell is included in the “Resources and additional links” section at the end of this document.

To optimize the performance of the SQL Server 2016 database workloads, the following parameters were utilized:

- Trace flag -T834 Use large pages
- “Lock pages in memory” right assigned to user running SQL Server service
- MAXDOP = 1 for OLTP database workloads
- Max Memory set 5-10% under physical maximum to prevent OS memory pressure

NUMA affinity was used to align different workload databases with specific NUMA nodes. Each NUMA node corresponds to one socket. Table 2 details the primary and secondary compute modules’ database NUMA affinity and TCP Port values (four socket platforms). Table 3 provides the NUMA affinity configuration and TCP port values for the two socket node running the secondary read-only workload. For more information on the relationship between the NUMA node address bitmask and the TCP Port configuration needed in SQL Server please see [Appendix A: NUMA Affinity bitmasks and ports](#).

Table 2. Four socket NUMA affinity

DATABASE NAME	SIZE	NUMA NODE AFFINITY	TCP PORT
OLTP80K_01	800GB	0	1501[1]
OLTP80K_02	800GB	1	1502[2]
OLTP50K_01	500GB	2	1503[4]
OLTP50K_02	500GB	3	1504[8]

Table 3. Two socket NUMA affinity

DATABASE NAME	SIZE	NUMA NODE AFFINITY	TCP PORT
OLTP50K_01	500GB	0	1501[1]
OLTP50K_02	500GB	1	1502[2]

Capacity and sizing

This reference configuration is designed to provide an OLTP platform capable of hosting multiple OLTP databases. The design load point of this entry configuration is that of an entry size database consolidation platform, and can be deployed using the Synergy D3940 Storage Module instead of a costlier SAN storage system. This configuration is ideal for a company or group consolidating databases, for example:

- Few mid-sized databases
- Several small-sized databases
- A combination of the above (this solution)

Deployments at this small load point typically have relatively low IOPS requirements. For this configuration we used the following databases and associated drive count and RAID settings:

Table 4. Database LUN configuration

SERVER	FORMATTED CAPACITY (GB)	LUN PURPOSE	DRIVE COUNT 450 GB SAS	RAID TYPE
HPE Synergy 680	3,680	Data	10	RAID 5
HPE Synergy 680	419	Logs	2	RAID 1
HPE Synergy 660	3,680	Data	10	RAID 5
HPE Synergy 660	419	Logs	2	RAID 1
HPE Synergy 480	2,040	Data	6	RAID 5
HPE Synergy 480	419	Logs	2	RAID 1

In addition to these LUNs the system storage module has 8 450GB mixed use SSD drives which can be used as either cache for regular drives or can be set up as dedicated SSD LUNs for database deployment. Enablement of SSD Smart Cache is performed using the Smart Array configuration utility. The SSD drives were equally allocated to both four-socket servers as shown in Table 5.

Table 5. Database LUN configuration

SERVER	FORMATTED CAPACITY (GB)	PURPOSE	DRIVE COUNT
HPE Synergy 680	1,090	Cache	4
HPE Synergy 660	1,090	Cache	4

The range of IOPS achievable with one storage module directly depends on the type of drives present in the enclosure. A single storage module full of SSD media is capable of exceeding 2 Million IOPS. A single Synergy Frame can host up to five storage modules. Performance data in this reference configuration explores only a configuration in which a single storage module is utilized with only 8 SSD disks. Larger SQL deployments are possible.

With the above configuration we tested the four database systems using three acceleration methods:

- No acceleration
- Four SSD drives allocated as Smart Cache to rotational SAS drive based volume
- Single database deployed entirely on RAID5 SSD

The figure below shows the batch requests per second achievable for each load with a latency under 17ms.

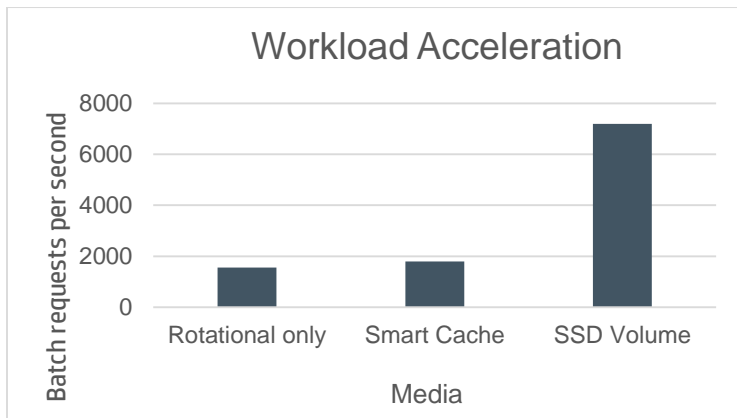


Figure 8. Workload performance using acceleration media

This shows that above a certain transaction throughput (and its base IOPS load), certain database loads benefit from solely SSD media over caching. Similarly, lower load databases benefit from partial acceleration.

Analysis and recommendations

HPE Synergy provides an ideal platform to support the range of deployment scenarios for SQL Server database workloads. From test and development to mission-critical scale-up environments, HPE Synergy provides a composable infrastructure with fluid resources to support the necessary SLA requirements for your database workloads. This white paper drilled into a specific workload use case to demonstrate the range of the HPE Synergy platform, highlighting multiple compute modules and leveraging the internal HPE Synergy D3940 Storage Modules. The following section provides additional detail around the benefits of the HPE Synergy platform for running DB workloads.

High availability

At its core, the HPE Synergy Frame has been built with redundancy and high availability by design and integrated through the frame, fabrics and HPE Composer management layer. The HPE Synergy Frame has two dedicated appliance management ports for either the HPE Composer or the HPE Image Streamer. Installed in pairs, the management components are completely redundant and can extend to manage the expansion of up to 20 frames.

OLTP Database workloads that have high availability requirements must be redundant at all levels of the system in order to achieve the desired uptime. For the deployment of SQL Server, the HPE Synergy platform can provide the degree of performance isolation, operational visibility and management for the highest levels of availability. This reference configuration describes a scenario using AlwaysOn Availability Groups to achieve an HA configuration for a physical server SQL DB deployment. To verify the HA design, testing was performed using a simulated workload and executed during a failover condition without issues. From a production perspective, additional degrees of isolation can be provided by extending to multiple D3940 Storage Modules to isolate DB replica copies and by extending the AlwaysOn Availability Group replicas to additional compute modules. This provides the granularity to configure the right level of the HA and DR capability for the right tier of database workloads.

While not in scope in this RC, another use case for SQL on HPE Synergy is to leverage a mix of virtualized and physical resources. A virtualized resource pool can be spun up from the composable compute, storage and fabric pools, providing a cluster for running SQL DB VMs. This supports a broader set of consolidation and DBaaS use cases and provides another layer of HA and workload mobility for the SQL databases.

Scaling

Scale-out performance increase can be achieved by directing read-only queries to a read-only replica. This allows the primary server to increase its core read/write workload and offload reports or backups to the replica without impacting primary server performance. This is a typical SQL Server use case when end-of-month reports and computationally intensive financial queries surge.

HPE Synergy provides a high bandwidth fabric that enables replication under load, and includes multiple 40 Gb/s external ports.

Scaling-up SQL Server with Synergy can be facilitated using HPE Composer profiles for compute modules to migrate database servers to larger compute modules. Storage can also scale up by increasing the number of storage modules and attaching to an external HPE 3PAR SAN.

In the same respect, but opposite direction, scaling-down is possible by releasing unused resources to their pools.

Microsoft SQL Server can capitalize on the composable resource provisioning for either scale or customized adjustments for additional compute modules, storage modules, or fabric resources.

Performance

Tests performed with the reference configuration hardware show the SQL databases able to work within expected latencies and utilizing the storage and compute modules efficiently. IOPS need to be considered in populating the storage modules since there are a number of media options that need to be considered depending on the workload rate (IOPS) and read/write characteristics.

For optimum performance of Microsoft SQL Server OLTP databases we recommend different combinations of media depending on workload intensity.

In low-to-mid transaction environments, a mix of rotational SAS and write-intensive SSD drives is recommended.

- In this configuration SSD drives can be utilized as dedicated volumes, or used as a larger cache for rotational disks (Smart Cache).
- Non-SSD cached rotational volumes can have their Smart Array controller cache tuned to favor read or write workloads.

In high transaction environments, a configuration using solely write-intensive SSD drives is recommended.

- In this configuration SSD drives are utilized as dedicated volumes

For example:

- In a low transaction scenario, use rotational SAS drives with Smart Array cache set to 100% writes for logs and use SSD Smart Cache on rotational SAS drive volumes for Data.
- In a mid-transaction scenario, use SSD Smart Cache on both Log and Data rotational SAS drive volumes.
- In a high transaction scenario, use write-intensive SAS SSD drives as volumes for Log and Data.

Operations

Setting up the different instances of SQL Server was straight forward. We used HPE Composer profiles for compute modules and allocate network and storage resources for each server. The storage modules provided direct attach storage to the different replicas.

Summary

The HPE Synergy platform provides a highly available and efficient solution capable of handling diverse SQL Server OLTP workloads. The composable infrastructure makes it easy to re-deploy SQL Servers to different compute modules when more or less processing is needed. The configuration tested in this document shows the flexibility of deployment options and scaling capabilities in terms of compute and storage.

Storage, Network, and Compute Module capabilities in Synergy provide Microsoft SQL Server 2016 (and earlier versions) databases with flexible and high performance storage, high bandwidth intra module connectivity, and advanced compute capabilities such as NUMA. This results in a well-balanced, scalable and efficient operational environment for SQL Server databases.

Appendix A: NUMA Affinity bitmasks and ports

NUMA Affinity is configured using a bit mask to reference the desired NUMA node(s):

Table 6. NUMA Affinity bitmask values

NUMA NODE (SOCKET)	BITMASK	DECIMAL VALUE
0	0001	1
1	0010	2
2	0100	4
3	1000	8

SQL Server Configuration Manager TCP/IP settings for server protocol are shown in Figure 9. Database connections established via the specified ports get affinity to the NUMA node(s) identified by the bitmask inside the square brackets, as shown in Figure 9

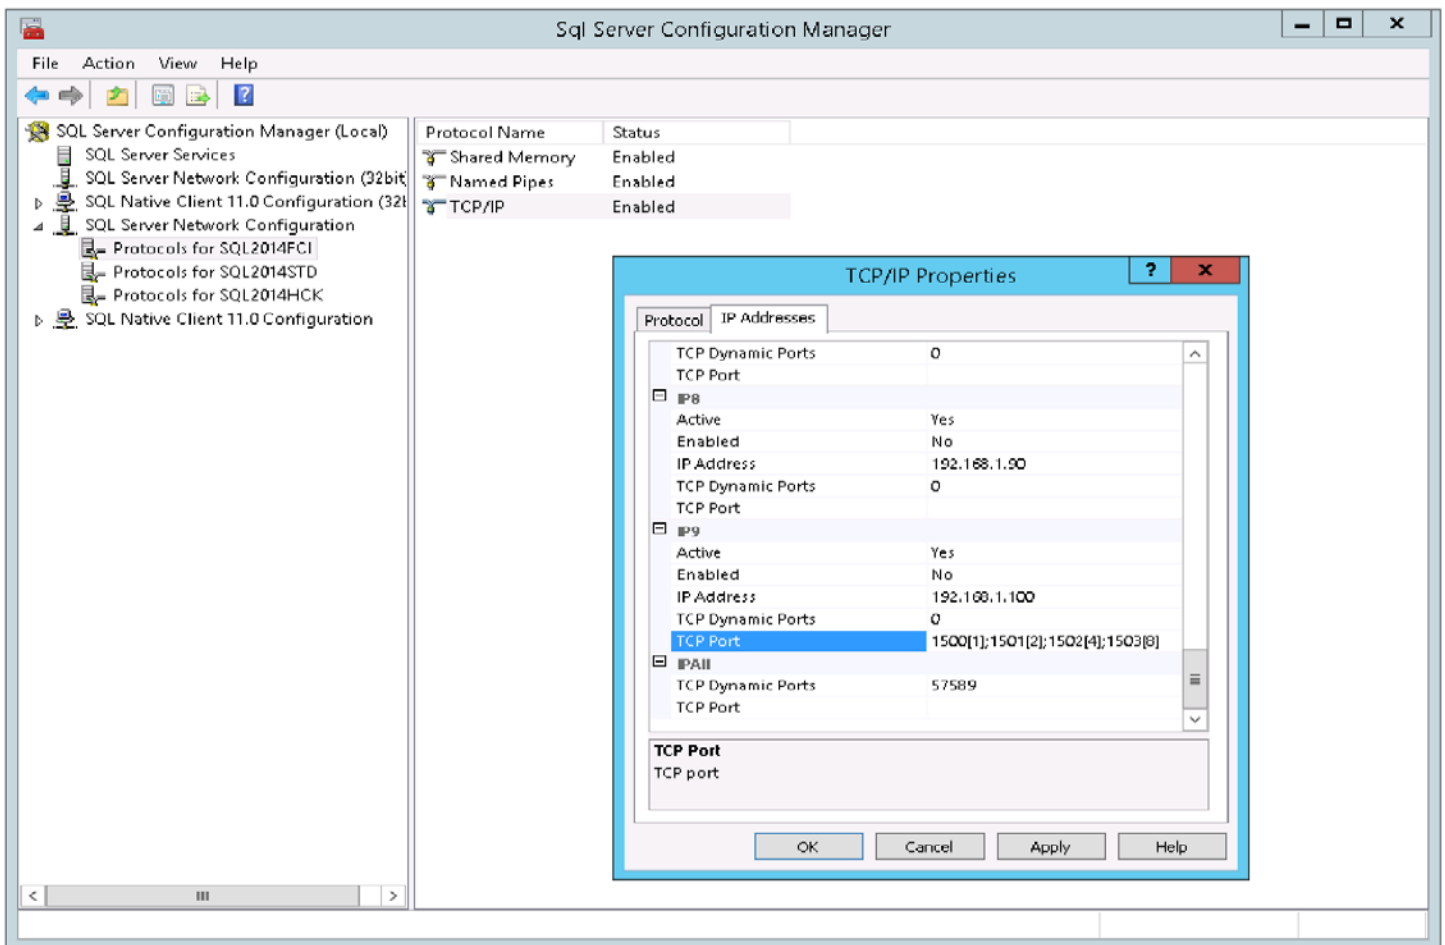


Figure 9. SQL Server TCP/IP port settings for NUMA Affinity

Resources and additional links

HPE Synergy

hpe.com/info/synergy

HPE Composable Infrastructure

hpe.com/info/composableinfrastructure

HPE Servers

hpe.com/servers

HPE Storage

hpe.com/storage

HPE Networking

hpe.com/networking

HPE Technology Consulting Services

hpe.com/us/en/services/consulting.html

Creating Microsoft SQL Server 2016 Availability Groups using PowerShell

<https://msdn.microsoft.com/en-us/library/gg492181.aspx>

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



Sign up for updates

★ Rate this document



© Copyright 2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for HPE products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows Server, and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

4AA5-8377ENW, March 2016