



REFERENCE ARCHITECTURE

Microsoft SQL Server 2016 Data Warehouse Fast Track

FlashStack™ 70TB Solution with Cisco UCS and Pure Storage FlashArray//X

FLASHSTACK REFERENCE ARCHITECTURE

September 2018





TABLE OF CONTENTS

MODERN DATA WAREHOUSE 3

RESOURCE UTILIZATION 3

WORKLOAD OPTIMIZATION 4

TEST CONFIGURATION 5

RESULTS.....7

SUMMARY 8

DWFT REFERENCE ARCHITECTURE 9





MODERN DATA WAREHOUSE

The data warehouse continues to grow. As data increases high bandwidth is essential to provide query results to users.

Microsoft SQL Server 2016 introduces several improvements to accommodate the challenges of storing and processing these large amounts of data.

SQL Server 2016 provides up to 10X compression of data utilizing columnstore index technologies. This columnstore technology provides the capability to scan only the columns needed while reducing I/O requirements and the memory required for a given number of rows from the source data warehouse. Efficient single-row lookup is improved by using additional (B tree) indices to columnstore-based tables.

Query optimization is also improved by operator pushdown which moves both filter and aggregation query operations closer to the data, reducing the volume of data which needs to be handled further on in query processing.

SQL Server 2016 provides Batch Mode Processing which processes many rows at a time rather than serially doing calculations on each individual row. These batch operations are further optimized by leveraging Single Instruction Multiple Data (SIMD) vector processing CPU instructions in the Intel® architectures.

RESOURCE UTILIZATION

The infrastructure that runs your SQL Server 2016 data warehouse needs to be balanced with regard to performance. Installing the fastest processors without considering the configuration as a whole, storage, memory, switch, etc. can lead to disappointing long-term results, as you need to scale. This testing on your own can be costly and time consuming.

The SQL Server 2016 data warehouse Fast Track (DWFT) program is a reference architecture designed to take the guessing out of building your data warehouse infrastructure. These reference architectures are already tested using bandwidth demanding workloads to meet specific query performance and to scale in size requirements designated by the Microsoft SQL Server performance team.


FlashStack™, a next-generation converged infrastructure solution from Cisco® and Pure Storage®, is designed with the philosophy to deliver meaningful performance with a focus on optimization of resources, not just a single benchmark score or metric. Nobody wants their expensive processors waiting around for work. Using the Cisco UCS® C460 M4 Rack Server for compute and the Pure Storage FlashArray//X product with 100 percent solid-state flash storage helps eliminate storage as the bottleneck and achieve a high utilization of all resources.

Together, Cisco and Pure Storage are delivering next generation converged infrastructure that enables flexible growth, streamlined operations, choice and transformation. FlashStack solutions validated by Cisco scale non-disruptively with business requirements, reduce risk, and lower costs, and transforms end-user experiences with simple on-demand IT. FlashStack solutions are available across a portfolio of mission critical data center workloads including Microsoft SQL Server, Microsoft Exchange, Microsoft Private Cloud with Hyper-V and more.





Resource utilization database environment



SOLUTION SUMMARY

Problem

Architectures based on a “build-your-own” method are time consuming and risky, often resulting in poor resource utilization, system imbalance, and scaling problems.

Solution

Cisco and Pure Storage FlashStack
Microsoft SQL Server Data Warehouse Fast Track reference architecture. Tested with SQL Server 2016 for row-store and columnstore workloads.

Results

A balanced system with efficient query processing and predictable performance.
70-TB-rated user capacity.
Row-store I/O throughput 7570 MB/s. Columnstore throughput 2248 queries/Hr/TB



WORKLOAD OPTIMIZATION



Pure Storage FlashArray//X

FlashArray//X is next-generation all-flash storage. Designed for the cloud era, FlashArray//X delivers the highest performance and enterprise reliability for Tier 1 applications. Microsecond latency and GBs of bandwidth – with all the rich data services, proven 99.9999 percent availability over two years (inclusive of maintenance and generational upgrades), and unrivaled Pure1® support that the FlashArray family is built on. With the FlashArray//X, organizations can dramatically reduce the complexity of storage to make IT more agile and efficient.





Cisco UCS C460 M4

The Cisco UCS C460 M4 Rack Server used in this reference architecture is a four-rack-unit (4RU) rack server supporting the Intel® Xeon® E7-4800/8800 v2, v3, and v4 processor families.

The Cisco UCS C460 M4 Rack Server offers industry leading performance and advanced reliability for the most demanding enterprise mission-critical workloads, large-scale virtualization, and database applications.



Cisco MDS 9148S

The Cisco MDS 9148S 16G Multilayer Fabric Switch is the next generation of the highly reliable Cisco MDS 9100 Series Switches. It includes up to 48 auto-sensing line-rate 16-Gbps Fibre Channel ports.

In all, the Cisco MDS 9148S is a powerful and flexible switch that delivers high performance and comprehensive Enterprise-class features.

TEST CONFIGURATION

COMPUTE	DESCRIPTION
Vendor	Cisco
Model	Cisco UCS C460 M4
Processor	Intel Xeon processors E7-8890 v4 24 cores Qty.(2)
DRAM	3 TB* tested with max 118 GB for scaling estimate purposes
QLogic QLE2672	QLogic QLE2672, 16Gb FC HBA, Qty. (4)
BIOS power	High-throughput-enabled



TEST CONFIGURATION

STORAGE	INFORMATION/PARAMETER
Model	FlashArray//X70 Qty.(4) 16Gbps ports
Operating environment	Purity 4.10.4
Total array capacity	47 TB
STORAGE	INFORMATION/PARAMETER
OS boot volume	1.04 TB (operating system)
Mount point anchor	20GB (mount point for volumes)
System	50GB (volume for system database files.)
Log	2TB (volume for database log files.)
Data and tempdb	16TB (volume for database data and tempdb files.)
Backup	12TB (volume for database backup files.)
NETWORK	INFORMATION/PARAMETER
Fabric switch	Cisco MDS 9148S Qty. (2)
WORKLOAD	INFORMATION/PARAMETER
OLAP type workload	Read/write ratio: 80/20






TEST CONFIGURATION

APPLICATION	INFORMATION/PARAMETER
Microsoft SQL Server Enterprise	2016 13.0.4446.0 (X64)
Trace flags	-E, -T1117, -T834
Minimum/maximum memory	118(MB)/118(GB)
Degree of parallelism	Row-store MA XDOP = 32 Columnstore MA XDOP = 96
OPERATING SYSTEM	INFORMATION/PARAMETER
Windows Server 2016 with updates	Build 10.0.14393
Power management	High-performance-enabled
Local security policy	Lock pages in memory: enabled
Perform volume maintenance tasks	Instant file initialization: enabled




Balanced predictable results

Rated user data capacity
70 TB





Row-store I/O throughput
7570 MB/s

Avg. CPU utilization



Columnstore throughput
2248 queries/Hr/TB

Avg. CPU utilization



RESULTS

Data warehouse workloads can range from loading large amounts of data to complex analytical processing of data for consumption. The I/O generated typically consists of concurrent streams of read operations. The SQL Server 2016 DWFT reference architecture with FlashStack, including Cisco solutions for compute and networking and Pure Storage solutions for fast solid-state flash storage, establishes an architecture ready to accommodate the most demanding and complex read-queries and scale requirements.

SUMMARY

Together, Cisco and Pure Storage are delivering next-generation converged infrastructure that delivers flexible growth, streamlined operations, choice, and transformation. A FlashStack SQL Server Data Warehouse Fast Track solution provides a reference architecture that meets the performance and scaling needs of today's modern data warehouse, as well as decreasing the time and implementation complexity for customers to create a balanced, reliable long-term solution for their data warehouse needs.



DWFT REFERENCE ARCHITECTURE

DWFT Certification: #2016 - 027	Pure Storage FlashArray//X70 and Cisco UCS C460 M4 DWFT Reference Architecture		Report Date: 8/31/2017
DWFT Rev. 5.4			
SYSTEM PROVIDER	SYSTEM NAME	PROCESSOR TYPE	MEMORY
	CISCO C460 M4	Intel Xeon E7-8890 v4 2.2 GHz (2/48/96)	1536 GB
OPERATING SYSTEM		SQL SERVER EDITION	
Windows Server 2016		SQL Server 2016 Enterprise Edition	
SYSTEM PROVIDER	STORAGE INFORMATION		
	FlashArray//X70 R2, 68.67 TB aggregated thin-provisioned RAID 3D for data, tempdb, and log 1 TB UCSC-MRAID C460 SCSI for OS boot		

PRIMARY METRICS			
Rated user data capacity ¹ (Tb)	Row store relative throughput ²	Columnstore relative throughput ³	Maximum user data capacity ¹ (Tb)
70	252	346	160

ROW STORE					
Relative throughput ²	Measured throughput (queries/Hr/TB)	Measured scan rate: physical (MB/sec)	Measured scan rate: logical (MB/sec)	Measured I/O throughput (MB/sec)	Measured CPU (avg.) (%)
252	330	6,635	8,506	7,570	83

COLUMNSTORE					
Relative Throughput ²	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)
346	2,248	3,328	N/A	N/A	95

The reference configuration is a 2-socket system rated for 25TB using SQL Server 2014 and the DWFT V4 methodology¹

¹ Assumes a data compression ratio of 5:1

² Percent ratio of the throughput to the row-store throughput of the reference configuration.

³ Percent ratio of the throughput to the columnstore throughput of the reference configuration.

*Reported metrics are based on the qualification that specifies the database size of the SQL Server memory.





© 2018 Pure Storage, Inc. and Cisco Systems, Inc. Pure Storage, the "P" Logo, Pure1 and FlashStack are trademarks or registered trademarks of Pure Storage, Inc., in the U.S. and other countries. Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries. All other trademarks are the property of their respective owners.

The Pure Storage product described in this documentation is distributed under a license agreement and may be used only in accordance with the terms of the agreement. The license agreement restricts its use, copying, distribution, decompilation, and reverse engineering. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Pure Storage, Inc. 650 Castro Street, Mountain View, CA 94041

PS-FS-RA-MSQL70TB-0618-0063v7