

Microsoft SQL Server 2016 Fast Track

55 TB Certified Data Warehouse

1.2 Petabyte Maximum SQL User Data

CISCO & Tegile Certified Systems Solution Review

CISCO B200 M4 Blade

CISCO UCS 5108 Chassis

Tegile HD T10KHD-300 All-Flash Storage Array

Author: Peter Kerr, Senior Director of Strategic Alliances



Table of Contents

Table of Contents.....2

Executive Summary:3

Introduction to Fast Track for SQL Server Data Warehouse3

Why Cisco UCS:.....3

Why Tegile:.....4

How are Cisco UCS and Tegile-based SQL solutions different?4

CISCO Blades and Features:6

Specifications at a Glance6

Tegile Intelligent Flash Array Features7

High Level Tegile Array Architecture:9

Data Warehouse Fast Track Certification: 10

Cisco B200 M4 Blade & Tegile T10KHD-300 Performance Overview 11

Solution Overview..... 12

Summary..... 13

Executive Summary:

The Microsoft SQL Server Data Warehouse Fast Track (DWFT) reference architecture is designed to eliminate the complexity of properly sizing hardware for your SQL Server-based Data Warehouse. This will streamline the decision-making process and reduce hardware and maintenance costs. The sizing technique used in SQL Server DWFT will properly size storage *and* servers, based on I/O throughput and CPU transaction capability. This consumption-based approach ensures that you can take full advantage of your hardware and software investment.

The Microsoft SQL Server Fast Track Data Warehouse Reference Architecture presented in this document is designed, tested, and validated by Microsoft, Cisco, and Tegile.

Target audience: This document is intended for CIOs, CTOs, decision makers, and managers, as well as individual (BI) Architects, DBAs, and developers. Tegile can also provide you step-by-step installation and configuration help. Request configuration support at www.tegile.com.

Introduction to Fast Track for SQL Server Data Warehouse

Data Warehouse Fast Track (DWFT) reference architectures provide tested and validated configurations and resources to help customers build the right environment for their data warehouse solutions. Following these best practices and guidelines will yield these tangible benefits:

- Accelerate data warehouse projects with pre-tested hardware and SQL Server configurations
- Reduce hardware and maintenance costs by purchasing a balanced hardware solution and optimizing it for a data warehouse workload
- Improve ROI by optimizing software assets
- Reduce planning and setup costs by leveraging certified reference architecture configurations
- Ensure predictable performance by properly configuring and tuning the system

The DWFT reference architectures provide a balanced server, memory, network, and storage hardware configuration, helping you avoid the risk of improperly designed and configured hardware systems. These guidelines will maximize the capability and throughput out of your storage and server investment.

Why Cisco UCS:

- (1) The Cisco's Unified Computing System (UCS) is the world's leading solution for blade-based systems
- (2) By incorporating unified fabric and fabric-centered technology, the Cisco Unified Computing System requires fewer physical components, less management, and less energy

-
- (3) By using a unified fabric, Cisco UCS can leverage multiple network protocols quickly and simultaneously

Why Tegile:

- (1) Tegile arrays are the leading storage solution for SQL Server deployments delivering raw capacities from 24TB to 480TB. Ideal for large analytics solutions while delivering aggressive cost per GB.
- (2) Tegile arrays deliver low-latency, SQL Server CPU offload features, and industry leading SQL management features
- (3) Tegile is the only storage company that lets you expand an all-flash array with additional all-flash, hybrid, or hard disk drive (HDD) storage. This allows customers to add more performance, more space, or both. The decision to expand can be deferred until you're sure what your SQL Server deployment needs to grow
- (4) Tegile delivers advanced data management, deduplication, and compression features, which increase your effective capacity and boost performance by minimizing your data footprint. Furthermore, all of Tegile's software features are included in the base price, at no additional cost. Everything comes in the box.
- (5) Tegile is the only storage company that offers a **FULL 5-year** performance and capacity warranty on all of its flash arrays. This includes the actual flash media (eMLC).

How are Cisco UCS and Tegile-based SQL solutions different?

Generally speaking, there are three types of data warehouse architectures and SQL deployments:

The traditional solution calls for separate servers and a huge allocation of spinning media on a traditional HDD-based storage array. This approach comes with all of the drawbacks associated with a 50-year-old architecture.

The converged solution calls for stuffing PCIe flash cards into a single server and creating a fast, but highly volatile solution that lacks enterprise features like failover, scale out, non-disruptive updates, etc. This approach combines computing and storage into one unit.

The modern solution is designed using enterprise deployment standards, with one or more servers, along with independent modern storage. Benefits include lightning-fast response times, high availability, and expandable all-flash and hybrid storage.

The traditional solution:

Legacy data warehouse reference architectures are well established in the industry. By now, we've figured out how to extract as much performance as possible from disk-based arrays. Unfortunately, the limitations of spinning media technology prevent these architectures from advancing much further. They also present other problems. For one, the high latency and lack of advanced features

create more work for DBAs and storage teams. To keep data warehouses operating well on spinning media and achieve a true sequential layout of data, DBAs must painstakingly load data sequentially into a data warehouse. Steps include single-threaded loading, the use of multiple staging tables, and the ordering of data between steps.

Traditional solutions trade overall loading speed for a physically sequential data layout in order to minimize fragmentation. This ETL process is tailored to the limitations of legacy spinning media. The intent was to try to minimize the movement of the read/write heads on the disks. Tegile IntelliFlash® storage arrays remove this requirement and allow the administrators to take the quickest path, with parallel loads and the ability to reload data at will.

The converged solution:

Converged solutions are designed to eliminate performance bottlenecks by placing storage as close to the CPU processing as possible—in this case, directly inside the server that's running the application. Some companies sell so-called "hyper-converged" systems that take on the role of compute, storage, and network at the same time. While PCIe flash and hyper-converged systems can seem somewhat speedy, they turn out to be very expensive over time and have many serious drawbacks.

First: In order to avoid catastrophic failure of the solution, the PCIe flash cards need to be doubled and mirrored inside the server, doubling the cost of the solution. Furthermore, PCIe flash failures are quite common due to poor ventilation of the server cases. They produce a tremendous amount of heat under pressure, which is the exact moment you need them the most.

Second: Using a single server creates a single point of failure. To overcome this, SQL Enterprise must be deployed in an always-on configuration. This will require another (ideally) identical server with approximately the same amount of flash (PCIe) storage crammed into it. Once again, this effectively doubles the cost of the deployment, both in licensing and storage. By doubling the storage in the server, and then doubling the storage in the fail over server, you end up purchasing 4 times the required storage. Furthermore, this flash can only be used by one solution and cannot be shared across the organization when it's no longer needed.

Third: There is no easy way to add storage performance and capacity. One must add a new server to increase the performance and/or capacity. While some DBAs will say sharding DBs is an option, it is complex and difficult to maintain for a reasonable period.

Fourth: "Hyper-converged" architectures use the same CPU, which must be licensed for SQL to perform complex and demanding storage and networking tasks. The amount of CPU time available to the SQL process can be very minimal, around 20%. Yet every hyper-converged box will have to be fully licensed (for each core), as if each core was dedicated 100% to the SQL Process. This is a very inefficient way to deploy SQL Server licenses.

The modern solution:

The modern solution is based on the proven strategy of separating highly available, fault-tolerant storage and computing resources. Cisco UCS, with its powerful blades and unified fabric, provides a

dense, highly elastic and adaptable environment. Tegile IntelliFlash™ arrays bring low latency flash storage to the UCS platform to create a properly balanced enterprise architecture. Tegile storage arrays are highly available, with two active-active controllers. They also protect all data with industry-standard RAID configurations. Tegile arrays deliver all-flash performance that can be expanded over time with all-flash, hybrid, or HDD-based shelves, depending on your needs. The Tegile IntelliFlash™ operating system (for hybrids) ensures that frequently accessed data is always stored in the fastest media, and less frequently used data is stored on more cost-effective media. Furthermore, the Tegile IntelliFlash operating system dynamically allocates the performance of the array/storage pool based on the requirements placed on each LUN or share.

CISCO Blades and Features:

The UCS B200 M4 delivers enterprise-class performance in a compact form factor. Optimized for data centers or the cloud, the Cisco UCS B200 M4 can quickly deploy stateless physical and virtual workloads, with the programmability of UCS Manager and the simplified server access of SingleConnect technology.

The UCS B200 M4 is built with the Intel® Xeon® E7-4370 v3 processor family, with up to 1.4 TB of memory (with 64 GB DIMMs), up to two drives, and up to 80 Gbps total bandwidth. It offers exceptional levels of performance, flexibility, and I/O throughput to run the most demanding applications.

In addition, Cisco UCS doesn't need to power and cool switches in each blade chassis. This reduces costs and provides an architectural advantage. With more power available for blades, the solution can use more capable blade servers with uncompromised expandability.

The UCS B200 M4 Blade Server delivers:

- Suitability for a wide range of applications and workload requirements
- Highest-performing CPU and memory options without constraints in configuration, power, or cooling
- Half-width form factor, offering industry-leading benefits
- Latest features of Cisco UCS Virtual Interface Cards (VICs)

Specifications at a Glance



CISCO UCS 5108 Chassis with 8x B200 M4 Servers

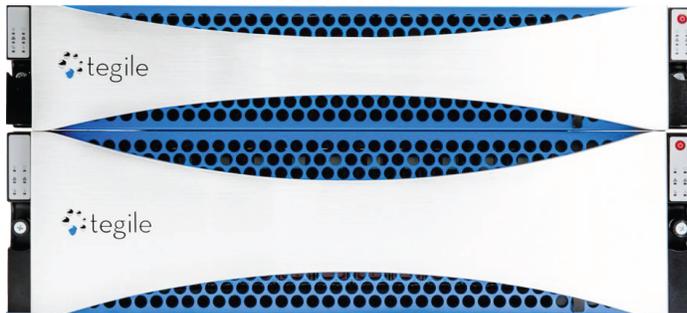


Cisco UCS 6248UP 48-Port Fabric Interconnect

Specifications at a Glance

Item	Specification
Form Factor	Half-width blade form factor
Processors	Either 1 or 2 Intel® Xeon® processor
Chipset	Intel C610 series
Memory	Up to 24 double-data-rate 4 (DDR4) dual in-line memory (DIMMs) at 2133 MHz speeds
Mezzanine adapter slots	2
Hard drives	Two optional, hot-pluggable, SAS, SATA hard disk drives (HDDs) or solid-state drives (SSDs)
Maximum internal storage	Up to 3.2 TB
Maximum number of blade servers	Eight (8) B200 M4 Blade Servers per blade chassis

Tegile Intelligent Flash Array Features



Tegile T10KHD-300 All-Flash Array (5U)



High IOPS and Low Latency: Tegile storage arrays are able to achieve very high IOPS by writing each I/O block into DRAM and flash on the array. Using advanced algorithms, the most important data is always kept at the fastest layer of the array. Tegile storage arrays can come equipped with several hundred terabytes of enterprise-grade (eMLC) solid-state drives. The actual usable space is even larger due to inline compression and de-duplication. These data reduction techniques free up cache space in DRAM and flash for faster reads and writes.



Compression: Tegile storage arrays use inline compression, which can remove multiple write operations and simplify read operations in real time. Compression also delivers space savings and eliminates multiple IOs, effectively increasing the throughput of the system. Compression is performed inline without compromising read or write performance. This allows the Tegile array to perform compression on the array side, instead of in the SQL Server Host, which means the CPU can focus on real SQL work.



Inline Deduplication: Tegile storage arrays come with low overhead inline deduplication. For most SQL Server workloads, this feature is only effective if the data is highly repetitive. However, this feature offers jaw-dropping results for virtualized servers, VDI, Exchange, and SQL deployments with multiple copies of the same databases.



@Rest Full Data Encryption: Tegile arrays exclusively use self-encrypting SSDs and hard disk drives (for hybrid models). This means there is no latency degradation and no CPU overhead on the SQL Host or the array controllers. Each drive has its own encrypt/decrypt hardware engine. So 24 drives will run 24 encryption engines, 100 drives will have 100, delivering a linearly scalable encryption solution. Tegile eliminates all encryption complexity by making encryption workloads and functions completely transparent to users, applications, and operating systems. Tegile encryption also meets the toughest healthcare and DoD standards. As such, encryption can be turned off on the SQL Server host, freeing up 20% of CPU cycles for other SQL tasks.



Snapshots and Clones: Tegile storage arrays can be used to present multiple copies of the same data from a **single source**. When copies are updated, only deltas will be saved. Using this feature can result in significant savings because development, testing, production, DW, and BI environments all use the same copy of the data. Snapshots (database/app-consistent) are created instantly, regardless of the size of the database. This makes them a great way to create instant replicas of the database and establish instant recovery points across complex application sets.

Note: While many other vendors have snapshots, they all face three core problems: (a) snapshot and clones are not zero-byte, (b) arrays are unable to deliver the required IOPS to serve all of the copies, or (c) the snapshots themselves are not database/application-consistent. Furthermore, many vendors are missing Microsoft VSS or VMWare VMAPI integration. Each of these problems render the feature effectively useless. In many such cases, the storage team will tell the data team to “simply” copy the data someplace else, causing very poor utilization of expensive storage resources and a significant amount of work for DBAs.



Replication: Tegile storage arrays have native asynchronous replication for both in-site and remote-site scenarios. The scheduling granularity is completely controllable by the customer and has zero impact on latency at the primary site. Replication can be used either alone, or in concert with SQL Server AlwaysOn to provide fast in-site and remote-site replication and high availability when business continuity is critical.

Replication can move complex applications, in addition to complex SQL Server Databases. Data is replicated after deduplication and compression, and it's based on snapshots. Because of this, only very small amounts of data are put onto the network, saving expensive bandwidth between remote sites.

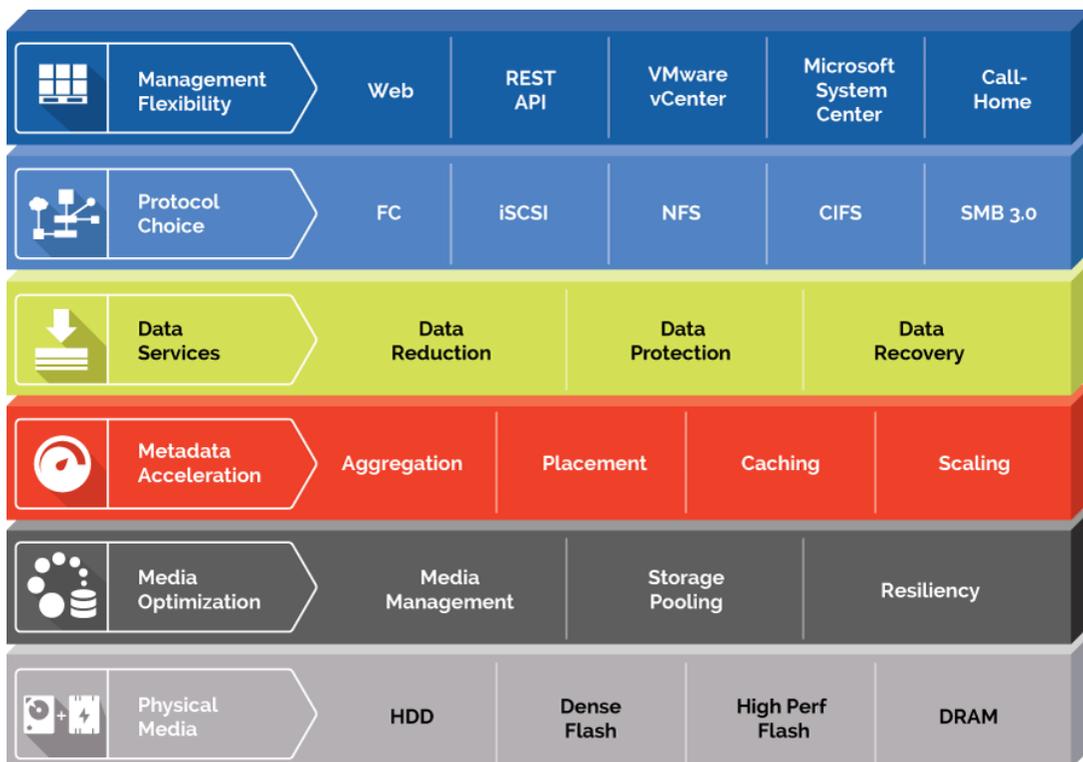


Cloud Analytics: Tegile IntelliCare™ is an industry-leading cloud analytics platform. IntelliCare delivers reporting, advanced troubleshooting, and proactive support from the Tegile support team.



Management: Tegile arrays are fully manageable using Microsoft System Center (SCVMM), VMware vCenter, or its native Tegile Web UI.

High Level Tegile Array Architecture:



IntelliFlash Architecture

Data Warehouse Fast Track Certification:

DWFT Certification #2016-008	CISCO UCS + Tegile T10KHD-300 DWFT Reference Architecture		Report Date: 8/15/2016		
DWFT Rev. 5.4					
System Provider	System Name	Processor Type	Memory		
	CISCO UCS + Tegile T10KHD-300	Intel Xeon E7-4370 2.1 GHz (2/24/48)	768 GB		
Operating System		SQL Server Edition			
Windows Server 2012 R2		SQL Server 2016 Enterprise Edition			
Storage Provider	Storage Information				
	24 x 2TB Metadata and Read and Write Cache 60 x 8TB Modules in (RAID 60)				
Primary Metrics					
Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹ (TB)		
55	235	228	1168		
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.) (%)
235	113	6,035	7,785	6,910	61
Column 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.) (%)
228	1,485	1,496	N/A	N/A	97
The reference configuration is a 2 socket system rated for 25TB using SQL Server 2014 and the DWFT V4 methodology					
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 column store throughput of the reference configuration.					
Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.					

Cisco B200 M4 Blade & Tegile T10KHD-300 Performance Overview

The testing and certification process was completed by running a Microsoft-designed and controlled benchmarking system that includes a full set of tests that simulate real world situations using terabytes of TPC-H-like data.

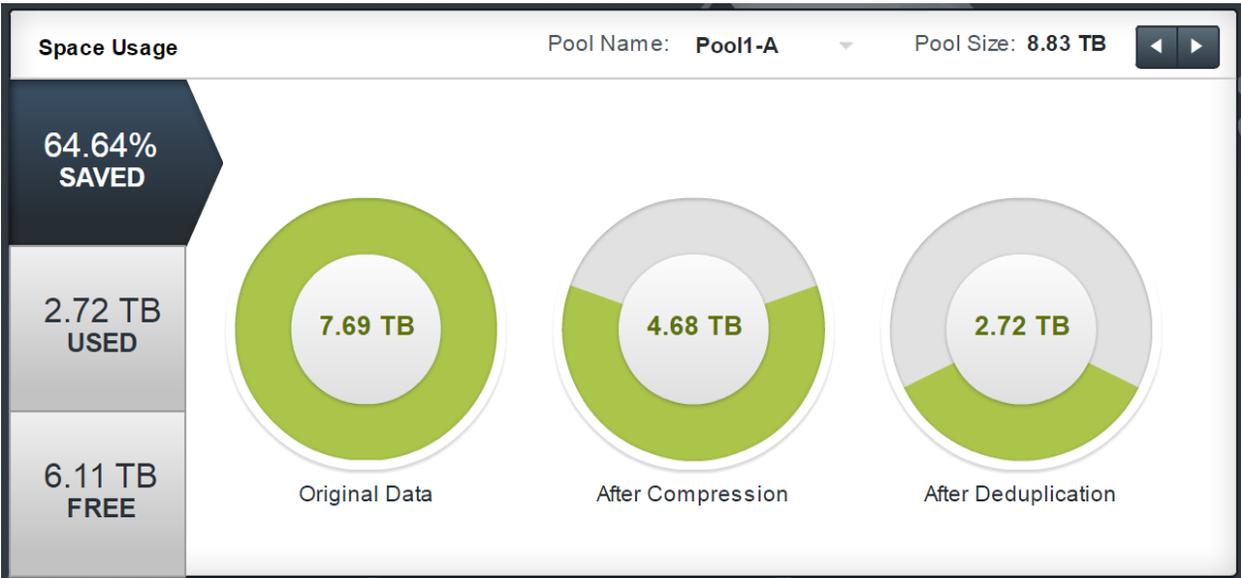
Cisco purposefully chose a B200, with total of only 24 cores and 128 GB of RAM, to demonstrate the power and efficiency of Tegile all-flash storage.

The certificate above essentially certifies the Tegile T10KHD-300 single (5U) unit which can support active data sets of 55TB and maximum user size of **1.2 Petabytes of SQL Data** in an all flash system.

It is important to point out that Microsoft considers *the total amount of data loaded* into the system as the metric used for sizing. Meaning a SQL Server DW running on top of a Tegile T10KHD-300 can fit up to 1.2PB of data.

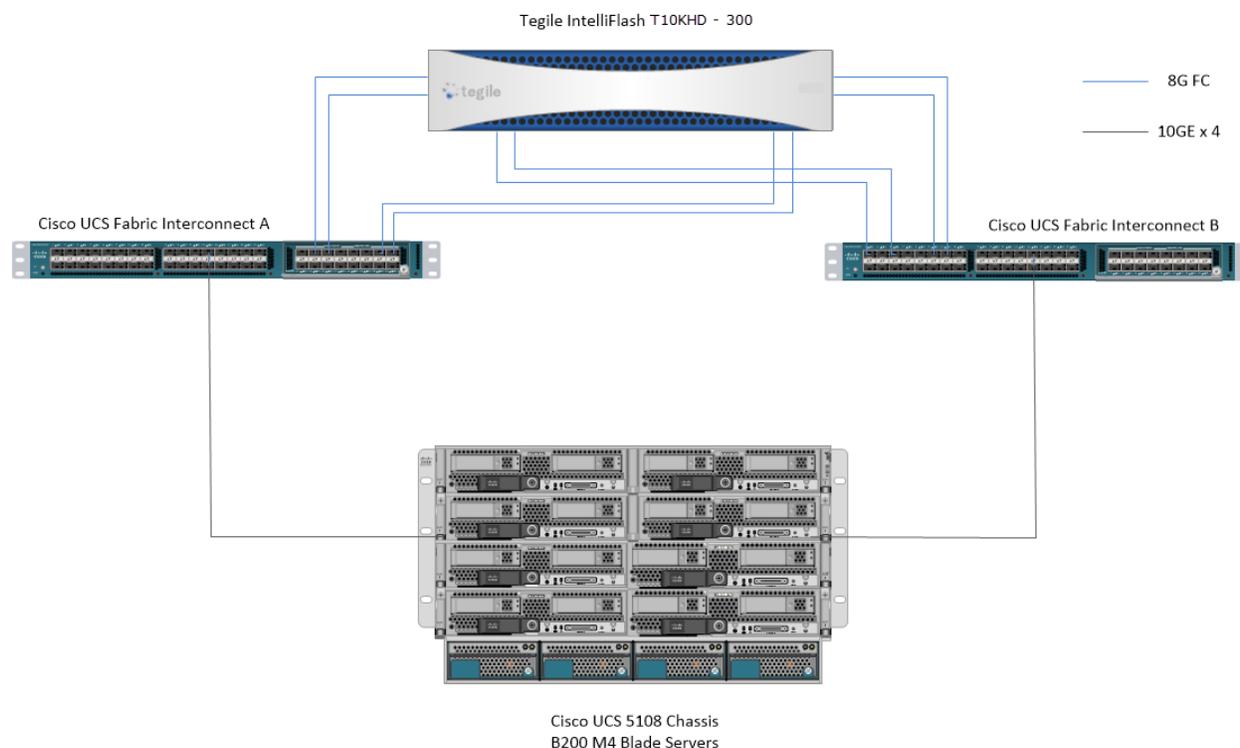
As these 1.2PB are loaded into SQL server, SQL Server applies data compression, reducing the amount of data the first time and resulting in at least 1.2PB of usable space. In addition, Tegile applies hardware level compression and de-duplication as the data is saved by SQL server onto the storage array resulting in additional savings. However, Tegile compression is not included in space the calculations and serves as an effective reserve for any configuration.

A real world example from the Tegile Array management UI is included below. In the example, we see a virtualized production SQL Server deployment. The 7.69TB is the data that was saved by SQL Server onto the array. Inline compression reduced that to 4.68TB, and inline deduplication reduced that to 2.72TB, for a total savings of 64%. This does *not* include any savings due to snapshots, which of course dramatically compound the space savings as well.



Solution Overview

The hardware chosen for this FT RA solution is highlighted below. A single B200 M4 blade was used for the testing. The UCS Chassis can house up to eight B200 blades, setting a precedent for power and density.



Cisco & Tegile Hardware setup

Component	Details
Server	Single B200 M4 Blade inside UCS 5108 Chassis
CPU	Intel Xeon E7-4370, 2.1GHz with Hyper-Threading
Total Sockets	2
Cores Per Socket	12
Total Physical Cores	24
Total Logical Cores	48
Memory	128 GB DDR3 (8 X 16GB DIMM's)
Internal Storage	2x Intel 520 Series SSD (200GB configured in RAID 1)
HBA	CISCO 6248UP IC (InterConnect)
Storage Array	Tegile IntelliFlash T10KHD - 300

Software	Details
Operating System	Windows Server 2012 R2 Standard
SQL Server	SQL Server 2016 Enterprise

Summary

The Cisco UCS Microsoft SQL Server 2016 Fast Track Reference Architecture powered by the Tegile All-Flash Storage Array balances price, performance, and ease of management for demanding data warehousing and BI scenarios. This Cisco + Tegile solution is certified for 55TB data warehouse(s) and can hold up to 1.2PB of data. Cisco and Tegile can deliver this with 100% flash solutions, all while preserving a proven enterprise architecture. The configuration is easy to expand in size or performance. It offers a simple way to accommodate single solutions or multi-tenant databases, along with “as-a-service” scenarios with complex billing. The infrastructure is exceptionally reliable and predictable, and with a total required space of 5U for storage, the system saves space and power compared to other solutions with the same performance parameters. The evolution of enterprise storage has placed Tegile in a strategic position to provide the most usable storage per dollar. Cisco and Tegile offer a wide array of all-flash and hybrid solutions that meet demanding performance requirements without breaking budgets.

- For more information, visit http://www.cisco.com/web/solutions/unified_comm/ucs/ms-sql-cisco-ucs.html and <http://www.tegile.com/solutions/database/sql/>
- To find out how to buy this solution, or something similar, contact sales at sales@tegile.com and mention SQL Data Warehouse Fast Track solution. We can also get you started on a proof of concept if you'd like to see the solution in action.



5 YEARS OF FRESH FLASH

Enterprise workloads require enterprise-grade SSDs, and we guarantee it for 5 years, even in high write I/O environments. Furthermore, we will upgrade your storage controllers at no additional charge after 5 years.

© Copyright 2016 Tegile Systems, Inc. The information contained herein is subject to change without notice. The only warranties for Tegile Systems 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. Tegile Systems shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows, and SQL Server are U.S. registered trademarks of Microsoft Corporation. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. All other trademarks and copyrights are property of their respective owners. All rights reserved.

Document: DW-5 Edit Version1.0