



Technical Report

Virtualized Microsoft SQL Server performance on NetApp AFF A-Series and C-Series systems

Ron Pratt, Manohar Vijay Kulkarni, Jeffrey Steiner | NetApp
July 2023 | TR-4976

Abstract

This technical report provides a performance summary of running Microsoft SQL Server on NetApp® AFF A-Series and AFF C-Series systems, plus guidance on how to select the right system for your storage needs based on workload usage.

NetApp provides high-performance flash (AFF or All SAN Array [ASA] A-Series), capacity flash (AFF C-Series), and hybrid (FAS-Series) systems in clusters of up to 24 nodes. With NetApp ONTAP® software, these systems offer performance acceleration, operation agility, industry-leading protection, and business continuity for database deployments.

Consult the [Interoperability Matrix Tool](#) (IMT) to determine whether the environment, configurations, and versions specified in this report support your environment.

TABLE OF CONTENTS

- Introduction 4**
- ONTAP data management 4**
- Microsoft SQL Server workload 4**
- ONTAP AFF platform 5**
- Database storage performance 5**
 - Read latency: Storage5
 - Read latency: Cache6
 - Write latency6
 - Bandwidth7
 - CPU7
- AFF performance: AFF A-Series and AFF C-Series 7**
 - 100% Read8
 - Read-Write9
 - Write latency9
- A-Series and C-Series compared 10**
 - Latency-sensitive workloads10
 - IOPS11
 - Durability11
 - Tiering11
 - Compression11
 - Spinning and hybrid array upgrades12
- Test configuration 12**
 - Database servers12
 - Microsoft SQL Server version13
 - Network configuration13
 - Storage layout13
 - Test methodology13
- Conclusion 13**
- Where to find additional information 13**
- Version history 14**

LIST OF FIGURES

Figure 1) A-Series and C-Series 100% reads.	8
Figure 2) A-Series and C-Series 70% reads, 30% writes.	9
Figure 3) A-Series and C-Series write latency comparison.	10
Figure 4) Microsoft SQL Server and AFF configuration.	12

Introduction

Data is the core of any business organization. Data continues to grow at an ever-increasing pace, so organizations are challenged to manage its volume and velocity. As database counts have increased along with data growth, organizations are looking to improve business agility and automate operations. Here are the other key considerations for organizations to manage databases:

- Optimize infrastructure costs by provisioning capacity to database workloads based on actual business needs, using modern technologies.
- Enable security teams to take advantage of storage capabilities to protect and secure large data volume from malicious external or internal attacks.
- Accelerate development by using storage features to quickly provision and manage data.
- Provide a reliable, high-performing storage solution that enables businesses to generate value and modernize SQL Server workloads.
- Harness the power of the hybrid cloud. Cloud-enable your IT services to get the benefits of on-premises storage with the flexibility of public cloud.

NetApp® ONTAP® and NetApp AFF storage provide high performance and reliable storage for mission-critical databases for data-driven enterprises, while modernizing infrastructure and unlocking the value of the data.

ONTAP data management

NetApp® ONTAP® is the data management platform that provides advanced data protection, inline efficiencies, and nondisruptive upgrades, as well as the ability to import a LUN from a non ONTAP storage array. ONTAP AFF clusters scale up to 24 nodes, simultaneously serving data through NFS, SMB, iSCSI, FC, and NVMe protocols. In addition, NetApp Snapshot™ technology is the basis for creating tens of thousands of online backups and fully operational database clones.

In addition to its rich feature set, ONTAP meets a wide variety of user requirements, including database size, performance requirements, and data protection needs. Deployments of NetApp storage include everything from a virtualized environment of approximately 5,000 databases running under VMware ESXi to a single-instance multi-TB data warehouse.

This document demonstrates and quantifies the differences and maximums of the AFF A-Series and AFF C-Series systems using a virtualized server infrastructure.

Microsoft SQL Server workload

SQL Server is the foundation of Microsoft's data platform, delivering breakthrough performance, availability, and manageability for mission-critical applications, whether on premises or in the cloud. The storage system is a key factor in the overall performance of a SQL Server database. NetApp offers several product features that allow your SQL Server database to deliver enterprise-class performance and help you to manage your environment.

The SQL Server database platform can support several applications. Before deploying SQL Server, you must understand the database workload requirements of the applications that your SQL Server instances support. Each application has different requirements for capacity, performance, and availability, and therefore each database should be designed to optimally support those requirements. Many organizations classify databases into multiple management tiers, using application requirements to define SLAs.

The following paragraphs describe SQL Server workloads.

OLTP databases are often the most critical databases in an organization. These databases usually back customer-facing applications and are essential to the company's core operations. Mission-critical OLTP databases and the applications they support often have SLAs that require high levels of performance and are sensitive to performance degradation and availability. They might also be candidates for Always On Failover Clusters or Always On Availability Groups. The I/O mix of these types of databases is usually characterized by 70% to 90% random read and 30% to 10% write.

Decision support system (DSS) databases (also known as data warehouses) are mission critical to many organizations that rely on analytics for their business. These databases are sensitive to CPU utilization and read operations from disk when queries are being run. In many organizations, DSS databases are the most crucial during the month, quarter, and year end. This workload typically has a 100% read I/O mix.

ONTAP AFF platform

As mentioned earlier, ONTAP data management software is the foundation for advanced data protection and management. However, ONTAP doesn't refer only to software. There are multiple ONTAP hardware- and software-defined platforms to choose from that rely on a variety of storage technologies, including flash media, spinning media, and virtualized storage. Nearly all databases deployed today use solid-state storage, and that trend is accelerating.

NetApp offers two AFF systems, the A-Series and the C-Series. Both are all-flash, solid-state storage solutions, but the A-Series targets ultra-latency-sensitive workloads, whereas the C-Series is aimed at solutions for which optimizing costs and capacity is a higher priority. The difference is almost entirely in the media. TLC flash media has emerged as the enterprise performance market leader in solid-state drive technology, whereas QLC costs significantly less.

A-Series and C-Series systems are based on the same controllers, allowing you to right-size the solution for your workloads in terms of CPU, host and network connectivity, and memory. You can mix A-Series and C-Series within the same cluster to create a tiered architecture. Finally, both the A-Series and C-Series use the same NVRAM and NetApp WAFL® technology in which write I/O is committed to mirrored NVRAM and written to media in full RAID stripes, yielding write latencies measured in microseconds.

Database storage performance

The most crucial step when choosing a storage platform for your database is to understand the application's actual needs. Many customers made the transition from spinning media to 100% solid-state storage the moment that flash became affordable, but not all of them experienced a clear benefit. Some databases were never limited by latency in the first place. They were limited by bandwidth, in which case solid-state and spinning media perform about the same. In other cases, databases weren't limited by storage performance to begin with.

The following sections explain considerations when choosing an AFF storage platform.

Read latency: Storage

Before affordable all-flash storage, storage latency was the #1 problem with database performance, for two reasons. First, reading a database block from spinning media required around 8ms to 10ms. If a database needed to serially read 1M individual blocks, read latency of 10ms each adds up to a lot of time.

The second reason that spinning media latency was a challenge was a result of the maximum number of I/O operations a single drive could service at a time—typically around 120 operations per second. Attempting to push more than 120 IOPS resulted in soaring latencies. The only solutions at the time were to either add more drives to the storage solution, or to use increasingly larger controllers with additional cache in the hope of avoiding drive I/O entirely.

NetApp AFF storage addresses both latency limitations:

- A-Series controllers with TLC media deliver high IOPS with submillisecond read latencies.
- C-Series controllers with QLC media deliver more IOPS with better latency than spinning media.

Whether you notice a difference between A-Series and C-Series performance depends on the type of workload. Although many database tasks require billions of individual reads performed serially, many databases are driven by user activity. For example, if you're waiting for a report that summarizes hundreds of thousands of individual bank transactions performed during the day, you'll probably see a benefit from the improved read latency of the A-Series controllers. If a database is hosting an online order entry system, it probably isn't latency sensitive. End users aren't going to notice a few extra milliseconds delay between the time they click the Submit button and the time they see the words "Order accepted!"

Read latency: Cache

The amount of RAM in a controller was often a critical consideration when sizing a storage solution based on spinning disk due to the limited IOPS and higher latency of spinning media I/O. Good performance depended on RAM offsetting the comparatively poor performance of spinning media and avoiding saturation.

With most all-flash storage solutions, the amount of RAM in a controller is rarely important because the service time of a read I/O from storage as compared to RAM is comparable. In this context, service time means the elapsed time from the moment the database issues an I/O and the moment it is received. The actual time required to read a block from RAM is obviously much lower than the time required to read a block from even an NVMe drive. However, when the time required by the host, protocol, and network layers is included, cache reads and drive reads show comparable latency.

This changes a little with C-Series due to the latency of the QLC media. Read I/O that can be serviced by cache on A-Series and C-Series can be serviced with latencies approaching 100µs, but the latencies on drive random reads can be noticeably higher with QLC.

Whether you see a difference between A-Series and C-Series sometimes depends on the chosen controller model. Databases often have very small working sets. For example, suppose that a NetApp customer with a nearly 10TB database has a working set of only few GBs. The result would be nearly equivalent performance on an AFF A400 and a C400 because the working set could reside in memory of the controller.

This is an important consideration when choosing between A-Series and C-Series: How big is the working set size of the database? It's extremely difficult to quantify, but often it can be estimated. For example, a call center database with 5 years of data that is 50TB in size nevertheless might have a working set of only 100GB because most of the activity is surrounding recent customer contacts. Customers rarely call in to inquire about bills that are more than a few months old, and if they do, additional latency in retrieving the required data is not likely to cause problems.

Write latency

Nothing damages database performance more than high write latency. Every time a change is committed by a database, one or more writes to the transaction logs must be completed and acknowledged by the storage system.

Fortunately, databases are almost never limited by write latency. The primary reason is that NetApp arrays, like most modern storage systems, do not commit writes directly to back-end media. With ONTAP, inbound writes are journaled into mirrored NVRAM and then acknowledged to the host. The update to the drives happens much later in the write process. This is why ONTAP delivers write latency measured in microseconds, even with spinning drive storage systems.

In addition, ONTAP WAFL technology avoids the RAID parity problem that affects writes with many competing storage systems. Without WAFL, you would never want to use a RAID 4/5/6 implementation

because of parity. Every database write would require multiple read I/Os from storage in order to recalculate parity. This was sometimes called "the RAID penalty" because completing a write required read I/O. ONTAP does not have this limitation because of WAFL. Inbound write I/O is journaled into NVRAM and then organized into full RAID stripes that can be written as a single unit.

Both the A-Series and C-Series use WAFL with the same NVRAM write technology, which means that both deliver the same ultra-low latency write I/O. Much of the legacy database documentation and guidelines surrounding the use of RAID, media, and other aspects of database storage sizing is not applicable to NetApp storage because ONTAP is immune to the problems that such recommendations are trying to solve.

Bandwidth

Some databases are bandwidth hungry. This is common among databases referred to as "data warehouses" and with tasks such as batch reports. The actual I/O pattern depends on many factors, but they include a lot of large-block sequential I/O operations. Despite the fact that such workloads are very I/O-intensive, the actual media type used rarely makes any difference. The reason is that large block sequential I/O operations are inherently very efficient. ONTAP storage systems can detect a sequential I/O operation in progress and proactively start reading drives and assembling required data before a host has even issued the request. Data is read more efficiently, in larger blocks.

The result is that these types of workloads should run about identically on A-Series and C-Series because these workloads are not latency sensitive. Databases that are known to be limited by sequential I/O performance usually have problems because of either host configuration errors or limitations in the network itself where ports are running at maximum line speed.

CPU

Nearly all database performance issues reported to the NetApp Support center are actually a result of database operations. In most cases, the root cause is CPU computation work on the database server itself. For example, if 85% of the database time is spent with database CPU work, there is rarely value in trying to optimize storage performance because improvements are not likely to be noticeable. Improving performance requires optimizing the SQL query itself to make it more efficient. In other cases, performance is limited by database contention; queries are delayed because they are waiting for other operations to release locks on data.

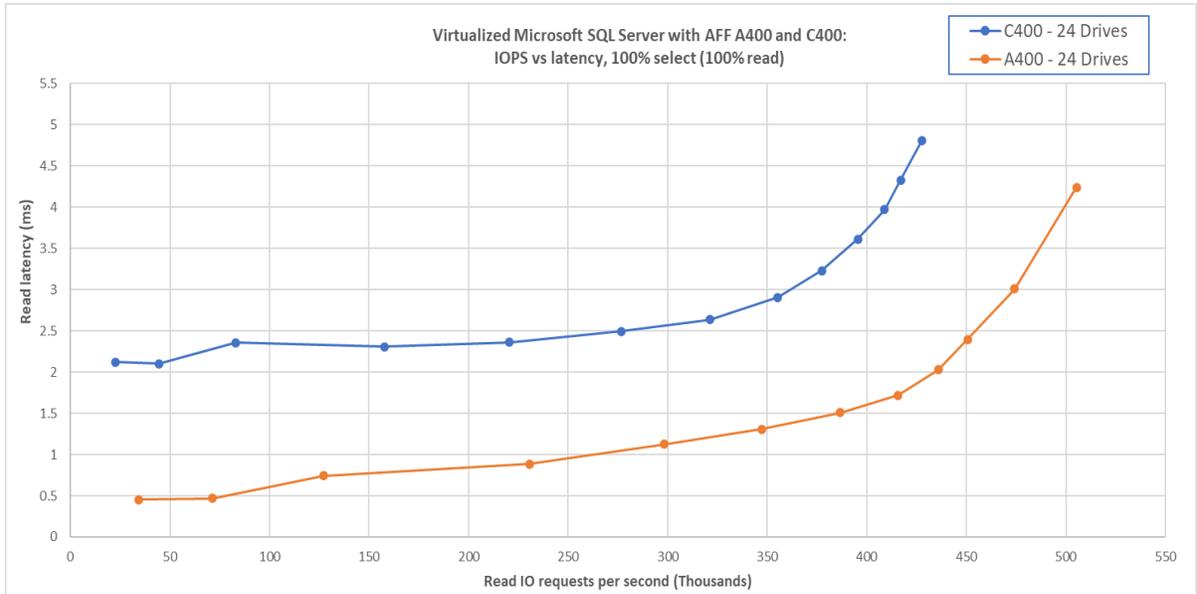
AFF performance: AFF A-Series and AFF C-Series

The following graphs show the performance capabilities of the AFF platforms in a number of configurations. These tests were set up with four Microsoft SQL Server instance running on four different ESXi hosts and were run on 2-node AFF A400 HA pairs and a 2-node C400 HA pair.

Note: We chose the FCP SAN protocol for these tests. Although other protocols are supported by various hypervisors, including NVMe/FCP, FCP is still the most popular.

100% Read

Figure 1) A-Series and C-Series 100% reads.



The following subsections describe the highlights of Figure 1.

Latencies

The A-Series latencies are mostly around 0.5ms to 1.5ms, depending on the controller model, and latencies peak as the controller nears saturation. This is a result of the faster TLC media. Most customers' database requirements are based on a latency cutoff of 1ms; the AFF A400 meets this requirement based on the performance curve. Customers also have the option to choose the next level of A-Series systems, the AFF A800, which can provide higher IOPS with lower latency. This is the reason that NetApp A-Series controllers usually eliminate storage as a database performance bottleneck. Databases become limited almost entirely by the query logic and CPU processing on the database server itself.

Latencies shown for the C-Series are higher due to the use of QLC media, but they are still far better than latencies for legacy spinning drive systems. Many advanced databases require the performance capabilities of an A-Series controller, but not all of them. Many of the largest databases in use today are performing the same functions they were performing 15 years ago, before the era of all-flash storage.

IOPS

The figures show the performance capabilities up to the point where the configuration saturates.

C-Series configurations reach a saturation point more gradually because the performance is mostly limited by the drives. Latency increases as the drives reach their limits. In contrast, A-Series configurations are usually limited by the controller CPU. The drives themselves usually offer greater performance potential than the controllers can extract. The result is much more consistent performance, right up to the point where the controller CPUs start reaching higher capacity utilization.

Read-Write

Figure 2) A-Series and C-Series 70% reads, 30% writes.

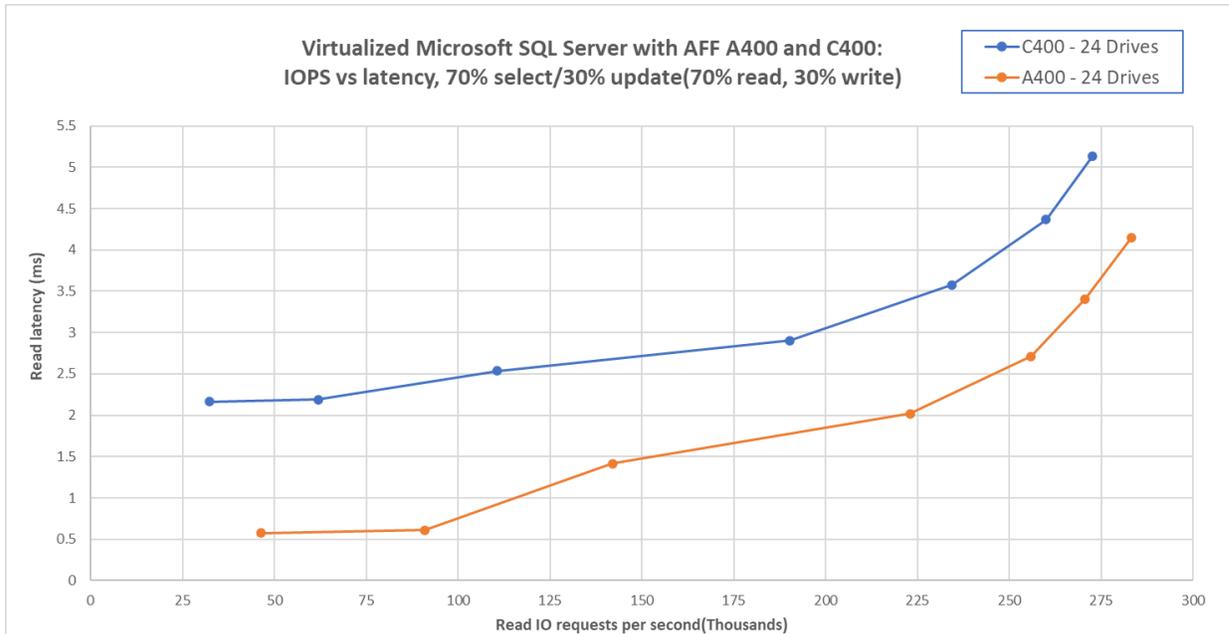


Figure 2 shows a 70% SELECT test, which could result in an 80% read ratio. There are two reasons for this result. One reason is that a test involving 30% update operations creates a read for each block updated, which slightly increases the read percentage. The other reason is that the Microsoft SQL Server engine looks for optimal way to fetch data as quickly as possible. SQL Server uses a read-ahead mechanism through which it tries to read pages in a specific order and thus copies more pages into the buffer cache in anticipation that those pages will be required to serve the request. The SQL Server read-ahead mechanism can use sequential read-ahead read or a random prefetching read-ahead mechanism.

Latencies

Latencies are slightly higher than the 100% read test, which is to be expected, because the write I/O is producing additional load on the controller. Both the A-Series and C-Series are equally affected by the write I/O because the write path is the same. The storage is still ONTAP, which means that the same write path is used when changes are committed to NVRAM, staged in RAM, and then written out to the drives.

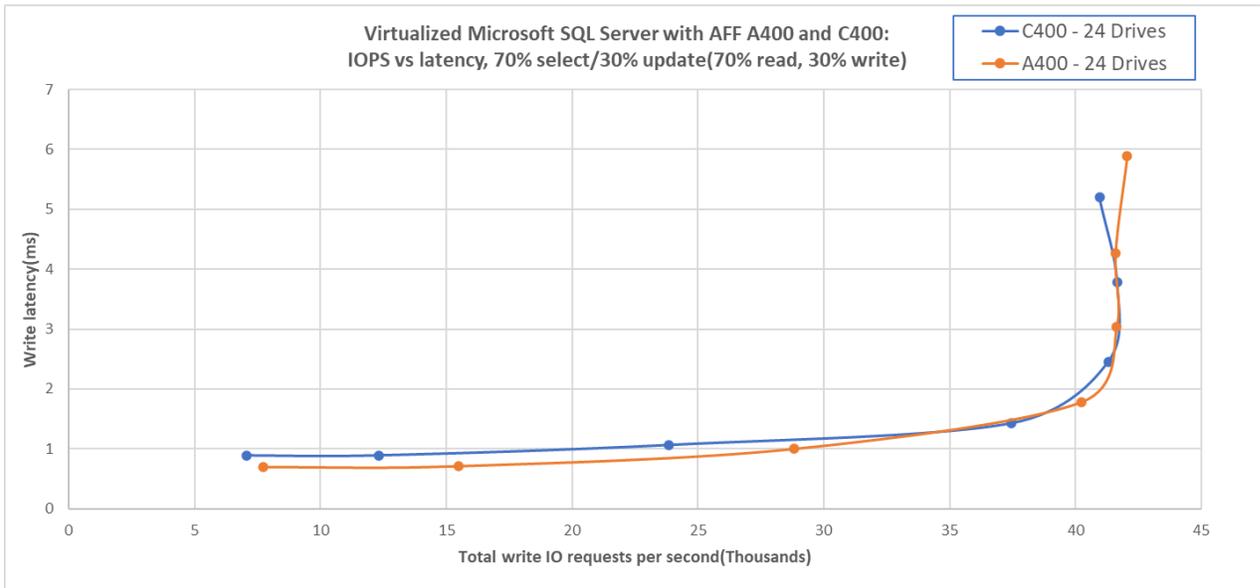
IOPS

The IOPS maximums are similarly affected on both platforms as a result of the increased write activity. Write I/O requires more CPU work to process, and drives are slightly slower with a mix of read and write I/O than with just read I/O.

Write latency

There are many ways to measure write latency at the database level.

Figure 3) A-Series and C-Series write latency comparison.



As explained earlier, the write path with A-Series and C-Series is the same. Whether you choose A-Series or C-Series, the write operation is complete from the host point of view once the I/O is journaled into NVRAM. Overall, write performance is the same on each platform.

A-Series and C-Series compared

With any database sizing effort, the most important question is “What do you need?”

For example, there is no reason to pay for IOPS capabilities or latency benefits if the workload does not require them. A proof of concept that is based on nothing more than which configuration shows the highest IOPS potential isn't useful unless you know that the limiting factor for a given business need is genuinely raw IOPS. In addition, a test of only IOPS ignores additional critical factors such as latency.

Not all workloads can be numerically quantified. Sometimes the only option is to take some time to understand the business need. For example, who is using the database? Are they just making updates, or are they running intensive reports that require millions of individual I/Os? Performance metrics can be collected by using the Windows performance counter and Microsoft SQL Server query statistics from dynamic management views. This data can provide insight into the IOPS requirements and current latencies and can identify whether storage is a contributing factor to performance issues.

Once you have an idea about what the requirements are, you can select the right controller.

Latency-sensitive workloads

The A-Series controllers are unquestionably the best option where random read latency is critical. Both the A-Series and C-Series deliver comparable write latencies, which means that critical processes are always running at peak performance. Similarly, reads of hot data that are already cached in memory yield the same ultra-low response times. The difference between A-Series and C-Series is in random read performance.

The difference between a less than 1ms read operation on A-Series and a 2ms read operation on C-Series might seem substantial. However, 2ms is a substantial improvement over the latency on spinning

disk solutions that are still used today by many mission-critical databases that require many terabytes of storage and high throughput and yet accept 8ms to 12ms of latency on random reads.

Memory management

As described previously, avoiding I/O altogether is preferable to optimizing I/O. Increasing memory allocated to the database effectively reduces latency by changing reads from storage into reads from server memory.

If a C-Series controller is not able to provide the required low latency for a specific database workload, increasing the memory allocated to the SQL Server instance can improve effective latency. Alternatively, it might be better to consider rehosting the SQL Server database on an A-Series system.

IOPS

The required IOPS level is mostly controlled by the storage controllers. NetApp sizing tools can help you select the right controller. The drive count with AFF systems does not affect the maximum IOPS as it did with spinning media, but there is some effect. As a general rule, AFF systems should have a minimum of 24 drives per HA pair (12 per controller).

Durability

A second difference (in addition to the latency characteristics) between QLC drives and other flash technology is *wearing*. The manufacturer specifications for QLC drives include a reduced overwrite capability as compared to TLC drives, but this has minimal effect on ONTAP storage systems. First, ONTAP RAID protects data against media failures and includes both double and triple parity drive options. Furthermore, ONTAP WAFL technology distributes inbound write data to free blocks across multiple drives. This minimizes overwrites of the individual cells within the drive, which maximizes the drive's useful life. Finally, NetApp support agreements that cover drive failures also include drive replacement for SSDs that have exhausted their write cycles.

Tiering

The choice between A-Series and C-Series is not an either/or decision. For example, you could build a 4-node cluster containing two AFF A400 controllers for latency-sensitive databases and two C400 controllers for others. As needs change, you can easily and nondisruptively migrate databases between tiers with vol move. And you could add C250 controllers to store Microsoft SQL Server backups.

Compression

Temperature-sensitive storage efficiency (TSSE) is available on A-Series and always enabled on C-Series. This feature improves storage efficiency by detecting inactive data and recompressing it with a larger compression block, further reducing storage requirements. In practice, there is usually enough routine database activity, including full table scans, backups, reindexing, upgrades, and other activities that touch most or all blocks in a database, to prevent TSSE from taking effect in the first place.

If TSSE with a 14-day cooling policy was enabled for a particular set of database volumes, and the database was shut down entirely for 30 days, the underlying blocks would be 100% recompressed by TSSE. If the database was then restarted and subject to heavy random read and write I/O, performance would initially be affected. Performance would eventually return to normal pre-TSSE levels.

This does not necessarily mean that the increased latency would be problematic, especially on a C-Series system where drive latency is already higher than A-Series drive latency. Furthermore, the increased savings on cold data might offset any such increase.

Spinning and hybrid array upgrades

If similar tests were performed on a legacy spinning-disk solution, the average latency would start at around 8ms because that is the time required for an electromagnetic drive head to seek to the right position and transfer a block. The C-Series delivers nearly 4 times improvement.

The saturation point problem also drives migrations from spinning disk solutions toward C-Series. Many workloads are not latency sensitive, but they still require high IOPS levels. A typical high-performance database in the days of spinning disk solutions would involve as many as 1,000 drives, because each drive could service only around 120 IOPS. 1,000 drives would yield around 120K IOPS.

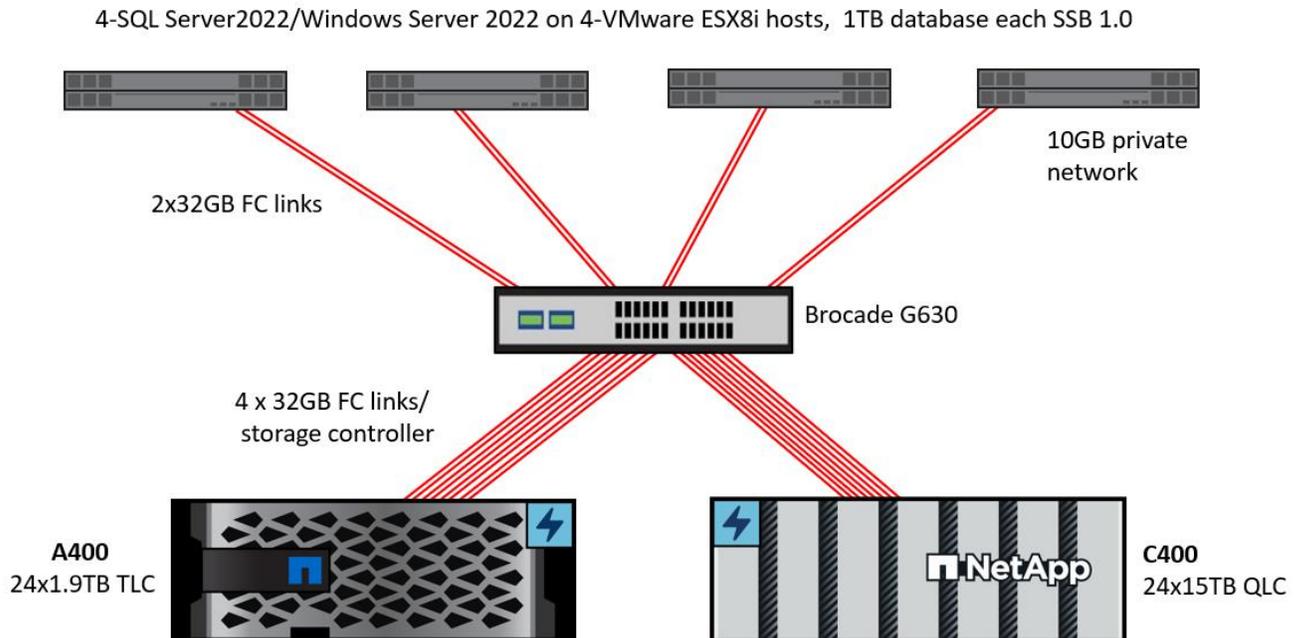
Such a solution is no longer viable. Even disregarding the decreased reliability, increased power consumption, and increased heat output of spinning drives, such solutions are impossible to build because the drives in the dwindling spinning disk markets are enormous in size. Storage solutions would reach their maximum capacity in terms of bytes with a small handful of drives. The IOPS limitation would be severe. A solution with 50 spinning drives would support only around 6000 IOPS.

C-Series solves this problem without the expense of the TLC media in the A-Series controllers. You get vastly more IOPS potential and increased reliability in a much smaller footprint with much lower power and cooling requirements. That alone would be a significant benefit, but you also get a 4 times (or better) improvement in latency.

Test configuration

This section describes how the tests were configured and performed.

Figure 4) Microsoft SQL Server and AFF configuration.



Database servers

Database servers running on four VMware ESXi 8 hosts, with one database server on each ESXi host.

Microsoft SQL Server version

Microsoft SQL Server 2022 on Windows 2022 with database size of 1TB on each instance per host.

Network configuration

The FC network was configured with two 32GB FC connections on each ESXi server and four 32GB FC connections for each storage controller. Because controllers are normally deployed in HA pairs, we used a total of eight FC connections for the A-Series system and eight FC connections for the C-Series controller. Testing was performed on either the A-Series or the C-Series systems (not simultaneously).

The IP network used 10GB NICs each for public and private network. Each ESXi is connected to public network for management and to a private network for Microsoft SQL Server communication.

Storage layout

Each ESXi host is running a single instance of Microsoft SQL Server and a total of four ESXi servers, Four Microsoft SQL Server instances are running. Each SQL Server instance has a database consisting of eight data files and a single log file that is spread across 9 VMDKs on separate datastores. Each datastore is hosted on its own LUN and each LUN on its own volume. A total of 18 storage volumes are laid out on each controller for a total of 36 storage volumes.

Test methodology

For our testing, we used a NetApp-internal SQL Server workload generator tool called SQL Storage Benchmark (SSB). SSB is a highly tunable workload generator that allows the user to test a variety of read/write ratios while using the SQL Server database engine to drive the workload. It's not a synthetic I/O generator. SSB allows the user to have any number of database user tables, allowing control over working set size, caching, random or sequential I/O, and read/write percentages. The data gathered provides a better measure of storage performance in terms of latency, IOPs, and throughput, as well as TPM, based on storage subsystem performance. The transactions being measured are a direct correlation of IOPs and throughput. SSB helps the user to understand the storage limits better because the database engine is driving a workload directly to storage.

- Using SSB generated a 100% read workload.
- Using SSB generated a 70% read workload.

Conclusion

This technical report compared performance between A-Series and C-Series systems running a Microsoft SQL Server workload in a virtualized environment. NetApp AFF A-Series has a proven history of data management with optimal performance and submillisecond latency for mission-critical applications. NetApp AFF C-Series offers higher-capacity data storage to achieve impressive return on investment. NetApp AFF C-Series is good fit for business databases that do not require submillisecond latency, or to store backup as secondary storage or for testing database workloads.

With the flexibility and scalability of these NetApp ONTAP AFF platforms, your organization can start with a modern IT infrastructure and can grow and adapt to evolving business requirements. With these benefits, your system can serve existing workloads while streamlining data management processes, reducing operational costs, and preparing for workloads of new business in the future.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp A-Series
<https://www.netapp.com/data-storage/aff-a-series/>
- NetApp C-Series
<https://www.netapp.com/data-storage/aff-c-series/>
- NetApp Product Documentation
<https://docs.netapp.com>
- Microsoft SQL Server best practice guide on ONTAP
<https://www.netapp.com/pdf.html?item=/media/8585-tr4590.pdf>
- Virtualized Oracle database performance on NetApp AFF A-Series and C-Series
<https://www.netapp.com/pdf.html?item=/media/85629-tr-4971.pdf>

Version history

Version	Date	Document version history
Version 1.0	July 2023	Initial release

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright information

Copyright © 2023 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data—Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.