



NetApp Verified Architecture

Quantum StorNext with NetApp E-Series Systems Design Guide

NVA Design

Ryan Rodine, NetApp
January 2021 | NVA-1150-DESIGN | Version 1.2

Abstract

This document provides details on how to design a StorNext parallel file system solution with NetApp® E-Series storage systems. This solution covers the NetApp EF280 all-flash array, the NetApp EF300 all-flash NVMe array, the EF600 all-flash NVMe array, and the NetApp E5760 hybrid system. It offers performance characterization based on Frametest benchmarking, a tool that is widely used for testing in the media and entertainment industry.

TABLE OF CONTENTS

Executive summary	3
Use case summary	3
Solution overview	3
Target audience.....	4
Solution technology	4
Technology requirements	9
Hardware used	9
Software requirements	10
Solution verification	10
Conclusion	16
Where to find additional information	16
Version history.....	17

LIST OF TABLES

Table 1) Software requirements.	10
--------------------------------------	----

LIST OF FIGURES

Figure 1) Technical components of the solution with an EF280 array.	5
Figure 2)) Technical components of the solution with an EF300 array.	6
Figure 3) Technical components of the solution with an NetApp EF600 all-flash array.	7
Figure 4) Technical components of the solution with an E5760 system.	8
Figure 5) Results from testing with an EF280 array.	11
Figure 6) I/O Response times with an EF280 array.....	11

Executive summary

As commercial video industry technologies continue to grow and improve, editing and storage requirements also continue to grow, along with the need to produce video in higher resolutions. With these changes come greater production demands and the need for massive bandwidth from the underlying production storage systems. These changes can overwhelm the proprietary media storage systems that are traditionally used in these environments.

NetApp® E-Series systems with StorNext provide broadcasters, media studios, and media content aggregators with a high-performance and high-capacity streaming tier that also optimizes rack space, power, and cooling usage.

The NetApp EF280 all-flash array provides an entry-level solution optimized for read performance of up to 9.3GBps of video throughput in a 2U building block, serving up to 8 simultaneous uncompressed 4K read streams with zero dropped frames.

The NetApp EF300 provides a midrange all-flash NVMe array which can provide even more performance of up to 19.8GBps of video throughput in a 2U building block, service up to 17 simultaneous uncompressed 4K read streams with zero dropped frames.

For the highest performance, the NetApp EF600 all-flash NVMe array, can provide up to 38.5GBps of video throughput in a 2U building block, serving up to 33 simultaneous uncompressed 4K read streams with zero dropped frames.

When \$/GB is a key metric, the E5760 configuration we tested provides capacity at lower cost with HDDs in 8U of rack space. It can also serve up to 9 simultaneous uncompressed 4K read streams with zero dropped frames.

Use case summary

This solution applies to the following use cases in a StorNext environment:

- Ingest and playout of high-bandwidth media streams
- The need for either a high-bandwidth streaming tier or a lower-cost, high-capacity streaming tier
- High-resolution media asset management
- Online media archiving
- Centralized storage to provision content creation workstations and to broadcast playout workstations

Solution overview

NetApp E-Series plus a file system such as StorNext dramatically streamlines workflow and improves productivity. The combination creates a shared repository that supports flexible, high-performance streaming, even with high-bit-rate media content. This repository also includes:

- A single namespace, for virtually limitless bandwidth or capacity
- Near-linear bandwidth scalability, for both scale-up and scale-out configurations
- Support for Linux, Microsoft Windows, and macOS clients

Growing media and entertainment companies are challenged to find storage tier solutions that both satisfy their high-density, high-bandwidth requirements and optimize rack space power and cooling. EF280, EF300, EF600, and E5760 systems fulfill these roles.

The architecture that is demonstrated in this design guide shows the capabilities of the four options. First, an entry-level all-flash option, using a single EF280 array. Second, a medium-tier option using a single EF300 array with high performing NVMe drives. Third, a high-performance building block using a single

EF600 array also with high performing NVMe drives. Fourth, an E5760 system using cost-effective SAS HDD drives. Additional similarly configured arrays can be added to the StorNext namespace to allow virtually unlimited scale-out.

Each option also demonstrates value. The EF280 is a low-cost entry into an all-flash storage array. The EF300 and EF600 have the ability to support a high number of high-resolution video streams with no dropped frames. The E5760 provides a lower-cost, large-capacity target in 8U of rack space. All options save the customer in footprint and power/cooling costs.

In addition, the EF300 and EF600 arrays support NVMe protocols, enabling the potential to upgrade to NVMe over Fabrics in the future.

Target audience

The target audience for the solution includes the following groups:

- Media content aggregators and distribution companies
- Broadcasters and networks
- Film studios
- Multimedia corporations
- Sports leagues

Solution technology

The StorNext solution in this design guide consists of at least four StorNext clients. Each configuration runs the same tests, but with a different target storage array in each case. All configurations use a StorNext metadata controller (MDC) with separate storage for the metadata and the journal volumes. In this example, a NetApp EF570 all-flash array is used for the MDC, but a StorNext MDC appliance can be used for the metadata and journal, as well.

The SAN consists of a 32Gb Fibre Channel (FC) fabric, facilitated with ATTO CTFC-324E adapters running ATTO MultiPath Director on the clients.

The EF280 has 24 flash drives and is provisioned into two 10+2 RAID 6 volume groups.

The EF300 and EF600 options each have 24 NVMe drives and are provisioned into two 10+2 RAID 6 volume groups.

The E5760 has 120 SAS HDDs and is provisioned into twelve 8+2 RAID 6 volume groups.

Figure 1 through Figure 4 show the technical components of the EF280, EF300, EF600, and E5760 solutions.

Figure 1) Technical components of the solution with an EF280 array.




Figure 2)) Technical components of the solution with an EF300 array.




Figure 3) Technical components of the solution with an NetApp EF600 all-flash array.

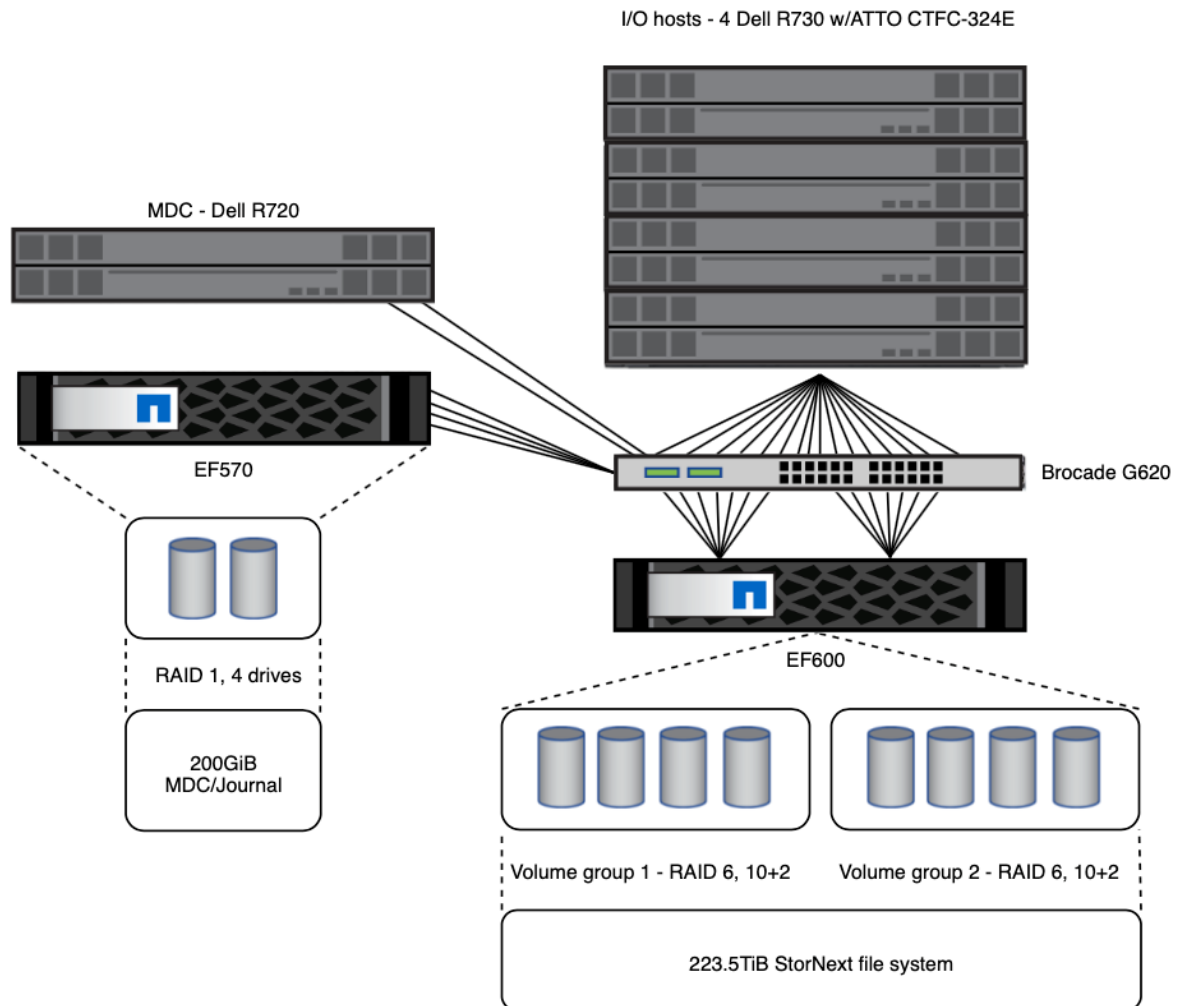



Figure 4) Technical components of the solution with an E5760 system.



Technology requirements

Hardware used

Servers for the NetApp EF600 array: One metadata controller (MDC) and four Dell R730 with ATTO Fibre Channel (FC) adapters are used:

- Two Intel Xeon E5-2670 v3 processors @ 2.3GHz
- 128GB RAM
- Client host bus adapters (HBAs): ATTO CTFC-324E
- MDC HBA: Broadcom LPe32002

Servers for the NetApp EF280, EF300, and NetApp E5760 configurations: One MDC and four Lenovo SR650 with the following specifications:

- Two Intel Xeon Gold 6136 CPUs @ 3.00GHz
- 96GB RAM
- Client HBAs: ATTO CTFC-324E
- MDC HBA: Broadcom LPe32002

Storage: NetApp E-Series arrays:

- Data volumes on the EF280 array:
 - 24 MZILS15THMLS-0G4 - 15TB SSD drives
 - Two 10+2 RAID 6 volume groups
 - Four 34.9TiB volumes per volume group with 512KiB segment size
- Data volumes on the EF300 array:
 - 24 MZWLJ7T6HALA-0G5 - 7TB NVMe drives
 - Two 10+2 RAID 6 volume groups
 - Four 15.7TiB volumes per volume group with 512KiB segment size
- Data volumes on the EF600 array:
 - 24 MZWLL15THMLA-0G5 15.3TB NVMe drives
 - Two 10+2 RAID 6 volume groups
 - Four 27.9TiB volumes per volume group with 512KiB segment size (capacity is 20% overprovisioned per NVMe drive guidelines)
- Data volumes on the E5760 system are comprised of:
 - 120 ST10000NM0096 10TB 7,200 RPM SAS drives
 - Twelve 8+2 RAID 6 volume groups
 - Twelve 71.2TiB volumes per volume group with 512KiB segment size
- MDC volumes on the EF570 array are comprised of:
 - Four PX02SMU080 800GiB SSDs
 - One 4-drive RAID 1 volume group
 - Two 200GiB volumes with 128KiB segment size

Switching: Brocade G620

Software requirements

This design includes StorNext file system version 6.1.0. The Quantum StorNext file system is a versatile, high-performance shared file system that offers heterogeneous, block-based access to a single storage system or striped across multiple external storage systems.

This example consists of tests that were performed on separate EF280, EF600, and E5760 storage targets running NetApp SANtricity® OS 11.60. The EF300 array was running SANtricity OS version 11.70.

The clients in this design run Red Hat Enterprise Linux (RHEL) 7.7, with ATTO MultiPath Director for failover support; however, a wide variety of host operating systems are supported. For a full list of supported operating system combinations, consult the [NetApp Interoperability Matrix Tool](#).

Table 1 lists the software components that are required to implement the solution. The software components that are used in any particular implementation of the solution might vary based on customer requirements.

Table 1) Software requirements.

Software	Version or Other Information
StorNext	6.1.0
SANtricity	11.60 and 11.70
MDC OS	RHEL 7.5
Client OS	RHEL 7.7
ATTO MultiPath Director	Driver 1.76 MP; firmware, March 26, 2019

Solution verification

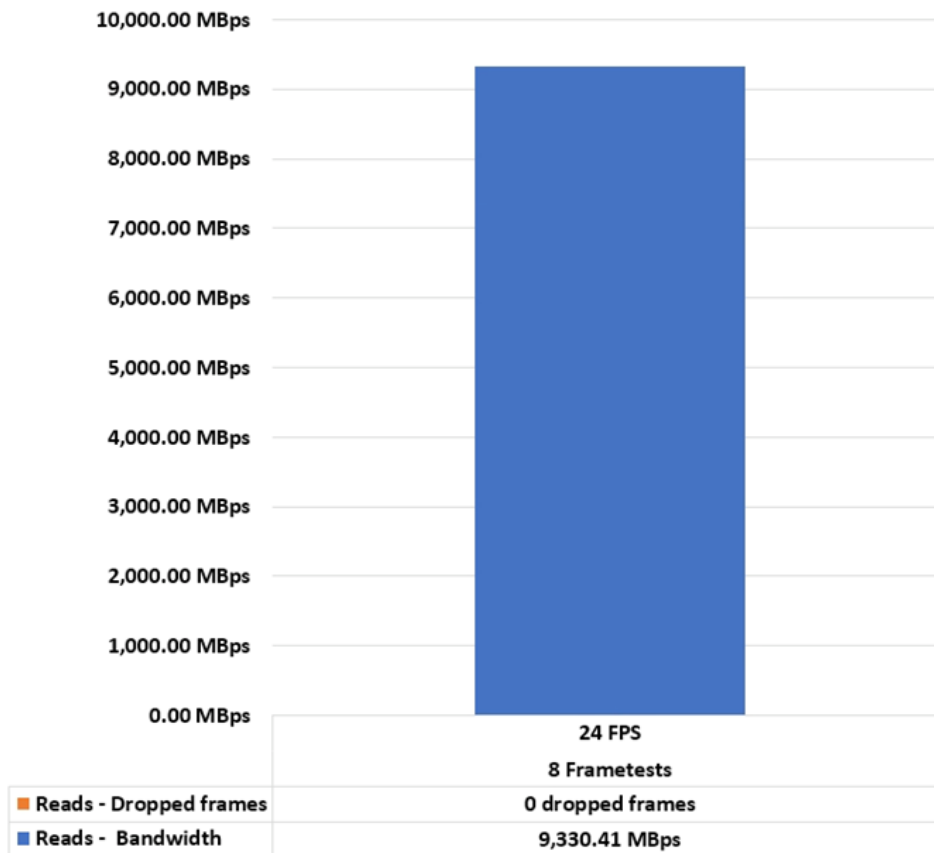
For this solution, the industry-standard Frametest benchmarking tool was used to generate uncompressed 4K streams at 24 frames per second (FPS). Multiple single-threaded Frametest streams were started on each host and were scaled up to see how many streams could run simultaneously without dropping frames.

To confirm consistency, the streams were intentionally run for a long period, 27 minutes each.

EF280 results

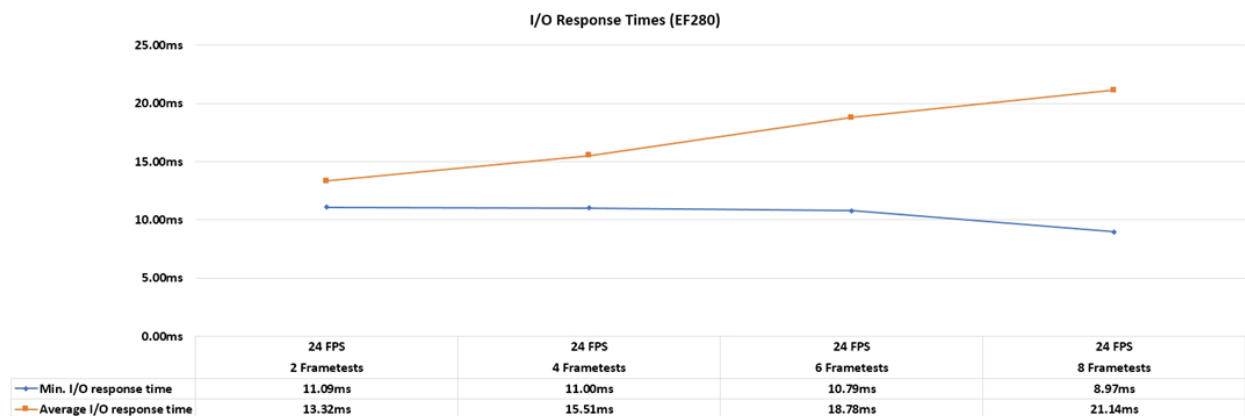
As Figure 5 shows, in testing the NetApp EF280 configuration achieved up to 8 simultaneous 4K read streams at a total of 9.3GBps, with zero dropped frames on 100% reads.

Figure 5) Results from testing with an EF280 array.



As Figure 6 shows, I/O response times were low, with average response times of 21.29ms at 9 streams.

Figure 6) I/O Response times with an EF280 array.



Some additional write and mixed workload testing was performed to see how many streams could be processed with no dropped frames.

At 100% writes, the EF280 array achieved up to two write streams with zero dropped frames, for a total of 2.3GBps writes. With a mixed workload, the EF280 array achieved up to six read streams at 6.9GBps while simultaneously writing one instance at 1.6GBps with zero dropped frames. See Table 2.

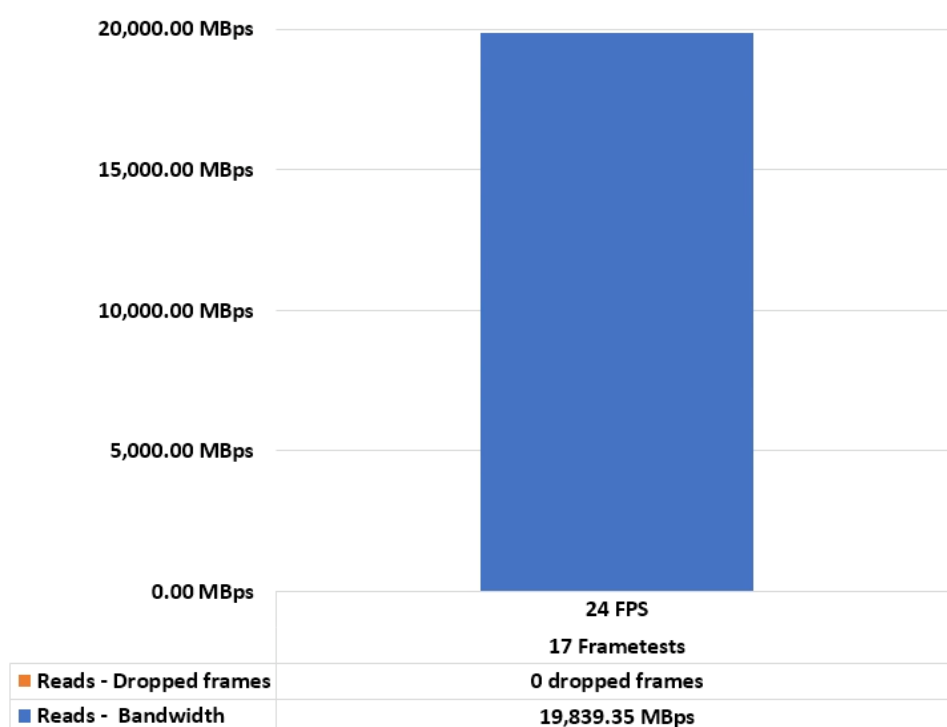
Table 2) Read, write, and mixed workload results for the EF280 array.

I/O Profile	Read Streams GBps	Write Streams GBps
100% reads	8 streams, 9.3GBps	–
100% writes	–	2 streams, 2.3GBps
Mixed reads/writes	6 streams, 6.9GBps	1 stream, 1.6GBps

EF300 results


As Figure 7 shows, in testing the NetApp EF300 configuration achieved up to 17 simultaneous 4K read streams at a total of 19.8GBps, with zero dropped frames on 100% reads.

Figure 7) Results from testing with an EF300 array.



As Figure 8 shows, I/O response times were low, with average response times of 27.63ms at 17 streams.

Figure 8) I/O Response times with an EF300 array.



Some additional write and mixed workload testing was performed to see how many streams could be processed with no dropped frames.

At 100% writes, the EF300 array achieved up to five write streams with zero dropped frames, for a total of 5.8GBps writes. With a mixed workload, the EF300 array achieved up to eight read streams at 9.3GBps while simultaneously writing two streams at 2.3GBps with zero dropped frames. See Table 3.

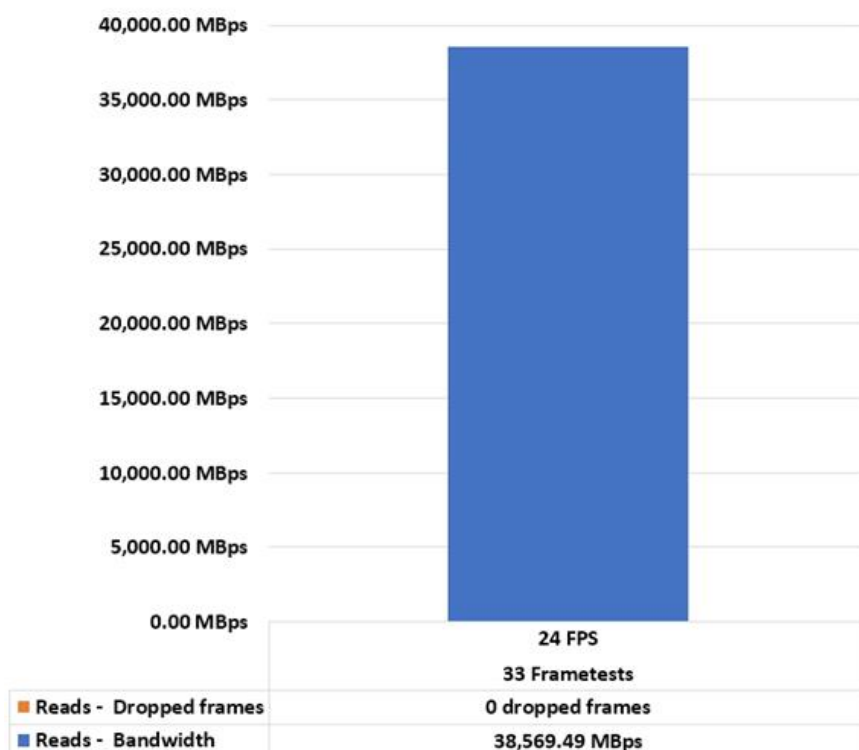
Table 3) Read, write, and mixed workload results for the EF280 array.

I/O Profile	Read Streams GBps	Write Streams GBps
100% reads	17 streams, 19.8GBps	—
100% writes	—	5 streams, 5.8GBps
Mixed reads/writes	8 streams, 9.3GBps	2 stream, 2.3GBps

EF600 results

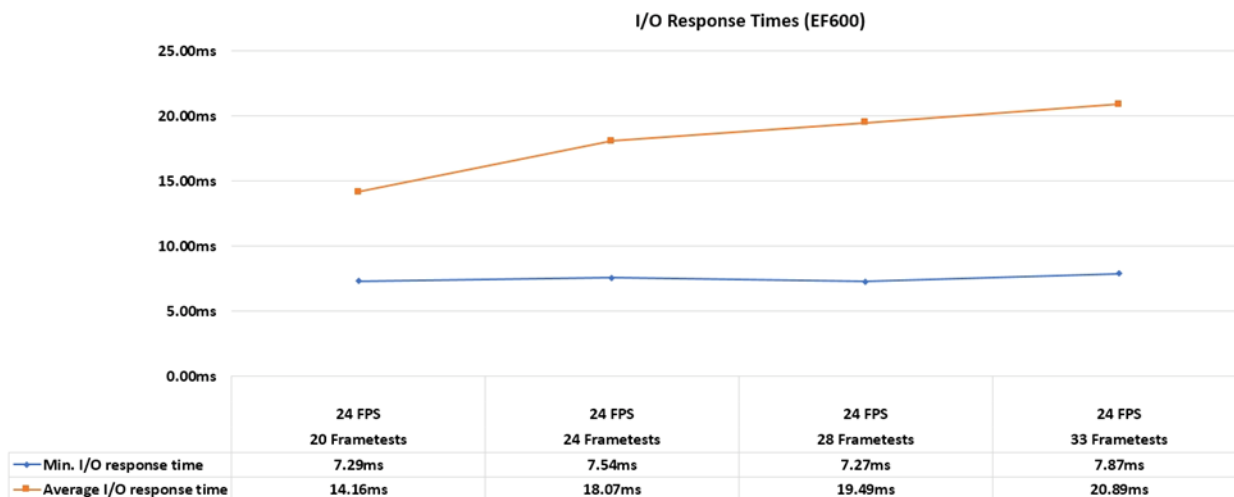
As Figure 9 shows, in testing, the NetApp EF600 configuration achieved up to 33 simultaneous 4K read streams at a total of 38.5GBps, with zero dropped frames on 100% reads.

Figure 9) Results from testing with an EF600 array.



In addition, I/O response times were low, with average response times of 20.89ms at 33 streams. Figure 10 illustrates the minimum and average response times as the number of streams increased.

Figure 10) I/O response times with an EF600 array.



Some additional write and mixed workload testing was performed to see how many streams could be processed with no dropped frames.

At 100% writes, the EF600 array achieved up to 9 write streams with zero dropped frames, for a total of 10.5GBps writes. With a mixed workload, weighted toward reads, the EF600 array achieved up to 16 read

streams at 18.7GBps while simultaneously writing 5 streams at 5.8GBps with zero dropped frames. See Table 4.

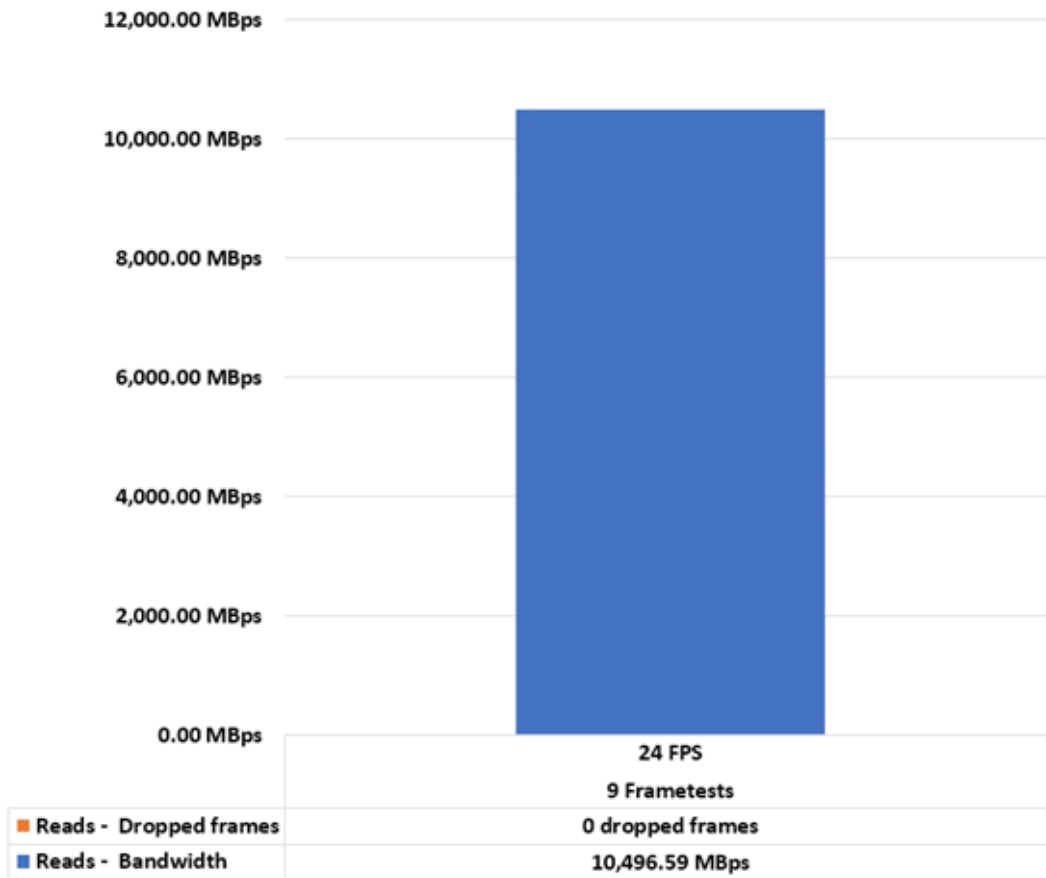
Table 4) Read, write, and mixed workload results for the EF600 array.

I/O Profile	Read Streams GBps	Write Streams GBps
100% reads	33 streams, 38.5GBps	–
100% writes	–	9 streams, 10.5GBps
Mixed reads/writes	16 streams, 18.7GBps	5 streams, 5.8GBps

E5760 results

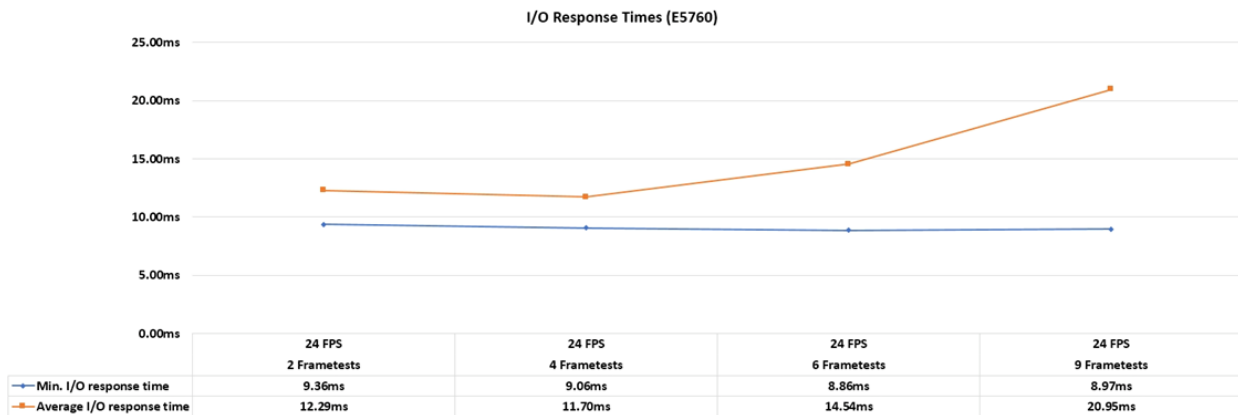
As Figure 11 shows, in testing, the NetApp E5760 configuration achieved up to nine 4K read streams at a total of 10.4GBps, with zero dropped frames on 100% reads.

Figure 11) Results from testing with an E5760 system.



I/O response times also remained low, up to nine streams. Figure 12 illustrates the minimum and average response times as the number of streams increased.

Figure 12) I/O response times with an E5760 system.



Some additional write and mixed workload testing was performed to see how many streams could be processed with no dropped frames.

At 100% writes, the E5760 system achieved up to seven write streams with zero dropped frames, for a total of 8.1GBps writes. With a mixed workload, weighted toward reads, the E5760 system achieved up to six read streams at 6.9GBps while simultaneously writing two streams at 2.3GBps with zero dropped frames. See Table 5.

Table 5) Read, write, and mixed workload results for the E5760 system.

I/O Profile	Read Streams GBps	Write Streams GBps
100% reads	9 streams, 10.4GBps	—
100% writes	—	7 streams, 8.1GBps
Mixed reads/writes	6 streams, 6.9GBps	2 streams, 2.3GBps

Conclusion

Using NetApp E-Series arrays in your StorNext solution can provide a variety of options.

The NetApp EF280 provides an entry-level all-flash option, which can support up to 8 simultaneous 4K uncompressed streams without dropping frames.

The NetApp EF300 array provides a midrange all-flash NVMe option, which can support up to 17 simultaneous 4k uncompressed streams without dropping frames.

Use of a NetApp EF600 array provides a high-performing shared storage system that is optimized for supporting many uncompressed 4K video streams. Testing shows that it can serve up to 33 simultaneous 4K uncompressed streams without dropping frames, in only 2U of rack space.

The NetApp E5760 system provides a lower-cost, high-capacity option in 8U of rack space and is efficient and cost-effective.

Where to find additional information

To learn more:

- NetApp Media and Entertainment Storage Solutions
<https://www.netapp.com/us/solutions/industry/media-entertainment-storage-solutions.aspx>

- E-Series for Media Solution Brief
<https://www.netapp.com/us/media/ds-e-series-media.pdf>
- NetApp Product Documentation
<https://docs.netapp.com>

Version history

Version	Date	Document Version History
Version 1.0	August 2020	Initial release.
Version 1.1	October 2020	Added reference to EF280.
Version 1.2	January 2021	Added reference to EF300.

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 © 2020–2021 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. 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).

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

NVA-1150-DESIGN-0121