



Technical Report

Best Practice Guide for Microsoft SQL Server with NetApp EF-Series

Mitch Blackburn, Pat Sinthusan, NetApp
March 2019 | TR-4764

Abstract

This best practice guide is intended for storage administrators and database administrators to help them successfully deploy Microsoft SQL Server on NetApp® EF-Series storage.

TABLE OF CONTENTS

1	Executive Summary	4
1.1	Intended Audience.....	4
1.2	Caveats.....	4
2	Introduction to EF-Series	5
2.1	EF-Series Hardware Overview.....	5
2.2	NetApp SANtricity OS.....	6
2.3	Provisioning the EF-Series.....	8
3	SQL Server Overview	8
3.1	SQL Server I/O Overview.....	9
3.2	Windows Volume Mount Points.....	10
3.3	Database Files and File Groups.....	10
3.4	SQL Server Database Files.....	11
3.5	Database Files for Tempdb.....	12
4	EF-Series Performance	13
5	High Availability	14
5.1	EF-Series and SANtricity OS.....	14
5.2	SQL Server HA Options.....	15
6	Sizing	15
6.1	EF-Series I/O Overview.....	16
6.2	Estimating I/O.....	16
7	EF-Series Performance Monitoring Using NetApp SANtricity System Manager	18
8	Summary	20
	Appendix A: Provisioning EF-Series with REST API	21
	Where to Find Additional Information	26
	Version History	26

LIST OF TABLES

Table 1)	EF-Series performance.....	5
Table 2)	Comparison of usable capacity for differing RAID levels.....	16

LIST OF FIGURES

Figure 1)	Dynamic Disk Pool components.....	7
-----------	-----------------------------------	---

Figure 2) Dynamic Disk Pool drive failure.....	7
Figure 3) Example of primary, secondary, and log files that locate in each drive.	10
Figure 4) Performance results for OLTP configuration.	13
Figure 5) Step to select E-Series Forward and reverse Sizer.....	17
Figure 6) E-Series reverse sizer input fields.....	18
Figure 7) Reverse sizing output.....	18
Figure 8) Example of Performance Monitor with logical view.	19
Figure 9) Example of Performance Monitor with physical view.	19
Figure 10) Example of Performance Monitor with applications and workloads view.	20

1 Executive Summary

In most OLTP systems, the processor, memory, and I/O subsystem in a server are well balanced and are not considered performance bottlenecks. The major source of performance issues in OLTP environments is typically related to the storage I/O activity. The speed of existing HDD-based storage systems does not match the processing capabilities of the servers.

As a result, a situation often occurs where a powerful processor sits idle, waiting for the storage I/O requests to complete, negatively affecting user and business productivity. The negative effect on productivity extends the time to return on investment (ROI) and increases overall TCO. Therefore, storage IOPS performance and latency become strategic considerations for business. It is critical to meet response time goals and optimize performance for other system resources (processor and memory).

The NetApp EF-Series flash array provides a robust platform for delivering exceptional performance with submillisecond response times to mission-critical applications. The EF-Series flash array uses the latest technology in solid-state drive (SSD) along with a strong heritage handling diverse workloads to provide superior business value through the acceleration of latency-sensitive and high-I/O environments.

Not only is the EF-Series flash array exceptional for OLTP databases but also for data warehouse environments. The entry-level EF-Series EF280 was recently certified under Microsoft's Fast Track Data Warehouse program, the certification can be found [here](#). The midrange EF-Series was also certified under the Fast Track Data Warehouse program and can be found [here](#).

The NetApp EF-Series flash array leads the market in delivering high performance and low latency. The EF570 is the \$/SPC-2 MBps leader¹. "To say the EF570 is quick is putting it mildly; it's a monster."²

NVMe has become the industry standard interface for Peripheral Component Interconnect Express (PCIe) solid-state drives (SSDs). With a streamlined protocol command set and fewer clock cycles per I/O, NVMe supports up to 64K queues and up to 64K commands per queue. These attributes make it more efficient than SCSI-based protocols like SAS and SATA.

The introduction of NVMe over Fabrics (NVMe-oF) makes NVMe more scalable without affecting the low latency and small overhead that are characteristic of the interface. NetApp EF-Series supports both NVMe over RoCE (NVMe/RoCE) and NVMe/InfiniBand.

The EF-Series, available with up to 1.8PB of raw solid-state drive (SSD) capacity, provides capacity and bullet-proof reliability to meet the requirements of the most demanding organizations. This technical report provides an overview of best practices for Microsoft SQL Server with NetApp EF-Series flash array.

1.1 Intended Audience

This technical report is intended for NetApp customers, partners, employees, and field personnel who are responsible for deploying a SQL Server database solution in a customer environment. It is assumed that the reader is familiar with the various components of the solution.

1.2 Caveats

This document assumes that the database is either being relocated to an EF-Series storage system or being created on an EF-Series storage system to achieve high performance.

This document also assumes that the user wants to improve performance of an OLTP application.

¹ [Top 10 SPC-2 by Price-Performance](#)

² [StorageReview.com](#)

2 Introduction to EF-Series

2.1 EF-Series Hardware Overview

The EF-Series flash array continues NetApp's long-standing heritage of delivering powerful solutions to meet specific business needs. With high IOPS and submillisecond response times, the EF-Series enables business-critical applications to deliver faster results and improve customer experience.

This combination of high IOPS and ultra-low latency makes an EF-Series flash array a great choice for database-driven applications requiring a dedicated extreme performance solution.

Table 1 provides an overview of the EF-Series flash array performance.

Table 1) EF-Series performance.

Performance Comparison	EF570	EF280
Sustained I/O rate: random disk reads (4KB) – DDP	950,000 @ 0.27ms 700,000 @ 0.18ms	300,000 IOPS @ 0.42ms 150,000 IOPS @ 0.24ms
Sustained I/O rate: random disk writes (4KB) – DDP	185,000 @ 0.17ms 100,000 @ 0.07ms	50,000 IOPS @ 0.32ms 40,000 IOPS @ 0.19ms
Sustained throughput: sequential disk reads (512KB) - DDP	21Gb/s	10Gb/s
Sustained throughput: sequential disk writes (cache mirror enabled 512KB) - DDP	9Gb/s	3.5Gb/s

The EF-Series flash array is a fully redundant 2U/24 system that supports up to four more 2U trays of SSDs for up to 1.8PB of raw capacity.

The EF-Series has dual active controllers, four 12Gb SAS disk expansion ports, and up to 64GB of battery-backed controller cache and runs on the performance-optimized and enterprise-proven NetApp SANtricity® software.

The EF-Series supports dual controller modules, power supplies, and fan units for hardware redundancy. The shelves are sized to support up to 24 2.5" disk drives in 2U of rack space.

Along with performance, the key to maximizing value is to maximize efficiency. Historically, companies have sacrificed efficiency to achieve extreme performance levels by overprovisioning their storage. But that is changing. The all-flash EF-Series helps customers balance performance and efficiency by eliminating overprovisioning, resulting in dramatically reduced costs.

With the performance of more than 1,000 traditional drives, a single EF-Series can meet extreme requirements with 95% less rack space, power, and cooling. This ability is a significant benefit to customers who are used to deploying partially filled disks to improve application performance.

In addition to cost efficiency, the EF-Series flash array provides application efficiency. By completing a higher volume of application operations, customers can become efficient and drive better results.

The EF-Series has a fully redundant I/O path that provides automated failover. Automated failover is surprisingly not in many of the flash products available today, but it is an absolute requirement for enterprises who want to implement this type of technology.

All management tasks are performed while the EF-Series remains online with complete read/write data access. Being online enables storage administrators to make configuration changes and conduct maintenance without disrupting application I/O.

The EF-Series flash array also offers advanced data protection common to enterprise storage to protect against data loss and downtime events, both locally with NetApp Snapshot™ technology and remotely with synchronous and asynchronous replication.

2.2 NetApp SANtricity OS

NetApp EF-Series flash array systems are managed by the SANtricity System Manager browser-based application. The SANtricity System Manager is embedded on the controller.

To create volume groups on the array, the first step is to assign a protection level during SANtricity configuration. This assignment is then applied to the disks selected to form the volume group. The EF-Series flash array storage systems support DDP and RAID levels 0, 1, 5, 6, and 10. DDP was used for all configurations described in this document.

To simplify the storage provisioning, NetApp SANtricity provides an automatic configuration feature. The configuration wizard analyzes the available disk capacity on the array. It then selects disks that maximize array performance and fault tolerance while meeting capacity requirements, hot spares, and any other criteria specified in the wizard.

For further information about the NetApp SANtricity® Unified Manager and the SANtricity System Manager, see the [E-Series Documentation Center](#).

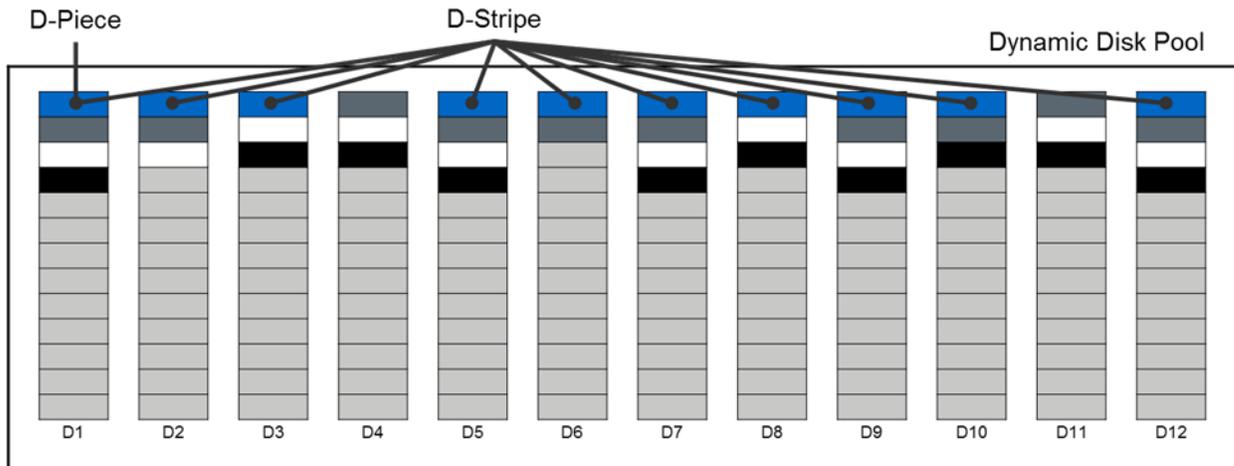
Dynamic Disk Pools

With seven patents pending, the DDP feature dynamically distributes data, spare capacity, and protection information across a pool of disk drives. These pools can range in size from a minimum of 11 drives to all the supported drives in a system. In addition to creating a single DDP, storage administrators can opt to mix traditional volume groups and DDP or even multiple DDPs, offering an unprecedented level of flexibility.

Dynamic Disk Pools are composed of several lower-level elements. The first of these elements is a D-piece. A D-piece consists of a contiguous 512MB section from a physical disk that contains 4,096 128KB segments. Within a pool, 10 D-pieces are selected using an intelligent optimization algorithm from selected drives within the pool. Together, the 10 associated D-pieces are considered a D-stripe, which is 4GB of usable capacity in size. Within the D-stripe, the contents are similar to a RAID 6 8+2 scenario. There, 8 of the underlying segments potentially contain user data, one segment contains parity (P) information calculated from the user data segments, and one segment contains the Q value as defined by RAID 6.

Volumes are then created from an aggregation of multiple 4GB D-stripes as required to satisfy the defined volume size up to the maximum allowable volume size within a DDP. Figure 1 shows the relationship between these data structures.

Figure 1) Dynamic Disk Pool components.

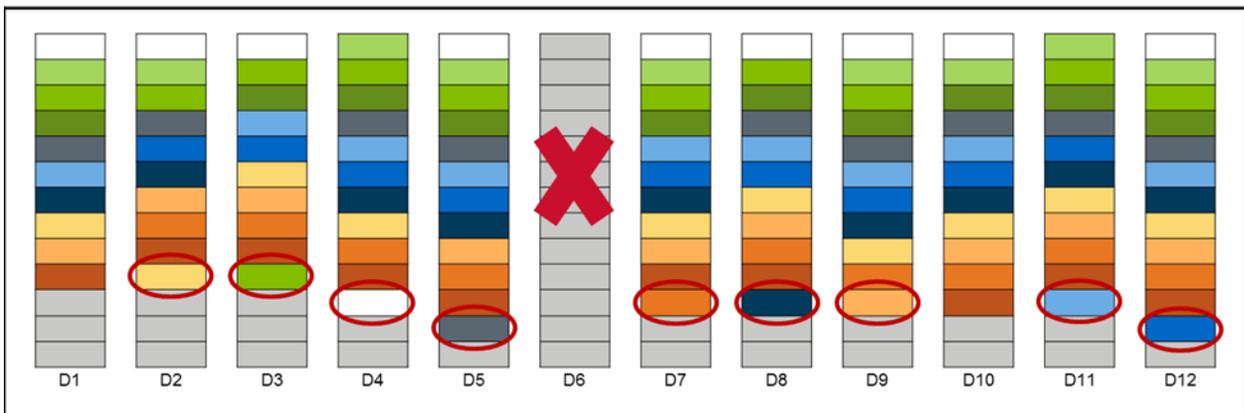


Another major benefit of a DDP is that, rather than using dedicated stranded hot spares, the pool contains integrated preservation capacity to provide rebuild locations for potential drive failures. This approach simplifies management, because individual hot spares no longer need to be planned or managed. The approach also greatly improves the time for rebuilds, if necessary, and enhances volume performance during a rebuild, as opposed to traditional hot spares.

When a drive in a DDP fails, the D-pieces from the failed drive are reconstructed to potentially all other drives in the pool using the same mechanism normally used by RAID 6. During this process, an algorithm internal to the controller framework verifies that no single drive contains two D-pieces from the same D-stripe. The individual D-pieces are reconstructed at the lowest available Logical Block Access (LBA) range on the selected disk drive.

In Figure 2, disk drive 6 (D6) is shown to have failed. Next, the D-pieces that previously resided on that disk are re-created simultaneously across several other drives in the pool. Because there are multiple disks participating in the effort, the overall performance effect of this situation is lessened, and the length of time needed to complete the operation is dramatically reduced.

Figure 2) Dynamic Disk Pool drive failure.



When multiple disk failures occur within a DDP, priority for reconstruction is given to any D-stripes missing two D-pieces to minimize data availability risk. After those critically affected D-stripes are reconstructed, the remainder of the necessary data is reconstructed.

From a controller resource allocation perspective, there are two user-modifiable reconstruction priorities within DDP:

- Degraded reconstruction priority is assigned to instances in which only a single D-piece must be rebuilt for the affected D-stripes; the default for this value is high.
- Critical reconstruction priority is assigned to instances in which a D-stripe has two missing D-pieces that need to be rebuilt; the default for this value is highest.

For large disk pools with two simultaneous disk failures, only a relatively small number of D-stripes are likely to encounter the critical situation in which two D-pieces must be reconstructed. As discussed previously, these critical D-pieces are identified and reconstructed initially at the highest priority. Doing so returns the DDP to a degraded state quickly so that further drive failures can be tolerated.

In addition to improving rebuild times and providing superior data protection, DDP can also greatly improve the performance of the base volume when under a failure condition compared to the performance of traditional volume groups.

For more information about DDP, see [TR-4652: SANtricity Dynamic Disk Pools - Feature Description and Best Practices](#).

2.3 Provisioning the EF-Series

SANtricity DDPs allow storage administrators to simplify RAID management, improve data protection, and maintain predictable performance under all conditions. DDP evenly distributes data, protection information, and spare capacity across the EF-Series entire pool of drives, simplifying setup and maximizing utilization. Its next-generation technology minimizes the performance effect of a drive failure and can return the system to optimal condition up to eight times more quickly than traditional RAID. With shorter rebuild times and patented prioritize reconstruction technology, DDP significantly reduces exposure to multiple disk failures, offering a level of data protection that simply can't be achieved with traditional RAID.

With SANtricity software, all management tasks can be performed while the storage remains online with complete read/write data access. This allows storage administrators to make configuration changes, conduct maintenance, or expand the storage capacity without disrupting I/O to attached hosts. SANtricity software's online capabilities include:

- Dynamic volume expansion allows administrators to expand the capacity of an existing volume.
- Dynamic segment size migration enables administrators to change the segment size of a given volume.
- Dynamic RAID-level migration changes the RAID level of a RAID group on the existing drives without requiring the relocation of data. Supported RAID levels are 0, 1, 3, 5, 6, and 10.
- Nondisruptive controller firmware upgrades (no interruption to data access) are supported

For the detailed procedures on provisioning EF-Series, see [Deploy SQL Server on EF-Series](#). If there is repeated volume creation and deployment, REST API commands can be used to automate this task. The example of creating volume, mount, initialize, and format disk using REST API is located in Appendix A.

3 SQL Server Overview

The SQL Server database platform can support a wide variety of 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 SLA (SLAs). SQL Server workload can be described as follows:

- OLTP (OLTP) databases are often also the most critical databases in an organization. These databases back customer-facing applications and are considered 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 clustering with Windows failover cluster or Always On availability groups. The I/O mix of these types of databases is characterized with 75% to 90% random read and 25% to 10% write.
- Decision support system (DSS) databases can be also referred to as data warehouses. These data warehouses are mission critical in 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 critical during month, quarter, and year end. This workload typically has a 100% read I/O mix.

3.1 SQL Server I/O Overview

SQL Server is sensitive to I/O latency issues due to the concurrent transactional nature of the SQL Server engine. SQL Server is built on a complicated system of row, page, extent, and table locks that provides transactional consistency throughout the SQL Server system. A poor I/O structure (for example, when I/O takes too long to respond) causes resources to be held longer than necessary, resulting in blocking within the system. When blocking occurs, it is typically not obvious that the I/O subsystem is the root cause.

- SQL Server reads. When reading data from SQL Server, the client first goes to the buffer cache. If the data is not in the buffer cache, SQL Server goes to the I/O subsystem to retrieve the data. The statement does not complete until 100% of the data is read; the user connection or process remains in an I/O wait state until completion.
- SQL Server writes. The user writes to the transaction log and the buffer cache. If the data to be modified is not already in the buffer cache, then it must be read into the buffer cache from the I/O subsystem. The buffer manager enables the transaction log to be written to first, before changes are written to the database. This is known as write-ahead logging (WAL). When the user makes the change and the commit is executed, a log write is displayed about the change that took place, allowing the commit to complete. After the commit is complete, the user process can continue to the next stage or command without waiting for the changes to be written to the disk. Rollback transaction follows the same process as the commit, but in reverse. The buffer manager moves the data from the cache to the disk. It tracks log sequence numbers (LSNs) for each log record.
- Transaction log. The SQL Server transaction log is a write-intensive operation that is sequential in nature. If the database or instance fails, the transaction log is used to provide recoverability of data.

The OLTP database system within the SQL Server environment is most dependent on getting the greatest number of transactions through the system in the least amount of time. Examples of different types of OLTP systems include web order systems and manufacturing tracking systems. OLTP systems can have large volumes of transactions per second (TPS), and for the OLTP system, it is all about throughput. For these transactions to take place, SQL Server relies on an efficient I/O subsystem. Based on a [Microsoft SQL Server best practices article](#), an OLTP transaction profile has the following attributes:

- OLTP processing is random in nature for both reads and writes issued against data files.
- I/O activity is approximately 80% read and 20% write.
- Usually, read activity is consistent and uses point queries; it does not consist of large time-consuming queries.
- Write activity to the data files occurs during checkpoint operations (frequency is determined by recovery interval settings).
- Log writes are sequential in nature with a varying size that depends on the nature of the workload (sector aligned up to 60KB).
- Log reads are sequential in nature (sector aligned up to 120KB).

3.2 Windows Volume Mount Points

NetApp storage solutions and Microsoft SQL Server 2005 and later support mount points. Mount points are directories on a volume that can be used to mount a different volume. Mounted volumes can be accessed by referencing the path of the mount point. Mount points eliminate the Windows 26-drive-letter limit and offer greater application transparency when moving data between LUNs, moving LUNs between hosts, and unmounting and mounting LUNs on the same host. This is because the underlying volumes can be moved around without changing the mount point path name.

Note: NetApp recommends using NTFS mount points instead of drive letters to surpass the 26-drive-letter limitation in Windows. When using volume mount points, the name given to the volume label and mount point must be the same.

3.3 Database Files and File Groups

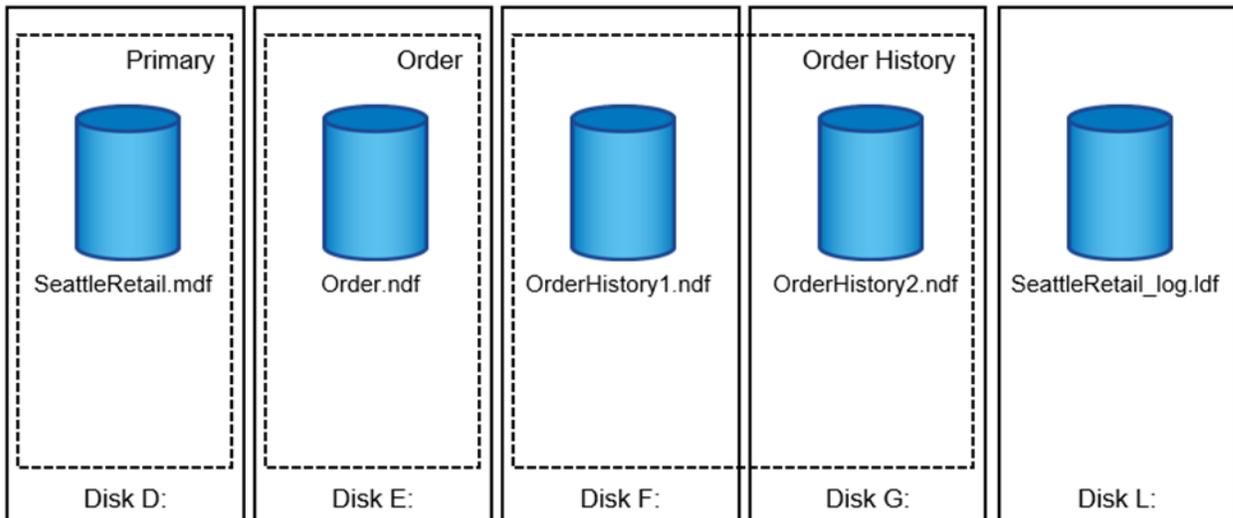
A SQL Server database is a collection of objects that allow one to store and manipulate data. In theory, SQL Server (64-bit) supports 32,767 databases per instance and 524,272TB of database size, although the typical installation usually has a few databases. However, the number of the databases SQL Server can handle depends on the load and hardware. It is not unusual to see SQL Server instances hosting dozens or even hundreds or thousands of small databases.

Each database consists of one or more data files and one or more transaction log files. The transaction log stores the information about database transactions and all data modifications made by each session. Every time the data is modified, SQL Server stores enough information in the transaction log to undo (roll back) or redo (replay) the action. SQL Server's transaction log is a part of SQL Server's reputation for data integrity and robustness. The transaction log is vital to the atomicity, consistency, isolation, and durability (ACID) capabilities of SQL Server. SQL Server writes to the transaction log when any change to the data page happens. Every Data Manipulation Language (DML) statement (for example, select, insert, update, or delete) is a complete transaction, and the transaction log confirms that the entire set-based operation takes place, enabling the atomicity of the transaction.

Each database has one primary data file, which, by default, has the .mdf extension. In addition, each database can have secondary database files. Those files, by default, have .ndf extensions. See Figure 3.

All database files are grouped into file groups. A file group is the logical unit, which simplifies database administration. They allow the separation between logical object placement and physical database files. When you create the database objects tables, you specify in what file group they should be placed without worrying about the underlying data file configuration.

Figure 3) Example of primary, secondary, and log files that locate in each drive.



The ability to put multiple data files inside the file group allows us to spread the load across different storage devices, which helps to improve I/O performance of the system. The transaction log, in contrast, does not benefit from the multiple files because SQL Server writes to the transaction log in a sequential manner.

The separation between logical object placement in the file groups and physical database files allows us to fine-tune the database file layout, getting the most from the storage subsystem. For example, independent software vendors (ISVs) who are deploying their products to different customers can adjust the number of database files based on underlying I/O configuration and expected amount of data during the deployment stage. Those changes are transparent to the application developers, who are placing the database objects in the file groups rather than database files.

NetApp recommends not using the primary file group for anything but system objects. Creating a separate file group or set of file groups for the user objects simplifies database administration and disaster recovery, especially in cases of large databases.

You can specify initial file size and autogrowth parameters at the time when you create the database or add new files to an existing database. SQL Server uses a proportional fill algorithm when choosing into what data file it should write data. It writes an amount of data proportionally to the free space available in the files. The more free space in the file, the more writes it handles.

Note: NetApp recommends that all files in the single file group have the same initial size and autogrowth parameters, with grow size defined in megabytes rather than percentages. This helps the proportional fill algorithm evenly balance write activities across data files.

3.4 SQL Server Database Files

Provisioning database files on EF-Series flash array can be done in two ways:

1. To create a database that has database files residing on an EF-Series LUN, you can use the following T-SQL script during database creation:

```
-- Assuming C:\MSSQL\Data and C:\MSSQL\Log is the mount points of EF-Series LUNs
USE master;
GO
CREATE DATABASE Sales
ON
( NAME = Sales_dat,
  FILENAME = 'C:\MSSQL\Data\saledat.mdf',
  SIZE = 10,
  MAXSIZE = 50,
  FILEGROWTH = 5 )
LOG ON
( NAME = Sales_log,
  FILENAME = 'C:\MSSQL\Log\salelog.ldf',
  SIZE = 5MB,
  MAXSIZE = 25MB,
  FILEGROWTH = 5MB ) ;
GO
```

2. To move the database files from non-EF-Series LUNs to EF-Series LUNs, you must detach the database files. After you detach the database files, you can copy the files to the path or mount points that reside on the EF-Series LUNs. After the files are copied, you can attach the database files at the new location.

The common best practice is to separate data, transaction logs, and tempdb files in separate logical LUNs. The origin of this recommendation lies with the separation of the types of workload between different physical storages. This is still a valid recommendation for environments where you can have that separation. However, it is common to see customers deploying SQL Server in a shared storage environment, where physical separation is much harder to achieve and usually is not even necessary for performance reasons.

It is still a good idea to maintain separation to help with manageability so that potential problems are easier to isolate. For example, separating tempdb onto its own logical disk means that you can pre-size it to fill the disk without worrying about space requirements for other files. The more separation you implement, the easier it is to correlate logical disk performance to specific database files.

3.5 Database Files for Tempdb

The tempdb system database is a global resource that is available to all users connected to the SQL Server instance, and it is used to hold the following:

- Temporary user objects that are explicitly created, such as global or local temporary tables, temporary stored procedures, table variables, or cursors.
- Internal objects that are created by the SQL Server Database Engine, such as work tables to store intermediate results for spools or sorting.
- Row versions that are generated by data modification transactions in a database that uses read-committed row versioning isolation or Snapshot copy isolation transactions.
- Row versions that are generated by data modification transactions for features, such as online index operations, multiple active result sets (MARS), and AFTER triggers.

Operations within tempdb are logged, enabling transactions to be rolled back. Tempdb is re-created every time SQL Server is started so that the system starts with a clean copy of the tempdb database. Temporary tables and stored procedures are dropped automatically on disconnect, and no connections are active when the system is shut down. Therefore, there is nothing in tempdb to be saved from one session of SQL Server to another. Backup and restore operations are not allowed on tempdb.

Because every SQL Server instance has a shared database named tempdb that is used by temporary objects, and there is only one tempdb per instance, it often causes a bottleneck for systems that use it heavily. Typically, this bottleneck happens because of PAGELATCH, an in-memory latch contention on the allocation bitmap pages inside the data files.

It is possible to reduce the contention on the in-memory pages by adding extra data files to tempdb with the same initial size and autogrowth configuration. This works because SQL Server uses a round-robin, proportional-fill algorithm to stripe the writes across the data files. When multiple data files exist for a database, all of the writes to the files are striped to those files, with the writes to any particular file based on the proportion of free space that the file has compared to the total free space across all of the files. This means that writes are proportionally distributed to the files according to their free space, regardless of their size, to make sure that they fill at the same time.

Microsoft recommends up to a one-to-one mapping between the number of files and logical CPUs. During testing of massive workloads, Microsoft has seen performance benefits, even with hundreds of data files.

A more pragmatic approach, however, is to have a one-to-one mapping between files and logical CPUs up to eight, and then add files if you continue to see allocation contention or if you must push the I/O subsystem harder.

Because the NetApp EF-Series flash array provides a robust platform for delivering exceptional performance, you can create multiple tempdb files and place them on the EF-Series flash array. You can accomplish this by using the following Transact SQL Server script.

```
select *
from sys.database_files

use master
go
-- Change logical tempdb file name first since SQL Server shipped with logical file name called
tempdev
alter database tempdb modify file (name = 'tempdev', newname = 'tempdev01');

-- Change location of tempdev01 and log file to C:\MSSQL\Tempdb path
```

```

alter database tempdb modify file (name = 'tempdev01', filename =
'C:\MSSQL\Tempdb\tempdev01.ndf');
alter database tempdb modify file (name = 'templog', filename = 'C:\MSSQL\Tempdb\templog.ldf');

-- Assign proper size for tempdev01
ALTER DATABASE [tempdb] MODIFY FILE ( NAME = N'tempdev01', SIZE = 2GB, FILEGROWTH = 100 MB );
ALTER DATABASE [tempdb] MODIFY FILE ( NAME = N'templog', SIZE = 4GB, FILEGROWTH = 100 MB );

-- Add more tempdb files
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev02', FILENAME =
N'C:\MSSQL\Tempdb\tempdev02.ndf' , SIZE = 2GB , FILEGROWTH = 100 MB);
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev03', FILENAME =
N'C:\MSSQL\Tempdb\tempdev03.ndf' , SIZE = 2GB , FILEGROWTH = 100 MB);
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev04', FILENAME =
N'C:\MSSQL\Tempdb\tempdev04.ndf' , SIZE = 2GB , FILEGROWTH = 100 MB);
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev05', FILENAME =
N'C:\MSSQL\Tempdb\tempdev05.ndf' , SIZE = 2GB , FILEGROWTH = 100 MB);
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev06', FILENAME =
N'C:\MSSQL\Tempdb\tempdev06.ndf' , SIZE = 2GB , FILEGROWTH = 100 MB);
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev07', FILENAME =
N'C:\MSSQL\Tempdb\tempdev07.ndf' , SIZE = 2GB , FILEGROWTH = 100 MB);
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdev08', FILENAME =
N'C:\MSSQL\Tempdb\tempdev08.ndf' , SIZE = 2GB , FILEGROWTH = 100 MB);

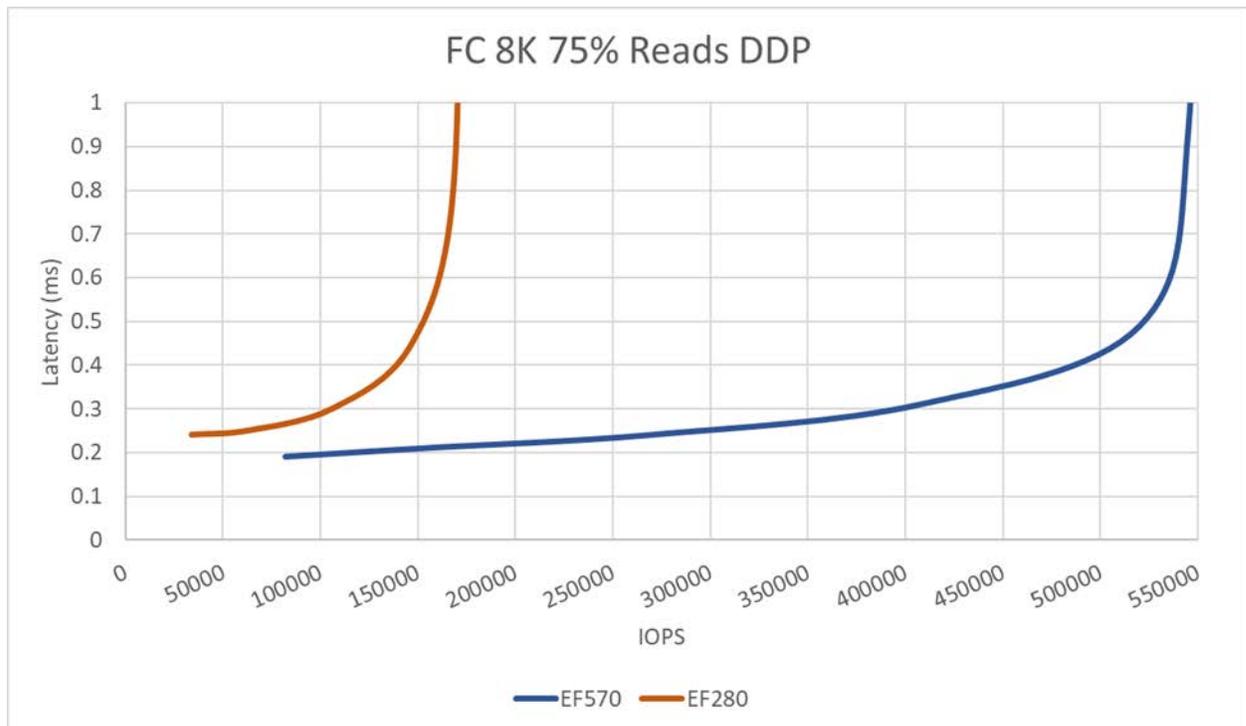
```

4 EF-Series Performance

Figure 4 shows EF-Series performance results for both the entry-level EF280 and mid-range EF570 configured for a typical OLTP workload. Testing was done with the array configured to use 24 solid-state drive (SSD) in a DDP, FC host protocol, and a workload of 75% read and 25% write.

At 500 μ s of latency the EF280 produces 150,000 I/Os per second, while the EF570 is producing nearly 525,000 IOPS.

Figure 4) Performance results for OLTP configuration.



5 High Availability

5.1 EF-Series and SANtricity OS

The NetApp EF-Series storage system has been architected for high reliability and HA with features such as:

- Dual active controller with automated I/O path failover
- RAID levels 0, 1, 5, 6, and 10 or DDP
- Redundant, hot-swappable controllers, disk drives, power supplies, and fans
- Automatic drive failover detection and rebuild using global hot spares
- Mirrored data cache with battery backup and destage to memory
- Nondisruptive controller firmware upgrades
- Proactive drive health monitoring
- Background media scan with autoparity check and correction

All components are fully redundant and might be swapped without powering off the system or even halting operation. This includes controllers, disk drives, power supplies, and fans. The EF-Series power supplies offer an 80-plus efficiency rating. The EF-Series flash array features several functions designed to protect data in every circumstance. Multiple RAID levels are available for use with varying levels of redundancy. Fail over from one path to another in the case of a lost connection is also automatically included with the system. Within the shelf, each drive has a connection to each controller so that even internal connection issues can be quickly overcome. Volumes on the system are available for host I/O from the moment they are created and can even have significant properties altered without stopping I/O.

Other features of the EF-Series flash array that protect data include mirroring and backing up controller cache. If power is lost to the system during operation, onboard batteries are used to destage the data from cache memory to internal controller flash so that it is available when power is restored. The RAID algorithms allow the system to re-create any lost data in the rare case of drive failure. You can also confirm data with RAID parity and even continue a rebuild if you hit an unreadable sector.

Behind the scenes, the system performs other tasks that protect data always. The optional media scan feature looks for inconsistencies even on sectors not currently accessed by any host. The EF-Series proactively tracks SSD wear and flags drives that are approaching the end of their expected life. All types of diagnostic data is constantly collected for use later by support, if necessary.

Not only does the EF-Series offer the detailed reliability and availability features already described, but by using the SANtricity software, it is also possible to maximize availability. For example, SANtricity software:

- Enables high-speed, high-efficiency Snapshot technology
- Protects data in seconds
- Reduces flash consumption by storing only changed blocks
- Provides robust disaster recovery (DR) protection
- Supports synchronous mirroring for no-data-loss protection of content
- Supports asynchronous mirroring for long-distance protection and compliance
- Maximized return on investment (ROI) with flexible protection
- Supports flash, near-line SAS (NL-SAS), or a mix of recovery targets based on cost and performance needs
- Delivers speed without breaking budgets

5.2 SQL Server HA Options

The high availability of SQL Server solutions masks the effects of a hardware or software failures and maintains the availability of applications to minimize the perceived downtime for users. SQL Server provides several HA solutions.

Log Shipping

Log shipping operates at the database level. It can maintain one or more warm standby databases (referred to as secondary databases) for a single production database that is referred to as the primary database. For more information about log shipping, refer to [About Log Shipping \(SQL Server\)](#).

Database Mirroring

Database mirroring increases database availability by supporting almost instantaneous failover. Database mirroring can be used to maintain a single standby database or mirror database for a corresponding production database that is referred to as the principal database. For more information, refer to [Database Mirroring \(SQL Server\)](#).

AlwaysOn Failover Cluster Instances

AlwaysOn Failover Cluster Instances applies Windows Server Failover Clustering (WSFC) functionality to provide local HA through redundancy at the server instance level: a failover cluster instance (FCI). An FCI is a single SQL Server instance that is installed across WSFC nodes and possibly across multiple subnets. On the network, an FCI appears to be an SQL Server instance running on a single computer, but the FCI provides failover from one WSFC node to another if the current node becomes unavailable. For more information, refer to [AlwaysOn Failover Cluster Instances \(SQL Server\)](#).

AlwaysOn Availability Groups

AlwaysOn Availability Groups is an enterprise-level high-availability and disaster recovery solution introduced in SQL Server 2012 to enable you to maximize availability for one or more user databases. AlwaysOn Availability Groups require that SQL Server instances reside on WSFC nodes. For more information, refer to [AlwaysOn Availability Groups \(SQL Server\)](#).

AlwaysOn Availability Groups support two availability modes, asynchronous-commit mode, and synchronous-commit mode:

- Asynchronous-commit mode is a disaster recovery solution that works well when the availability replicas are distributed over considerable distances. For more information, refer to [Asynchronous-Commit Availability Mode](#).
- Synchronous-commit mode emphasizes HA over performance, at the cost of increased transaction latency. In synchronous-commit mode, transactions wait to send the transaction confirmation to the client until the secondary replica has hardened the log to disk. For more information, refer to [Synchronous-Commit Availability Mode](#).

Automatic failover provides high availability by making sure that the database is quickly made available again after the loss of the primary replica. To configure an availability group for automatic failover, you must set both the current primary replica and one secondary replica to synchronous-commit mode with automatic failover.

6 Sizing

SQL Server performance has been centered on I/O. Traditionally, this performance was improved by either increasing the number of spindles or making the spindles go more quickly. With the advent of the EF-Series flash array, performance improvement can be achieved by using SSD drives.

6.1 EF-Series I/O Overview

There are several factors that can affect the overall performance of an EF-Series storage system, including physical components, such as networking infrastructure, and the configuration of the underlying storage itself. Generically, storage system performance tuning can be defined as following a 40/30/30 rule, where 40% of the tuning and configuration is at the storage system level, 30% is at the file system level, and the final 30% is at the application level. The following sections describe the 40% related to storage system specifics. At a high level, some of the considerations at the file system and application level include:

- I/O size. EF-Series storage systems are largely responsive systems that require a host to request an I/O operation to complete that operation. The I/O size of the individual requests from the host can have a significant effect on either the number of IOPS or throughput (described in terms of megabytes per second (MBps) or gigabytes per second (GBps)). Larger I/Os typically lead to lower numbers of IOPS and larger MBps, and the opposite is true as well. This relationship is defined with the equation $\text{Throughput} = \text{IOPS} \times \text{I/O size}$.
- Read versus write requests. In addition to the I/O size, the percentage of read versus write I/O requests processed at the storage system level also has a potential effect on the storage system and should be considered when designing a solution.
- Sequentiality or randomness of the data stream. Whether the host requests to the underlying disk media logical block addresses (LBAs) are sequential or random has a significant effect on performance at the storage system level. The sequence affects the ability of physical media to respond effectively to the request with minimal latency and the effectiveness of the storage system's caching algorithms. An exception to increased latency of random requests is for SSDs, which do not have mechanically invoked latency.
- Number of concurrent I/O operations. The number of outstanding I/O operations applied to a given volume can vary based on several factors, including whether the file system uses raw, buffered, or direct I/O. Generally, most volumes in an EF-Series storage system are striped across several drives. Providing a minimal amount of outstanding I/O to each individual disk can cause underutilization of the resources in the storage system, resulting in less than desired performance characteristics.

For those SQL Server customers new to NetApp EF-Series, it might be helpful to review the differences between RAID 10, RAID 5, RAID 6, and DDP technology. Table 2 compares the usable capacity for differing RAID levels. For completeness, all RAID levels supported by NetApp EF-Series are shown.

Table 2) Comparison of usable capacity for differing RAID levels.

Desired Feature	RAID 0	RAID 10	RAID 5	RAID 6	DDP
Usable capacity	100%	50%	$(N-1) \div N$ where N is the selected drive count in the volume group	$(N-2) \div N$ where N is the selected drive count in the volume group	80% minus selected preservation capacity

6.2 Estimating I/O

Estimating the number of I/O operations required for a system is crucial when sizing a database. This exercise helps the administrator understand how to keep the database instance performing within acceptable limits. You must estimate I/O when you are unable to get the actual physical I/O numbers for the system. This is typically the case in new systems that are in the process of being constructed. The following sections provide formulas for estimating I/O.

New OLTP Database System

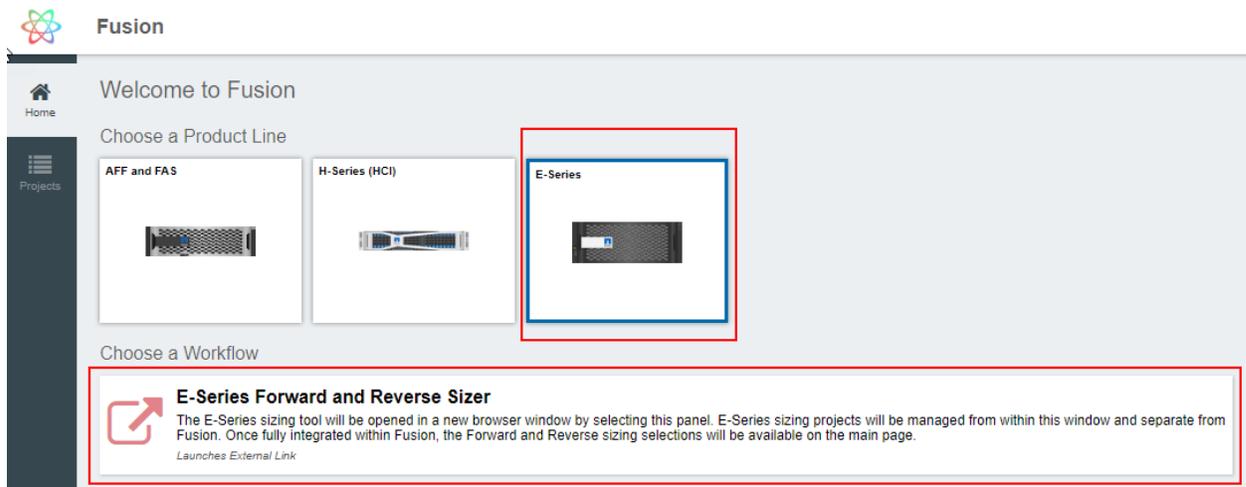
The following items must be considered:

- Business transaction
- Duration of business transaction
- Acceptable latency for the application
- Approximate ratio of transaction and system transactions (I/O)

The easiest way to estimate sizing for SQL Server with EF is by following this example:

1. Estimate the number of business transactions. For example, estimate the number of business transactions as 600,000,000 per day. The business runs 24/7.
2. Assume that one business transaction creates 25 I/O operations. Therefore, for our example, the I/O operations per day are $600,000,000 \times 25 = 15,000,000,000$ I/O/day.
3. To use standard storage sizing tools, we need IOPS information. From the example, this is $(15,000,000,000) \div 86,400 = 173,612$ IOPS.
4. Acceptable latency for the application should be considered. For example, the application might want less than 1 millisecond of latency.
5. Input this number to the [NetApp Fusion](#) tool and use the E-Series Forward and Reverse Sizer, see Figure 5.

Figure 5) Step to select E-Series Forward and reverse Sizer.



6. To use the sizer, you must also estimate the percentage of reads and writes. 80% reads and 20% writes is typical of an OLTP database.

Existing OLTP Database System

When sizing for an existing database environment, understanding the type of workload and interpreting the statistical data is helpful. It is important to gather statistics during periods of peak stress on the system. PerfMon displays the high-water marks for the time frame in which you monitor the system.

After either IOPS or throughput (MBps) of the system is captured, it can be entered into the E-Series Performance Sizing tool. Go to [NetApp Fusion](#) and select the E-Series product line for an E-Series sizing. Figure 6 shows the input fields for a reverse sizing. From our example, 173,612 IOPS was required. As shown in Figure 7, one EF280 with 24 SSD drives can provide over 184,000 IOPS with under 1ms of latency.

Figure 6) E-Series reverse sizer input fields.

Figure 7) Reverse sizing output.

System Count	System Shelf Count	System Drive Count	System Rack Space (U)
1	1	24 (Conf) 24 (Actual)	2

Commodity	Description	Part Number	Quantity
Controller	EF280	E2800A-8GB	2
Shelf	DE224C	E-X5721A-QS	1
Disk	1.6TB, SSD	E-X4095A	24
Interface	12Gb; SAS	X-56027-00-0E-C	2

System Capacity (TIB)	Usable Capacity (TIB)	Reserve Capacity (TIB)	Spare Capacity (TIB)*
34.93	24.98	2.91	0.00

System Throughput (MB/sec)	System IOPS (Conf)	System IOPS (Actual)	System Latency (ms)
1439	184227	184227	0.70

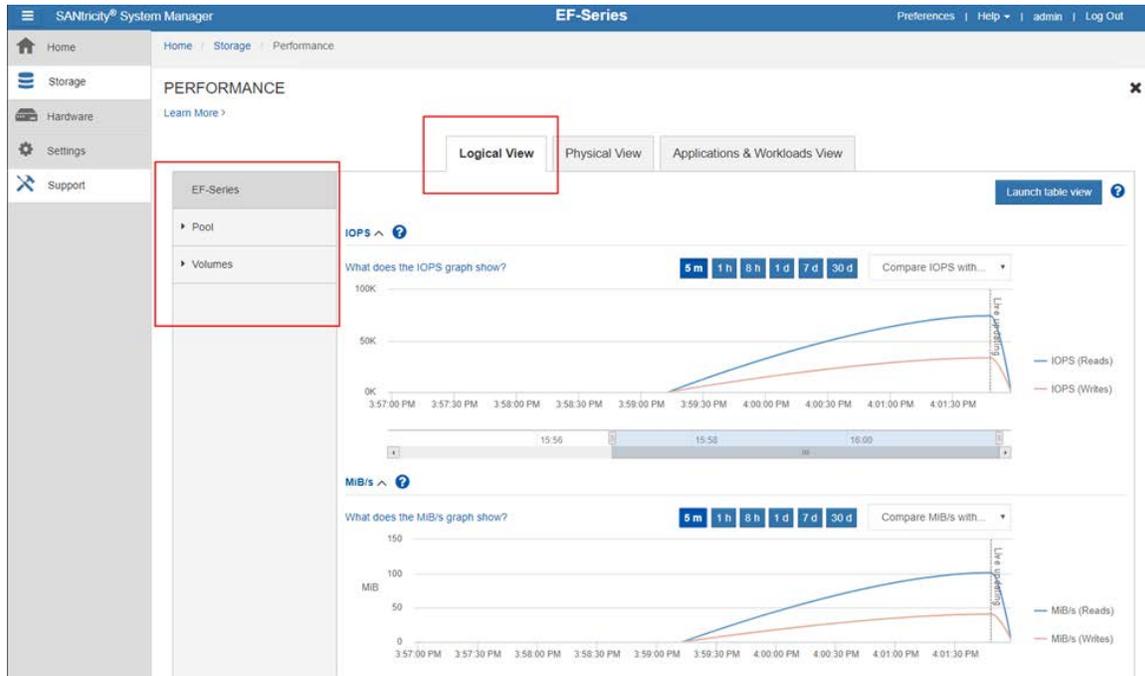
7 EF-Series Performance Monitoring Using NetApp SANtricity System Manager

During storage system operations, it can be useful to monitor storage system performance. Using SANtricity System Manager (SSM), you can view EF-Series performance data in both textual and graphical dashboard formats.

SSM provides exceptional performance monitoring enabling users to capture logical, physical, and application level workloads. It adds new functionality, including application and workload tagging, enhanced performance data, embedded monitor, and a graphical view of the volume usage. The Performance Monitor can be observed by selecting View Performance Details in SSM. You can observe the performance of three main categories:

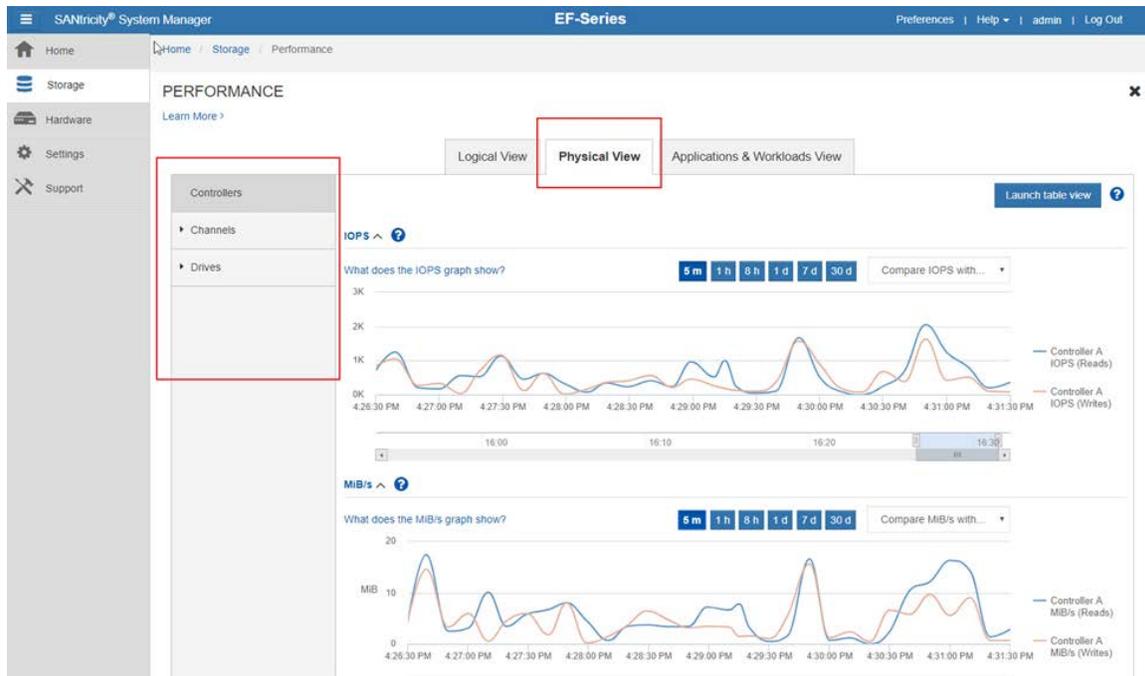
1. The Logical View enables you to filter information based on pools and volume groups or volumes, see Figure 8.

Figure 8) Example of Performance Monitor with logical view.



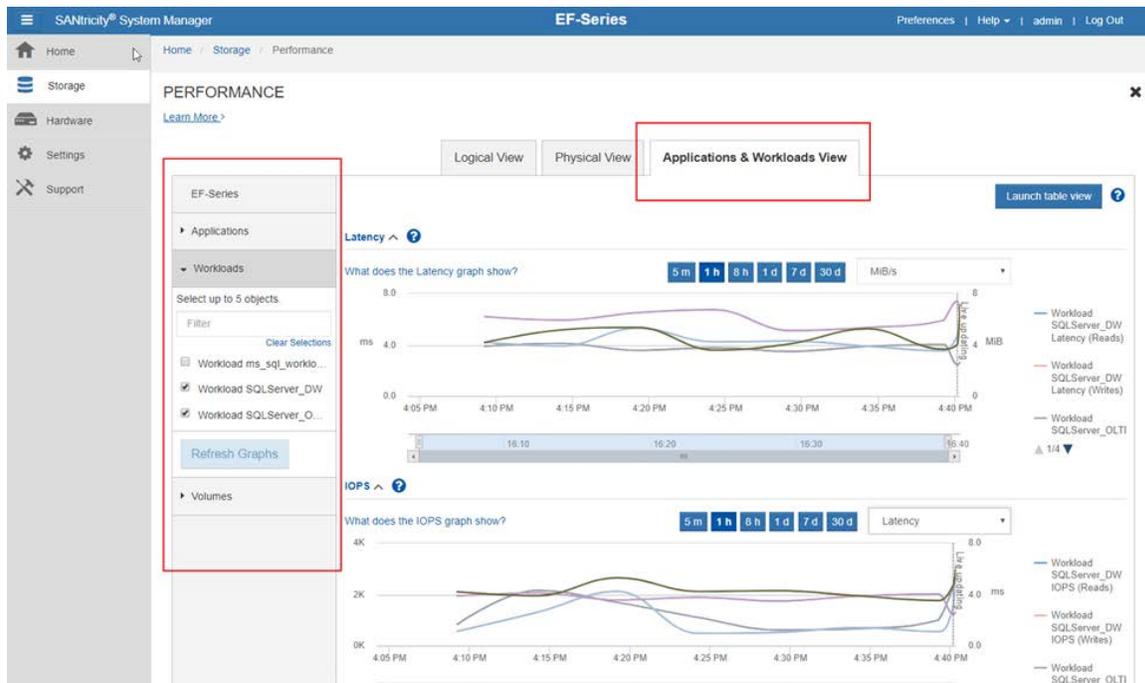
2. The Physical View enables you to filter information based on host channels and drives, see Figure 9.

Figure 9) Example of Performance Monitor with physical view.



3. The Application and Workload View enables you to filter information based on applications, workloads, and volumes within the workloads, see Figure 10.

Figure 10) Example of Performance Monitor with applications and workloads view.



For more detail on Performance Monitor, see the video [Monitor SQL Server Performance with EF-Series](#).

8 Summary

The NetApp EF-Series flash array is a market leader in delivering high performance, consistent low latency, and advanced high-availability features. The importance of this varies depending on business needs.

The EF-Series array is easy to provision with the embedded System Manager as shown in the video [Deploy SQL Server on EF-Series](#). It can also be provisioned using the REST API for users who have many systems to prepare, see Appendix A.

The EF-Series Flash array also has robust built-in monitoring capabilities that enable users to troubleshoot performance issues at the logical level, physical level, and application level, as shown in the video [Monitor SQL Server Performance with EF-Series](#).

In addition to solving the extreme latency requirements, the EF-Series also resolves SQL Server database challenges in the following ways:

- Dramatically boosts the performance of existing applications and lowers the cost per IOPS ratio without rearchitecting the application
- Increases SQL Server performance with RAID10 and DDPs
- Increases user productivity with better response times, improving business efficiency
- SQL Server AlwaysOn Availability Groups provide an alternate solution for HA and DR

You can find more technical information about EF-Series Flash Storage Arrays [here](#).

Though the EF-Series array can be configured in many different ways, the best practice is to configure a single DDP with all the SSDs participating in the pool. Volumes are carved out of this pool regardless of the purpose of the volume (such as a log or a database). Single DDP results in extreme performance at either an entry-level or at an enterprise level for an 80/20 read/write 8K workload. Only in extreme cases

does the EF-Series flash array need to be configured in a RAID 10 where IOPS increases and latency decreases. The objective should be to keep it as simple and fast as possible.

Appendix A: Provisioning EF-Series with REST API

```
# EF-Series System Manager RESTAPI Volume Provisioning

# =====
# Globals
# =====

# bypass untrusted certificates -- only do this for sources you trust!
add-type @"
    using System.Net;
    using System.Security.Cryptography.X509Certificates;
    public class TrustAllCertsPolicy : ICertificatePolicy {
        public bool CheckValidationResult(
            ServicePoint srvPoint, X509Certificate certificate,
            WebRequest request, int certificateProblem) {
            return true;
        }
    }
"@
[System.Net.ServicePointManager]::CertificatePolicy = New-Object TrustAllCertsPolicy
[Net.ServicePointManager]::SecurityProtocol = "Tls12, Tls11, Tls, Ssl3"

# IPs and credentials
$EFSeriesC2A = "https://10.251.229.52:8443/devmgr/"
$username = "admin"
$password = "Infini1"

# persist web session after logging in
$session = New-Object Microsoft.PowerShell.Commands.WebRequestSession

# headers
$headers = @{}
$headers.Add('accept','application/json')
$headers.Add('Content-Type','application/json')

# misc
$storageSystemId = "1"
```

```
# =====
# Initiate Session
# =====
# as alternative, can also store creds as plaintext in url:
# Invoke-RestMethod -Method Get -Uri
($EFSeriesC2A+"devmgr/utils/login?uid=admin&pwd=myPassword123&xsr=false&onlycheck=false") -
Headers $headers

Function LoginAPI ($username, $pass) {
    $credentialsJSON = '
    {
        "userId": "' + $username + '",
        "password": "' + $pass + '"
    },'

    $restParams = @{
        Method      = 'Post'
        Uri         = $EFSeriesC2A+"utils/login"
        Headers     = $headers
        Body        = $credentialsJSON
        WebSession  = $session
    }
}
```

```

    Invoke-RestMethod @restParams
}
# LoginAPI $username $pass

```

```

# =====
# GET Functions
# =====
# Ex endpoints:
# Invoke-RestMethod -Method Get -Uri ($EFSeriesC2A+"devmgr/utils/about") -Headers $headers
# Invoke-RestMethod -Method Get -Uri ($EFSeriesC2A+"devmgr/utils/buildinfo") -Headers $headers

Function GetStorageSystems {
    $restParams = @{
        Method      = 'Get'
        Uri         = $EFSeriesC2A+"v2/storage-systems"
        Headers     = $headers
        WebSession  = $session
    }

    try {
        Invoke-RestMethod @restParams
    }
    catch {
        CatchException $restParams
    }
}
# GetStorageSystems

Function GetEvents {
    $restParams = @{
        Method      = 'Get'
        Uri         = $EFSeriesC2A+"v2/events"
        Headers     = $headers
        WebSession  = $session
    }

    try {
        Invoke-RestMethod @restParams
    }
    catch {
        CatchException $restParams
    }
}
# GetEvents

Function GetVolumes ($getIdsOnly) {
    $restParams = @{
        Method      = 'Get'
        Uri         = $EFSeriesC2A+"v2/storage-systems/"+$storageSystemId+"/volumes"
        Headers     = $headers
        WebSession  = $session
    }

    try {
        if ($getIdsOnly -eq $true) {
            (Invoke-RestMethod @restParams).id
        }
        else {
            (Invoke-RestMethod @restParams)
        }
    }
    catch {
        CatchException $restParams
    }
}

```

```

}
# GetVolumes $true # '$false' returns all volume details

Function GetControllerId {
    $restParams = @{
        Method      = 'Get'
        Uri         = $EFSeriesC2A+"v2/storage-systems/"+$storageSystemId+"/controllers"
        Headers     = $headers
        WebSession  = $session
    }

    try {
        (Invoke-RestMethod @restParams).controllerRef[0]
    }
    catch {
        CatchException $restParams
    }
}

```

```

        else {
            (Invoke-RestMethod @restParams)
        }
    }
    catch {
        CatchException $restParams
    }
}
# GetDriveIds $true # '$false' returns all drive details

Function GetSnapshots {
    $restParams = @{
        Method      = 'Get'
        Uri         = $EFSeriesC2A+"v2/storage-systems/"+$storageSystemId+"/snapshot-images"
        Headers     = $headers
        WebSession  = $session
    }

    try {
        (Invoke-RestMethod @restParams)
    }
    catch {
        CatchException $restParams
    }
}
# GetSnapshots}
# GetControllerId

Function GetStoragePoolId {
    $restParams = @{
        Method      = 'Get'
        Uri         = $EFSeriesC2A+"v2/storage-systems/"+$storageSystemId+"/storage-pools"
        Headers     = $headers
        WebSession  = $session
    }

    try {
        (Invoke-RestMethod @restParams).volumeGroupRef
    }
    catch {
        CatchException $restParams
    }
}
# GetStoragePoolId

Function GetHosts {
    $restParams = @{
        Method      = 'Get'
        Uri         = $EFSeriesC2A+"v2/storage-systems/"+$storageSystemId+"/hosts"
        Headers     = $headers
    }
}

```

```

    WebSession = $session
}

try {
    Invoke-RestMethod @restParams
}
catch {
    CatchException $restParams
}
}
# GetHosts

Function GetDriveIds ($getIdsOnly) {
    $restParams = @{
        Method      = 'Get'
        Uri         = $EFSeriesC2A+"v2/storage-systems/"+$storageSystemId+"/drives"
        Headers     = $headers
        WebSession  = $session
    }

    try {
        if ($getIdsOnly -eq $true) {
            (Invoke-RestMethod @restParams).id
        }
    }
}

```

```

# =====
# POST Functions
# =====

Function CreateVolume ($volName, $size) {
    $controllerId = GetControllerId
    $poolId = GetStoragePoolId
    $volumeDetails = '
    {
        "poolId": "' + $poolId + '",
        "name": "' + $volName + '",
        "sizeUnit": "gb",
        "size": "' + $size + '",
        "segSize": 0,
        "dataAssuranceEnabled": false,
        "owningControllerId": "' + $controllerId + '",
        "metaTags": [
            {
                "key": "string",
                "value": "string"
            }
        ]
    }'

    $restParams = @{
        Method      = 'Post'
        Uri         = $EFSeriesC2A+"v2/storage-systems/"+$storageSystemId+"/volumes"
        Headers     = $headers
        Body        = $volumeDetails
        WebSession  = $session
    }

    try {
        Invoke-RestMethod @restParams
    }
    catch {
        CatchException $restParams
    }
}

# CreateVolume delThisVolume 25

Function MapVolume ($targetHostId, $targetVolId) {

```

```

$mappingDetails = '
{
  "mappableObjectId": ' + $targetVolId + ',
  "targetId": ' + $targetHostId + '
}'

$restParams = @{
  Method      = 'Post'
  Uri         = $EFSeriesC2A+"v2/storage-systems/"+"$storageSystemId+"/volume-mappings"
  Headers     = $headers
  Body        = $mappingDetails
  WebSession  = $session
}

try {
  Invoke-RestMethod @restParams
}
catch {
  CatchException $restParams
}
}
# MapVolume "84000000600A098000BF6F310030060D5C3C58CA" "02000000600A098000BF85F300001E555C41E60F"

```

```

# =====
# DELETE Functions
# =====

Function DeleteVolume ($targetVolId){
  $restParams = @{
    Method      = 'Delete'
    Uri         = $EFSeriesC2A+"v2/storage-systems/"+"$storageSystemId+"/volumes/"+"$targetVolId
    Headers     = $headers
    Body        = $mappingDetails
    WebSession  = $session
  }

  try {
    Invoke-RestMethod @restParams
  }
  catch {
    CatchException $restParams
  }
}
# DeleteVolume 02000000600A098000BF85F300001E615C45CC56

# =====
# Exception Handler
# =====

Function CatchException ($restParams) {
  if ($_.Exception.Message -like "**401*") {
    Write-Host "Not logged in -- attempting log in now . . ."
    LoginAPI $username $pass
    Invoke-RestMethod @restParams
  }
  else {
    Write-Host ($_.Exception.Message)
  }
}
}

```

```

# =====

```

```
# Function Calls
# =====

# new session:
$session = New-Object Microsoft.PowerShell.Commands.WebRequestSession

GetHosts
```

Where to Find Additional Information

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

- Microsoft SQL Server Product Documentation
<https://docs.microsoft.com>
- NetApp Product Documentation
<https://docs.netapp.com>

Version History

Version	Date	Document Version History
Version 1.0	March 2019	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 © 2019 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.