

SAN Storage Evaluation Guide Including Solid State / Flash Storage

December 2019



Evaluator Group

Enabling you to make the best technology decisions

Table of Contents

- Evaluation Guide Overview..... 2**
- What is SAN Storage..... 3**
 - General Overview of SAN Storage..... 4
- Critical Characteristics for SAN Storage Systems 5**
 - Operating Environment 5
 - Storage System Characteristics 6
 - Physical Environment 7
 - Connectivity 7
 - Performance 8
 - Advanced Features 8
 - Installation and Management 9
- Evaluation Questions..... 9**
- Evaluator Group *EvaluScale*™ Requirements 10**
- Summary 11**



Evaluation Guide Overview

Evaluator Group's SAN Storage Evaluation Guide is part of a series of guides designed to help IT professionals evaluate storage technology alternatives. This Evaluation Guide and the accompanying workbook are designed to assist potential buyers understand the options and products available and to help match requirements to the available technology choices.

What sets Evaluator Group's Evaluation Guide series apart from vendor sponsored whitepapers is the lack of vendor bias. Our Evaluation Guides are not sponsored by vendors and are written for IT managers seeking a vendor neutral discussion of the design considerations behind new products, technologies, and trends.

What is SAN Storage

In this Evaluation Guide, SAN storage is defined as storage connected to servers by use of a storage area network or SAN for block access to data. This is distinguishable from connection over Ethernet as a NAS or Network Attached Storage for access to files or object storage object access. Many SAN storage systems also support file access as a NAS system as well. The NAS function is covered in the NAS Evaluation Guide while this guide focuses on the block access.

SAN storage may include solid state technology as well, in the form of SSDs or Solid-State Devices that mimic Hard Disk Drives (HDDs) in form and protocol or with custom flash memory implementations. Systems that support both solid state and disk drives are covered in this Evaluation Guide.

Included in SAN storage is the use of NVMe (Non-Volatile Memory express), a protocol that is memory-based, enabling reduction in overhead when using solid state technology. NVMe is used in attaching solid state devices to storage controllers and servers as well as a protocol over a storage network called NVMe over Fabric (NVMeoF).

Solid State Technology or Flash Storage in its various forms has become an integral part in the planning for Information Technology storage professionals in addressing storage needs. With a number of different implementations for storage with the use of solid-state technology, there can be confusion over the value of the solutions in different environments. The different approaches can change the thinking for storage professionals on how the implementations are applied in their environments. In planning for using the value brought with solid state in the evaluation, understanding the system level characteristics and performance is required. The many variations, each needing to be understood for value and capabilities. Evaluator Group divides the different implementations for solid state into two groups: new designs and adapted designs. New designs are storage systems designed specifically for solid state technology. Adapted designs were originally designed for spinning disk devices but were modified to take advantage of solid-state technology. Both of these are generally referred to as "all flash arrays" by some vendors. Even though there are all flash NAS systems, the majority of systems are for block storage and included in the discussion for SAN storage.

All SAN storage systems can be direct connected to servers as well. The information in this Evaluation Guide is applicable but storage devices that may only be direct connected (such as JBODs – Just a Bunch of Disks) are not included in the discussion.

SAN attached storage is seen as the most critical type of storage in enterprise IT currently. There are many competing products and different feature capabilities. The different products require a detailed understanding so an effective evaluation can be performed.

General Overview of SAN Storage

Storage Area Networks (or SANs) allow storage systems to be shared by multiple computer servers over a network specifically designed to carry storage data. This flexibility provides storage administrators the ability to more easily adapt to changing business requirements.

The term SAN is often used to mean transporting the SCSI protocol over a Fibre Channel network. The SCSI over FC protocol is known specifically as FCP, and is what is typically implied by IT personnel when they talk about SAN storage. This is a common example, particularly for open system environments. However, a SAN may also refer to mainframe systems connected via a FICON network to storage systems.

Another protocol, iSCSI, allows clients to send SCSI commands over an IP network; allowing users to leverage their existing IP networks. Fibre Channel over IP (FCIP) enables the transmission of FC information over IP networks for longer data transmissions. The protocol known as Fibre Channel over Ethernet (FCoE) maps Fibre Channel natively over an enhanced version of Ethernet. A nascent protocol is NVMe over Fabric with different fabric implementations. All of these technologies may accurately be described as SAN storage connectivity.

NVMeoF, NVMe over Fabric, uses the NVMe protocol over storage networks which may be Fibre Channel, Ethernet, or InfiniBand. Many if not most vendors that have Fibre Channel attached storage have added support for NVMeoF, which lowers the protocol overhead significant. Using NVMeoF with Ethernet has more than one solution and does eliminate the necessity of the addition iSCSI software layer. Ethernet implementations are RoCE (RDMA over Converged Ethernet) which requires RDMA capable NICs and Data Center Bridging switches, iWARP which require RDMA capable NICs but uses standard switches, and NVMe over TCP which uses standard NICs and switches. InfiniBand natively supports RDMA for the protocol reduction.

An early promise of SAN storage was the ability to share data between computers connected to the storage network. However, sharing of data requires coordination between applications, operating systems on the compute servers or between the file systems. SANs provide access to a common pool of storage, resulting in a reduction in the amount of under-utilized and wasted storage capacity.

Additionally, storage connected using FC, iSCSI, or FCoE (and NVMe over Fabric) deliver data with limited overhead, resulting in high speed access. The overhead associated with file data can impact applications

that require fast access to data, and as a result transaction processing data is typically used in a database are often stored on SAN connected storage systems.

The largest issues with SANs are typically cost and implementation complexity. A SAN solution consists of storage, computer systems, the network including networking switches, and management software. All of these components add cost and complexity to the system, requiring more effort to deploy and maintain SAN connected storage.

The SAN Storage Systems have advanced features that enable IT to better protect systems and improve operations. These advanced features were once the distinguishing characteristics of systems offered in the high-end enterprise market segment but now are available even in entry level systems. The complexity or completeness of these features can vary and need to be included in the evaluation consideration.

The features include the ability to internally tier data between different performance and cost devices. The devices may be HDDs, SSDs, or caching elements (DRAM, SSDs, Persistent Memory/Storage Class Memory – PM/SCM – or flash cards). Other distinguishing features include support of virtualization hypervisors and integration with external software through APIs.

Critical Characteristics for SAN Storage Systems

SAN storage, when equated to block storage access, has a number of very well-understood areas where the characteristics are evaluated when making a purchasing decision. These areas may have different weight depending on the usage needs for the particular environment. SAN storage is usually regarded as primary storage in an IT operation but may also be used as secondary storage. In many cases, a primary storage system that is replaced with a better performing system may be relegated to a secondary storage role.

The evaluation of a SAN storage system needs to include consideration of the following areas:

Operating Environment

The different operating systems or hypervisors that will utilize the storage resources from a SAN system need to be understood. The storage system may not support every operating system or hypervisor or may not support a particular level. The storage system vendors typically have a compatibility guide on their website for each product that identifies the software and version number supported.

There may be advanced features that will improve operations between the host and the storage system that would also be of value to understand. These may be in the form of advanced features such as VMware vSphere APIs for Array Integration (VAAI) for example. Knowing what the operating environment will be and the storage system capabilities will help make a knowledgeable choice and avoid a big surprise later.

Storage System Characteristics

The SAN storage system has certain hardware capabilities that are important in capability and overall operation. The number of controllers in the system can affect the performance and the failover operation. Multiple controllers also may be the determining factor of the ability of the system to scale both performance and capacity. Increases in capacity accomplished by adding more devices – HDDs and SSDs are needed for planning for future capacity demands. The types of devices affect both performance and cost.

The amount of DRAM or PM/SCM cache can affect the system performance and typically will require battery hold up or a super capacitor to retain data in the case of power loss. Extended cache using flash technology in the form of SSDs or PCIe flash cards can also be a factor for performance.

The sparing of physical devices such as disk drives or SSDs and the ability to replace while online is another availability issue and can be a critical factor in the selection.

Currently, the solid-state technology of interest is NAND Flash because of the declining price. The low price is driven by usage of flash in consumer electronics such as cameras and tablets. Much of the focus has been on the transition from Single Level Cell technology to Multi-Level Cell technology to 3D or stacked NAND flash, some with Triple Level Cell technology. The change or transition to denser technology is underway and will continue with the commensurate price declines. Characteristics of managing the different number of bits per cell and the issues in handling the limited number of erase cycles available in the technology are changing areas as well and may be interesting for discussion of design details, but from a storage system standpoint, the important issues are how the solid-state technology is applied in a complex system environment.

The most often cited values for Flash usage as storage include:

- High performance with no rotational latency or actuator movement as there is with hard disk drives.
- Rugged devices with no moving parts.
- High reliability with an MBTF greater than six times that of hard disk drives.
- Low power requirements with typically less than one quarter of the power required for hard disk drives.

There are different areas where NAND Flash has been applied for storage of information in Information Technology. Each of these areas needs to be explored when used in a system as to the implications for performance and costs. The general usage areas are:

- Solid State Devices or SSDs which also may stand for Solid State Disks or Solid-State Drive, depending on the vendor marketing, are usually NAND Flash devices with the form factor of a hard disk drive (3.5" or 2.5") and have a disk drive interface such as SAS or SATA. Some SSDs

may also have a PCIe connection using NVMe (Non-Volatile Memory extended) to eliminate the overhead of mapping a memory technology to spinning disk protocols.

- Flash memory cards with NAND Flash technology on a PCIe card and are plugged into servers as storage accelerators or plugged into storage controllers and primarily used as cache.
- Storage systems enhanced with internal tiering using software to move data to solid state technology based on detected patterns of access for performance acceleration.
- Storage systems that use solid state technology for caching of highly active data. The cache may be complementing DRAM-based caching that exists in most storage system with a larger capacity, less-expensive caching technology.
- Storage systems that consist entirely of solid-state technology. With many different implementations, the all solid-state storage systems can vary in performance and cost.

A trend developing for solid state technology used in storage systems is to capitalize on the performance aspects to enable inline data reduction using compression and data deduplication to effectively multiply the capacity of solid state. With data reduction the effective capacity varies depending on the type of data but five times the capacity is very typical in the systems available currently. Increasing the efficiency of the solid-state storage in the amount of data that can be stored effectively reduces the cost by the multiplier amount.

Physical Environment

SAN storage systems may be quite large – both in the capacity supported and in the physical space required. The evaluation of the storage needs to include all the physical details so that the planning can be accomplished. Power, cooling, and space are usually the first things thought of when the physical environment is mentioned. But it also needs to include other areas such as what phase of power is required, how much does the system weigh, what plug connection is needed, whether there needs to be two independent AC sources and the oft forgotten independent ground plane connections.

Connectivity

The SAN connectivity is one of the basic areas for consideration and would seem simple but can become very complex when architecting a fault tolerant system. The host connections include the type, speed, and number of connections. Failover capabilities lie in different areas including the host driver, switch, and the storage system.

Typically, higher performance interfaces such as Fibre Channel and InfiniBand are used to get the greatest value out of the storage performance. For Fibre Channel, there is a significant difference in the amount of overhead in the protocol for a Gen 6 or 7 implementation at 32/64 Gb/s compared to the earlier Fibre Channel versions. The overhead is visible in the reduced response time seen from the Fibre Channel adapters and switches.

Use of Ethernet protocols such as iSCSI and FCoE will have the added time spent in protocol encapsulation and conversion visible in the response time for storage access. Storage systems targeted at less performance demanding environments will typically be where Ethernet-based protocols are deployed.

Performance

The performance characteristic is usually believed to be the most important decision point for SAN storage. This is usually because storage is the bottleneck for most applications and systems and the faster the storage system, the better the applications will operate. As with an all solid-state storage system, there is a tendency to look at the aggregate number of IOPs supported and the total bandwidth as being the measure of goodness for SAN storage but that can be very misleading. A faster response time does mean a greater number of I/Os but, more importantly, it means that the application can proceed to the next operation quicker resulting in more important measures such as increased number of transactions per second. Aggregate bandwidth may be very important depending on the applications in use.

Performance numbers should also include industry standard benchmarks such as SPC and IOMARK to be relevant. Vendor supplied numbers may not be very useful for comparison and may not reflect true independence.

Advanced Features

The advanced features in SAN storage systems have evolved over a very long time and, for the most part, are mature features that have proven to be highly valuable for IT. The features themselves may be complicated and require additional effort for management by storage administrators.

There can be considerable differences between the different storage systems in how the features are implemented. The differences and the changes in operational procedures to take advantage of the features may make changing storage systems a very difficult task. The advanced features of importance to look for include:

- Snapshots – a large number with redirect-on-write or copy-on-write implementations and support for read/write usage of the snapshot copy.
- Remote replication – both synchronous and asynchronous replication with the capability to support consistency groups to maintain write-order integrity. Periodic synchronous replication where differences are replicated at points in time may be very useful for some types of data and help minimize the bandwidth expenses.
- Hypervisor features – VAAI support for VMware and ODX support of Hyper-V.
- Local clones – ability to make local copies of volumes
- Tiering between devices of different performance and cost
- Tiering to cloud/object storage with LUNs assigned to clouds storage pools for data protection (snapshots) or storage expansion.

- Encryption of data at rest with either internal key management or exported keys
- Stretched cluster support in order to create a business continuous environment across local distances
- Security capabilities
- Data integrity capabilities to continuously verify data accuracy
- NAS support for some systems that can operate as unified SAN and NAS storage
- Data reduction – support for greater capacity by using data reduction techniques such as data compression or data deduplication. There can be a wide difference in implementation by vendors.
- Direct backup by instructing the storage system to make a copy of data (usually a LUN) to another storage system attached to the SAN.

Installation and Management

Storage systems have a basic element manager but many are now integrated with top-level management software such as vCenter. The ability to work with top level management tools and reporting software such as Storage Resource Management is highly valuable for managing the storage environment.

The complexity of the installation may also be an indicator of how complex the system is to administer. The installation ease including the time required, the amount of configuration, and whether it can be installed by the storage administrator or requires a specialist should be understood both from the time expectation of getting the system into production and the management complexity that will be ongoing.

Evaluation Questions

There are many questions to start with in the discussion of SAN storage given the wide understanding of the systems and usage. Understanding the real requirements and what is needed from the storage system is important so as to focus on them and not the marketed product distinctions. Starting with the set of requirements should lead to a set of useful questions for the evaluation. The Evaluator Group SAN Storage Workbook should be useful in prompting for some of the information. A few of the questions that should be considered follow:

- What is the planned usage for the storage?
- What are the applications that will be accessing data on the system?
- How will data be migrated to the system?
- What is the capacity growth expected for SAN storage?
- What type of security and compliance requirements must be met?
- What advanced features will be required?

Evaluator Group *Evaluscale*™ Requirements

Working with many IT clients, Evaluator Group has developed a list of the most important criteria for making product selections. These criteria and the associated requirements comprise the *Evaluscale*. For each product, Evaluator Group publishes a Product Brief that includes an *Evaluscale* showing how the product measures up.

Requirements do vary depending on usage and IT environments, variations that generally follow a segmentation of high-end enterprise, mid-range or entry-level. The *Evaluscale* incorporates these differences into each requirement and orders the criteria based on Evaluator Group's opinion and information gathered from IT client engagements.

Criteria for SAN Storage Systems

	Criteria	Description	Requirement
1	Capacity	Current capacity of system to meet demand	Must have enough capacity to meet current demand and have ability to scale-up – adding more capacity up to a practical limit.
2	Price – including data reduction	Cost of system. This includes data reduction effect – compression/deduplication	Must be competitive with other leading solutions in this space meaning prices have no more than 20% variance from an average of the other solutions. This includes the effect of data reduction according to the Evaluator Group Data Reduction Estimator tool.
3	Performance	Latency IOPs Bandwidth	The performance requirement can vary based on high-end, mid-tier or entry-level usage. Requirements are that the system be within 10% of the top numbers for products in this class using SPC or IOMARK benchmarks.
4	Scaling – performance and capacity	Ability to increase to meet future demands	Scale-out means scaling both performance and capacity to meet demands up to a practical limit – more capacity without sacrificing performance by crippling the access density.
5	Security – encryption	Data at rest encryption and key management	Data encryption is a perceived requirement across different usage categories for many industries and applications. The high-end enterprise requires an external key manager. A key manager also is a positive factor for the mid—tier.

6	Data protection	Snapshots – R/W, large number Synchronous replication Asynchronous replication Tiering to clouds	Read/write snapshots are a requirement with a number that roughly equals the number of LUNs supported. The high-end enterprise requirement for remote replication is for both synchronous and asynchronous technology while mid-tier and entry usage require asynchronous. Tiering to clouds may be a benefit to all segments but would only be a requirement in the high-end enterprise.
7	Business continuity	Active-active stretched clusters	The high-end enterprise generally requires active-active stretched clusters. Some mid-tier environments will need that as well but would not be a current requirement but could be a differentiator.
8	Economic considerations	Warranty Evergreen updating Environmentals – power & space Simplicity for admin	The overall environmental footprint being roughly on par with other leading systems in this area is the requirement measure. An extended warranty period for devices and an evergreen program for the controllers in the case of an all-flash system is now a requirement given the competitive nature. 100% data availability guarantees and “as-a-service” offerings are also positive factors.
9	Storage technology	Use of solid-state technology as storage devices and for caching	The requirement is for the system to support solid-state technology for at-rest storage and to have an extension cache using solid-state storage technology. End-to-end NVMe and SCM devices are strong positive factors. S
10	Application/system integration	VMware VAAI, VASA, SRM, etc. and Microsoft ODX, SCOM integration, Container support	Given the pervasiveness of hypervisor use in server virtualization, the integration for storage primitives and management is a requirement. As container usage increases, the need for container support such as CSI is becoming more critical as well.

Summary

SAN Storage systems are well known, mature systems with many advanced features that provide value to customers. The technology used in the systems continues to evolve and the systems advance in capabilities such that each new purchase requires an evaluation. The diligence involved in the evaluation will result in better selections of SAN storage to meet the requirements. The Evaluator Group Evaluation

Guide, Workbook and Comparison Matrix for SAN Storage Systems will be very valuable in these evaluations.

About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and **education**. www.evaluatorgroup.com Follow us on Twitter @evaluator_group*

Copyright 2019 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.