



SAN Sizing Guide

The Essential Guide to sizing a SAN. Selecting the right number, size and speed of disks, RAID levels and SAN performance.

OVERVIEW

The purpose of this guide is to help determine the correct disks for use in a SAN. The basic choices when selecting disks include the variables: disk speed, disk size, number of spindles. The challenge is to ensure the correct IOPs (i/o's per second) are attainable when many clients are connecting to shared storage on the SAN. Over sizing disks for performance can lead to excessive costs and management headaches, while under sizing disks can lead to degradation in an applications performance and unexpected expenditures.

A common scenario may be to use slower speed drives (7.5K rpm SATA) for files servers and less critical applications, and faster drives (15K rpm SAS) for critical apps. The number of drives in a RAID group also affect performance, which may dictate a minimum number of drives required to achieve the necessary performance.

SAN SIZING

When your existing storage infrastructure is no longer performing as required or management of storage capacity and backups is out of control it's time for a change. Whether you already own a SAN which is out of date or end of life or you are investing in your first SAN you need to make some critical decisions in terms of SAN sizing and performance. Let's assume you've read the "Why iSCSI Guide" and agree that iSCSI is a contender for your new SAN. Now it's a matter of correct sizing, and picking things like SAS drives, SATA drives or a combination of SAS and SATA.

You should be concerned with ensuring you get enough disk capacity, performance and scalability to satisfy your applications today and for the life span of the new SAN. One advantage of enterprise class SANs is the ability to perform "thin provisioning" allowing you to allocate more storage to each of your application servers than physically exists on the SAN; thus you only buy new disks as your applications data grows, which saves you from guessing how much disk you will need 2-3 years from now.

The better prepared you are as a customer before talking to vendors, the less you rely on sales reps and vendors "selling you" on products that may not be the best fit for your budget or technical requirements. It's common in the storage business for SAN sales people to over size, both in terms of disk and performance. This is costly for you the customer. A little bit of knowledge will empower you as a SAN consumer to select the right storage solution at the best price.

APPLICATION PERFORMANCE REQUIREMENTS

Let's start by looking at the data throughput on each of your servers. Keep in mind the each application is going to be unique, for example some applications will spend more time reading data than writing data (and vice versa), and sequential vs. random reads or writes will affect the performance requirements.

Create a simple spread sheet to track each servers stats, this will be useful to hand to your storage vendors so they can come up with their own analysis and recommendations for your unique environment.

For each physical server you want to connect to the SAN you will need to know the following. If the server is currently in a virtual machine or you intend to use a virtual machine in the future for this server you would require everything below except the physical NIC ports and slots would be shared on the physical host server:

1. How many free NIC ports? Ideally 2 free NIC ports are required for multi-pathing / high availability when using iSCSI.
2. Is utility VLAN used or only production LAN traffic?
3. Current disks (both internal and external). Number of spindles. Disk type SATA, SCSI, SAS.
 - a) Is the server currently SAN connected?
4. RAID level used.
5. List existing HBAs (iSCSI/FC), and the type of slot on the server.
6. How many free PCI-X and PCI-E slots? What speed are they i.e. 133MHz, 66MHz, etc.
7. List your server manufacture and model.
8. List the operating system and version, along with relevant patches.
9. List the applications running on the server.
10. The OS and application information will help vendors determine support for things such as HBAs, software Initiators, multi-pathing, and application aware snapshots.
11. List any other relevant facts, such as servers set up in clusters or unusual performance demands such as video streaming or real time data capture servers.

Now that you have a spread sheet listing the above information you can easily determine some basic requirements for SAN performance, capacity and reliability without getting too deep into complicated capacity planning analysis or running lots of benchmark tests.

HOW MANY DRIVES PER RAID GROUP?

Let's determine how many drives you need in each RAID group in order to achieve the performance required for your applications, such as Exchange, SQL Server or file servers. Let's also determine when to use SAS vs SATA drives. Some considerations with RAID groups involve the number of drives that can fail (hot spares) before data loss occurs along with the RAID type.

RAID is essentially a group of disks where data is spread across the disk group to allow a combination of performance (striping) and redundancy (parity). Striping increases disk performance by increasing the number of drives that are spinning at one time. Parity provides redundant blocks of data spread across multiple disks. The choice over RAID types provide varying levels of performance, availability and disk overhead. In terms of capacity, you really only need to consider the overhead of the RAID type, for example for each raw TeraByte of RAID 5, you can expect to have 800GB of usable data (20% overhead). If your RAID array allows global hot spare disks then you will have less usable space, but more availability.

You should include your data base and/or email administrators in discussions to assist you with sizing requirements. If you are implementing a SAN for backup to disk (raw disk or VTL), then consult with your IT admin that administers backups. SAN performance will depend on whether your data access consists of sequential reads, random reads, sequential writes, random writes or any combinations of these.

When in doubt, and for smaller IT shops you can get away with selecting disks at the same speed (or higher) as what you currently have in your servers as direct attached disks. So if your Exchange server has 5 X 146GB SAS, 10K drives in a RAID 5 config, you could safely assume going with an identical drive config, or even faster 15K drives will do the job.

If you have any doubts, or your applications are currently not performing as fast you should run the Microsoft utility, PerfMon, to help isolate the problems. Taking the effort to run PerfMon will help eliminate bottle necks with your SAN. For assistance running PerfMon contact your local sales rep who will be able to coordinate this with a sales engineer.

CHOOSING THE TYPE OF DRIVES

By categorizing your applications into three categories you can determine whether to go with 15k SAS, 10k SAS, or 7.2k SATA drives. The three categories are applications with the highest IOPS requirements i.e. Exchange, SQL Server, followed by ap-

plications with mid-to-high disk i/o requirements, i.e. backup to disk (or VTL), smaller departmental applications. The third category consists of basic utility applications, i.e. web servers, file/print, archiving.

Ultimately, you will need to determine the number of drives to achieve the correct IPOS. As previously mentioned this can be "eye-balled" based on a servers existing direct attached disks. As an example you may need to consider whether five 300GB 15k drives will do the job, or if you need more IOPS would ten 146GB 15k drives be a better choice. The obvious trade off is with more drives you get more IOPS, but you lose expansion capabilities, hence possibly incurring additional costs for extra disk shelves/arrays.

To summarize, 15k SAS drives should be used for your highest performing applications, consider 10k SAS drives for your mid-tier applications, and 7.2k SATA drives for disk backup or basic file/print/web services.

WHICH RAID TYPE?

Let us get into the depths of RAID with an overview on selecting the RAID type best suited to your applications. RAID is a method of providing a level of redundancy (parity) to a storage system, protecting data in the event of drive failures. Performance can also be increased via striping (RAID 0) data across disks. Typical RAID levels allow for a combination of both parity and striping enabling additional redundancy and performance across a disk system. Below is a summary of the most common RAID types used on a SAN.

RAID 5: Provides more capacity than RAID 1, while ensuring high reliability, improving read performance since more disks are working at the same time. RAID 5 also provides good performance for random data access for applications such as streaming media, disk backup or replication for disaster recovery. RAID 5 requires a minimum of three disks.

RAID 10: Provides block level striping of RAID 0 with the parity of RAID 1. RAID 10 is commonly used for high load databases, since there is no parity to calculate it provides faster write speeds.

RAID 50: Provides block level striping of RAID 0 with the parity of RAID 5, requiring a minimum of 6 drives. RAID 50 provides better performance than RAID 5, especially during writes, providing a higher level of redundancy. RAID 50 improves upon the performance of RAID 5 particularly during writes, and provides better fault tolerance than a single RAID level does.

SELECTING A RAID TYPE

Choosing a RAID type is a balance between capacity, performance and availability. Once you have decided upon which RAID types each of your applications require you can move onto the step of mapping out where to place your SAN volumes (LUN). For example will you create each LUN on a separate RAID group, or place multiple LUNs on a single RAID group? As long as you take into consideration the performance demands applications will place on each LUN, and the IOPs that each RAID group is capable of it becomes a simple matter of how much availability with the disk subsystem do you want to build into your SAN design.

SAN OPTIMIZATIONS FOR PERFORMANCE

Here we summarize a few key points to keep in mind when designing your SAN infrastructure. These points apply to any type of SAN protocol (FC or iSCSI):

1. Keep data intensive applications on separate RAID groups. This will ensure dedicated drive spindles for each of your mission critical applications.
2. Plan for performance first, and redundancy second. The theory here is that If the performance is sub standard then additional costs of redundancy should be used to improve the performance.
3. Ensure your SAN switch configuration can handle the IOPS required. Essentially this translates into using enough ports on the storage server and on each of your application servers to sustain the required i/o throughput. You may need to add an additional NIC or HBA (single or dual port), not an huge expense, but there is the question of available slots and unplanned downtime.
4. Use Gigabit Ethernet for the entire SAN fabric. If you try to leverage older 10/100 switches and NICs don't expect much in terms of performance, although for small test/dev environments or backup to disk for small environments you might be able to get away with 10/100, it's still worth the effort to upgrade to Gigabit Ethernet.
5. Use the iSCSI command line interface (ISCSICLI) command to configure a persistent login target or the iSCSI Initiator Control Panel tool to make your volumes persistent.

6. Use the ISCSICLI command to bind persistent volumes or the iSCSI Initiator Control Panel tool to permit the iSCSI service to configure the list of persistent volumes.

SUMMARY

By taking the effort to plan out your new SAN's design, taking into consideration all the factors which affect performance and availability of data you will ensure your new SAN is a worry free investment as part of your data protection strategy. The correct number of disk spindles and RAID level, along with choosing the most appropriate speed and type of drive (SAS, SATA) will prevent unexpected performance degradation and the additional costs of fighting fires. Implementing a SAN is the best way to reduce IT administration costs, improve application uptime with shared storage and gain enterprise level features for protecting data, such as snapshots and replication. A SAN is a key ingredient as part of a disaster recovery plan, with an affordable cost of entry.

ABOUT STARWIND SOFTWARE INC.

StarWind Software is a leading global provider of storage and iSCSI SAN software. StarWind Server converts any Windows server into a reliable, powerful and scalable SAN and is designed for use with solutions such as Microsoft Hyper-V Server 2008, VMware ESX and VMware ESXi and Microsoft Windows Server clusters for database applications such as Microsoft SQL Server, Microsoft Exchange and Microsoft SharePoint Server. StarWind is an affordable iSCSI SAN solution and combines ease of use with enterprise-class features such as Mirroring and Replication, CDP and Snapshots, Thin Provisioning and Virtual Tape Library (VTL). Since 2003, StarWind has pioneered the iSCSI SAN industry and has been the solution of choice for thousands of global customers in over 50 countries, from small and mid-size businesses, to governments, and Fortune 1000 clients.

RELATED LINKS

To download a fully functional FREE edition of the StarWind iSCSI Server for Microsoft Windows or for more information from StarWind Software about the solution outlined in this paper visit: www.starwindsoftware.com.



Turn Any Server Into a SAN.
Enterprise Features.
SMB Price.™

www.starwindsoftware.com

ABOUT STARWIND

Since 2003, StarWind has been the storage solution of choice for thousands of global customers in over 50 countries, from SMBs, to governments, and to Fortune 1000 clients. StarWind has pioneered the iSCSI / IP SAN industry with its storage virtualization software that converts any Windows Server into a reliable and scalable shared storage.

©2009, StarWind Software Inc. All rights reserved. Reproduction in any manner whatsoever without the express written permission of StarWind Software, Inc. is strictly forbidden. For more information, contact StarWind. Information in this document is subject to change without notice. StarWind Enterprise Server is a registered trademark of StarWind Software.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.