StarWind iSCSI Target for Microsoft Windows:

"Best Practices Guide"

Rocket Division Software
www.RocketDivision.com

INTRODUCTION

StarWind is an advanced, full-featured software-only iSCSI Target for Microsoft Windows. StarWind enables just anyone to quickly install and configure IP SAN solution that delivers immediate benefits allowing storage to be consolidated, virtualized and centrally managed. That's why it's an ideal choice for any sized business.

StarWind is a hardware-independent, cost-effective and scalable storage management solution that enables the benefits of Storage Area Networking such as enhanced disaster recovery and increased regulatory compliance and does not consume any extra resources on application servers. It does not require any additional hardware either, perfectly utilizing pre-installed Ethernet and IP network infrastructure, leveraging existing storage hardware.

With StarWind's volume snapshot feature it is simple and easy to perform scheduled or on-demand snapshots of critical data for backup, and if necessary retrieve all backed up data, or just retrieve individual files. With StarWind's unique ability to "virtualize" and share physical storage hardware there is no need to have an expensive DLT tape or high-capacity Blu-Ray or HD-DVD burner attached to every machine.

Operating System

StarWind designed to work with the following Microsoft Windows operating system environments:

- Microsoft Windows 2000 SP4
- Microsoft Windows XP SP2
- Microsoft Windows Server 2003 SP1
- Microsoft Windows Vista
- Microsoft Windows Longhorn

Both 32-bit and 64-bit native binaries are available. As StarWind designed to deal with large amounts of data, it performs much better on a server class OS such us Windows Server 2003.

Processor and Physical Memory

As a software-only iSCSI Target StarWind completely depends on the installed hardware it is running on. The most critical elements of the system are:

- Processor: StarWind requires at least Intel Pentium4 class CPU but best performance achieved with Intel Xeon based machine.
- Physical Memory: StarWind needs at least half gigabyte of RAM but for getting top-notch performance, we recommend to have at least two gigabytes installed. All available memory would be used for iSCSI volume disk I/O caching and buffering.

Network Performance

Whole iSCSI works on the top of existing TCP/IP network. TCP/IP performance measured in two ways: throughput and latency. Throughput is the amount of data transferred in a given time period. Latency defined as the amount of time between request of the data and actual data availability.

Optimal TCP/IP and thus iSCSI performance is obtained by using a Gigabit Ethernet network infrastructure. It has low latency while providing maximum throughput. Gigabit Ethernet latency is typically less than the seek time of today's even 15K rpm hard drives, while it's throughput (1,000 Mbps or 130 megabytes per second in single direction) is more than adequate for most existing drive technology.

For low-performance tasks like iSCSI-based replication, tape drives and high-capacity optical media devices sharing it is possible to utilize existing 100 Mbps hardware and Ethernet wiring. As it is destination storage hardware being a performance bottleneck and not network itself.

Getting more and more mainstream 10 Gigabit Ethernet from the other side is both performance and price Fibre Channel killer. StarWind works quite well with for example Neterion 10 Gigabit Ethernet controllers. Latest performance records are as high as 500+ megabytes per second in single direction (true wire speed expected soon).

Network Adapter

Network adapters can increase overall system performance by supporting hardware offload of Transmission Control Protocol/Internet Protocol (TCP/IP) CPU-intensive tasks. For example, with TCP/IP checksum offload alone, up to a 30 percent performance gain in CPU utilization been seen in testing at Microsoft.

Alacritech iSCSI offload adapters. Unlike iSCSI HBA (Host Bus Adapters) from for example Adaptec and QLogic, Alacritech network interface controllers do iSCSI offload (CRC and digest calculation) and they are still generic IP tasks capable. QLogic and Adaptec iSCSI controllers do nice job in iSCSI processing but they are unusable for other tasks and they are client-side only (no iSCSI Target acceleration).

Rocket Division Software recommends the use of network adapters with hardware offload (at least TCP offload). These features are typical in most server adapters available today.

Network Switch

A network switch is a small hardware device that joins multiple computers together within one local area network (LAN) segment. Technically, network switches operate at layer two (Data Link Layer) of the OSI multi-layer network stack model.

Network switches appear nearly identical to network hubs, but a switch generally contains more "intelligence" (and a slightly higher price tag) than a hub. Unlike hubs, network switches are capable of inspecting data packets as they received, determining the source and destination device of that packet, and forwarding it appropriately. By delivering each message only to the connected device it originally intended for, a network switch conserves network bandwidth and offers generally better performance than a hub.

Rocket Division Software recommends using network switches rather then hubs when using StarWind (to say truth it is nearly impossible to find GbE and up hub these days, they are only switches now).

Network Topology

When deploying StarWind in a IP SAN environment, separate the network segment from other traffic in the enterprise. This helps to minimize network congestion. Since the segments separated, network packets destined for the StarWind nodes do not have to compete with other corporate traffic for bandwidth. Another advantage of this type of network segmentation is the ease in security and network management. Troubleshooting errors, detecting network failures, and auditing security is simpler since problems can be easily isolated to the SAN segment. See Figure 1 below.

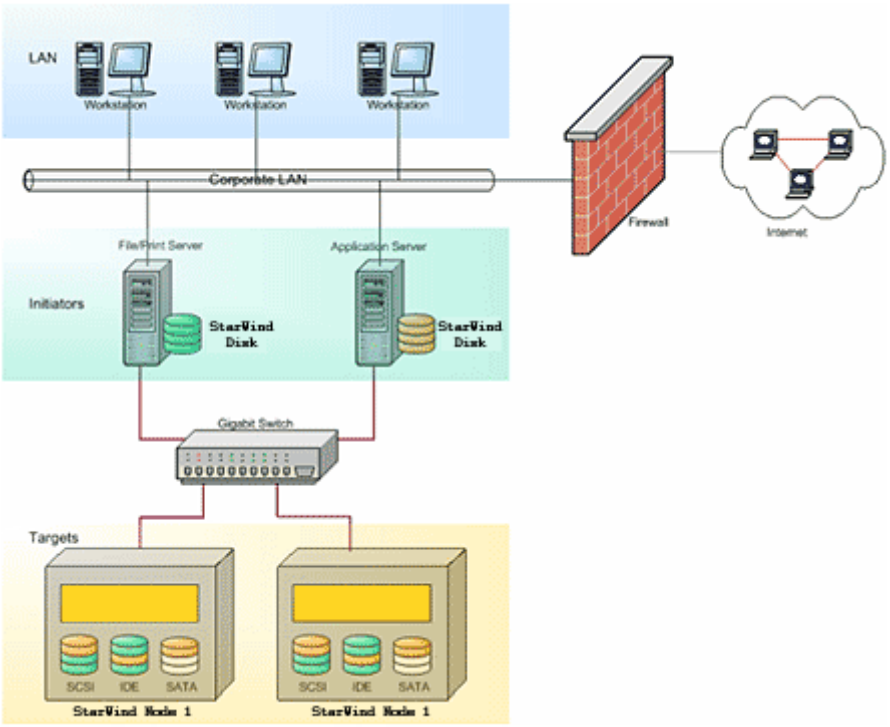

Figure 1 - Network Segmentation.

iSNS

The IP SAN (iSCSI) could be set up to use iSNS. iSNS is analogous to DNS. Just as DNS provides name service for servers and workstations in a LAN, an iSNS server provides name service for initiators and targets in iSCSI. This makes the task of managing the storage network easier because the data centralized on a server. From the other side iSNS is purely optional and could be simply skipped for comparably small SAN installations or for the ones using static IP addresses.

Disks and RAID Volumes

Hard disks use limited number of bus types to attach to computer. They are Serial and Parallel ATA and SCSI, Fibre Channel, USB and FireWire.

To improve disk performance the disks can be set up in a RAID configuration to provide striping.

RAID allows StarWind to see many hard disks as one single disk, or to perform replication and mirroring.

In computing, the acronym RAID (originally "redundant array of inexpensive disks", now also known as "redundant array of independent disks") refers to a data storage scheme using multiple hard drives to share or replicate data among the drives. Depending on the configuration of the RAID (typically referred to as the RAID level), the benefit of RAID is one or more of increased data integrity, fault-tolerance, throughput or capacity compared to single drives. In its original implementations its key advantage was the ability to combine multiple low-cost devices using older technology into an array that offered greater capacity, reliability, speed, or a combination of these things, than was affordably available in a single device using the newest technology.

The original RAID specification suggested a number of prototype "RAID levels", or combinations of disks. Each had theoretical advantages and disadvantages. Basic is RAID-0 configuration (striping).

At this very simplest level, RAID combines multiple hard drives into a single logical unit. Thus, instead of seeing several different hard drives, the operating system sees only one. RAID is typically used on server computers, and is usually (but not necessarily) implemented with identically sized disk drives. With decreases in hard drive prices and wider availability of RAID options built into motherboard chipsets, RAID also found and offered as an option in more advanced personal computers. This is especially true in computers dedicated to storage-intensive tasks, such as video and audio editing or huge database processing.

The easiest way to provide fault-tolerance to important application data is by using RAID-1 (mirroring) functionalities. StarWind is fully capable of handling any RAID configurations supported by Windows (software RAID) or dedicated hardware RAID controller. When StarWind Disks created using mirrored volumes, the Disks are available for use by clients such as Exchange Server with virtually zero downtime. The Windows Logical Disk Manager ensures that StarWind continues to read and write to the disk even when one of the mirrors fails. The mirror is re-established when the failed disk is replaced making the volume fault-

tolerant once again. In this scenario, the system incurred no downtime and no data loss. To gain both the performance benefits of RAID-0 and the fault-tolerance of RAID-1, configure StarWind with RAID-5 volumes, see Figure 2.

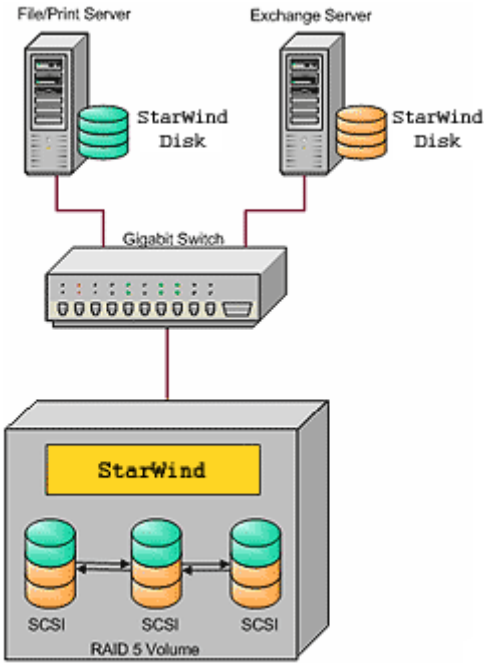

Figure 2 - StarWind Disk backed by RAID 5 volume.

To minimize downtime in the event of a hardware problem, deploy StarWind on a server in combination with hot-swappable backplane hardware. In the unlikely event of a hard disk failure for instance, it is possible to replace failed disk without having to shut the whole system down. Such procedures do not have to involve scheduling downtime for the system.

StarWind can perfectly maintain RAID-1 (mirror) over the different physical hard disks attached. There is really no need in stand-alone expensive RAID controller or dedicated software RAID volume configuration.

CONCLUSION

As this document mentions, there are many factors contributing to the overall performance of StarWind, such as network switches, RAID configurations and choice of operating system. Following are some recommendations for best performance of StarWind:

Operating System	Microsoft Windows 2003 Server SP1 (either 32-bit or 64-bit)
CPU	Intel Xeon 2 GHz and up
RAM	2 GB and more
HDD/RAID	15K rpm SCSI hard disk drives attached to hardware SCSI RAID controller for best performance, 10K rpm SATA hard disk drives in software RAID for best value for money.
NETWORK ADAPTER	Virtually any PCI-X attached GbE NIC with TCP offload capabilities for entry level, Alacritech iSCSI offload controllers or Neterion 10 GbE network adapters for true "iSCSI hardcore".

Contacts

Support E-mail: **support@RocketDivision.com**

Online Support Forum: **www.RocketDivision.com/forum**

Support FAQ: **www.RocketDivision.com/faq.html**

Sales E-mail: **sales@RocketDivision.com**

Phone: +38 067 532 2233
Fax: +38 044 414 8717

Rocket Division Software
04210, Kyiv, Ukraine,
Geroiv Stalingrada Str. 18B/Office 21
www.RocketDivision.com