

# Improving Your NetWare® Backup Performance

[www.novell.com](http://www.novell.com)

---

WHITE PAPER



**Novell®**

Introduction .....	3
The whole system .....	3
Novell Findings .....	3
Storage System .....	3
Controller Performance: .....	4
Disk Drives: .....	4
SAN storage Arrays: .....	4
Controller Drivers:.....	5
File System .....	5
File Sizes: .....	5
File System Tuning:.....	5
Backup Agents .....	5
Remote Agents: .....	6
3 <sup>rd</sup> Party Agents: .....	6
Backup Engine .....	6
Backup Medium .....	6
The worst thing: .....	7
Faster than Tape: .....	7
Conclusions .....	7

## INTRODUCTION

In recent years, storage capacities as well as the amount of generated corporate data has exploded, which has created a significant problem for corporations trying to protect their data within a “backup window” of time. Backup times have significantly encroached upon production time, creating tenuous situations for IT departments.

Three years ago, Novell® embarked on an engineering research exercise to look at the backup performance problem from a complete system viewpoint. This paper articulates Novell findings and recommends best practices.

## THE WHOLE SYSTEM

Backup system performance can only be as fast as the slowest system component. So the key to improving backup performance is finding the slowest components and focusing there. The system components involved:

- Storage system (disks, storage arrays, controllers, drivers)
- File System (file system drivers, cache settings, file sizes/types)
- Backup Agents (the agent library and helper [accelerator] modules)
- Backup Engine (the backup utility that moves data between agents and backup medium)
- Backup Medium (the tape library & drives, SCSI/FC connections and tape drivers, or VTL [virtual tape library])

## NOVELL FINDINGS

Our research discovered many different problems in each area, many of which have been addressed in recent products. As a result, Novell has demonstrated:

- One Terabyte of data backed up in 3 hours<sup>1</sup>
- SYS volume backup in 25 Seconds

So what can you do to improve your backup performance? The remainder of this document outlines what Novell has learned and what you can do. Tools available for doing your own analysis: TSATEST.NLM, and Technical Information Document (TID) number 2962002

## STORAGE SYSTEM

While investigating backup performance Novell found that it is possible to build storage systems that are optimized for backup, which means that it is also possible to build storage systems that are not optimized for backup. The problem is that storage systems are purchased, plugged-in and volumes are carved out using the most obvious strategies: Carve across as many disks as are necessary to create the required capacity and use as few racks of disks as are possible for a given requirement. This practice puts the activity on as few physical SCSI buses as is possible which does not aggregate disk throughput. We found as much as 600% difference in performance between comparable storage systems; tuned and un-tuned. Note: ALL applications and services will benefit from tuning your storage system, not just backup/restore. When Novell built the system for BrainShare 2003 to demonstrate

---

<sup>1</sup> Single stream backup to a non-tape device using an HP DL380 server with dual 2.8 GHz Intel Xeon processors running NetWare 6.5 connected to the Fibre Channel SAN via a 64-bit 100MHz PCI-X HBA. Fastest commercial tape drives (SDLT600 - as of 1H2004) can run at up to 36 Megabytes per second raw speed, or 72 MB per second at 2:1 compression ratio of data. This is 1TB in 7.7 hours raw, or near 4 hours for easily compressible data (plus tape change time). Aggregation can improve these times.

backup performance we used 18 SCSI disks on four SCSI buses. We used two 66MHz, 64 bit dual channel RAID controllers. We used a server with multiple independent PCI buses and put each controller on their own PCI bus. We used 8 physical processors. The storage configuration was four RAID-0 arrays in hardware (one per SCSI channel per controller). We then used software RAID-0 striping on the host across the four hardware logical devices. Note that all of this was older equipment; 3-4 years old. We demonstrated 5.3GB/min at BrainShare 2003 and, most interestingly, were able to backup the entire SYS: volume on the host in 25 seconds! The performance obtained derives from tuning the system comprised of eight CPUs, two PCI buses, two controllers, four controller channels (2 per controller), and 18 physical disks. TSATEST is designed to investigate cases like this, following the instructions in the Technical Information Document. Since the time when Novell obtained these results in 2003, we have achieved even higher performance levels from more modern equipment. We appreciate that real customers will not use software RAID-0 systems on top of hardware RAID-0 systems but a good configuration is software RAID-0 across multiple hardware RAID-5 logical devices. Your design goal should be to use RAID configurations to achieve aggregation of disk throughput.

Another item worth noting is about Hyper-Threading CPUs. Hyper-Threading can have adverse affects on storage system performance when the storage system is under heavy load. This has a simple cause. Intel quotes 30% benefit from HT. So, a 1GHz HT processor would be like two processors with a combined capacity of 1.3GHz. This is like two 650MHz processors. The reality is that individual thread performance will vary between zero and 1GHz with the sum of two threads being an effective 1.3GHz. This is fine for word-processor or computational use but it is very poor for I/O. The reason is very simple. I/O drivers and related layers must interact with hardware devices which inherently are slower than today's fast CPUs. This results in CPU thread execution delays waiting on hardware which are counter productive to the hyper-threading architecture. In NetWare 6.5, Novell has enhanced this type of operation so that it is less invasive on HT CPUs. This is accomplished by not spreading hardware interrupts across the HT virtual processors but only across the physical processors in a multi-processor machine. Even with these changes, HT is still not that beneficial for I/O intensive operations/systems.

#### **Controller Performance:**

The controllers Novell used had 166MHz processors in them and 128MB of battery backed-up RAM. They are optimized for a single task but are still built with general purpose devices. There's no possibility that a 166MHz processor and 128MB of RAM can offer better performance than multiple 3GHz processors and multiple Gigabytes of host RAM. Our tests as well as reports from server hardware vendors have shown that software striping and RAID-5 on-host is noticeably faster than hardware striping or RAID-5 on the controller. The battery backed-up RAM of hardware RAID controllers is valuable for preserving data in light of system failures.

#### **Disk Drives:**

The head seek time will always be the bottleneck in disk drives. It's taken the industry 10 years to reduce the head seek times from 10 milliseconds to 4 milliseconds. Seek times are typically a couple of orders of magnitude slower than the actual read times. This means that a 15,000 RPM drive will not be twice as fast as a 7,500 RPM drive. Aggregating drives in various RAID configurations is the best way to reduce latency which is caused by head seek time. SAN Storage arrays build this into their designs in order to achieve greater performance than direct attached disks.

#### **SAN storage Arrays:**

Novell's testing and investigations discovered that some Storage arrays are not optimally designed for Intel Architecture systems. The IBM PC BIOS limits the maximum number of sectors in SCSI disk geometry to 63 sectors. Some SAN arrays expect an alignment of 64 sectors, so PC Servers attaching to these arrays cause mis-alignment. The solution is to align the sector alignment

on a factor of 64 such as 32. Novell has a new SCSIHD.CDM that supports 32 sectors alignment geometry, and we have seen DOUBLING of performance on the storage arrays that have 64 sectors alignment with this new SCSIHD.CDM. You will want to check with your storage array vendor to determine if this is a possible issue (not all storage arrays have these restrictions).

#### Controller Drivers:

Novell has found that at least one vendor that had shipped Fibre Channel drivers that cause storage systems to perform very badly. TSATEST was used to discover this problem. In the case of the driver in question the hardware had a 16 entry request queue and the driver only used 8 entries in it. Once the vendor discovered this problem and corrected it, the performance of the driver close to doubled.

## FILE SYSTEM

Our tests showed that performance is mostly affected by file sizes. Novell uses the Storage Networking Industry Associations (SNIA) standard backup performance test data set for our backup performance testing.

#### File Sizes:

A backup of a great number of small files will be much slower than a backup of a few large files. For backups of files smaller than about 8K, the system spends most of the time obtaining the file metadata, as opposed to the file data. This becomes especially important with respect to backup engine operations (see Backup Engine discussion).

#### File System Tuning:

File system cache and file system read-ahead can be tuned for specific work loads to improve performance. For larger files, Novell's testing discovered that setting the read-ahead parameter to 16 (from the default of 2 which equals 2 blocks or 8K) gave a large file performance improvement of about 500% with a slower disk subsystem, and 30 to 50% improvement on very fast storage subsystems. For small files, the read-ahead does not give any performance improvement. Caution should be exercised when tuning the read-ahead. If the volume has medium to large files, and few users/processes will be accessing that volume, a larger read-ahead value will result in dramatic performance improvement. However, a system that has many users accessing unique files on the volume simultaneously, may suffer performance from increasing the read-ahead value. In such an environment, an increased read-ahead value will end up causing file cache memory thrashing or poisoning and hamper system performance, resulting in slower performance than the default value of 2. Adjustable read-ahead is only available for the NetWare NSS file system in NetWare 6.5.

## BACKUP AGENTS

Novell's investigation revealed that the Target Service Agent (TSA) architecture used to access Novell file systems and services is an excellent architecture, but that backup engines were not exploiting its architecture. So, for NetWare6.5, Novell made some design enhancements to the file system TSA to both help with older backup engines and to allow more modern backup engines exploit its capabilities for performance. Modern backup engines are capable of multiple streams (to aggregate tape devices), implement an asynchronous read/write engine (to stream data), and submit one job with file include/exclude lists to accomplish the whole backup. The one job per backup feature of modern backup engines allows the new TSA to perform a great deal of pre-processing to significantly improve system performance. The new TSA is called TSAFS.NLM. This is the TSA that was used in our performance tests. Novell attributes up to 800% performance increases to TSAFS, but most importantly, it provides the infrastructure for backup engines to be able to achieve much faster performance than they currently are able to support. In other

words, some backup engines (which need redesign) may only see 20-30% performance improvement with TSAFS. But other modern backup engines can achieve up to 800% performance improvement just by using TSAFS.

### Remote Agents:

The Novell Storage Management Data Requester (SMDR) provides two functions:

- 1) the local system interface to the TSA, and
- 2) a Remote Procedure Call (RPC) interface for remote backup operations.

An RPC implementation is simply a request/reply interface. This does not scale, so Novell is currently implementing a "streaming" SMDR using an Inter Process Communications (IPC) system to upgrade this capability. Initial prototypes have shown theoretical gigabit Ethernet LAN bandwidth limits, and more importantly, it removes LAN latency effects. We have prototyped 1 TB data backed up over a gigabit LAN in under 4 hours (to a non-tape backup medium). This new feature will be delivered in a future release. Please note: the current RPC SMDR is capable today of saturating a 100Mb Ethernet link, so the streaming SMDR will not result in improved remote backup performance on 100Mb Ethernet. The 100 Mb Ethernet is your slowest system in this case.

### 3<sup>rd</sup> Party Agents:

Most 3<sup>rd</sup> party backup applications install "accelerator" agents on the NetWare server for remote backup solutions. These agents talk directly to the local SMDR interface to obtain the data, format the data into the backup vendor's format and send it across the network to the backup engine. Most of these agents also incorporate some form of compression such that the data on the wire is compressed for greater throughput.

## BACKUP ENGINE

Novell's investigation revealed that most of the backup engines in the market are not properly implemented for performance. There are two main areas of improvement needed in backup engines. Some engines have properly dealt with at least one of the issues.

1. Single stream – request/reply architecture. Most of the engines with this implementation deficiency were designed and built over a decade ago. They cannot scale to the speeds of today's storage and tape library systems. They should be re-designed to support asymmetric read/write and multiple streams (threads) to aggregate tape devices
2. Backing up each file as independent jobs. In this case, for a full backup of a volume that has 100,000 files, the backup engine would submit 100,000 jobs. This is bad as it does not allow the TSA to do any pre-processing or pre-fetching of data for the backup engine. Hence, any gains from TSAFS pre-processing features cannot be achieved, not to mention the extra overhead incurred in processing 100,000 jobs as opposed to simply 1 job. Engines should submit one job for each backup, specifying include/exclude lists as necessary in order to be more efficient. How backup engines submit jobs to the TSA makes the greatest performance differences on the overall system.

Novell's SBICON backup engine has been modified as part of the performance research to support asymmetric read/write, and this resulted in a 60% performance improvement with the older TSA600.

## BACKUP MEDIUM

One of the problems traditional tape libraries have had is that their capacity and speed has not kept up with the growth in disk capacity and speed. While significant improvements have been made, low cost IDE or ATA disk drives have become less expensive

per gigabyte than similar capacity tape cartridges. This has created a new market for Virtual Tape Libraries (VTL). This is a low cost IDE or Serial ATA storage array that appears as a tape library to existing backup applications. The performance is orders of magnitude greater than traditional tape, and the reliability (mean time between failures) is much greater. However, VTLs may not be for every corporation, mostly based on corporate procedures and policies. If tapes are required to be physically transported offsite (removable media), a traditional tape library is still required. In this case, a multi-tier data protection architecture may be the solution needed: data is backed up to a middle tier hard disk storage array, and then data is archived from the middle tier to tape medium for moving off-site as per corporate requirements.

The goal of any performance tuning practice will be to ensure that the tape system is the bottleneck in the performance of your overall backup system. Tape should be the slowest component.

### The worst thing:

The worst situation to ever have in your backup environment is known as buffer under-runs. In this condition, the tape device is not fed data fast enough to keep it running at full speed. The result is a situation where the tape must stop, rewind, and resynchronize to the place where it ran out of data. This is called in the industry "shoe-shining". It causes extremely slow backup performance, and causes an extreme amount of wear on the tape and tape drive. This results in significantly reduced life of the tapes and drives. Newer tape drives have the ability to vary their speed to avoid "shoe-shining", but they cannot completely avoid the possibility. Tuning the rest of your system to ensure that it has to wait on the tape drive will ensure the tape is always running at full speed and never "shoe-shines".

### Faster than Tape:

More and more backup utility vendors are supporting tape aggregation. This allows for the backup data stream to be multiplexed across more than one tape device simultaneously. This allows the system to aggregate the performance of multiple tape devices to achieve faster backup speeds than would be possible with a single tape device. The only drawback is that unless the data is placed in a RAID style configuration across the tape devices, if one fails, all the backup data is lost. It is through aggregation that vendors claim 1TB an hour (some demonstrations have shown as much as 10 TB an hour using massive aggregation of many tape devices).

## CONCLUSIONS

Vast improvements in backup performance require finding the bottlenecks within the system, looking from the storage system through to the backup medium. Eliminating one bottleneck may reveal another that is only marginally faster than the first.

But more importantly, eliminating bottlenecks in current backup systems still may not meet the needs of the rapid growth of data that needs to be protected. New or different policies for data protection and archive may need to be explored. First, the data protection requirements that are being solved with current backup solutions need to be identified. Typically this breaks down into 3 areas:

1. Disaster recovery – data protection and Business Continuation.
2. Archive – Historical data, including off-site storage.
3. Recovery of user mistakes – file by file restores.

Disaster Recovery can range from current tape backup methods to full replicated data centers to achieve business continuation. Technologies in this area are rapidly improving and costs are rapidly decreasing.

Archive typically is the result of regulatory requirements. Technologies such as Hierarchical Storage Management (HSM), Information Life Cycle (ILM) are also maturing and should be explored.

User mistake recovery should be put in the hands of the user to handle as opposed to requiring IT to manage. Novell's Salvage capability in the NSS file system has provided a means for this for some time. In NetWare 6.5, the File Versioning and Archive server lets users retrieve older versions of files themselves in-case they accidentally delete or overwrite data they did not intend to lose.

Other effective approaches to manage data protection are becoming available which allow companies to remain compliant to regulatory and company policies without requiring backup/restore from tape. Check with your data protection vendor for the latest technologies.