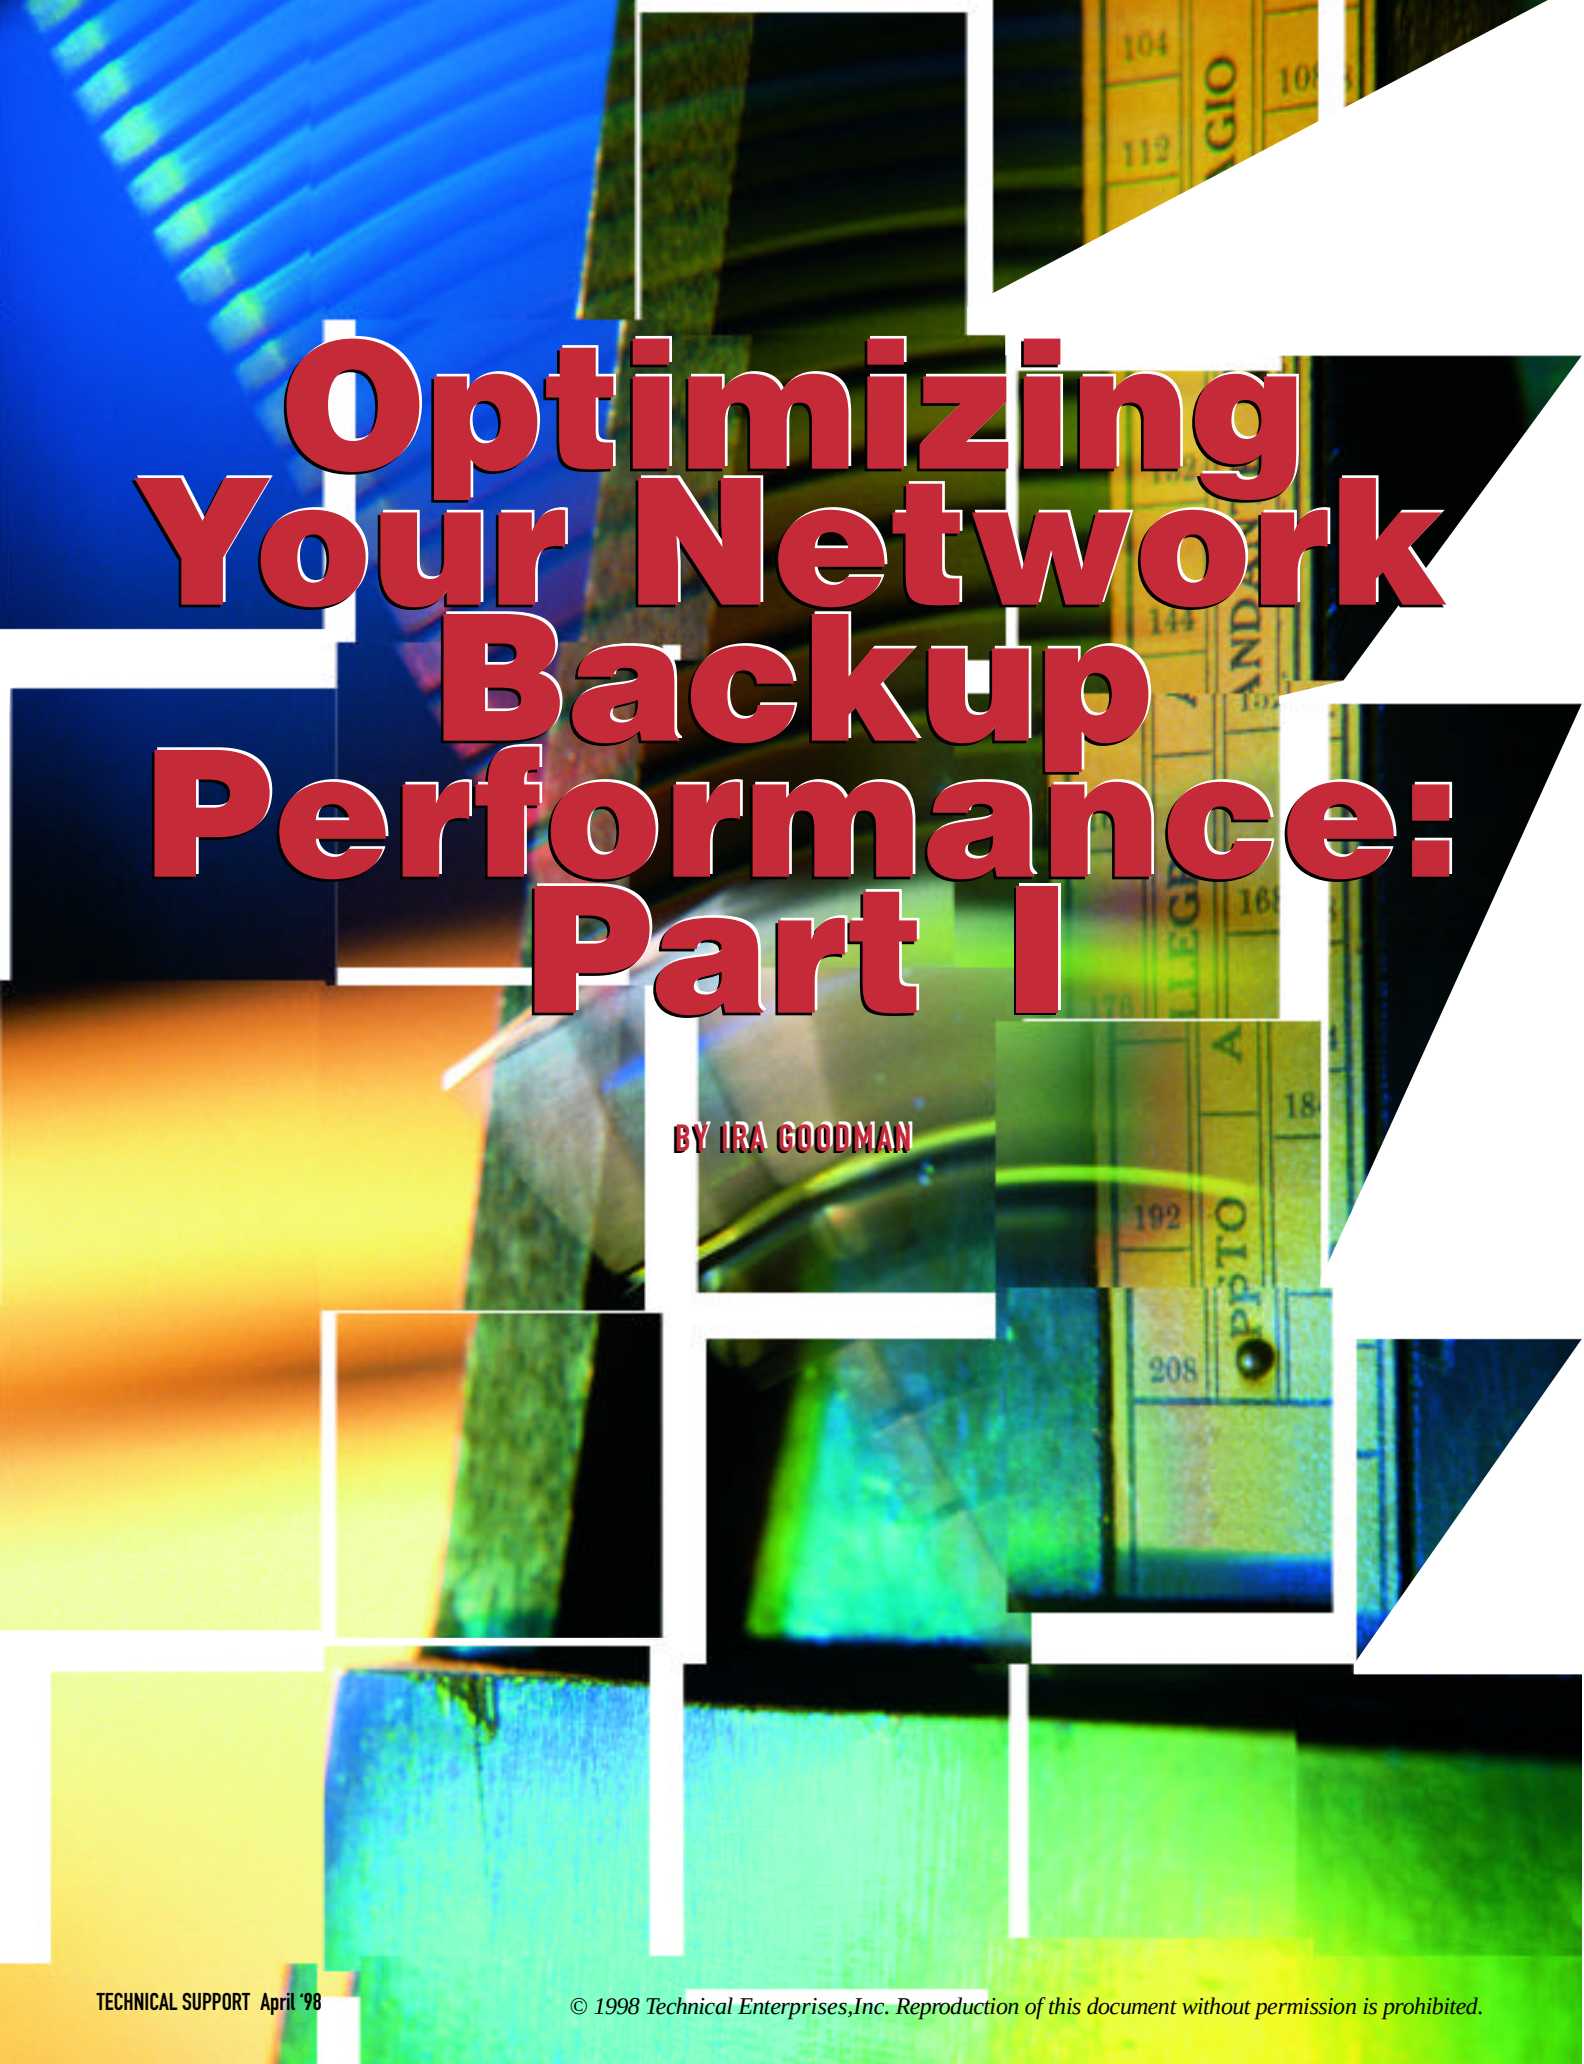

Optimizing Your Network Backup Performance: Part I

BY IRA GOODMAN

Increasing amounts of mission-critical data are now stored on networked computers. The ability to back up and restore this data completely and quickly with as little administrative hassle as possible has become extremely important. Ideally, backup and restore jobs should run at the limits of today's high-speed tape drives and SCSI devices. However, this is often not the case in the real world.

Tuning network backups and trouble-shooting performance problems are extremely complex tasks because of all the variables involved. Gone are the days when a backup administrator dealt with only one or two vendors and compatibility among components was assured. Because of this complexity, I will take a close look at the basic set of variables you need to focus on when optimizing a network backup. The concluding article in this series will describe a few simple utility programs that can measure the performance of some of the variables. Additionally, software settings that have the most marked affect on backup performance will be discussed.

VARIABLES THAT AFFECT PERFORMANCE

To tune a backup, you must understand all the variables that can affect backup performance. The following is a list of the major categories:

- network bandwidth
- hardware
- software
- backup servers
- amount of data to be backed up, its location and type

Any of the variables in this list can make backup jobs run far more slowly than they theoretically should and can degrade overall backup performance.

WHY BANDWIDTH IS A COMMON PROBLEM

Network bandwidth is the most likely backup bottleneck you will encounter if you are not backing up to locally-attached (distributed) backup devices exclusively. The easiest way to understand bandwidth bottlenecks is to think of the network as a pipe with water flowing through it. The bandwidth is comparable to the physical size of the pipe, and the data is the water flowing through it. There is a finite amount of data that can travel across a network, just as there is a finite amount of water that can flow through a specific size of pipe. Too much data leads to both slow backups and almost complete network paralysis.

Network bandwidth is the most likely backup bottleneck you will encounter if you are not backing up to locally-attached (distributed) backup devices exclusively.

Two of the most common networks are 10BaseT, which allows 1MB of data to flow per second, and FDDI,a fiber optic network that allows 10MB per second of data to flow. If you are lucky enough to have the larger FDDI network (600MB per minute, or 36GB per hour), bandwidth might seem sufficient at first glance. Unfortunately, it's extremely unlikely that backup data will be the only data traveling over the network during backup processing. Even late at night, backup processing won't have the entire bandwidth to itself because of other applications running overnight.

SOLVING BANDWIDTH PROBLEMS

First, the following questions need to be asked:

- Do I have enough bandwidth now?
- How fast is the data on the network growing?
- What else is running on the system during backup processing?
- Can I make any scheduling changes to allow the backup data more room to flow?

The second question is probably the most critical. If your data is growing rapidly and your backup is centralized, network bandwidth is the first culprit to consider.

The earliest commercial backup products, many of which are still in use today, were all centralized. They required most of the data in a backup job to travel across the network to reach a cluster of backup devices, usually tape drives, that were attached to a single "centralized" backup server. These centralized backups worked fine in the early days of networking because the amount of data on the networks was relatively small.

Today's backup products are distributed and are able to cope with large amounts of data efficiently, eliminating bandwidth problems completely. A distributed backup product allows you to "distribute" backup devices across your network and attach them wherever they are needed. The data flows to the locally-attached devices and not across the network.

BREAKING BANDWIDTH BOTTLENECKS

Network bandwidth is a physical limitation, and there are essentially three solutions:

- reschedule other work to free up more bandwidth for backup
- invest in a network with a larger bandwidth
- change to a distributed backup product

The first solution may or may not be available to you. The second, a network with more capacity, is a major investment and can be a major disruption when it is installed. The third alternative, a distributed backup, is probably the most cost-effective solution, especially if you are expecting the data on your network to grow rapidly.

HARD FACTS ABOUT HARDWARE

The second "variable" you need to consider when optimizing backup performance is hardware. Whether your backup is distributed or not, there are four basic components:

- the disk on which the data to be backed up is held or on which it is being processed for delivery to one or more backup devices
- the SCSI card that channels data from the computer drive to the backup device(s)
- the cable that connects the SCSI card to the backup device(s)
- the device or devices (usually tape devices) to which the data is actually backed up

Any of these four components can be your bottleneck because each will have a physical limitation that you will not be able to work around.

Each disk drive has a maximum speed at which it can read and write data, and some can process data at a rate of 10MB per second. A SCSI card will have a maximum throughput, and the cable used in the SCSI connection will also have a throughput limit. Finally, the backup devices have both a maximum speed and a maximum capacity.

When choosing a SCSI card, you have two alternatives,SCSI-II and SCSI-III. The differences between the two are the number of devices that can be attached to them and the length of the connection. The first device is actually attached to the SCSI, and all the others are daisy-chained to each other. A single-ended SCSI bus allows a maximum distance of 18 feet between the card and the last device in the chain, while a differential SCSI bus allows a distance of 82 feet. The number of devices each type can handle varies.

HARDWARE RULES OF THUMB

Follow these rules of thumb when dealing with hardware:

- Always make sure that capacities are compatible. If you have a SCSI card that is rated at 20MB per second and a cable rated at 20MB per second and a backup device that can process data at a rate of 20MB per second, don't connect them all to a disk drive that can only read and write data at a rate of 5MB per second. Your maximum throughput will be 5MB per second, and you will have wasted your money on capacity that can't be utilized.

- Never push your components to their limits. Always leave a comfortable cushion of excess capacity. Even if your SCSI card allows a maximum distance of 82 feet, never try to connect a device at that maximum length. Pushing components to their limits can cause all sorts of exotic and very hard to identify, performance problems.

- Be suspicious of everything. It's not uncommon, for example, to have performance problems because of I/O errors caused by an inferior cable or improper SCSI terminations when everything else is working fine.

Because four hardware components are involved, the bottleneck combinations can seem endless. Let's say you have three DLT 7000 devices, and these devices can read and write data at a rate of 5MB to 10MB per second compressed. If you have a cable that is rated at 10MB per second, the cable cannot transport data from more than two devices at a time, even if your SCSI card technically supports more. And, of course, the SCSI has to be able to handle a throughput of at least 10MB per second, and the disk drive must be able to sustain that rate also.

The tape format you choose for your backup devices can make a significant difference in backup performance in two important areas: capacity and throughput.

CHOOSING A TAPE FORMAT

The tape format you choose for your backup devices can make a significant difference in backup performance in two important areas: capacity and throughput.

A simple comparison will make this clear. DDS format, which has been on the market for several years, is available in standard and high capacity. The maximum native (uncompressed) capacity of standard DDS media is about 7GB per tape, and its throughput is a maximum of 0.5MB per second. A DDS high capacity tape jumps to 20GB maximum, and throughput to 3MB per second maximum.

The difference between standard and high capacity is already considerable. However, if you choose devices that use DLT format, you can use DLT 7000, which has a maximum capacity of 35GB native, and a throughput maximum of 5MB per second. Compression doubles both capacity and throughput.

WHY CAPACITY AND THROUGHPUT ARE IMPORTANT

Why is this important? The idea in any backup system is to write the backup image as quickly as possible, and with as little manual intervention as possible. If you use a tape format with a maximum throughput of 0.5MB per second native, you can never write to your backup devices any faster unless you compress the data. If you compress, you will get better throughput, but you also have to do more work (and compression is definitely work).

A second problem is that you can have the most finely tuned backup system possible, but if your operator is on a coffee break when it's time to change a tape, you are going to lose 10 minutes or so, regardless of how fast your system writes to tape. You want to minimize tape handling as much as possible since even rewind and unload speed on many devices today is comparatively slow.

In addition, higher capacity devices in automated libraries give you the ability to run "lights out," that is, without any human intervention whatsoever. Today's high capacity devices make this both feasible and cost-effective.

MORE FACTORS TO CONSIDER

Unfortunately, there is still a host of other factors to consider. One is sustained data rate vs. bursts of data,both of which depend on the size of the file being processed and the size of the block being used to transport the data. You will want to have plenty of capacity for these "bursts."

Hardware compression can also be a factor. Not all files are equally compressible, and this affects throughput too. Small files without a lot of repeating characters are not very compressible while a large file with a lot of blanks will be highly compressible. This compression ratio can be very important in predicting performance.

START WITH VANILLA

Does all this seem complicated? It is, and the best solution is to try and set up as "vanilla" a situation as possible. The simplest set up is a single disk drive and SCSI card connected to a single backup device with a high quality cable. And all these components must have a comfortable cushion of capacity among them. This eliminates all sources of contention, including contention from the network.

From this "ideal" you can go on to build more complex combinations, as long as you always keep in mind the capacity of the lowest rated component. Capacity mismatches cause not only slow throughput but also data loss and propagation problems.

SOFTWARE: WHO'S DRIVING?

Once you have all the hardware variables under control (at least temporarily), it's time to turn to two related software components:

- operating system levels
- device drivers

Upgrading to a higher operating system level can often improve backup performance substantially because the operating system is delivered with new, improved device drivers. These drivers are usually far more sophisticated than those delivered with the previous release. And, of course, there will be drivers for devices that were not previously supported.

What do you do if your operating system lacks drivers for new backup devices? Recently, this was a serious problem for some Windows NT users. Windows NT version 3.5.1 wasn't delivered with drivers for DLT tape devices, but Windows NT version 4.0 is. If you buy a new backup device but are forced to stay at an older operating system level, try calling your backup software vendor. "Pass-through" drivers are usually available from the vendor.

Tabletop devices commonly have operating system support. You can usually take these devices out of the box, attach them to a backup server, and everything works fine. However, if you find yourself with a choice among drivers, it's best to use the driver provided with the operating system, although "pass-through" drivers are usually more than adequate.

The major exceptions are robotic devices, such as automated tape libraries (ATLs), where the backup vendors are expected to provide the library drivers. Since libraries usually contain drives already available as tabletop devices, drivers for the individual devices have already been written by the operating system vendors. The backup software vendors base their library drivers on these "tabletop" device drivers.

The right driver can make a dramatic difference in device performance and can sometimes break a hardware bottleneck for you. Always try to use the newest driver available, preferably supplied by the operating system vendor because, of course, the developers of the operating system know the internals of their operating system best.

SUMMARY

In this article, many of the basic variables that affect network backup performance have been reviewed. These include network bandwidth, hardware components, such as SCSI cards, and software device drivers. Part II will discuss guidelines for choosing backup servers, some simple programs that allow you to test variables in isolation, and the backup software settings that most often affect backup performance. ts

Ira Goodman is a software services manager at Syncsort, Inc. He has 24 years of experience in information processing, and has managed backup product support for almost a decade. He currently manages support for Backup Express, a backup product for networks running UNIX, NT and NetWare, and advises customers on the setup and maintenance of distributed backup systems and automated libraries.