

By Dan Osborne

January/February 2003

Systems Manager for Top Manufacturer Decreases Back Up Time by 80 percent

Are nightly backups taking longer to complete than the time available? This situation faces a surprising number of companies as they struggle to stay on top of critical maintenance routines. At one East Coast manufacturer, for example, servers were taking 16 hours to backup.

“In some cases, it took around 14 to 16 hours to backup the entire server and data each night,” said Bonnie Manley, open systems manager at Reynolds Metals’ Richmond, Virginia-based data center. “There was no time for maintenance, virus scan or anything else.”

A couple of the slowest servers couldn’t be backed up before the start of the working day and had to be aborted. These machines, therefore, could only be backed up at the weekend,

placing corporate data at risk. Additionally, the backup window had expanded to such an extent that IT struggled to cope with other routine maintenance actions.

Eventually, she resolved the situation, reducing nightly backups to within a four-hour slot, and did so while also improving system reliability and increasing performance. How? The company instituted a defragmentation schedule thus eliminating a primary source of slow or aborted backups – disk fragmentation. This alone decreased the time consumed by backup by 80 percent on several servers.

“Regular defragmentation makes the system much more manageable within a shorter maintenance time frame,” said Manley. “The time I now save on backups each night has

exceeded any of my expectations.”

Head-Ache

It’s fairly easy to grasp the mechanics behind why fragmentation multiplies the amount of time needed to complete a routine backup. If each file to be backed up is stored on disk contiguously (in one piece) the head requires only one motion to be in position to read its entire contents. If files are fragmented, though, the head has to locate and gather together numerous fragments before they can be consolidated for back up. Thus it is quite common for IT departments to report backup times taking many hours longer than they used to – out of proportion with any increase in data being backed up – and even the inability to backup corporate data over the course of one night.

This situation has been exacerbated by the fact that the window of opportunity to conduct system backups has been steadily shrinking for the past decade. IT departments used to have fourteen or more hours available for backup and maintenance, as well as the entire weekend. Today, however, more and more businesses are operating globally, with IT consolidated onto one worldwide network. And with the maturation of e-commerce, a 24/7 presence is a virtual necessity. Backup and maintenance staff is left to cope with two almost irreconcilable factors – the amount of data which grows exponentially while the backup window gets smaller and smaller.

Take the example of Reynolds Metals Company, part of the Alcoa Inc. manufacturing empire. Reynolds Metals is a global manufacturer, distributor and marketer of aluminum products. It is the third-largest aluminum company in the world, with annual revenues of around \$6 billion. 18,000 employees at more than 100 locations in 24 countries produce numerous packaging and consumer products and take care of approximately one million metric tons of melting capacity, bauxite reserves and two petroleum coke plants.

To safeguard organizational data, Reynolds conducts daily backups on all of its eighty servers. These are done using

Backup Exec by Veritas Software Corp. of Mountain View, CA. The company's 10 Windows NT and Windows 2000 file servers, for example, each utilize direct-attached DLT 7000 tape drives – one for each server. Although this amounted to only about 200 GB of data, the speed of nightly backups slowed to a crawl.

Manley traced these backup delays to the impact of disk fragmentation. Server files are subjected to heavy fragmentation within a very short time period due to the large number of writes, deletes and temporary folders created by multiple applications and users.

What occurred at Reynolds is not uncommon. According to a survey of large corporations by American Business Research Corporation (Irvine, CA), half reported extensive fragmentation on Windows NT/2000 server files. This ranged from 2000 to 10,000 fragments. Another third of the respondents found files with between 10,001 and 95,000 pieces. Similar results were found on workstations. 56 percent of Windows NT/2000 workstations had files containing between 1050 and 8102 pieces, and a quarter had files ranging from 10,000 to 51,222 fragments.

To combat this situation, Reynolds Metals initiated a daily defragmentation schedule

on all machines using Diskeeper by Executive Software (Burbank, CA). Standardizing with the Diskeeper utility resolved the backup delays.

“I installed Diskeeper on each server and saw a big improvement,” said Manley. “Now my nightly backups take a maximum of four hours to complete.”

The company's nightly maintenance schedule has been reorganized. Backup is done between 10:00 PM and 2:00 AM. Defragmentation is then scheduled to run automatically, and as needed, in a two-hour slot between 2:00 AM to 4:00 AM.

“Now that the drives stay defragmented, it doesn't take long to run each night,” she said. “Backup runs and finishes, Diskeeper runs and finishes, and then we do a complete scan all before 6:00 AM.”

System Reliability/ Stability

Faster backups weren't the only benefit reported. The company also noticed major improvement in performance and system overall stability. Calls about system reliability issues have fallen sharply since instituting regular defragmentation in the datacenter.

“I don't have the system complaints that I had before,” said Manley.

Having all program and data files stored contiguously is actually a crucial factor in keeping systems stable. File being broken into pieces and scattered across a drive, opens the door to a host of stability and reliability issues. Just one or two key files in a fragmented state can lead to crashes, conflicts and errors. If this situation is not understood or addressed, the IT burden escalates. Troubleshooting, for example, takes up more and more time yet fails to isolate the correct reasons for instability. In frustration, system administrators attempt to resolve the situation by reinstalling software, re-imaging hard drives or replacing hardware. Yet unacceptable levels of downtime continue, IT budgets rise and user productivity is adversely affected.

As well as degraded performance and slow backups, some of the common problems caused by file fragmentation are slow boot times, file corruption, data loss, crashes, errors in programs, memory slows, and hard drive failure. The technical reasons why fragmentation impacts these areas are fully documented in a white paper entitled, "Identifying Common Reliability/Stability Problems Caused by File Fragmentation." (See www.execsoft.com). The report refers to a series of Microsoft Knowledge Base articles that highlight the severe influence of fragmentation on system stability and reliability.

Shrinking Windows?

IT departments used to manage relatively small amounts of data and have twelve to sixteen hours available to back it up. These days, however, they must

cope with an explosion in storage demands and an ever-shrinking backup window. If this situation isn't resolved, backups can take so long that there is little time for other routine maintenance tasks. In extreme cases, backups may even have to be aborted, putting data at risk.

One of the easiest and most cost-effective ways to stay on top of this dilemma is to institute a routine defragmentation, automatically scheduled across all servers and workstations on a network. By consolidating files into one piece before backing them up, a much shorter backup window is required.

"To maintain optimal system performance, companies need to schedule disk defragmentation on a regular basis for all their servers and workstation," said Steve Widen, analyst at International Data Corp (IDC) of Framingham, MA. "Otherwise files can take 10 to 15 times longer to access, boot time can be tripled and nightly backups can take hours longer."