

ExaGrid Helps U.S. Federal Government Agencies Reduce Backup Windows, Improve Data Protection

Customer Overview

Federal government agencies, both civilian and military, are steadily moving towards a point where the reliance of tape backup is minimal. The U.S. Government has been the largest user of tape-based backup systems since the 1970s. Most agencies have begun to deploy disk storage and disk-based backup because of reliability and security concerns.

Tape Issues Impact Backup Times, Data Security

Federal government agencies struggle with the ever-increasing amount of data that is generated on a daily basis. Current regulations and legislation require the need for government organizations to backup or archive virtually everything.

Magnetic tape as a medium for backing up data is cumbersome, time consuming, and unsecure. As terabytes become petabytes and regulatory demands become more stringent, traditional tape solutions cannot keep up - resulting in longer backup windows, unreliable backups and restores and increased costs. Tape may have a low initial cost, but long-term costs in lost and compromised data and production hours can cripple an agency.

Over the last decade, a range of technologies has emerged that makes it feasible for disk to replace tape. Disk-based solutions now offer the benefits that only tape once offered, such as infinite capacity, portability and manageability.

Replacing Tape for Backups Makes Better Use of Agency Resources

The use of disk storage for augmenting tape, or of disk storage and deduplication either augmenting or eliminating tape, is becoming a more logical investment for agencies. Scarce agency resources once used to deliver "just" data protection can be repurposed for carrying out strategic initiatives, disaster recovery and continuity of agency services.

Using backup alternatives to tape offers:

- Less IT staff time spent on backups, resulting in time to focus on other valuable IT initiatives
- Faster and more reliable backups
- Faster and more reliable restores
- Ability to meet all financial, regulatory and legal retention requirements
- Achieve all of the above without making any major changes to the current environment that could create work, risk or change

More Agencies Are Adopting Disk-Based Backup

As agencies determine that it is now economically feasible to move from a tape-based to disk-based backup approach, there have been a large number of vendors emerging to offer varying approaches to replacing tape with disk-based systems. This has caused a great amount of confusion for IT managers looking to adopt a disk-based backup system for their agency.

Overview of Data Deduplication Techniques

One of the few remaining arguments for tape is that tape libraries will technically never "run out of retention capacity." As soon as a tape cartridge fills up, it can be replaced with another tape cartridge and the full cartridges can be stored. When writing to disk, storing the same amount of data that is stored on tape would require a massive amount of disk, resulting in high cost.

"The system was more cost-effective than some of the other solutions we looked at. It was less expensive to acquire, and we were able to use it along with our existing backup application and tape drive. It was also more advantageous to purchase the ExaGrid system in terms of licensing and maintenance. We were able to set the system up ourselves, and it's easy to manage and administer."

Dana McCutcheon
IT Specialist
Federal Mediation &
Conciliation Service



However, if you could use a fraction of the space required to store the data on disk and bring the cost of disk storage close to the cost of tape, then disk is clearly the better alternative. From week to week, only about 2% of the bytes change. However, with tape backup 98% of the unchanged data is backed up repeatedly, resulting in saving the identical data dozens and even hundreds of times. With disk storage for backups, deduplication software can intelligently save only the 2% of the data that changes from week to week, saving just the changed data. The net result of using disk storage and data deduplication together is you only need 1/20th to 1/50th of the storage you would need on tape.

Since tape costs about 1/20th the price of disk per terabyte of usable capacity, using data deduplication effectively neutralizes the price gap between tape and disk by using far less disk space than is required to store the same data on tape.

Data Deduplication Methods

There are many methods of data deduplication available in commercial offerings including:

- Fixed data block (64KB to 128KB) - used in backup software applications
- Changed storage blocks - used in primary storage SNAPS
- Byte level - used in target side appliances
- Data block with variable content splitting - used in target side appliances
- Zone-level - used in target side appliances


All of these methods reduce redundant data in backups to some degree but differ on the specific details of their implementation and achievable deduplication ratios.

Without Deduplication



In general, the higher the deduplication ratios, the better. A higher deduplication ratio uses less disk space over time and needs far less WAN bandwidth to replicate data to the offsite disaster recovery site.

With Deduplication



Factors Impacting Deduplication Results

The deduplication approach selected impacts the amount of storage savings that will result.

- 64KB to 128KB fixed block deduplication will average about 7:1 reduction in data storage
- Byte, segment-block and zone-level deduplication will average from about 20:1 to 50:1 reduction in data storage

Deduplication Reduces Storage over Time

The deduplication ratio can range from 10:1 to as much as 50:1, depending on the mix of data types being backed up. Databases can get very high deduplication ratios of over 100:1. Unstructured file data will see an average ratio of 7:1 to 10:1. Deduplicating compressed or encrypted files does not yield a high ratio or significant space savings.

Data Mix Affects Deduplication Results

Retention Period Affects Deduplication Results

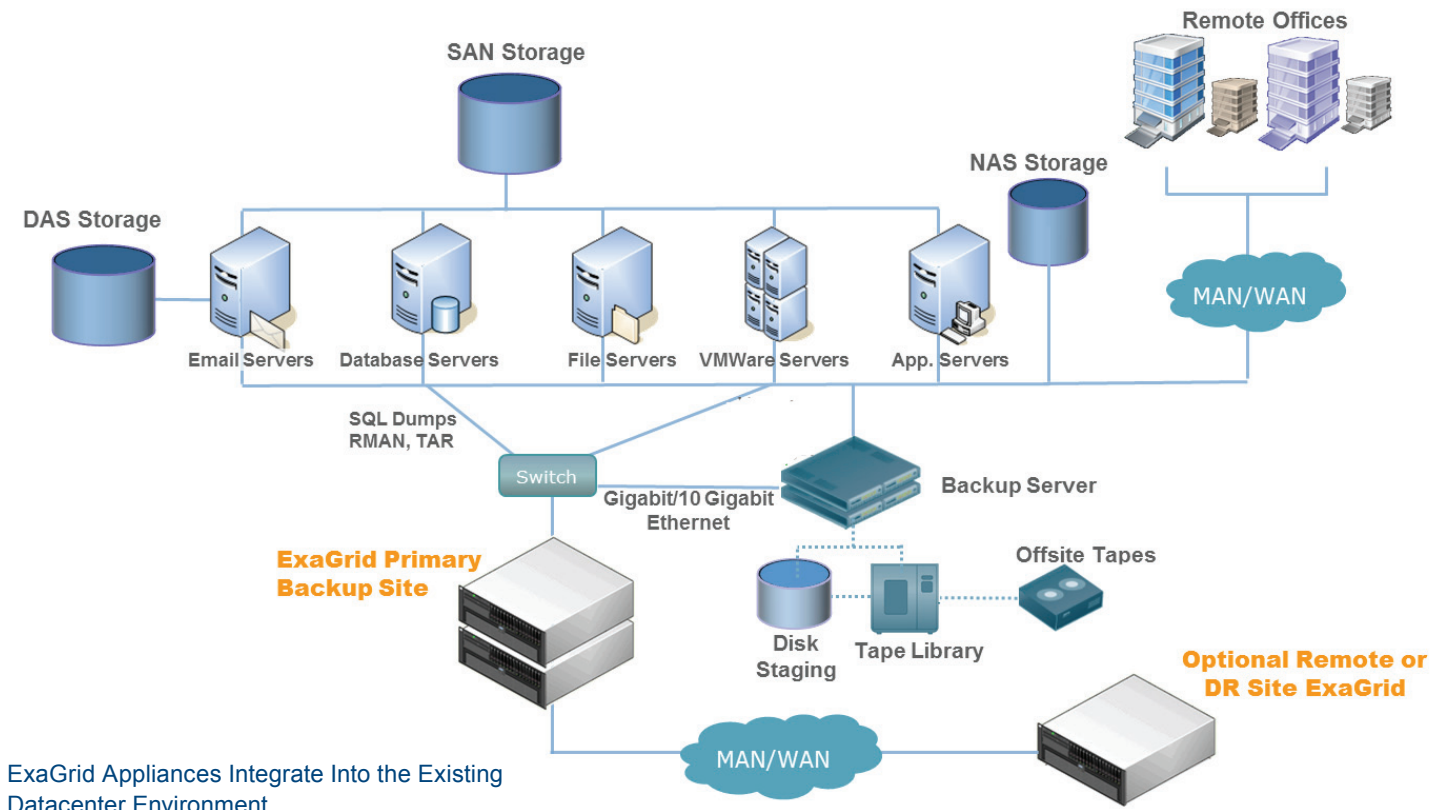
The longer the retention period, the higher the deduplication ratio will be. Given the trend toward longer retention periods for data in general, both for industry and in the public sector, retention policies can have a large impact on the deduplication results.

Getting the Best Deduplication Results

The net is that not all deduplication approaches achieve the same results, and deduplication ratios are clearly impacted by data types and retention periods. All of these factors need to be taken into consideration when choosing the proper disk backup approach.

Disk-Based Backup with Data Deduplication Ensures Reliability and Low Cost

After evaluating various solutions, many federal government agencies select ExaGrid for reasons of performance and cost-effectiveness. ExaGrid is the only purpose-built target-side appliance that combines last backup compression along with data deduplication to reduce the disk space required for backups by a range of 10:1 to 50:1 or more. ExaGrid delivers extremely fast backup performance because data is written directly to disk, and data deduplication is performed post-process after the data is stored to reduce data. When used to replicate data between two or more sites, ExaGrid systems are extremely efficient because only the byte-level changes are moved across the WAN, so only about 1/50th of the data needs to traverse the WAN.



ExaGrid Appliances Integrate Into the Existing Datacenter Environment

Leveraging Existing Backup Application Saves Time and Money

The ExaGrid system works in conjunction with a federal government agency's existing backup application. Because ExaGrid is simply a disk-based target for the backup application, the agency does not incur any additional license or maintenance costs, and there is no need to learn a new backup application. Additional site licenses are not required for second site ExaGrid systems.

Scalability to Meet with Future Backup Requirements

ExaGrid's GRID computing software makes the system highly scalable to meet increased storage and regulatory demands over time. When plugged into a switch, different-sized configurations can be mixed and matched into a GRID system with capacities of up to 130TB. Once virtualized, they appear as a single system to the backup server, and load balancing of all data across servers is automatic.

Sample List of ExaGrid Federal Government Customers

- | | |
|--------------------------------------|--|
| Army Fleet Support | U.S. Army MEDDAC |
| Army National Guard | U.S. Dept of Interior |
| Defense Logistics Agency | U.S. Federal Courts |
| Department of Commerce | U.S. General Services Administration |
| Department of Justice | U.S. Holocaust Memorial Museum |
| Department of Veterans Affairs | U.S. House of Representatives |
| Library of Congress | U.S. Navy Advanced Information Systems |
| U.S. Environmental Protection Agency | U.S. Senate |
| U.S. Air National Guard Arizona | |

ExaGrid Systems Available Through Government Wide Acquisition Contracts

GSA Contract #: GS35F4342D

ExaGrid is available through distribution on the Promark Technology GSA Schedule. Promark is authorized to provide the ExaGrid product to federal reseller partners nationwide. This allows federal agencies to purchase from the Value Added Reseller of choice.

NETCENTS Contract #: FA8771-04-D-0003

Contract Line Item Number (CLIN) via the Air Force sponsored NETCENTS contract through one of its prime vendors, Harris Corporation.

SEWP Contract #: NAS5-02146

Contract Line Item Number (CLIN) via the NASA Scientific & Engineering Workstation Procurement IV contract vehicle (SEWP) through one of its prime integrators, Intelligent Decisions.

ECS Contract #: 263-03-D-0538

Contract Line Item Number (CLIN) via The Electronic Commodity Store III contract vehicle (ECS) through one of its prime integrators, Intelligent Decisions.

US Courts LAN BPS Contract #: USCA-08-B0010

The ExaGrid System is a featured Contract Line Item Number (CLIN) via the US Courts LAN Blanket Purchase Agreement (BPS) contract through one of its prime integrators, Intelligent Decisions.



ABOUT EXAGRID

ExaGrid is the leader in cost-effective disk-based backup solutions. A highly scalable system that works with existing backup applications, the ExaGrid system is ideal for companies looking to quickly eliminate the hassles of tape backup while reducing their existing backup windows. ExaGrid's innovative approach minimizes the amount of data to be stored by providing standard data compression for the most recent backups along with zone-level data deduplication technology for all previous backups. Customers can deploy ExaGrid at primary and secondary sites to supplement or eliminate offsite tapes with live data repositories or for disaster recovery.

ExaGrid Systems, Inc.
2000 West Park Drive
Westborough, MA 01581

1 800.868.6985
www.ExaGrid.com

