

Storage Savings Analysis: Storage Savings with Deduplication and Acronis[®] Backup & Recovery[™] 10



Table of contents


Executive Summary	3
The Importance of Deduplication	4
Approach	5
Analysis Details	5
• The Inputs.....	6
• Results.....	6
What is Acronis Backup & Recovery 10?	8
How is Acronis deduplication better?	8
Network Traffic & Remote Office Backup.....	9
Customer Scenarios	9
• Technology Consulting Company.....	9
Civilian Government.....	10
Conclusion.....	11
• More Information	11

Executive Summary

The use of deduplication technologies are becoming very prevalent when backing up workstation and servers containing large volumes of redundant data. Acronis® Backup & Recovery™ 10 introduces a deduplication technology that applies to file backups and to full system images. Acronis Backup & Recovery 10 is a unique solution that offers data-protection, system recovery and deduplication in one integrated package.

In conjunction with industry leading compression ratios, Acronis is able to provide the high levels of storage efficiency, saving up to 90% (or more) of the storage capacity required to provide fast and reliable data protection.

This paper outlines the storage savings in terms of both storage volume and actual IT budget saved, based upon a typical customer scenario. At the end of this paper we have included customer results based on a recent Acronis survey. This information is intended to provide realistic expectations on what Acronis® Backup & Recovery™ 10 Deduplication can do for your company.



Deduplication is a data consolidated technique that stores redundant data one time only

The Importance of Deduplication

Almost every organization is looking at the escalating cost of storage and is trying to ensure that it is controlled as much as possible. Trends suggest that in the typical organization, storage increases 50-100% every year. According to a 2008 IDC report, the amount of global digital data created and stored on a worldwide basis has increased over 3,000% in just three years! The true costs associated with storage are elusive, but many analysts agree that the cost of managing storage can be as much as four times the pure hardware, software and other tangible costs. Storage is expensive.

Balancing the costs of storage with meeting an organization's recovery time objective (RTO) can be challenging. When most workstations and servers are considered critical, how do you balance the costs of providing fast recovery times and controlling storage costs?

Deduplication on both an image and file backup can drastically reduce the overall storage capacity requirements for maintaining data protection levels. Deduplication is a data consolidation technique that searches for large blocks of redundant information (typically 4KB and larger), and stores them only once, even if multiple copies exist. A "pointer" references the original blocks of data in a way that is seamless to the user, who continues to use a file as if it were a discrete copy.

Certainly, at its most basic level, deduplication reduces the amount of storage required for data protection, in some cases up to 90% or more. Depending upon the deduplication approach deployed, a reduction in network traffic may also be achieved.

This paper will detail the savings that can be possible from deploying data deduplication. A 90% reduction in storage is huge. By reducing overall storage capacity needs and costs, organizations can afford to take more disk images or retain the same data for longer periods of time, without increasing capacity.

Deduplication may be able to save enough on overall storage costs to justify instituting a real backup and recovery strategy for PCs and workstations, not just for mission critical services. Personal computers typically hold 40-80%¹ of mission critical data, yet few organizations have a comprehensive plan in place to ensure that data is not lost.

1: *Online Backup Market Research: MegaNet Online Data Backup Market Research* <http://www.meganet.net/pdfs/onlinebkresearch.pdf>

Approach

The basis of this paper is the online deduplication calculator available on Acronis’ website at <http://www.acronis.com/backup-recovery/deduplication-roi-calculator.html>


The deduplication ratios are based on Acronis R&D tests and verified against early survey trends. This analysis is a guideline for potential results; each environment will have its uniqueness and results may vary . The Acronis calculator is flexible able to accommodate your storage scenarios to provide an estimate of the amount of storage that your organization could save using Acronis® Backup & Recovery™ 10. Additionally, the Customer Scenarios section contains generalized data figures to compare how deduplication could affect your storage requirements.

Analysis Customer Specifics:	
50 Workstations	25 Servers
Average stored data per workstation: 10 GB	Average stored data per server: 50 GB
Data 90%+ Office docs	Data: 100% File & Print, Email Servers
Backup Methodology: 1 full weekly backup, additional backups incremental	Backup Methodology: 1 full weekly backup, additional backups incremental
Retention: 8 week retention of full backups and incrementals are replaced weekly	Retention: 8 week retention of full backups and incrementals are replaced weekly

Analysis Details

To keep the analysis simple, it was assumed that the organization is already performing image-based backups. This simplifies the calculations because of the many sunken costs - the servers or tapes for storage, the administrative time to manage, and the network traffic are already known. All savings that are achieved with deduplication can be calculated and quantified with a simple figure – the fully loaded costs of storage. In the above sample customer scenario, the results of deduplication show as 8500 GB of storage without data deduplication, 1200 GB with data deduplication, yielding a storage capacity savings of 7400 GB.

This is a **87% storage savings** with Acronis data deduplication.



The Inputs

This section details the exact values and assumptions made for this analysis.

	Workstations	Servers
Source Data Total (GB)	500	1250
Daily change % rate	1%	1%
Data retention (weeks)	8	8
Cost of Backup Storage (GB)	\$10	\$15

How does your environment compare?

Visit <http://www.acronis.com/backup-recovery/deduplication-roi-calculator.html> and use the interactive calculator to see how much you could save.

As already noted, this sample set is a basic storage scenario. It includes 50 workstations with an average of 10 GB of storage each, totaling 500 GB of source data; and 25 servers averaging 50 GB of storage each, totaling 1250 GB of source data.

Daily change rate: This calculation specifically feeds into the amount of data included in the incremental backups. Industry studies show about 5% average daily change rate, although our customer data reflects smaller percentages of change, closer to 1%.

Data retention: The range on data retention is from 1 to 52 weeks. An average retention rate of 8 weeks was used for this analysis, although our customer survey data averages 26 weeks. More retention weeks translate to more storage needed; however, it also can provide additional deduplication savings. A 26 week retention period in this case would yield a storage savings of 90%.

Cost of storage: This number is fairly elusive. It should include all storage related costs including floor space, electricity, hardware and software maintenance and overall administration. Industry averages range between \$15-\$30 per GB per year. For the purpose of this analysis \$10 (workstations) and \$15 (servers) were used.

Results

This section details the results using Acronis Backup & Recovery 10.

Description	Value
Total Source Data	1800 GB
Storage capacity without deduplication	8500 GB
deduplication	1200 GB
Cost of Backup Storage (GB)	\$10

Total source data: The total source data entry simply reflects the total source data in the analysis.

Storage capacity without deduplication: This calculation assumes that on a weekly basis, the first backup instance is full and that the other weekly instances are incremental. The weekly totals are multiplied by the number of retention weeks. For retention, it is assumed that the full backup is kept and that the incremental is replaced weekly. The total of this storage space is then cut by 50% to accommodate for the compression achieved in the backup process.

Storage capacity with Acronis Backup & Recovery 10 Deduplication

The same assumptions apply to the storage content: one full image backup and the additional instances are incremental backups. It is also assumed that the full images are retained over 8 weeks.

This analysis assumes a reduction in storage capacity required for the base images of 90%. Acronis Backup & Recovery 10's compression technology alone can achieve 50-60% reduction in storage requirements. Adding deduplication, depending on the data type, can reduce resulting storage on its own by 90%. Thus overall savings of 90% is achievable. Tests using real customer data sets (file and print and email servers) are achieving close to 94% storage capacity savings.

The assumptions for the incremental changes are not as important because they consume only a small percentage of the total storage. To be conservative, it is assumed that the lower end of the basic compression ratio, with the assistance of deduplication, is achieved, resulting in a 50% savings.

The size of the full image is obviously the most important factor in determining the ultimate cost of storage. *The resulting Acronis backups can be 5-20% smaller than those created using other deduplication technologies.*

Savings with Acronis Deduplication	
Storage Capacity Saved with Acronis deduplication	1200 GB
deduplication	\$78,535

A savings of **87%** of storage capacity and resulting costs could be achieved with Acronis Acronis® Backup & Recovery™ 10.

The resulting storage capacity savings is 1200 GB; storage capacity was reduced to 87% of the original storage required. Translating that into IT budget dollars, \$78,535 was saved in the first year.

The 90% claims are aggressive but it is reflective of the data mix. File, print and email content is heavily duplicative. Email attachments are a primary example of how big savings can be achieved. One email with an attached presentation sent to 5 people will cancel out about 500 unique no attachment emails in storage capacity. The deduplication benefits can add up quickly

What is Acronis Backup & Recovery 10?

Based on Acronis' patented disk imaging and bare metal restore technologies, Acronis Backup & Recovery 10 is the next generation of the Acronis True Image disaster recovery product family for physical and virtual environments. It simplifies and automates backup and restore processes across Windows and Linux environments, while having the ability to manage geographically distributed PCs, workstations, and servers and scale from one to thousands of machines.

With Acronis Backup & Recovery 10, the operating system, applications and all data can be recovered in minutes, instead of hours or days. This can be contrasted with the traditional approach which requires the full re-install of the OS, all applications, service packs, and installation of the backup client, finally followed by the recovery of data. Using the Acronis Backup & Recovery 10 Universal Restore option, systems and data can even be restored to dissimilar hardware, including virtual machines, maximizing flexibility and reducing downtime, therefore saving time and money.

How is Acronis® Backup & Recovery™ 10 Deduplication better?

There are many deduplication technologies available. Some are paired with hardware. Some are designed for legacy backup technologies only. Many have limited functionality and are paired with lower compression ratios, resulting in larger storage requirements.

Acronis Backup and Recovery 10 introduces deduplication technology that, in addition not only applies to file backups, also works with full system images. The storage savings and benefits derived from data deduplication can now be applied at any level of data protection and recovery - from routine file-level backups to full system images.

Powerful

What if you could recover any data, any file, or any full system in minutes?

Acronis Backup & Recovery 10 with powerful deduplication meets or exceeds RTOs at a value that justifies deploying on any workstation or server. The superior compression ratios and technological deduplication innovation make it an unparalleled disaster recovery solution.

Acronis Backup & Recovery 10 Deduplication power:

- Deduplication supports both file and image backup, maximizing the storage efficiencies of either type of data protection.
- Deduplication can be done either at the source or the target, placing the processing and overhead where it fits the need.
- Source deduplication employs a patent-pending technology that reduces network traffic.
- Deduplication can be applied to full, incremental and differential backups, fully leveraging the storage efficiencies.

Best Value

Acronis Backup & Recovery 10 value:

- Advanced and flexible deduplication technology, cutting up to 90% or more of storage capacity requirements.
- Straight forward backup and recovery procedures enable fast installs and minimal required training and ongoing administration.
- Simplify the overall data protection strategy and protect more data in the highest category of RTO. With the cost-effectiveness of Acronis Backup & Recovery 10, more data can be made available for recovery in minutes, simplifying data protection management as well as minimizing the impact of any outage.

Network Traffic & Remote Office/ Branch Office Backup

The majority of this paper has been concerned with storage savings that are simple and easy to quantify. Network traffic is another distinct area of savings accomplished with deduplication. In some organizations, the savings in network traffic will be transparent because the traffic is on internal networks running primarily in off hours. The biggest impact that deduplication could have on network traffic is in remote backup scenarios. Previously, remote locations would find it easier to complete backups at remote sites and keep the backup instance in the same location as the source workstation or server. However, disparate backups, in multiple locations are difficult and expensive to manage remotely. Certainly, any further deduplication benefits between site data is lost by storing the images separately, additionally this approach requires each remote location have its own destination for the local backups. In many cases, with smaller remote sites, the backups are not stored in high-density configurations and their resulting costs are much higher than centralized backups.

With source-level deduplication, the size of the backup becomes quite small . It is now possible to reliably transmit backups over a remote connection to a centralized backup location. Centralizing backups has many benefits, from ease of management to the highest storage densities and overall lowest costs.

Customer Scenario Examples

Deduplication, and the resulting benefits, certainly vary on an organizational basis. Early results and trends have been noted and may help determine a reasonable amount of expected storage efficiencies. The biggest determinant of savings is the data mix; the more user-specific data, the higher the potential savings from deduplication.

Technology Consulting Company

A technology consulting firm, heavy with customer data, proposals and internal documents, had very high rates of data duplication within the organization

Company profile:

- 100 Servers, each with 50- 400 GB of data
- 50 file & print servers, 10 email servers, 40 application and operational servers (not candidates for deduplication)
- 90% duplication ratio on file & print servers, 94% duplication ratio on email servers (based on storage capacity not on a number of files)

	File & Print Servers	Email Servers
Source Data Total (GB)	2500	4000
Number backups per week	1	5
Daily change % rate	1%	1%
Data retention (weeks)	8	8
Cost of Backup Storage (GB)	\$9	\$20

The company treated the data types with two different data protection methods. Email was considered the highest priority for the organization. Email was backed up daily and was stored on more reliable hardware.

The results after applying data duplication were dramatic. File and print data had 90% duplicated data and email data was even more redundant nearing 94% duplication.

Data	
Total Source Data	6500 GB
Storage capacity without deduplication	5200 GB
deduplication	5200 GB

After deploying Acronis deduplication, the weekly storage results decreased by over 90%, providing a very large capacity and storage budgetary savings.

Deduplication	
Storage Capacity Saved with Acronis deduplication	46899 GB
deduplication	\$596,640

For organizations with longer retention requirements, the costs of storage go up dramatically. The same organization would save 94% of storage capacity if the retention time frame was increased to 16 weeks.

Civilian Government

A civilian government agency has strict guidelines for data retention and storage and constituent accessibility, with high rates of duplicated data .

Profile:

- 60 Servers, each with 20- 500 GB of data
- 20 file & print servers, 10 email servers, 20 report repositories
- 70% duplication ratio on file & print servers, 95% duplication ratio on email servers, 99% duplication ration amongst report repositories.

	File & Print Servers	Email Servers	Report Servers
Source Data Total (GB)	400	5000	10000
Number backups per week	1	5	5
Daily change % rate	1%	1%	1%
Data retention (weeks)	4	4	8
Cost of Backup Storage (GB)	\$9	\$15	\$20

The agency treated the data types with different data protection approaches. While the report servers had few changes, it was necessary to back them up daily, retaining full backups for 8 weeks on very expensive hardware. The results of the savings for the reporting servers alone was 36,864 GB, and over \$900,000.

Conclusion

Data growth is driving the need to increase storage and network capacity. Organizations must find better ways to control storage costs. At the same time, pressure for fast recovery times disk imaging and deduplication technologies to quickly restore some or all of a company's stored data.

Acronis Backup & Recovery 10 has many innovations to help any organization maximize storage efficiency, lower the overall cost of data protection. Data deduplication, specifically, addresses the runaway issue in most organizations - user-specific data. Data stored on file and email servers are the most prone to duplication. With Acronis Backup & Recovery 10 deduplication organizations can often save tremendous amounts of storage and network costs.

More Information

Please take advantage of the many available resources for Acronis® Backup & Recovery™ 10:

- Try a free 15 day trial of Acronis® Backup & Recovery™ 10 at <http://www.acronis.com/enterprise/download/backup-recovery/advanced-server/>
- For a more complete discussion on the deduplication implementation in Acronis® Backup & Recovery™ 10, please read <http://www.acronis.com/backup-recovery/advanced-server/deduplication.html>
- To find out more about potential savings in your storage environment, use the deduplication calculator at <http://www.acronis.com/backup-recovery/deduplication-roi-calculator.html>
- Visit the Acronis website for more information on deduplication and Acronis® Backup & Recovery™ 10. <http://www.acronis.com/backup-recovery/>
- For more information, please contact your reseller or send an email to sales@acronis.com.



For additional information, please visit <http://www.acronis.com>

Enterprise/SMB sales:
 Email: sales@acronis.com
 Call +1 877 669-9749

OEM inquiries:
 Email: oem@acronis.com
 Call +1 650 875-7593