

WHITE PAPER

The ROI of Backup Redesign Using Deduplication: An EMC Data Domain User Research Study

Sponsored by: EMC

Richard L. Villars

Eric Hatcher

March 2010

EXECUTIVE SUMMARY

IDC conducted in-depth interviews with nine enterprise companies currently using EMC Data Domain deduplication storage systems (note: Data Domain was acquired by EMC in July, 2009). The interviews provided valuable insight into the customers' backup practices and yielded a large amount of quantitative data that IDC entered into its ROI analysis. This IDC White Paper takes a look at the specific business value associated with using Data Domain systems as part of rapidly evolving backup and data protection practices. The paper also examines how the use of Data Domain replication capabilities to move data between systems further reduced costs for backup/recovery and ongoing management, improved IT staff productivity, and increased space management efficiency.

Based upon the averages determined in this study, the highlights of customer benefits realized from Data Domain can be summarized as follows:

- ☒ Customers saved over \$1,000,000 annually.
- ☒ Customers achieved a return on investment over three years of 264%, and payback in 6.6 months.
- ☒ Customers avoided \$200,000 in tape drive and library acquisition costs in the first year after deployment.
- ☒ Customers saved \$156,555 annually in maintenance for existing tape drive/library hardware.
- ☒ Customers avoided over \$80,000 in new server costs annually.
- ☒ Customers reduced restore time from 4.6 hours to approximately 35 minutes.
- ☒ Customers saved IT staff time equal to 1.6 full-time equivalents (FTEs) over three years.

THREE KEYS TO SOUND IT INVESTMENT

Effective collection and use of information are key requirements for organizations in today's challenging business environment. With the right information, delivered at the right time and in the right place, companies can better manage costs, make smarter decisions, and react more quickly to changing conditions.

It is critical for executives across a wide range of industries to make wise investments in the IT systems that support an organization's business objectives. When making IT investment decisions, executives need to judge offerings based on three business requirements. Does the proposed investment:

- Help reduce or control increases in the cost of doing business (Deliver a strong ROI)?
- Ensure the integrity of the business in the face of system and site level failures (Enable more effective disaster recovery [DR])?
- Support increasingly diverse information management requirements (Deliver capabilities that work across a growing range of applications, information types, and data retention requirements)?

Controlling the Cost of Doing Business

Unforeseen or uncontrollable increases in the costs of IT systems are not acceptable especially during times when companies need to tightly control capital and operating expenses. Senior IT executives must reduce spending on IT equipment and administrative resources, and a key component of these efforts is the consolidation of IT assets to ensure optimal use as requirements grow. Technologies like server virtualization — which allow IT departments to consolidate proliferating server assets, boost server utilization levels, and be more responsive to new server requirements — are seen as key elements in improving IT asset utilization.

However, advancing IT operational efficiency through consolidation is not a one-time effort. As companies continue to expand, the resources required to sustain the business can quickly begin to erode margins and inhibit future growth. Consolidation must be seen as an ongoing, iterative process.

For example, the gains associated with server consolidation can quickly be overwhelmed by increases in the storage and management required to back up the growing pool of virtualized servers.

Therefore, IT organizations must deploy storage solutions that are part of ongoing efforts to optimize datacenter efficiency. These efforts include reducing the costs associated with legacy tape environments and minimizing highly redundant data sets generated in virtualized server environments.

Sustaining the Business

The best information in the world is useless if customers, employees, or business partners can't access it when they need it. Concerns about the availability and integrity of information can directly affect revenues and profits. Data management issues related to corporate governance are raised that can influence reputations and have significant legal implications.

The issues of business continuity and information availability have never been more critical for business and IT executives as they make investment decisions. They must deal with an expanding set of business-critical applications, a shrinking window for acceptable time to application recovery, and a heightened awareness of the costs associated with lost or corrupted data.

Improving recovery and data availability standards through more extensive and frequent backups can often prove impractical or overly expensive. The need to operate 24 x 7 means that backup windows are dramatically shrinking while the data to be backed up is rapidly growing. Simply throwing more of the same resources at the problem is financially and technically unfeasible.

In this new world, IT organizations must invest in solutions that allow them to consolidate and automate data protection and recovery processes across the entire organization. IT executives must lower costs and improve manageability for local recovery while also leveraging data replication for backup across remote location to boost overall application continuity and enterprise wide DR.

Accommodating Diverse Information Needs

Increasing IT efficiency purely to reduce the operational costs is acceptable, if an organization's only goal is to enhance profits. Business success, however, is rarely based purely on profit margins. Companies must continually seek out new customers and new revenue streams through innovation in products and services. One of the most important ways to develop and deliver innovative new offerings, today, is through the more effective use of more diverse information assets. These include:

- Relying on email, collaboration tools, and Web sites to communicate and conduct business with customers and business partners
- Collecting, storing, and analyzing more information about products, customers, and transactions
- Digitizing records, design documents, videos, and other types of unstructured data to boost efficiency, offer new services, and comply with evolving government regulations

The continued expansion in the creation and archiving of information is already having a dramatic impact on organizations' storage environments, with IT executives coming to the realization that the cost of storing and managing diverse data repositories will become a major concern for many datacenter administrators.

One clear example of the scope of the issue is the problem of dealing with aging CAD files. In conversations with a number of design and engineering companies, IT administrators told us that the active life for most CAD content ranges from 30 days to a year. A quick scan of their file servers and NAS systems reveals that greater than 80% and often more than 90% of the CAD data on the system is duplicated and may not have been accessed in months or years.

Today, many of these growing content pools (like the above mentioned CAD files) over-consume expensive primary storage capacity and scarce backup resources (time and capacity). In this world of exploding, yet sometimes unpredictable information growth, IT organizations need to deploy storage solutions that can accommodate a wide range of information assets and use cases.

The remainder of this white paper examines the many ways that EMC Data Domain deduplication storage solutions help companies better address the current challenges associated with companies' backup, recovery, long term storage and DR processes. It also provides a detailed assessment of the business value and return on investment that companies obtained from using Data Domain deduplication storage systems.

BUSINESS VALUE OF DATA DOMAIN SYSTEMS

To ascertain and quantify the benefits from deploying Data Domain deduplication storage systems, IDC interviewed IT executives and managers at nine companies that had incorporated Data Domain solutions into their storage infrastructures, including a major financial services company, Canadian National Railway (CN) which operates North America's only transcontinental rail network, an international media company, and a manufacturer headquartered in Canada with a global market and additional manufacturing facilities in the United States, Europe, and Asia. The interviews focused on the experiences of these end users and the actual savings associated with data backup since implementing Data Domain systems. The number of employees per company ranged from 5,000 to over 100,000 and the IT staff size ranged from 12 to over 700. Within the IT organizations, an average number of five staff worked with Data Domain systems directly. Study demographics are presented in Table 1.

TABLE 1

Demographics

Average number of employees	10,067
Average number of IT staff	285
Average number of Data Domain users	5

Source: IDC, September 2009

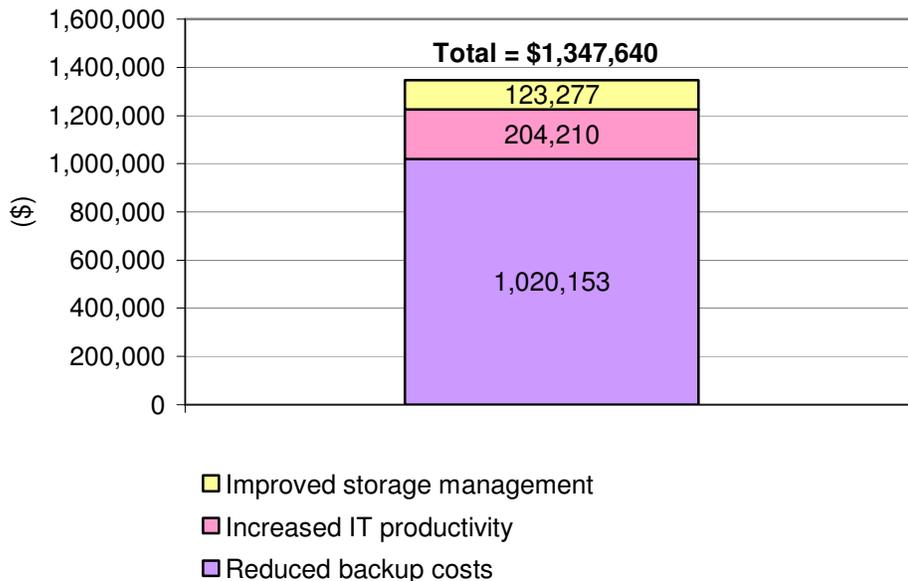
Benefits

All nine companies interviewed by IDC have realized significant benefits from deploying their Data Domain solutions. Reductions in backup costs have provided the greatest savings, averaging over \$1 million annually for each company. Annual savings from increased IT productivity have averaged \$204,210, while savings from simplified and improved storage management have contributed an average further benefit of \$123,277 annually. Averaged over the nine companies, the savings realized since deploying the Data Domain solutions have amounted to \$1,347,640 annually.

The proportion of these benefits and the average annual savings is shown in Figure 1.

FIGURE 1

Average Annual Benefits of Data Domain



Source: IDC, September 2009

Reduced Backup Costs

Much of the cost reduction stemmed from not having to upgrade tape libraries and drives and from savings in tape drive maintenance. Customers found themselves having to choose between increasing the size of an old and inefficient tape environment and upgrading their means of backup to Data Domain systems. Customers compared their expenses since the Data Domain deployment versus what they would have spent on tape and, on average, customers enjoyed savings totaling \$429,264 annually (see Figure 2). Savings on tape transportation have contributed a further annual savings of \$71,510, on average, while reductions in the number of contractors needed have added average savings of \$77,571 annually.

The following are summaries of the narrative discussions with customers that deployed Data Domain systems and realized backup cost reduction benefits.

Before deploying the Data Domain solution, a financial services company had one or two LTO tape drives that cost roughly \$11K each, at all of about 80 locations. "Before, we had a lot of tapes and we spent plenty of resource-hours swapping them," the manager said. "Now we are tapeless. We're also saving on tape transportation — and about 30 minutes a day spent swapping tapes at each location. With deduplication, we are saving on disk space and probably avoiding \$25,000 a year in new disks."

Prior to its Data Domain deployment, CN Railway had tape libraries at 10 remote sites and tape drives at 17 other remote locations. Kevin Whelan, Manager, Operating Systems and Storage said, "Now we have a Data Domain unit at each location. We use them to replicate data to our primary datacenter. Before, we needed someone at each office to take the tape out of the drive each morning and mount the new one." The company estimates it is saving at least 20 minutes per day at its 27 facilities, which equals nine staff hours of savings per day. This company also uses three Data Domain units to replicate data to its DR site. "We're saving time and transportation costs because before, we would load the tapes at one site and a third party service provider would drive them about 20 hours to our DR site so that we could mount them into a tape library."

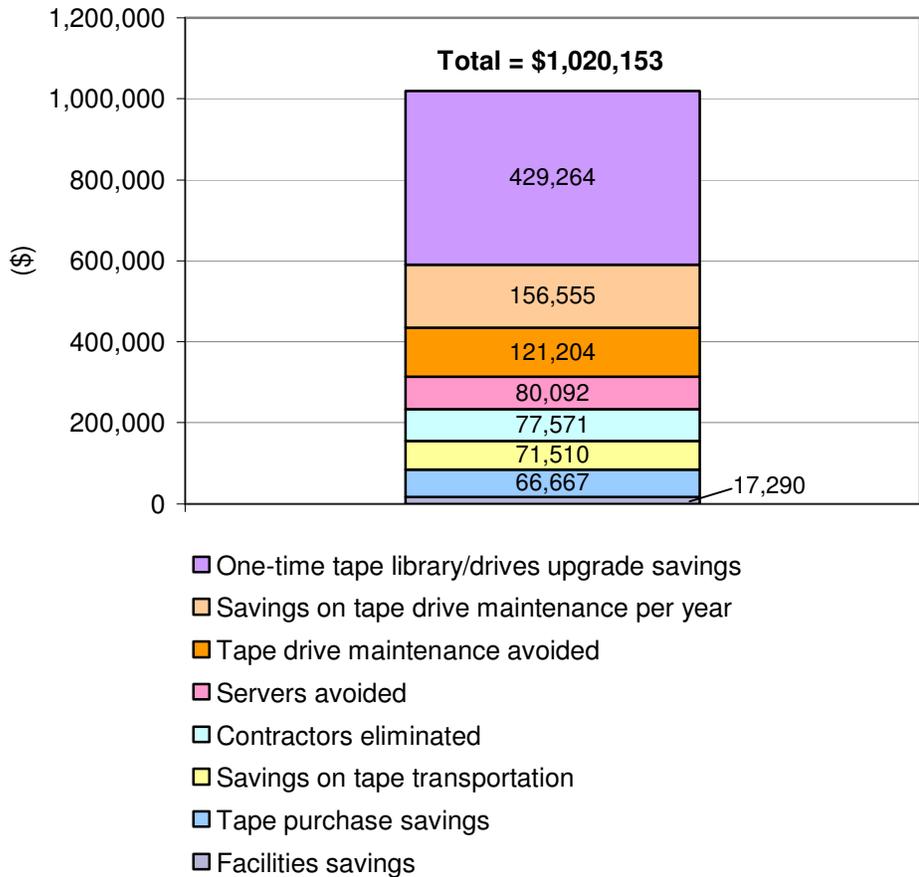
By using the bandwidth throttling capabilities of the Data Domain Replicator software to control the data transmission rate between locations, companies are also able to save on bandwidth costs. One manager said, "We transmit at a low speed during certain busy periods of the day and at a far greater speed during quieter periods. Without this feature, I would have needed twice the bandwidth at some facilities, so it's saving from \$20,000 to \$30,000 a month on bandwidth."

By deploying the Data Domain solution, a media company was able to eliminate tape libraries in its DR site and avoid an upgrade at its primary datacenter. "We had an antiquated tape library at one site, and the local backup solution was problematic," the manager said. "We eliminated the tape library and replaced it with the Data Domain system. Now we do electronic vaulting with two West Coast facilities and replicate the data to our New York office."

Before its Data Domain deployment, one company was using an LTO format, that was no longer supported by the tape drive manufacturer. "We had about 10,000 tapes, so we were looking at a major investment in migrating away from that media," the manager said. "We opted for the Data Domain solution, which paid for itself very quickly since we were going through about 1,000 tapes every three months. We don't write to tape anymore so the savings continue to add up." The company has also been able to eliminate the operational bottleneck caused by tape backups. Prior to the implementation, CPU utilization of the backup servers was restricted, but now the utilization rate is above 80%, which allows the company to avoid buying additional servers.

FIGURE 2

Annual Reduced Backup Costs



Source: IDC, September 2009

Increased IT Productivity

With the Data Domain solution, each of the nine companies needed fewer full time equivalents (FTEs) for managing backup, auditing, and other functions. Also, on average, the time spent on backup was more than halved and restore times were reduced by 53%. Averaged over the nine companies, the savings from increased IT staff productivity have amounted to \$204,210 annually (see Figure 3).

The reduction in backup hours has accounted for \$75,989 in average benefits annually. Savings in FTEs to manage backup have contributed average annual benefits of \$32,838, while the need for fewer FTEs for auditing has provided an average of \$10,343 in additional annual savings.

The following are summaries of the narrative discussions with customers that deployed Data Domain systems and realized increased IT productivity benefits.

A financial services firm was running out of time for backup before it deployed the Data Domain deduplication systems. A manager said in the interview, "Our backup

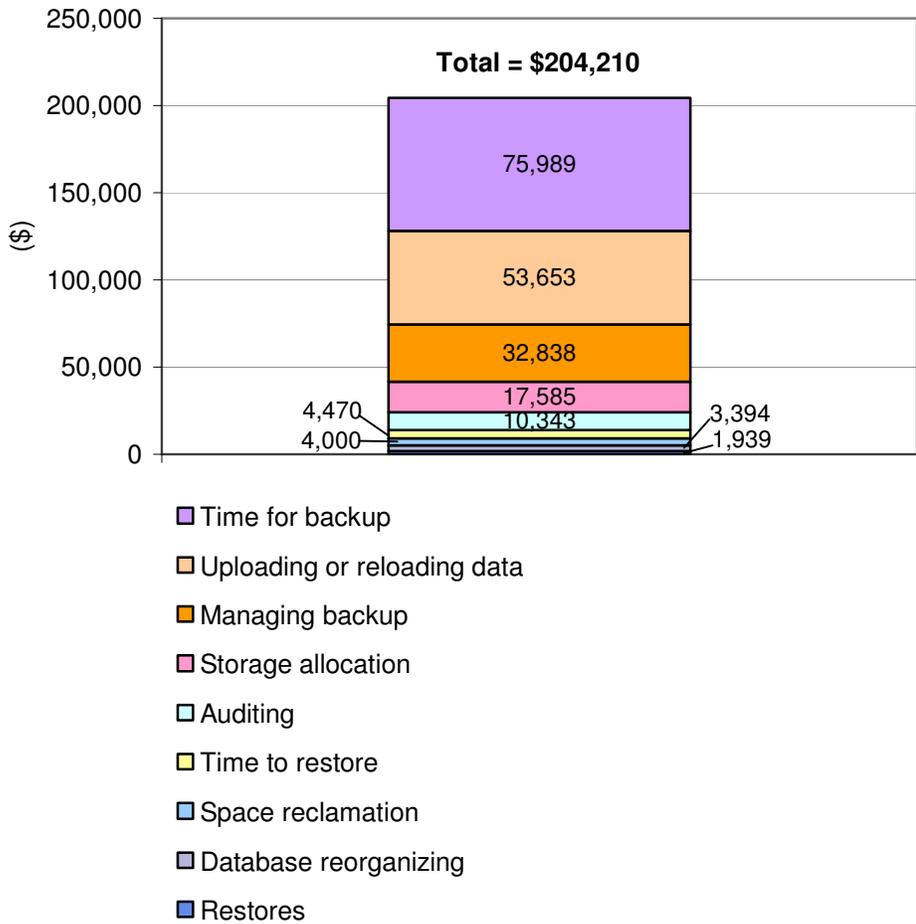
window is from 5:00 p.m. to 6:00 a.m., and our servers were still crunching data backups when they needed to be available for our users. Now it's all being done on a real-time basis. Also, we don't need people to swap tapes, which saves us 30 minutes a day at each of our locations."

One Data Domain customer is now saving approximately 47 hours a month, eight hours per week in back-office IT labor because the Data Domain solution is easier to administer and an additional five to 10 hours a month on managing storage allocation and then five more hours a month on space reclamation.

Data Domain systems have allowed customers to greatly reduce restore times. One customer was spending 24 hours a month managing backups at its main facility and now requires only four hours for this task. The company's more consistent backup strategy has also helped with audits. A manager mentioned, "We do audits every quarter, and it used to take nine days to complete. Now I can do it in one."

FIGURE 3

Increased Productivity



Source: IDC, September 2009

Improved Storage Management

By shortening the restore time and reducing data storage errors and failed backups, the Data Domain solution has helped the nine companies realize average benefits of \$123,277 annually.

The following are summaries of the narrative discussions with customers that deployed Data Domain systems and realized improved storage management benefits.

A manufacturing company went from 25 errors a month with its tape system to zero errors with the Data Domain solution. "We would lose an average of two hours each time an error occurred on our tape backups," the manager said. "We also lost data when a tape failed. That has never happened with Data Domain. Also, the four or five restores we do each week now take minutes each. Before, if the tape was onsite, it might take an hour. If the tape was offsite, it could take up to eight hours."

With its tape system, CN Railway would do 20 or 30 restores a week. "If the tape was onsite, a restore would take three to four hours. If it was offsite, and there was a 50-50 chance it would be, a restore could take two days," manager Kevin Whelan said. "Doing restores quicker is definitely business-enabling. Before, if a user accidentally deleted a report and had to wait two days, he or she would simply redo the report, wasting valuable time."

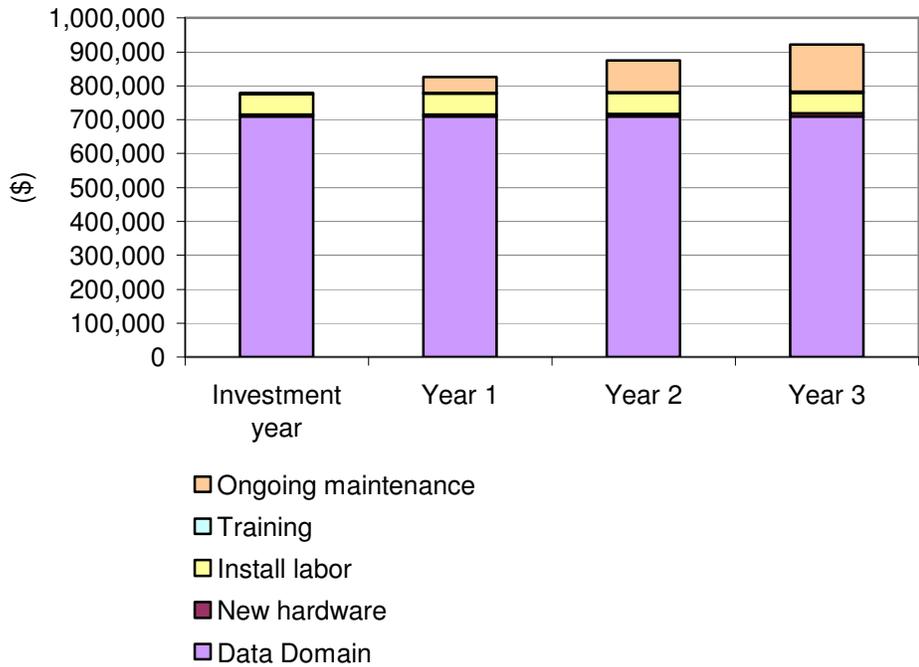
The media company has reduced the number of system errors significantly by replacing its tape system. "Recovery time is down, too, averaging about five minutes compared with 30 minutes before," the manager said. "Also, we save time on restores because they are instantaneous now."

Benefits, Investment, and Cumulative Cash Flow Summary

Figure 4 shows the three-year investment in the Data Domain solution. As is common in IDC ROI research, investments are highest in the deployment year, the year Data Domain systems are purchased. Once the initial cost is accounted for, investment in the following years declines and levels off over time. Ongoing investment is based on standard hardware turnover rates and IT time required to maintain Data Domain systems.

FIGURE 4

Three-Year Investment in Data Domain Systems

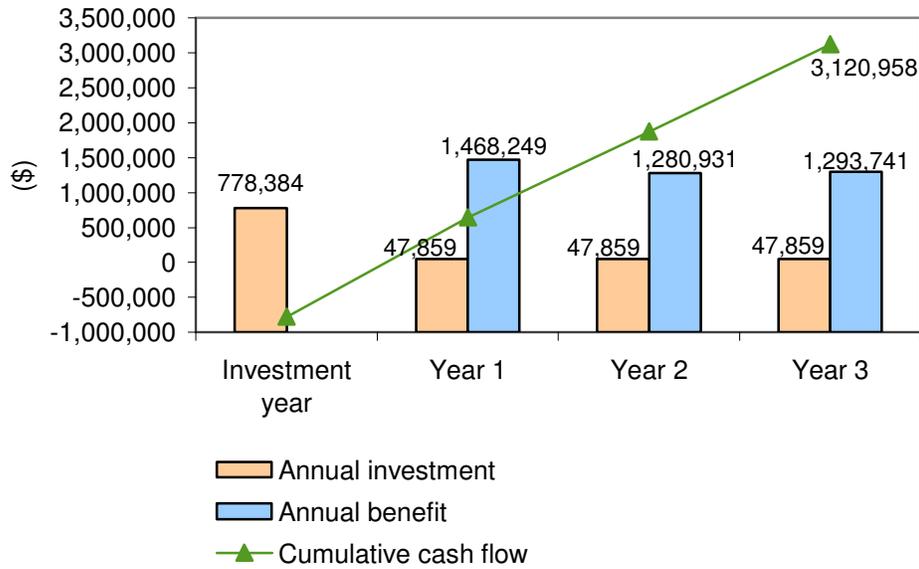


Source: IDC, September 2009

Figure 5 illustrates the aggregated annual benefits, investment, and cumulative cash flow over three years for the customers IDC interviewed in this study. Benefits are highest in year one because customers avoided tape purchases immediately after the deployment. Ongoing savings related to tapes include avoided tape upgrades and avoided drive maintenance.

FIGURE 5

Investment, Benefits, and Cash Flow



Source: IDC, September 2009

Return on Investment

The three-year IDC ROI analysis on Data Domain systems is based on initial and annual investments compared with the benefits over the three years. Based on the data gathered from the customers, this solution offers an ROI of 264% and payback occurs in 6.6 months. Table 2 displays the ROI results.

TABLE 2

Three-Year ROI Analysis

Benefits (discounted)	\$3,252,946
Investments (discounted)	\$893,334
NPV	\$2,359,611
ROI	264%
Payback	6.6 months
Discount rate	12%

Source: IDC, September 2009

IDC ROI METHODOLOGY

IDC's ROI methodology measures the efficiency of solutions uses the findings to calculate ROI for the deployed systems. The method includes four steps:

1. Evaluate the internal and external costs of administering the systems before deploying the solution.
2. Ascertain the investment in the purchase, implementation, and deployment of the solution. It is important to estimate not only the initial purchase cost but also the required implementation, integration, consulting and/or training costs. To measure the total deployment investment required, IDC includes questions that cover both the cost of purchase, setup and integration as well as ongoing software fees and IT maintenance time.
3. Measure the cost savings and gains in productivity, availability, and efficiency achieved using the solution. Portions of the interviews are dedicated to the discovery of cost reductions, including both "hard" IT costs, such as savings in server and back-up tape purchases and "soft" costs, including IT staff productivity, IT management efficiency and end-user productivity. For this study, we have modified our standard category labels to what is found in the parenthesis:
 - Cost reduction (reduced backup costs): IDC asks about what costs have been avoided or reduced for servers, back-up tapes, bandwidth, licensing fees and avoided travel. Savings are reported in terms of dollars per unit saved or annual reduction in spend.
 - IT staff productivity (increased IT productivity): To measure changes in IT productivity, IDC specifically asks about the reduction in time to restore databases or mission-critical business applications, the reduction in back-up window time, time spent auditing and uploading/reloading data.
 - Improved management of space (improved storage management): IDC asks the customers specific questions about the number of data storage errors per week, recovery times, failed backups per week, restore times and user productivity lost during restore.
4. Calculate the payback period and ROI for the deployed solution. Based on the aggregated interview data, IDC calculates the payback period and rate of return based on the overall cost savings resulting from the investments in Data Domain systems.

ROI and Payback Period Calculation Assumptions

IDC bases the payback period and ROI calculations on a number of assumptions, which are summarized below:

- Time values are multiplied by burdened salary (salary + 28% for benefits and overhead) to quantify efficiency and manager productivity savings.
- Downtime values are a product of the reduction in downtime hours multiplied by the number of users affected and their hourly rate.

- ☒ Because not every hour of downtime equates to a lost hour of productivity, IDC specifically asks about the percentage impact of an hour of downtime and attributes a fraction of the hourly result to the dollar savings.
- ☒ All IT solutions require a deployment period. The full benefits of the solution are not available during deployment. To capture this reality, IDC prorates the benefits on a monthly basis based on the average deployment term.
- ☒ The net present value of the three-year savings is calculated by subtracting the discounted three-year investment from the discounted three-year benefit. IDC uses a 12% discount rate to account for potential outlays made at the time of deployment and interest on that expense.

CHALLENGES AND OPPORTUNITIES

The insights EMC Data Domain customers provided clearly show that the use of disk-based data deduplication and replication services can have a dramatic impact on the costs and problems associated with backup/recovery and DR processes today. Organizations' data protection, DR, and long-term data management requirements will only increase in coming years; therefore, EMC will need to further extend the scale and scope of its Data Domain solutions.

One key challenge for companies is the need to accelerate the recovery times associated with large numbers of virtual machine images. As companies shift to a primarily virtualized server environment (and make broader use of desktop virtualization), they are increasingly concerned about potential "boot storms" triggered during full datacenter recoveries or start-of-day desktop boots. The Data Domain deduplication storage systems play an important role in helping manage this process, but the company must also ensure that it offers solutions that meet the intense data recovery performance required in a broad spectrum datacenter recovery.

Another key challenge for companies is the looming "archive problem." The digitization of many business processes (records, images, videos) is driving an explosion in the creation and storage of rich content. Concurrently, changing corporate governance requirements associated with regulatory compliance and eDiscovery are forcing companies to store this content for ever longer periods. These pools of rich content can quickly overwhelm backup and archive systems geared to support more traditional applications. EMC has the opportunity to help companies better manage fast-growing content archives through further enhancements in its Data Domain solutions (e.g., more content-specific services) as well as greater links to its broader storage and information management portfolios.

FINAL THOUGHTS AND ESSENTIAL GUIDANCE

IT managers at large and midsize enterprises are wrestling with many of the same challenges as they seek to boost IT asset utilization and improve business continuity/availability. They are looking for storage solutions that improve the usefulness of other IT assets while significantly reducing both capital and operational expenditures. Meeting these objectives, however, takes more than just identifying

the right use case and selecting the right solution. They also need to make wise implementation decisions that shield them from costly system reconfigurations and storage migrations.

IT managers who acquire storage products must evaluate suppliers such as EMC, as well as their business partners, based on more than just which company is providing the cheapest or highest-performing systems. IT managers need partners that can help them overcome or avoid the traditional shortcomings:

- Under use of installed assets
- Less than optimal data/application availability
- Excessive administrative overhead

IT managers should judge suppliers based on how well their solutions allow them to optimize the use of IT investments now and for an extended period. They also need to look for business partners that leverage emerging solutions to address specific application and business challenges while delivering faster, more consistent implementations with minimal risk of disruption to applications, processes, and business operations.

APPENDIX: DATA DOMAIN CUSTOMER CASE STUDIES

Case Study on Canadian National Railway

Overview

Canadian National Railway (CN) deployed EMC Data Domain deduplication storage systems in 2007 to improve backup and restore functionality using the solution's ability to deduplicate. As Kevin Whelan, Manager, Operating Systems and Storage, said, "We could see great advantages to Data Domain's approach. And we investigated a number of companies when we looked at our storage needs."

CN has over 22,000 employees with 725 staff supporting them. The IT organization has three dedicated staff managing Data Domain. The primary datacenter is based at the Montreal headquarters. The company's storage demand was constantly growing across all of its 27 North American offices. There is now a Data Domain appliance at each of these locations and one central system that replicates to the primary datacenter.

CN is using Data Domain in two ways: First, the IT organization replicates data from its many satellite locations back to the Montreal office. Second, the company uses a dedicated Data Domain system in Philadelphia for disaster recovery (DR). Prior to the implementation, CN would transport all of the disaster recovery tapes to Philadelphia by truck.

Reduced Backup Costs

The Data Domain deployment was a good fit for CN's business needs as the company's current tape drives were at the end of their life cycle and some form of storage upgrade was a necessity.

CN discontinued tape services and maintenance contracts, which saved the company approximately \$100,000 per year. The company has also avoided repairing tape hardware failures — resulting in an additional estimated savings of \$175,000 annually.

The company has avoided purchasing between 50 and 60 new tape drives total while meeting the continually growing demand for new storage. Company managers estimate they are saving over \$100,000 on new tapes annually, while storage demand grows at an average rate of 25% per year.

As Whelan said, "We are saving on bandwidth because of Data Domain systems. Some of our locations were running at very slow speeds on 256K circuits. If we continued transmitting entire backups over that link at the same time that my users are online — everyone would suffer." To improve transfer rates, the company estimates it would have to spend up to \$30,000 more per month — that cost is avoided due to Data Domain replication efficiency.

Increased IT Productivity

Before the Data Domain deployment, all of the CN sites were using tape backups. As Whelan said, "Each location would have someone that would go in the morning, take a tape out of the drive, and then mount the next tape. That took 15–20 minutes by the time they went to the server, took out the tape, filed it away, put in the new one." But since the implementation, those tasks have been completely eliminated. The company is saving an estimated nine hours a day across all of its offices.

The IT organization does 20–30 data restores, on average, per week. Before the Data Domain deployment, if a tape was onsite, then the job would require three to four hours of work. But if a tape was offsite, restore time quickly increased to two days. With Data Domain replication now in place, the company is still well protected from site disaster, and restores are completed in two hours since recoverable data is always onsite. In a data recovery situation, CN would avoid approximately 20 hours of vendor transport time because tapes are no longer moved between Montreal and Philadelphia.

Improved Storage Management

CN has approximately 22,000 users that benefit from the Data Domain deployment. In the past, if a report was accidentally deleted, the user would have to recreate it from scratch. But with Data Domain the report can be restored in less than two hours, allowing the user to avoid re-inputting that data. As Whelan said, "If a user needed a report the same afternoon, they would just start redoing their work because the old restore process was too long for them to wait. Now, they don't panic because they know all the data is going to be there for them."

Data Domain includes a bandwidth throttling capability which allows the IT organization to determine the data transfer rate. Obviously, the organization does not want to perform a large backup during peak usage. Using this feature, CN starts backups at the end of a business day and completes them before 8:00 a.m. the next morning. During typical work hours, the staff slows the Data Domain transfer rate so that performance is not adversely affected. Whelan said, "The ability to control the only transmission rate was a tremendous feature for us."

CN has saved a significant amount of disk space using Data Domain. In the past, the IT organization would write the same piece of data to tape every day. But since deploying Data Domain deduplication, data is written to only once. As Whelan said, "If the data has already been saved once, it will not be saved again."

Offsite data management has been improved since CN has moved away from using an outside vendor. On any given day, IT productivity could be negatively impacted by simple logistical issues. As mentioned in the interview, "One day we asked for the vendor to bring us an offsite tape so we could do a restore. He said, 'As soon as they finish changing the oil in my car. The tape is in the trunk, and it's up on a lift right now.' With Data Domain, we don't run into issues like that."

Case Study on Global Communications Corporation

Overview

A global communications firm deployed EMC Data Domain deduplication storage systems because of the system's highly efficient means of data backup. The company's demand for storage was getting so high that the IT organization simply could not afford to continue backing up to tape. As a Senior Technical Director, said in the IDC interview, "We chose Data Domain because of the way that they compress data and the way that they deduplicate data. It has the ability to take a gigabit at line rate for the duration of the backup window. The algorithm that Data Domain runs lends itself to very high throughput numbers right to their published specification. Their specifications are 200 megabytes/second for the DD500 Series, and those machines will do that all night long."

The company has been using Data Domain for over a year and is in the process of deploying more systems at numerous facilities.

Reduced Backup Costs

Data Domain is now the primary backup target, and the company has significantly reduced the amount of tape and tape hardware used across the organization. As the Technical Director said, "In most cases, Data Domain completely replaced tape. In our five new locations where we started with Data Domain in the first place, we never even deployed tape."

The company estimates that tape and the associated hardware savings are approximately \$385,000 per site annually. As he said, "When you look at those savings over our 27 sites in this deployment, this has been a good value." This is a savings of \$10.4M annually — \$30M savings over a standard three-year hardware life cycle.

By using Data Domain, the company purchased smaller backup servers than it otherwise would have. The director said, "We have not avoided new servers altogether, but the ones we did buy were smaller, cheaper units than we would have normally purchased. We are saving about \$35,000 per server, per site." This is equal to an additional \$945K in annual savings.

The firm has also been able to reduce potential risks during auditing. Since deploying Data Domain, it is possible to vault and ship data offsite, which keeps the company in compliance.

Increased IT Productivity

The company's staff managing backups has avoided time spent studying trends on tape usage, replenishing tapes, labeling, and then storing the backups. The company has three dedicated staff working with tapes, but their time spent on these tasks has been reduced, "by 15% — and they can work on other tasks for us now."

Data backups take less time than in the past because the data transfer rate of Data Domain systems is faster than that of tape. The Company estimates that in an average week, 22 staff members are saving at least 10 hours loading data. In the past, restores would take between four and six hours. Since the deployment, that time has been reduced to less than one hour.

The Data Domain deployment required only about two hours to install the storage units and needed no further maintenance time. This is more efficient than other solutions. As the director said, "With other storage, we would allocate and continue to allocate based on backups and backup servers. A lot of the benefit of Data Domain is that it takes many of the tasks multiple teams were doing in the past and consolidates them so that one person is able to do it." The company estimates the IT staff is saving 15 to 20 hours per Data Domain system, per year in storage allocation tasks.

Improved Storage Management

The director summarized, "Data is now more portable. It can be backed up quicker, and that actually has a dramatic impact on the organization. Data Domain has changed the dynamics of how our teams work. We are starting to see a shift in personnel as we take tape libraries off of the floor and reduce pressures on that team. It is all because of electronic vaulting — my team can now handle all of the storage based on the efficiencies that Data Domain brings."

Case Study on Financial Services Corporation

Overview

The financial services corporation, managing more than \$180B in education loans, first deployed EMC Data Domain deduplication storage systems in May 2008 — a decision that was driven by the company's need to reduce costs by eliminating its dependence on tape backup. At its central location, the company had three people managing backups.

Reduced Backup Costs

The two greatest factors in reducing costs were savings on tapes and savings on the associated maintenance. The company's Technical Architect estimates that the corporation was able to save over \$800,000 in tape expenses and \$100,000 in maintenance costs in the first year after deployment. The company has 12 remote sites that purchased new tapes regularly and handled their own separate tape backup environment. Since implementing Data Domain, expenses including tape purchase, drive maintenance, and transportation have been completely eliminated.

Because the corporation uses Data Domain, it has avoided a tape refresh that would have cost over \$500,000. The company has been able to collapse all of its tape libraries or eliminate them altogether. The 12 remote sites no longer use any tape at all — that data is now housed in the main datacenter.

The company has been able to avoid purchasing at least three media servers per year, which has resulted in a savings of over \$30,000 annually. In addition, the company no longer faces annual software renewals and maintenance renewals related to its old backup system. The Technical Architect estimates those savings total \$25,000 annually.

Increased IT Productivity

At each site, the company has increased its IT productivity by one full time equivalent (FTE). If that time savings is carried out to all 12 sites, the benefit is over \$700,000 in staff time.

As the Technical Architect said, "We are no longer constantly repairing bad tape drives and mechanics. We have reduced our transportation contract and are saving our administrators three to four hours per day for not having tape."

In order to protect against potential tape failures and minimize the amount of time spent waiting for restores, the company duplicated all of its tapes and kept one set in-house and another offsite. The company would back up to these duplicate tape sets three or four times a month in the past, but it avoids the process altogether since the Data Domain deployment. In addition, there has been a significant improvement in the amount of time required to perform a restore — that time has been reduced from one hour to just 30 minutes.

Before the Data Domain deployment, it was common for the company to experience one tape failure per day. Repairing the average failure required about three hours of an internal technician's time. Plus, a company administrator would have to coordinate with a vendor to help solve the problem, which would typically add another three hours to the process. As the Technical Architect said, "We had many, many tape failures. Those would result in loss of staff time and loss of data on that tape. About 5–10% of the time, the tape was just unrecoverable, and we would have to look for another copy. And if we had to go offsite for that, it was about a half a day to find the tape."

Data Domain's capacity allows the company to recover data faster and to points in time earlier than using tape drives. As mentioned in the interview, "One time some tapes fell down an elevator shaft. The tapes were being brought out by the disaster recovery team — about five of them fell all the way down the shaft. And one of those tapes was critical and not backed up."

The Technical Architect also said, "With the number of restores we have, Data Domain has gained a good reputation with us. It is known to be a reliable restore solution. People are happy when they see that their data is on a Data Domain appliance, as opposed to a tape appliance."

Case Study on Global Equipment Supplier

Overview

A global equipment supplier for the plastics industry manufactures custom injection molding machines for its clients. Customers use these custom molds for a variety of products, including automobiles, cell phones, toys, and home storage units.

The company was at a crossroads in 2005 when it selected EMC Data Domain deduplication storage systems. Storage requirements dictated either refreshing the entire tape environment or finding another technology that could serve its growing data needs. A challenge facing the company was to either purchase new tapes in LTO-1 format, which was no longer supported, or find a new solution altogether.

Deploying Data Domain required a minimal amount of time — the time between being taken out of the box and the staff performing backups was only about eight hours. The company spent one additional day validating the data prior to putting the Data Domain systems into production.

In the first three months after deploying Data Domain, the firm was able to pull approximately 1,000 tapes out of circulation. The company keeps one of its tape libraries active for data restores from legacy tapes only, and as an IT Systems Manager said, "We are just letting those old tapes expire naturally. We don't write anything to tape anymore."

Reduced Backup Costs

The company has avoided purchasing new tapes and new tape drives since the Data Domain deployment. The company estimates that over the past five years, it has avoided purchasing at least 12 tape drives that would have cost \$11,000 each.

The company has also avoided server expenses and increased its current servers' utilization rates. The IT Systems Manager mentioned, "We have avoided buying servers and we also find our backup time is faster and servers are running a lot harder." Prior to the Data Domain deployment, the tape drives were causing a bottleneck and server utilization rates averaged 20% to 25%. But since the implementation, backup times have been improved and the average utilization rate has been increased to roughly 80%.

Data Domain systems have created greater data management efficiencies, which have allowed the company to avoid hiring new IT staff. The manager said, "It has allowed us to keep up with the workload. Prior to bringing Data Domain in, my group was restructured and two people were moved to another team. So, for me, the question was 'How am I going to make up their hours?' But after the install, we managed with no problem — even after losing those two."

Increased IT Productivity

Since the Data Domain deployment, the company is saving IT staff hours when performing backups — instead of spending about 40 hours per month managing backups, the staff now requires only four hours per month. The IT staff is now able to pursue more business-enabling tasks. As the IT manager said, "Basically, since we are not spending all of our time doing backups, we can spend about 10 more hours per week on business projects and system maintenance."

The company experiences fewer data, system, and tape errors since the implementation. The company estimates that one in every 100 tapes completely failed, and in the past, about 25 basic tape errors occurred per month. "To address an event, it could have taken a number of actions to completely repair. If we had a physical drive problem, we could spend 8–12 hours for just one event." Since the Data Domain deployment, the company is saving at least two hours per month on average just by avoiding errors related to tape.

The company now has more consistent backup windows and reduced restore times. Before, when writing to tape, backups could take as long as six hours to run and the restore time was between 12 and 14 hours. Since deploying Data Domain, backups and restore time have been reduced to only four hours each. As the IT manager said, "There was a huge difference in our restore time, especially on things like databases. Our file systems run in about half the time."

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2010 IDC. Reproduction without written permission is completely forbidden.