IDC
Analyze the Future

WHITE PAPER

Demonstrating the Business Value of StorNext

Sponsored by: Quantum

Randy Perry Ashish Nadkarni

August 2012

EXECUTIVE SUMMARY

As the data created and consumed by users worldwide continues to double every two to four years, the role of scale-out file systems and policy-based data management solutions such as Quantum's StorNext is becoming critical to managing this growth and streamlining data workflows.

Quantum recently commissioned IDC to ascertain the financial benefits of deploying a StorNext solution in environments where unstructured and semistructured (machine-generated) data management has historically been a daunting task. IDC interviewed 10 organizations (based in North America and Western Europe) using Quantum's StorNext data management solutions. Based on the data gathered via interviews, IDC found that organizations that deployed the Quantum StorNext solution stood to gain a significant return on their investment.

Business Value Highlights

Over three years, Quantum StorNext customers interviewed for this study enjoyed benefits of over $423,000 per 100 users, or returns of $6.47 for every $1 invested. Additional analysis includes:

- Three-year ROI: 547%
- Payback period: 4.2 months
- IT productivity increase: 69%
- Service disruptions decrease: 82%
- Operations productivity increase: 14%

Highlights of the financial analysis include:

- Reducing service disruptions/outages, backup and disk failures, and capacity issues improved end-user productivity by 82%, saving companies over $141,000 per 100 users annually.
- Speeding up data flow and enabling multiple users to access data simultaneously, organizations using StorNext increased their operations productivity by 14%, gaining over $50,000 per 100 users annually.
- Optimizing the utilization of data storage, companies were able to save an average of $43,423 per 100 users annually.
- Increasing IT storage staff productivity by 69%, companies saved $2,520 per 100 users annually.

Over three years, Quantum StorNext customers interviewed for this study enjoyed benefits of over $423,000 per 100 users, or returns of $6.47 for every $1 invested, and a payback in just over four months.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

SITUATION OVERVIEW

The file-based storage market is undergoing a silent change. This market was originally dominated by monolithic dual-controller NAS appliances and by file servers (servers sharing their local or SAN file systems). The market soon accommodated a new breed of scale-out file-based storage solutions — hardware and software solutions that were powered by distributed file systems. Such distributed file systems were even available as multiplatform software solutions.

As the data created and consumed by users worldwide continues to grow year over year, the role of scale-out or distributed file systems such as StorNext is becoming critical. With data management and automated tiering capabilities in heterogeneous client environments, file systems such as StorNext are taking aim squarely at organizations with burgeoning data growth, especially those in data-intensive industries such as media/entertainment, natural resource exploration, surveillance, life sciences, and scientific research.

In its recent 2012–2016 forecast on file-based storage (see Figure 1), IDC expects the revenue for scale-out file-based software to grow at an average annual rate of 50% and account for 55% of the file-based storage market segment by 2016.

Organizations have grasped the benefits of scale-out solutions. Scale-out solutions support data growth by making it easy to add new storage "nodes" when extra capacity and performance are needed. Because the multiple nodes in scale-out solutions behave like a single larger file system, known as a single global namespace, organizations can scale more quickly using fewer resources.

FIGURE 1

Worldwide File-Based Storage Revenue, 2007–2016

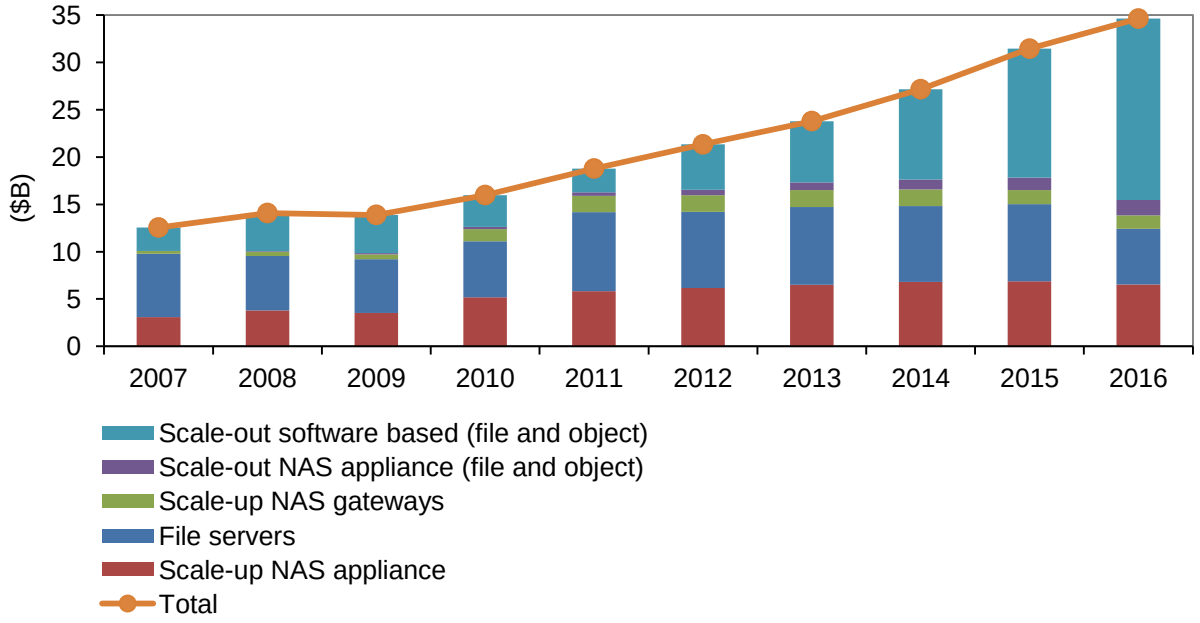

Source: IDC, 2012

 ©2012 IDC

QUANTUM'S STORNEXT PORTFOLIO

As mentioned in the previous section, Quantum's StorNext scale-out file system offerings are targeted specifically for organizations in data-intensive industries. Scale-out solutions, as the name suggests, are designed to handle huge data files, which implies that they need to handle scale from a capacity and performance perspective. IT organizations use products like StorNext not only to store data for data-intensive applications but also to leverage the products' capabilities to perform time-sensitive processing on this data. This means that the distributed file system needs to have an efficient and optimized mechanism for data location, organization, and shared access. Additionally, it needs to support a robust data management and tiering framework that includes built-in data protection and archiving capabilities.

Quantum has designed StorNext as an enterprise-grade solution keeping these requirements in mind. StorNext's portfolio features:

- StorNext File System software delivers shared high-performance file access in heterogeneous environments.
- StorNext Storage Manager provides policy-based tiered storage and archiving capabilities.
- StorNext appliances provide a packaged approach for high-performance workflow productivity.

StorNext software consists of two core components: StorNext File System and StorNext Storage Manager. The StorNext File System is enabled by a SAN client that resides on a Windows, Linux, Unix, or Mac OS X platform that is connected to a SAN. The centralized storage can also be shared out over Ethernet through gateway nodes to clients running StorNext distributed LAN software. A StorNext Metadata Appliance or Metadata Controller (MDC) runs on a Windows or Linux server and acts as the "traffic cop" for shared heterogeneous server access to the data pool on the shared storage. Additionally, gateway servers can allow additional clients high-speed access to a federated namespace across LAN networks using Quantum's proprietary Distributed LAN Client.

StorNext appliances include Metadata Appliances for high-performance file sharing, StorNext Archive Enabled Library (AEL) Archives for cost-effective nearline storage and offline data retention, StorNext Gateway Appliances to enable fast and highly available access to the central storage over Ethernet, and StorNext Q-Series Storage for high-speed primary storage in big data environments.

As pointed out in *The Future for Namespaces in File-Based Storage* (IDC #236010, July 2012), a key component of scale-out file systems is namespace federation. The StorNext File System, enabled by the MDC, simultaneously shares a common namespace and data from the centralized storage pool with the applications across different server platforms. This multiplatform support enables support for a diverse portfolio of applications, giving organizations the freedom to select a platform that best suits their application needs.

The StorNext Storage Manager, which runs on the MDC, is a policy-based data mover that automatically and transparently copies or moves data to another tier of storage. Files can be truncated in the primary file system, leaving behind a pointer, or

can remain whole with a copy in a lower storage tier. Since this movement is all done within the StorNext File System, the data appears local to the application accessing it. If an archived file is requested, the Storage Manager fulfils the request and presents the file back to the application as if it were resident in the primary file system. Policy-based tiering and archiving reduce the overhead on IT staff who do not have to manually identify and move data from one tier to another as its heat index changes (i.e., as the data goes from being "hot" [always accessed] to "cold" [hardly accessed]).

BUSINESS VALUE

Study Demographics

In the spring of 2012, IDC interviewed 10 organizations that had deployed Quantum's StorNext solutions. The organizations were small to medium sized with an average of 1,571 employees (see Table 1). Most organizations had high IT user to employee ratios (94%), and the universities included in the study had 4.4 students per employee so that the average number of IT users exceeded the average employee count (1,833). The study base represents a broad range of experiences from organizations across North America and Europe and from verticals including media, research, education, life sciences, satellite imagery, and government. The interviews were designed to elicit quantifiable information as well as anecdotes so that IDC could interpret the full impact of StorNext solutions on the organization.

TABLE 1

Quantum StorNext Customer Study Demographics

Organizational Unit	
Average number of employees	1,571
Average number of IT end users	1,414
Average number of terabytes of storage in use (total IT environment)	5,198
Annual growth rate for storage (%)	220
Average number of IT staff	49
Industries	Media, government, research, education, life sciences, and satellite imagery

Source: IDC, 2012

These organizations were remarkable in their intense data requirements and the central role data played in their operations. On average, each organization had storage assets of 5.2PB, or 3.3TB per employee, which is 1,000 times the industry average. These organizations selected Quantum's StorNext to solve data issues common to such data-intensive operations. Three of the customers mentioned data growth and the need

 ©2012 IDC

to scale. Two of the organizations needed to streamline their workflows. Many talked about dealing with the complexity of their operations or IT environments:

- Running multiple data types on the same system
- Needing access to multiple dissimilar operating systems — Linux, Unix, Mac, and Windows — at the same time
- Running simultaneous operations with multiple users and types of data content (text, audio and video, etc.)

FINANCIAL BENEFITS ANALYSIS

Overall, the organizations recognized financial benefits from the following areas (see Figure 2):

- **Optimized IT staff productivity.** By reducing the time IT staff spent dealing with storage management for servers, media management, and backup activities by 69%, organizations were able to save $2,520 per 100 users annually in IT labor costs.
- **Enhanced end-user productivity.** The most significant benefit was the improvement in the data-driven business operations. End users benefited from reduced service disruptions, increased speed of operations, increased capacity, and reduced recovery and restore times as well as better server performance, resulting in savings of $141,098 per 100 users annually.
- **Reduced storage environment costs.** StorNext enabled a more efficient storage environment, lowering capex and opex costs directly associated with storage by $43,423 per 100 users annually.

FIGURE 2

Quantum StorNext Customer Average Annual Financial Benefits

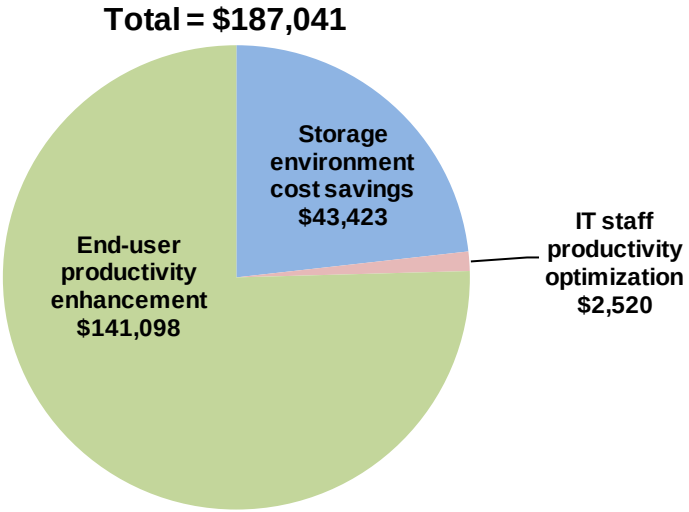

Note: Data is per 100 users.

Source: IDC, 2012

Optimized IT Staff Productivity

IT organizations today run very lean — really to the bone. The organizations in our study are data-heavy operations with significant dependence on using data resources to deliver products and services. Storage requirements ranged from 100TB to 22.5PB, with an average of 4.4TB per employee. Furthermore, these organizations were doubling their data requirements every six months. Yet, thanks to StorNext, they were able to manage their storage operations with 1.15 IT staff (10% of total IT staff).

StorNext sped up the analysis and enabled backup to be done in real time through StorNext's continuous data protection policies. This automated the day-to-day functions so that, as one organization put it, "once it is set up it's done." Table 2 shows that organizations interviewed for this study were able to reduce the time IT staff spent on management and support tasks by 69%.

Overall, among the organizations interviewed, the improvement in IT efficiency associated with better management of data resources contributed $2,520 per 100 users in annual savings.

TABLE 2

Quantum StorNext Customer IT Staff Productivity Optimization (%)

Management/Support Task	
Uploading or reloading data	98
Planning/purchasing	91
Storage provisioning/allocation	82
Managing backup	61
Space reclamation	51
Managing capacity	26
Total	69

Note: Data represents reduced time IT staff spent on task based on hours per 100 users.

Source: IDC, 2012

Enhanced End-User Productivity

As mentioned previously, the organizations interviewed were heavy data users, so improving data management had its greatest impact on the day-to-day business operations. Those interviewed, on average, realized a gain of nearly 56.7 hours per year for each end user through productivity enhancements that fell into two categories: reducing disruptions to operations and increasing productive output.

In the first category, fewer disruptions increased user productivity by providing more operating time and increased confidence in the operation and included the following:

- **Reduced service disruptions/outages.** Limiting the primary causes for storage-oriented outages such as backup and disk failures and capacity issues by 82% added 27.1 hours.
- **Reduced capacity shortfalls.** Eliminating planned outages required to move or add capacity added 3.1 hours.
- **Reduced disaster recovery time.** Faster operations, including backup and disaster recovery, added 8.5 hours.

The second category, increasing productive output, is aligned with the increasing speed of operations. StorNext's high-throughput content sharing enabled multiple users to access data simultaneously, increasing overall dataflow. This, in turn, increased their operations productivity by an average of 14%, adding 17.9 hours. Customer observations include:

- **Government.** "The dataflow is faster — much faster. The analysts now have much faster access to the data. Before, they could spend 10–15 minutes sometimes, waiting for large amounts of this data to load up into their computers. Well, now a lot of it is fiber based, and quick access and all of that, with the SAN … now they can pull stuff up in a matter of seconds. I'd say that we're getting at least a 10–20% efficiency increase."
- **Life sciences.** "It [dataflow] is faster now. Before, we only had one entry to the file system. Now, with this clustering system, we have 10. So we have more bandwidth; we have more access to the data. A lot of the stuff that we do with this is statistical analysis with the data that we have. Let's say in a specific country, instead of picking a random 10 data points, we can do 1,000. So that's enabling us to do more accurate research."
- **Media.** "We can do twice as many projects. The way that our business functions is that projects are based on a timeline. So before, a motion capture, or something like that, used to take a month. Now we can now do it in two weeks. So we have reduced the time by 50%."
- **Instrumentation manufacturer.** "It [dataflow] is lots faster. When we did the first evaluation, the engineer who was testing it, said, 'This can't be working, it must be broken …' I said, 'Why?' He said, 'Because it's too fast.' We measure dataflow from time to time, comparing StorNext with our other three systems that we can benchmark. Everything from small fiber arrays to drives you could carry around in your briefcase. StorNext is easily like 50 times faster than the drive in the briefcase, and maybe 10 times as fast as the small fiber array, because it is like having 10 small fiber arrays all glued together."
- **Research.** "Dataflow rate: that was vastly improved. We changed the network. The system itself is faster than the previous system. Anywhere from 25–50% faster … that's communicating with the CIFS and NFS gateways. We also deployed Quantum's Distributed LAN Client for StorNext, which allows us to connect directly to the StorNext host and bypass all of the inefficiencies of NFS and CIFS file system overhead."

- **Media.** "So I'm thinking that the time savings there is like about 10–20%. But it's really the capabilities that it provides us. Previous to StorNext, we just didn't have any shared storage for the video guys. We came from a place where we had one video editor, so shared storage didn't matter. However, even when we got to five editors, they were still working in that model. They were basically storing things locally. So if editor A started working on a project, it's on his computer, and he was the only one that could get to it because we didn't have shared storage. They could copy things up to a shared volume on the server, which anyone in the office would be familiar with. But they couldn't edit from it. Now, the guy's working on the file and the file is sitting on the StorNext, and the next guy that comes to work on the file can do that without switching LUNs and switching disks and all of that. The users save time because it looks like a local drive … a little silver icon. This guy is copying it here. The other guy is copying it there. Not only does StorNext save them time, but they wouldn't have changed to the current process, because it would take so much time without that shared storage … that intelligence that StorNext provides."

With Quantum's StorNext solutions, total enhancements to end-user productivity returned an average of over 39,000 person hours of productive time annually, saving $141,098 per 100 users (see Figure 3).

FIGURE 3

Quantum StorNext Customer End-User Productivity Enhancements

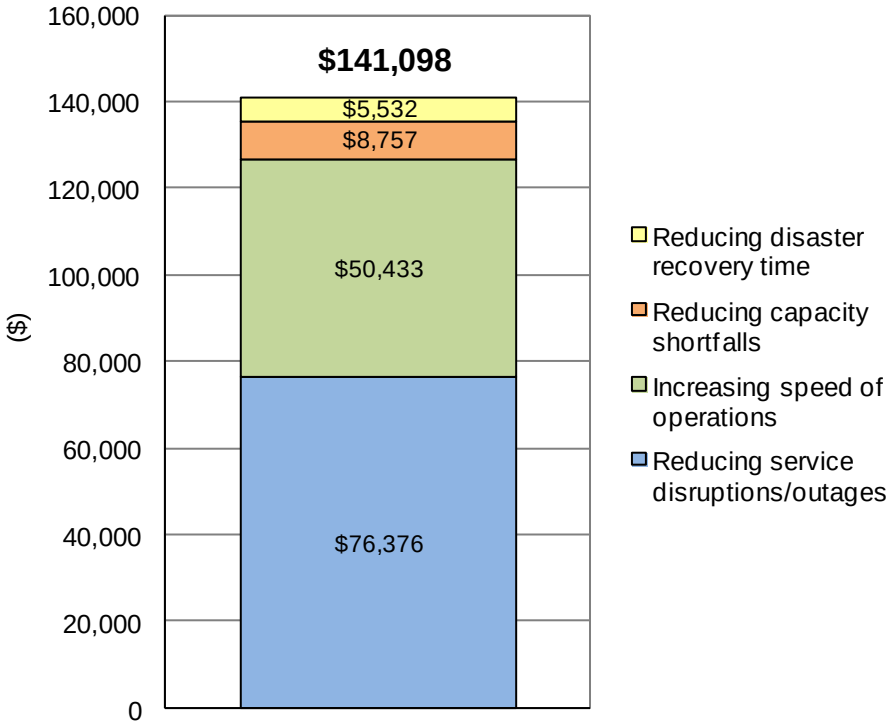

Note: Data is per 100 users.

Source: IDC, 2012

 ©2012 IDC

Reduced Storage Environment Costs

The organizations interviewed were adding 50TB to 1.0PB annually, increasing the requirements to manage and store their data assets by 50–900%. StorNext allowed them to consolidate storage in some cases and to slow the increase in storage related costs (see Figure 4). In more detail:

- **Storage management consolidation.** StorNext changed the way the organizations architected their storage environments. StorNext allowed IT to consolidate storage management, eliminating the "islands" of departments and connecting everyone to one system. StorNext's greater flexibility also contributed to consolidation. As one company stated: "The reason that we went with StorNext is the protocol [it's protocol neutral]. It allows flexibility from a vendor perspective, as far as who makes the storage. So StorNext gave us the flexibility to pick and choose a vendor that fit in a price point for us … and performed in a way that we needed."

- **Disk savings.** Storage requirements are unavoidably growing, so the key is growing efficiently. Most of the companies interviewed for this study used both disk and tape. StorNext extended the use of tape for backup and archive, reducing the need to purchase disk and saving these companies 30–40%. One company estimated that it did not have to add 100TB of storage last year because it was not using its capacity to the fullest.

- **Server and services reductions.** Because organizations were able to exploit storage hardware, they needed fewer application servers — in some cases, half as many. StorNext Storage Manager also enabled data protection and archive to be done on the same file system, speeding up the process, which enabled some of the organizations to drop third-party backup services.

- **IT staff savings.** More efficient storage management meant IT staff did not have to grow at the pace of storage growth.

- **Facilities savings.** Reducing the physical storage requirements resulted in corresponding reductions in datacenter racks, storage libraries, and power requirements. Some of these reductions were quite substantial, as one company revealed: "So we were able to consolidate our physical library infrastructure into one system for two primary applications. Then I didn't have to have two different libraries. I avoided buying a separate library and maintaining a separate library, which saved about $500,000."

- **Reduced storage capacity requirement with dedupe.** One organization estimated it was able to reduce its required backup and archive capacity by as much as 93% because of StorNext's deduplication capabilities.

FIGURE 4

Quantum StorNext Customer Average Annual Cost Reduction Benefits

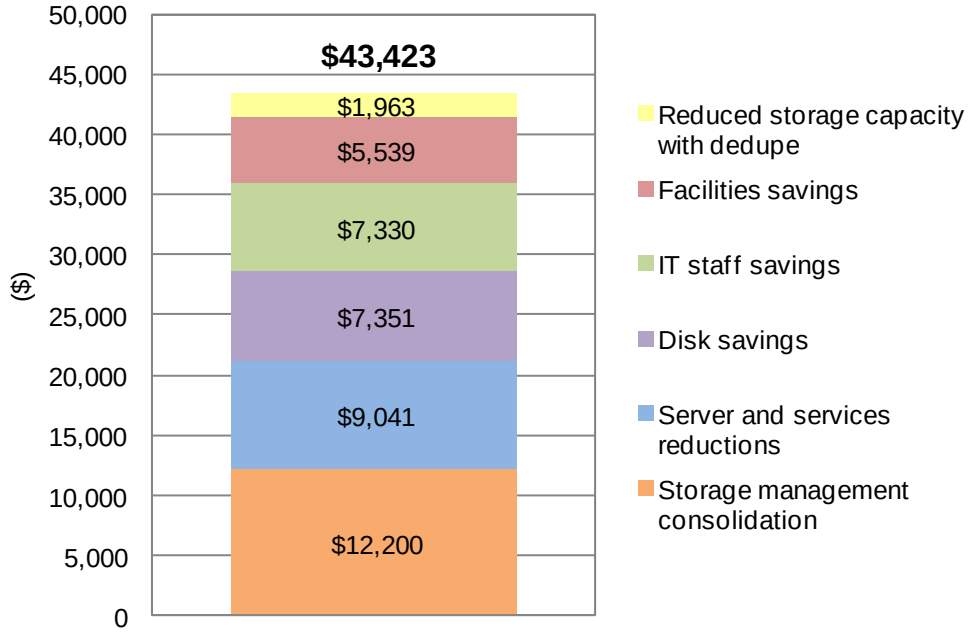

Note: Data is per 100 users.

Source: IDC, 2012

ROI ANALYSIS

IDC reviewed the cash flows of the financial benefits compared with the investment or total costs of the IT initiatives over a three-year time period to assess the return-on-investment (ROI). The initial investment included the average total costs to purchase and deploy the new systems, both hardware and software; consulting and other third-party services required to architect and install the new systems; the costs and lost productivity associated with IT training; and the IT labor required for installation and migration of applications. Annual costs were for support and upgrades. Annual benefits included infrastructure savings, reduced IT labor support, increased end-user productivity, and added income.

Figure 5 shows the undiscounted cash flow analysis. Organizations interviewed for this study made an initial investment of $53,465 per 100 users. Based on that investment, the organizations realized average annual benefits of $187,041 per 100 users. Over a three-year period, this represented an average cumulative net savings for each organization of over $464,147 per 100 users.

 ©2012 IDC

FIGURE 5

Quantum StorNext Customer Three-Year Cash Flow Analysis

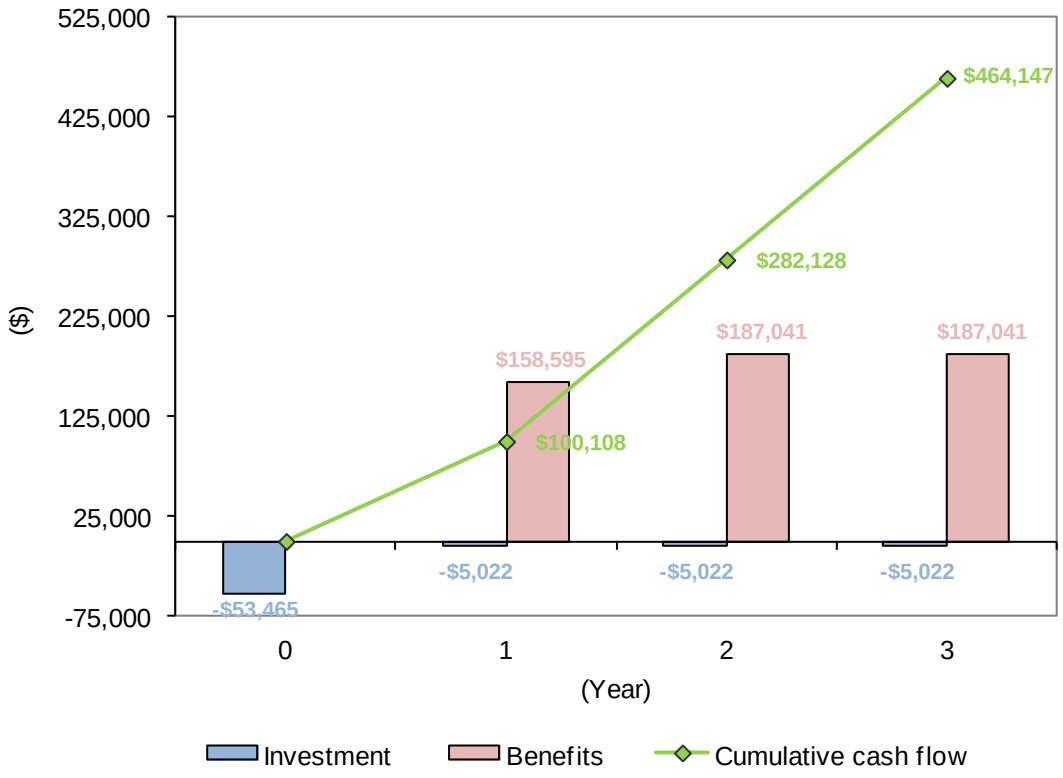

Note: Data is per 100 users.

Source: IDC, 2012

Table 3 constitutes a three-year view of the financial benefits of Quantum's StorNext solutions on a per-100-users basis. IDC uses a 12% cost of capital to discount cash flows.

Each of the organizations interviewed for this study experienced strong returns on its investment. The three-year ROI analysis shows that, on average, the organizations spent **$**65,473 (discounted) per 100 users on technology refresh and received $423,844 (discounted) per 100 users in benefits, for a net present value (NPV) of $358,370. On average, the organizations saw a payback in just over four months (after deployment) and an exceptional ROI of 547%.

TABLE 3

Quantum StorNext Three-Year ROI Analysis

Organizational Metric	
Benefit ($M discounted)	423,844
Investment ($M discounted)	65,473
NPV ($M)	358,370
ROI (NPV/Investment) (%)	547
Payback (months)	4.2
Discount rate (%)	12

Note: Data is per 100 users.

Source: IDC, 2012

CHALLENGES

Quantum's StorNext offers a key set of benefits for organizations that struggle with explosive unstructured data growth. Often the biggest barrier to adoption of new solutions such as StorNext is in understanding the quantitative benefits the solutions offer. As organizations seek innovative approaches to data management, they will consider solutions such as StorNext as long as they can articulate the value of the investment vis-à-vis the challenges they face with data management as well as the challenges posed with integrating the solution into their existing infrastructure. In this respect, Quantum may find it easier to get StorNext inserted into the stack for greenfield deployments as opposed to changing existing deployments over.

StorNext is a highly adaptable and efficient platform; however, Quantum may do it some amount of disservice by positioning part of its product line as a one-size-fits-all platform. Quantum should stay laser focused on the key performance and scalability benefits offered by StorNext in unstructured environments.

CONCLUSION

The convergence of exploding data growth and the need to control costs confront IT system decision makers. Nowhere is this data growth more profound than in the area of unstructured and semistructured (machine-generated) data — which is largely driven by the proliferation of video and high-resolution data sets used in various industries beyond just media and entertainment. Geographic dispersion and the need to have information at one's fingertips further add to the challenge of managing growth and providing disruption-free service.

 ©2012 IDC

IT continues to be pressured to reduce spending on new equipment and administrative resources while increasing operational efficiencies and meeting the demands of this data growth and availability. Organizations must deploy storage solutions that are part of ongoing efforts to optimize operations and improve IT asset utilization. Challenges can be met and goals achieved through informed implementation decisions that avoid costly storage migrations.

To that end, Quantum's StorNext solution alleviates these challenges by creating a global namespace for unstructured data management. As noted in this white paper, businesses deploying StorNext can benefit from increased storage efficiency and resource utilization and lower administrative overhead. This drives down IT costs while ensuring that the infrastructure is rightsized.

Quantum, a long-established player in data protection and data management, is well positioned to address today's data management challenges with StorNext. The company's broad range of open systems, scalable StorNext software, and appliance platforms support cost control, risk mitigation, and asset utilization. They are designed for a variety of environments and use cases to meet customer demand for comprehensive and flexible data management solutions.

APPENDIX

Methodology

IDC utilized its standard ROI methodology for this project. This methodology is based on gathering data from current users of the technology as the foundation for the model. Based on these interviews, IDC performs a three-step process to calculate the ROI and payback period:

- Measure the savings from reduced IT costs (staff, hardware, software, maintenance, and IT support), increased user productivity, and improved revenue over the term of the deployment.
- Ascertain the investment made in deploying the solution and the associated training and support costs.
- Project the costs and savings over a three-year period and calculate the ROI and payback for the deployed solution.

IDC uses the net present value of the savings and increased revenue over three years in calculating the ROI and payback period for the deployment. The NPV of the savings is determined by subtracting the amount that would have been earned by investing the original sum in an instrument yielding a 12% return (to allow for the missed opportunity cost that could have been realized using that capital).

IDC bases the payback period and ROI calculations on a number of assumptions:

- 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 number of hours of downtime multiplied by the number of users affected.
- The impact of unplanned downtime is quantified in terms of impaired end-user productivity and lost revenue.
- Lost productivity is a product of downtime multiplied by burdened salary.
- Lost revenue is a product of downtime multiplied by the average revenue generated per hour.
- The net present value of the three-year savings is calculated by subtracting the amount that would have been realized by investing the original sum in an instrument yielding a 12% return to allow for the missed opportunity cost. This accounts for both the assumed cost of money and the assumed rate of return.

Because every hour of downtime does not equate to a lost hour of productivity or revenue generation, IDC attributes only a fraction of the result to savings. As part of our assessment, we asked each company what fraction of downtime hours to use in calculating productivity savings and the reduction in lost revenue. IDC then taxes the revenue at that rate.

Further, because 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 and then subtracts the deployment time from the first-year savings.

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 2012 IDC. Reproduction without written permission is completely forbidden.