



Nirvanix White Paper



## Total Cost of Ownership Analysis of Alternatives for Managing Media-Centric Content

*The explosion of user-generated, media-centric content and applications, fueled by the Web 2.0 phenomenon, is opening a vast field of opportunity. With its enormous file sizes and unstructured characteristics, media-centric content is placing unprecedented demands on data storage infrastructures.*

*This paper examines the different storage alternatives that a company has to meet the data management challenge while capitalizing on the revenue-generating possibilities. Each storage alternative has different cost-of-ownership factors to consider, including acquisition and operating costs, as well as time-to-market factors and risks. Because of the critical long-term business implications, selecting the correct alternative demands a clear understanding of the advantages and disadvantages of each one.*



# Table of Contents

- The 3 Options .....3
- Option 1: Extending the Traditional Storage Infrastructure.....3
- Option 2: Building with Clustered/Virtualized File Systems .....4
- Option 3: Partnering with Storage Delivery Network Providers .....5
- Evaluating the Total Cost of Ownership.....6
  - One-Time Equipment & Infrastructure Costs.....7
  - One-Time Labor Costs.....8
  - Non-Labor Operating Cost .....8
  - Ongoing Labor Cost Beyond First Year.....9
- Total Costs of Different Alternatives.....9
  - Option 1: Extend Traditional Storage .....9
  - Option 2: Build With Clustered/Virtualized File System Technology ..... 10
  - Option 3: Storage Delivery Service Provider ..... 10
- Business Case Scenarios ..... 11
  - Scenario 1: High Capacity Upfront with Low/No Growth ..... 12
  - Scenario 2: Low Capacity Upfront with High Growth ..... 12
  - Scenario 3: Minimize Upfront Cost (i.e. Year 1)..... 13
  - Scenario 4: Minimize Total Cost of Ownership..... 13
  - Scenario 5: Quickest Time to Market} ..... 14
  - Scenario 6: Most Flexibility in Adding Capa ..... 15
- The Strategy for Pursuing the Opportunity ..... 15

## The 3 Options

The storage alternatives to consider are:

- Extend the current, traditional storage infrastructure.
- Build a next-generation infrastructure
- Partner with a service provider of Internet storage, also known as Storage Delivery Network provider (SDN).

As elaborated in this paper, the first option is best suited to an organization that mainly seeks to respond to internal storage pressures – as opposed to generating revenue in Internet storage. The second option can enable a company to pursue market opportunities, but with time-to-market constraints and a high implementation risk.

	Extend Traditional Storage Infrastructure	Build with New Clustered /Virtualized File System Technology	Partner with Storage Delivery Network Provider
Acquisition Cost	Low to Medium	Medium to High	None
Operating Cost	Medium to High	Medium to High	Low to Medium
Time to Market	Slow	Slow	Fast
Implementation Risk	Medium	High	Low
Technology Risk	High	Low	Low

Table 1

The approach with the most strategic long-term value, the one best suited to meeting both internal storage demands and responding to revenue possibilities, is partnering with an Storage Delivery Network provider. Within this approach, however, there are choices. This paper reviews the recommended criteria for selecting a Internet storage provider.

### Option #1:

#### Extending the Traditional Storage Infrastructure

For many established companies seeking to extend their current products and services to include online storage offerings, the obvious path to meeting the resulting challenge is to extend their existing storage infrastructure. This option allows the addition of capacity with minimal disruption of the current architecture. An established company can leverage their existing IT investment as well as the skills and knowledge of their staff.

While appealing because of its apparent simplicity, extending the traditional infrastructure is not optimal for companies seeking to create entirely new revenue streams based on Web 2.0 services. Traditional storage systems were designed for traditional data, which is mostly structured, text-based, and stored in databases. Media-

centric content, in contrast, is mostly unstructured data stored in content management systems. It consists of images, audio, video, PDFs, and graphics documents.

When it comes to handling media-centric content, even the newest Network Attached Storage (NAS) and Storage Area Network (SAN) systems have inherent architectural limitations. Also, the capacity needed for media-centric files is costly for traditional systems. Beyond certain capacity levels, traditional systems are unable to handle the complexity. The systems also experience serious performance bottlenecks due to built-in performance limitations.

A “one-size-fits-all” approach enables a company to react in the short term to the impact of media-centric content on the storage infrastructure. But the approach is not optimal for organizations ready to make a strategic commitment to capitalizing on Web 2.0 services.

## Option #2

### Building With Clustered/Virtualized File Systems

There are ways that companies can overcome the limitations of traditional NAS/ SAN architectures. One option is to build an in-house storage infrastructure incorporating next-generation technologies optimized for media-centric content. A number of solutions on the market today use clustered/virtualized file systems technology, offering the ability to handle the specific characteristics of mediacentric content. These newer solutions offer such capabilities as:

- Scalable file system (single name space for 100+TB)
- High performance (high throughput that scales linearly, non-disruptively)
- Easy to use and manage (cluster management)

Proprietary solutions with these capabilities include Isilon, Ibrix, and NetApp/ Spinnaker. Additionally, there are Open Source systems and tools from vendors such as MogileFS, Lustre, GlusterFS, MySQL, Gfarm, dCache, and FUSE.

Open Source solutions generally cost less than proprietary solutions, in terms of direct capital expenses, because there are no software license fees, and because inexpensive, commodity storage hardware is an option. But there are downsides, such as the lack of proven Open Source deployments, lack of guaranteed performance levels, and lack of recognized benchmarks. Also, unlike with proprietary solutions, Open Source tools are unsupported. As a result, the organization needs a full staff of internal support expertise for development and ongoing operations.

Whether proprietary or Open Source, solutions using clustered/virtualized file systems technology are far more effective than traditional NAS/SAN architectures with regard to media-centric content. The newer systems can handle streamed media as well as input/output-intensive applications characterized by frequent peak loads of simultaneous access. Energy/geophysical applications and some financial applications require this capability.

The build-it-yourself option can be an attractive choice for enterprises with an inhouse oriented culture and a large IT staff. This option may also be appealing to up-and-coming, high-growth start-ups looking to create their entire infrastructure in-house. The build-it-yourself option, however, involves considerable capital outlay (i.e. equipment) and medium-to-high ongoing operating expenses (i.e. support and maintenance). Also, there is a potentially longer time-to-market for the storage infrastructure. If Internet storage is meant to be an strategic revenue stream, the build-it-yourself option could seriously hamper a company's efforts to compete and build up that revenue stream.

There is another issue, in addition to operating expenses, capital expenses, and time-to-market. For most companies, building Internet storage is not a core competency. Product managers and software developers bringing new content services to the market face considerable competitive pressures. Most of them do not have core competencies in storage infrastructure, and they prefer to focus their efforts on the customer application or service. As a result, time spent on hosting, email management, cooling, power, and floor space will dilute a company's focus and divert attention from core business activities.

### Option #3

#### Partnering with Storage Delivery Network Providers

Whether they build it themselves or not, companies need a reliable, scalable and robust infrastructure in order to deliver Web 2.0-related services. Storage Delivery Network providers offer a compelling alternative to the build-it-yourself option. They provide a pay-as-you-use service model that enables companies to focus on their core business and accelerates time-to-market for revenue-generating products. Easy to use and quick to implement, this service model generally offers a broad range of features, from simple to fully supported, turnkey solutions.

Companies seeking to immediately capitalize on the growth of media-centric content should seriously consider this service model. With little or no initial capital required, this model suits up-and-coming and high-growth startups. At the same time, larger enterprises benefit from the ability to focus on core business activities such as internet-scale media storage, global network delivery, and multi-million user account management and security. While usually not requiring a capital outlay, partnering with storage delivery service provider requires a commitment to a service model. With a small number of vendors to choose from, companies need to make an informed decision. They need to identify a provider that can create a solution for them that is easy to integrate, easy to scale, and easy to control.

As they evaluate different storage delivery providers, companies need to consider the ease with which they can integrate their application, scale their media content, and control their storage infrastructure. In this paper we want to make an additional distinction inside of this category. We differentiate Basic Internet storage providers from Enhanced Storage Delivery Network providers because their offers vary in feature

set. An enhanced storage delivery network provider exposes a file system-like API via web services, so that developers looking to access storage find a very intuitive interface that looks and feels like a standard file system. Also, built-in management of end-user accounts and built-in media processing abilities allow applications to be built much faster.

**Basic Internet Storage:** just online storage in buckets, accessed over the Internet

**Enhanced Storage Delivery Network:**

### Instant Integration

- Rich Web Services API (REST & SOAP) & reference apps
- Advanced media processing functionality
- Full-featured NMFS with multi-tenant support

### Unlimited Scalability

- On-demand Internet storage with global scalability
- Clustered file nodes support millions of users & billions of files
- Dynamic load-balancing and file routing across nodes

### Effortless Control

- Real-time management portal and reporting
- Control costs with usage limit policies
- Guaranteed SLA and 24x7 support

For a more detailed comparison, check out the overview here:  
<http://www.nirvanix.com/comparison.aspx>.

## Evaluating the Total Cost of Ownership

The strategic alternatives described above have different TCO considerations which may include any or all of the following four basic components, as summarized in Table 2.

Cost Structure Component	Examples
One-Time Equipment & Infrastructure Costs	Storage arrays, software license, data center startup costs
One-Time Labor Costs	Storage development and application integration costs
Non-Labor Operating Costs	Hosting fees and bandwidth costs; online storage service costs, if applicable
Ongoing Labor Costs beyond first year	Online storage infrastructure and application maintenance costs

Table 2

## One-Time Equipment & Infrastructure Costs

The one-time equipment and infrastructure costs only apply to the options (a) to “extend traditional storage” or (b) “build with clustered / virtualized file systems”.

Detailed assumptions are spelled out in Table 3.

Key Components	Cost Range	Comments
Data center start up cost	\$90,000 to \$150,000 for 1.5PB scenario	“Build with proprietary clustered storage” option will have lower costs because of built-in capabilities. These scale linearly with storage size.
Storage system	\$1,000 to ~\$3,500 per TB	“Build with open source” option will have lower storage hardware cost than either “Extending traditional high-end NAS” or “Build proprietary clustered storage”.
Separate software / switches	~\$16,000 per node and \$7,600 per 24-port switch	Applies to “Build proprietary clustered storage” only.
Other software costs	\$15,000 for Open Source software support;	
\$200,000 / PB for Virtualization Software	“Build with open source” will have some software support costs while “Extending traditional storage” will require additional virtualization software, which scales linearly with storage size.	

Table 3

## One-Time Labor Costs

The one-time labor costs are applicable across all alternatives and vary greatly depending on the type of approach taken. Detailed assumptions are spelled out in Table 4.

Key Components	Cost Range	Comments
Storage Design	3 to 6 storage IT staff @ \$120,000 each	Only applies to the options to extend traditional storage or build with clustered / virtualized file systems.
Application Integration	0.4 to 15 application developers @ \$150,000 each	Enhanced Storage Delivery Network has the least cost while build with open source has the highest cost.

Table 4

## Non-Labor Operating Cost

For the extend / build alternatives, the non-labor operating costs consist of ongoing data center costs while for the storage delivery service providers, they consist of on-demand storage and bandwidth costs. Detailed assumptions are spelled out in Table 5.

Key Components	Cost Range	Comments
Data center rack space	\$360 per TB	Only applies to the options to extend traditional storage or build with clustered / virtualized file systems.
Data center bandwidth	\$90 per TB / month	Only applies to the options to extend traditional storage or build with clustered / virtualized file systems. 135 TB / year is assumed.
Online storage service	\$150 per TB / month	Only applies to storage delivery services.
Online storage transaction cost	\$0.01 per thousand transaction; assume 100K average file size and 20% of files for upload/download	Only applies to basic storage delivery services.
Online storage bandwidth	\$150 to \$180 per TB / month	Only applies to storage delivery services. 135 TB / year is assumed.

Table 5

## Ongoing Labor Cost Beyond First Year

The ongoing labor costs beyond the first year are applicable across all alternatives and vary greatly depending on the type of approach taken. Detailed assumptions are spelled out in Table 6.

Key Components	Cost Range	Comments
Storage maintenance	1 to 3 storage IT staff @ \$120,000 each	Only applies to the options to extend traditional storage or build with clustered /virtualized file systems.
Application maintenance	0.5 to 5 application developers @ \$150,000 each	Enhanced Storage Delivery Network has the least cost while build with open source has the highest cost.

Table 6

## Total Costs of Different Alternatives

The alternative a company chooses has long-term, strategic consequences. Therefore, it is critical to examine the cost structure of each alternative. The following scenarios provide an idea of what to expect based on 1.5 PB of storage. This 1.5PB comes from a fictional case study, where a company has an audience of 20 million end-users (either current customers off-line or target customers), plans include online backup offering for MyDocument files (average of 2Gb per end-user), possibly picking up 4% of this target audience (800K end-users).

Taking previous assumptions, we have assembled theoretical storage systems to address the requirements of the case study. Once built, we added all the costs, again based on assumptions above. Below are the results, in costs, of the different alternatives.

### Option #1: Extend Traditional Storage

The first year cost is \$7.3M for extending the traditional storage infrastructure. The largest cost in the first year is one-time infrastructure and equipment, which accounts for 78% of the total cost. Traditional high-end NAS or low-end SAN has a cost of \$3 to \$4 per gigabyte.

To make an “apples-to-apples” comparison, the calculations here include \$300K for virtualization software. It is necessary for avoiding the management of disparate main spaces across multiple filers.

Ongoing cost beyond the first year is \$1.1 million. Total Net Present Value at 7% is \$11 million for a five-year period.

## Option #2: Build With Clustered / Virtualized File System Technology

In the do-it-yourself option using clustered/virtualized file system technology the first-year cost ranges from \$5.3 million for Open Source to \$11.4 million for proprietary clustered storage. The biggest cost in the first year is storage/ application development expenses for Open Source, which accounts for 56%, or the one-time infrastructure and equipment for a proprietary solution, which accounts for 86%.

Beyond the first year, the ongoing cost beyond is \$1.6M for Open Source vs \$890K for a proprietary solution.

Total Net Present Value at 7% ranges from \$11M for Open Source to \$14M for proprietary over the five-year period. Open Source is less costly but more laborintensive for both storage development and application integration.

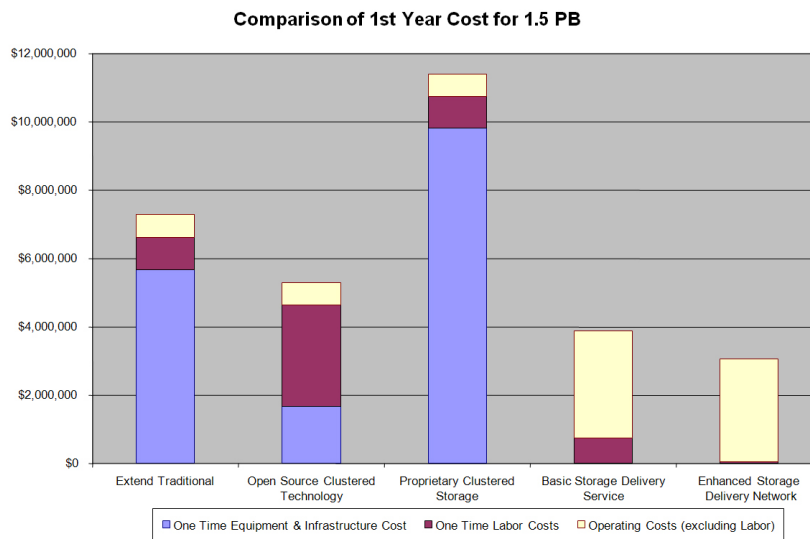
## Option #3: Storage Delivery Network Provider


Choosing to partner with a Storage Delivery Network provider entails a first-year cost ranging from \$3M for Enhanced to \$3.9M for Proprietary Clustered Storage. The difference is in the additional resources needed for application integration in the Basic Service. Service costs are similar and based on per GB per month plus the bandwidth.

Ongoing cost beyond 1st year is \$3M for Enhanced vs \$3.4M for Basic Service.

Total NPV at 7% ranges from \$13M for Enhanced to \$15M for Basic Service over the five-year period.

Figures 1 and 2 summarize the first year cost and ongoing costs, respectively, of the five different approaches showing an example of a 1.5 PB deployment.





## Business Case Scenarios

The sensible way to determine which alternative makes the most sense is examine various usage requirements that emphasize only one aspect of the end result. The best approach, then, would be the one that supports the strategic business objectives.

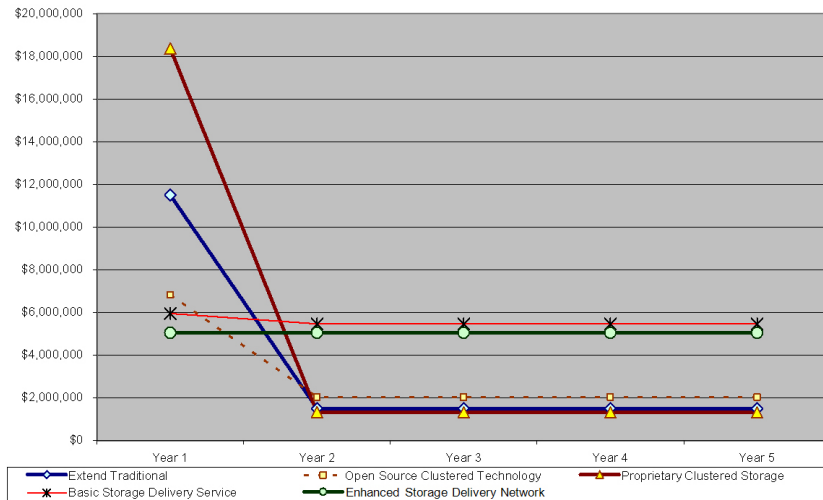
Possible cost-of-ownership usage requirements or scenarios are:

1. High upfront capacity with low / no growth (for example, when migrating an existing system to a new platform)
2. Low capacity upfront with high growth (like a new initiative or a start-up offering)
3. Minimize upfront cost (for example a yet-to-fund strategic initiative, or a frugal start-up)
4. Minimize total cost of ownership (most cases this is desired)
5. Quickest time to market (start-up opportunities, or competitive initiatives where being first to market matters)
6. Most flexibility in adding capacity (most cases this is desired, but specially when adoption rates are unknown).

## Scenario #1: High Capacity Upfront with Low / No Growth

Storage Delivery Networks provide the ability to add capacity quickly and without incurring high upfront storage costs in year 1. The higher the capacity needed upfront, the more attractive it is to go with a Storage Delivery Network provider.


First 5-Yr Cost for 2.5 PB Fixed Capacity of Different Alternatives



## Scenario #2: Low Capacity Upfront with High Growth


Storage Delivery Networks are the best alternatives for low capacity upfront with high growth, given the high upfront cost of building the in-house storage infrastructure.

First 5-Yr Cost for 500 TB Growth Per Year of Different Alternatives



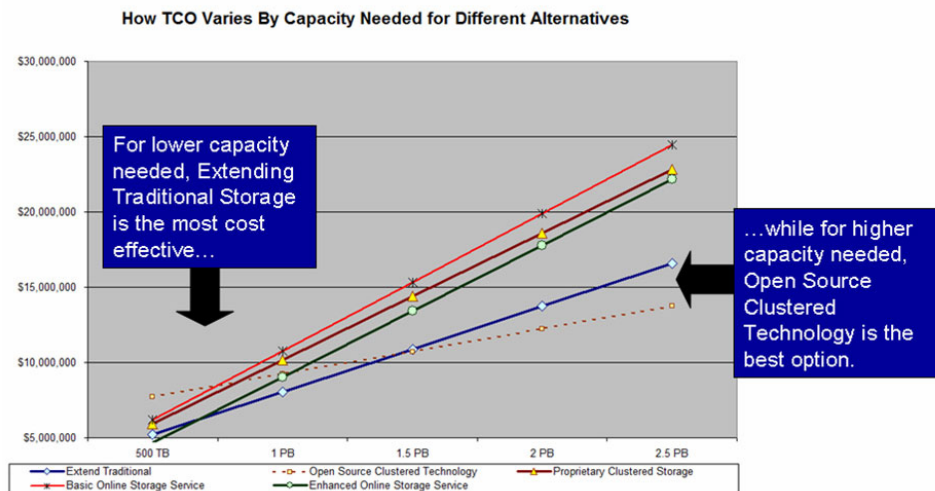
### Scenario #3: Minimize upfront cost (i.e. year 1)

Storage Delivery Networks are the best alternatives to minimize upfront costs in year one, given that the upfront cost of building the in-house storage infrastructure grows dramatically as capacity increases.



### Scenario #4: Minimize total cost of ownership

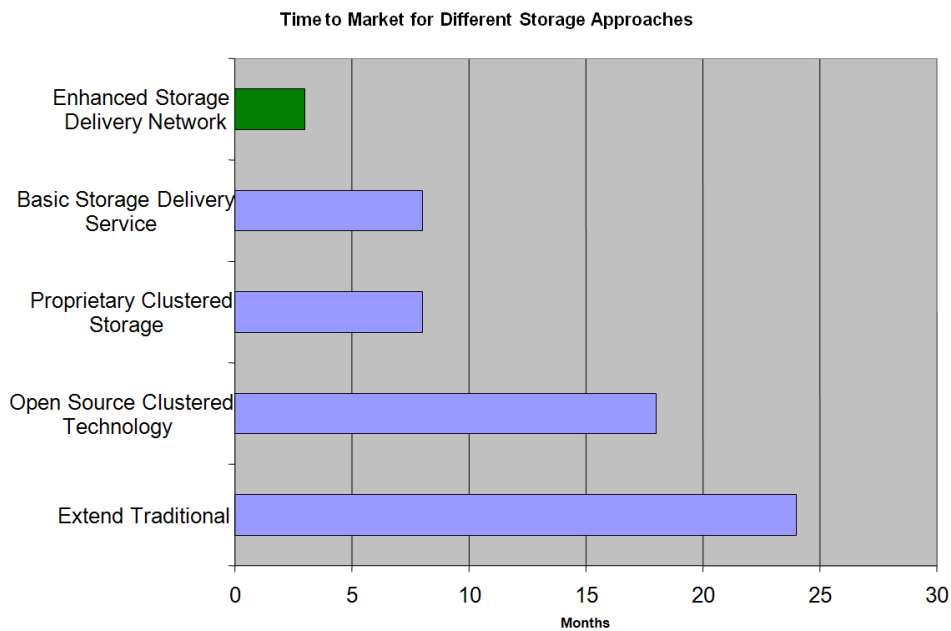
When lower capacity is sufficient, Extending Traditional Storage is the most cost effective choice. For higher capacity, Open Source Clustered Technology is the best option.



Note that these results may be misleading, as they are based on pure theory and math. In reality, choosing an Open Source solution dilutes the focus and skill set away from the company's core competencies and require special skills that may be hard to find, recruit, and retain. Also, rarely does one look at one requirement in the complete absence of others (like time-to-market) – A lower cost of ownership may be the worst option of all if time-to-market is important at all (coming to market 3 years later may cancel out all revenue, killing any possibility of having a return on this investment). Last, Open Source code may need significant overhaul work, or need a non-trivial re-architecture (you become the product designer and maintainer).

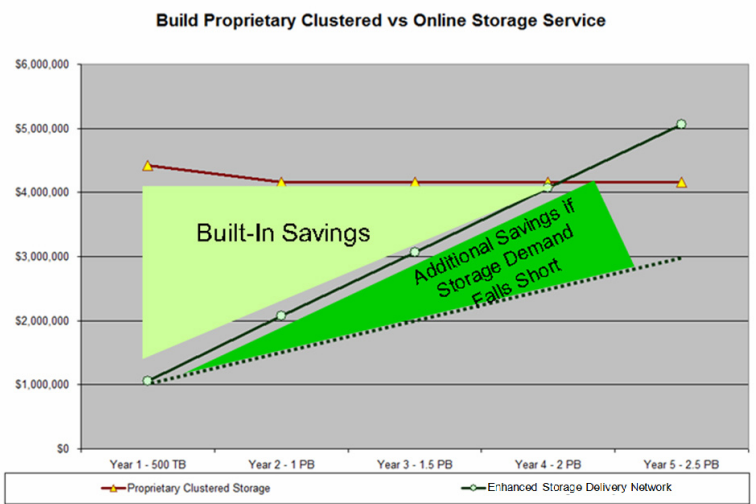
## Scenario #5: Quickest Time to Market

Enhanced Storage Delivery Network is the best for achieving fast time-to-market.



### Scenario #6: Most Flexibility in Adding Capacity

Enhanced Storage Delivery Network provides the maximum flexibility in adding capacity if demand falls short and/or ramps up more slowly.



### The Strategy for Pursuing the Opportunity

Clearly, the emergence of Web 2.0 services has created compelling market opportunities as well as significant storage challenges. In response, there are three main strategies that a company can pursue. Each strategy has different cost-of-ownership, time-to-market, and risk-level considerations. A thorough examination of the data shows that partnering with a Storage Delivery Network provider offers the greatest ability to seize the opportunities and meet the challenges.

For more details on this analysis of Total Cost of Ownership, to prepare a custom version for your business, or to learn more about Nirvanix, please visit us at [www.nirvanix.com](http://www.nirvanix.com).



525 B St. Suite 1000 | San Diego, CA 92101

Tel 1.619.764.5650 | Fax 1.619.764.5654 | [info@nirvanix.com](mailto:info@nirvanix.com) | [www.nirvanix.com](http://www.nirvanix.com)

To learn more about how Nirvanix can help your business grow contact us at 1.619.764.5650 or visit our web site at [www.nirvanix.com](http://www.nirvanix.com).

© Copyright 2008 Nirvanix, Inc. All rights reserved. #NWP--TCOMediaAlternatives-0708