August 25, 2011

File Storage Costs Less
In The Cloud Than In-House

by Andrew Reichman
for Infrastructure & Operations Professionals

August 25, 2011

File Storage Costs Less In The Cloud Than In-House
We Calculate A 74% Cost Reduction, But It's Hard To Compare Apples To Apples

by Andrew Reichman
with Robert Whiteley and Eric Chi

EXECUTIVE SUMMARY

There's a great deal of hype surrounding cloud delivery models of all flavors, with the assumption that it will usher in a new era of cheaper, better IT. But most I&O organizations struggle to compare the capital, operational, and staffing cost differences between internal and cloud storage. This document will break down the cost differences by modeling a common workload — file storage — built and deployed traditionally versus consumed through the public cloud. Our models reveal a significant cost difference, with the cloud-based model coming in 74% less expensive than I&O running it in-house. However, we encourage I&O leaders to use this document and the associated "Cost Analysis Tool: Cloud Versus Internal File Storage" Excel workbook as a tool and methodology for comparison. Customize the tool with numbers from your own environment and make a better decision as to whether cloud file storage makes sense in your business technology environment.

TABLE OF CONTENTS

NOTES & RESOURCES

Forrester interviewed multiple vendor and user companies, including Amazon, Ctera Networks, F5 Networks, Nasuni, Nirvanix, Riverbed Technology, Scality, StorSimple, TwinStrata, and others.

Related Research Documents

"The Three Stages Of Cloud Economics"
April 28, 2011

"Cloud Storage Comes Down To Earth"
August 5, 2010

"Controlling Storage Cost Amid High Growth"
February 3, 2010

CLOUD STORAGE IS WORKLOAD-SPECIFIC

You can fantasize about tearing down your data center and riding into the sunset on a cloud, but in reality that cloud is just another way to deliver IT services and it has limitations, both operational and financial. For cloud to be viable for a given workload, it has to deliver a performance and security experience that is as good as or better than what can be deployed locally. Latency is a big problem when it comes to storage, because the public cloud is distant from servers in your physical data center. Workloads that are very sensitive to latency aren't good candidates for cloud deployment.[1]

Many I&O professionals have recently been tasked with exploring the financial impacts of moving to cloud delivery models. In looking at the impacts within storage, I&O managers should start by looking at three storage workloads and how they relate to the cloud:

- **Whole applications that are storage-centric.** Getting rid of the owned infrastructure for a business application and moving it to the cloud could streamline operations, but it really isn't cloud storage; rather, it's software-as-a-service (SaaS). Because both server and storage are colocated in the cloud, the latency introduced by cloud delivery is limited to client-to-host traffic, not host-to-storage, so this is a likely delivery scenario for effective cloud deployment by application vendors. Going down this path would change the amount of storage that a given enterprise needs internally, so it's worth thinking about as you plan your storage strategy. It's not a good candidate for comparing internal storage costs with cloud storage costs though, since it has many more variables than just cost of storage.

- **Backup and archive data that are less sensitive to latency.** Backup is a secondary copy of data, and archive is data that isn't being frequently accessed, so these are workloads that are generally insensitive to latency. Sending incremental pieces of data sets over the WAN to the cloud can happen throughout the day without negative impact to the applications or processes, unlike primary database data that has slim latency tolerances. Bringing the entire data set back for recovery or archive discovery can be a challenge across the WAN, and cloud providers of backup and archiving are working on ways to enable this. Some even use trucks and ship physical systems to accommodate the large data transfer.

- **Files that are discrete packets of data and generally less performance critical.** Files are also called "unstructured data," referring to the fact that there is no need for an application to make sense of the information. Instead, files have their own metadata and can be independently read anywhere, regardless of who or what created it. With general departmental file sharing, performance is not a critical factor unless it's a specialized use case such as CAD/CAM, financial simulation, virtual server configurations, or database data. But for most files, their relatively low performance requirements make them viable for cloud storage, and Forrester sees this as one of the most likely and productive use cases to move data to the cloud.

Given the characteristics of the above workloads, file storage is the best candidate for cloud storage, and this document will explore the economics of moving storage in particular to the cloud. Our cost comparison will focus on cloud file storage for primary data compared with traditional, internal storage of that data.

QUICK STORAGE COMPARISONS DON'T MEASURE STORAGE COSTS EFFECTIVELY

Many people look at the cost of buying storage and compare that with the cost of procuring public cloud storage from a third-party vendor. That comparison usually looks something like this: For 100 TB of data stored, you would estimate your $/GB price of procurement and annualize that over the expected life of the storage you buy and run in your own data center. The cloud side of the equation would involve paying for that same quantity of storage for a year at prevailing market rates (see Figure 1).

This comparison is not accurate, because it leaves out significant costs on both sides. A quick comparison can be a good way to get a feel for the dynamics of a decision, but this is too quick and too simplistic. Both sides have significant gaps in the associated costs when measured in such a cursory fashion. A more detailed I&O business case needs to consider that:

- **Operating costs are high for internal storage, and you need more capacity than data.** When you run traditional infrastructure, you have to pay for the operations and the redundancy, both of which are included in cloud infrastructure services. Operational expense for storage is roughly estimated at 100% of the annualized cost of buying that storage, so it is significant. What's more, when you have a given quantity of data, you must buy much more raw storage capacity to deliver enterprise-class reliability and performance and to allow for growth buffers.

- **Cloud storage includes almost everything, but not WAN and transaction charges.** A huge benefit of cloud storage is that operating costs and capacity redundancy are included in the service, which makes measuring the cost easy and the results predictable. Or does it? There are a few gotchas when it comes to measuring the real cost of cloud storage, including: 1) the impact to network links from sending more data across the WAN; 2) transactional charges levied by providers that charge you for writing data into the service and reading or modifying it; and 3) variability in redundancy schemes or performance capabilities that can require you to evaluate and tier the service you get. This can make measuring the real cost of cloud storage more complicated than it may initially seem, but the variability is less significant than with internal storage. What's more, some cloud capacity vendors like Nirvanix and gateway vendors like Nasuni are moving toward simpler and more predictable pricing schemes.

- **Cloud options don't typically have native support for data access as you need it.** Most cloud storage repositories are built with object repositories as the back end. Object storage allows for massive scalability, custom metadata, geographic dispersion, and low cost, which is attractive for

cloud vendors.[2] But you would have to modify your data or business processes or code to vendor Application Programming Interfaces (APIs) to use object stores directly. These aren't trivial tasks, and the costs associated with the effort is significant and hard to predict. Part of the cloud gateway value proposition is to do this data transformation, taking your files and translating them to object stores. If you want to keep accessing files as you do today, only in the cloud, then the cost of a gateway or your own custom integration may need to be added to the base cloud storage price tag.

Figure 1 It's Easy To Get An Incomplete And Incorrect Comparison Of Cloud And Internal Storage

Cloud storage		Internal storage	
100 TB of actual data			
Assumptions		Assumptions	
$/GB/month cloud charge	$0.118	4	Years expected lifespan of storage
Months/year	12	$4	$/usable GB purchase price
GB/TB	1,000		
Calculations		Calculations	
Monthly cost of cloud storage	$11,800	$400,000	Acquisition cost of base amount of storage
Annual cost of cloud storage	$141,600	$100,000	Annualized over life of storage
Simple annual cost of cloud storage	**$141,600**	**$100,000**	**Simple annual cost of internal storage**

57696   Source: Forrester Research, Inc.

AN ACCURATE COMPARISON REQUIRES A MORE DETAILED MODEL

To accurately measure the difference between getting storage capacity from the cloud and building it yourself requires a model that takes more factors into consideration. Keep in mind, though, that measuring cost differences between the two scenarios is intended to simplify the analysis and take a pragmatic approach, so it's not necessary to look exhaustively at every possible cost factor. This type of modeling is known as relative cost of operations. Total cost of operations is a good idea but often hard to do in practice, so isolating only the changing factors and focusing on them while ignoring the rest can suffice.[3]

The Cost Of Storing 100 TB Of Data Internally Is Higher Than You Think

You might want to sit down for this: The cost of storing 100 TBs of file data is close to $1 million per year, with a significant front-loading of the capital costs. Our embedded model for the cost of

internal storage uses assumptions based on Forrester's experience talking with many companies that deploy internal storage (see Figure 2). You can download our Excel workbook, the "Cost Analysis Tool: Cloud Versus Internal Storage Deployment," and plug in your own numbers. If you have higher use, lower cost of storage, or more efficient staff, then you might find that your cost of storage is lower than $1 million, but it's still probably higher than you realize — or want. I&O teams should familiarize themselves with the internal storage calculations for this comparison. Here are some of the key factors that you should understand and measure for storing 100 TB of primary file data:

- **Cost varies widely based on the number of years you use internal storage.** Firms have different philosophies on how long to keep their storage arrays in service. Some I&O shops lease storage and get new gear every three years like clockwork. Others run their arrays until the vendor discontinues support, sometimes for seven years or more. Since our model annualizes storage costs, there is a great deal of variation based on the lifespan of the storage. The acquisition and migration costs go up with short life cycles, while the cost of maintenance goes up with longer life cycles. There are hidden costs of longer life cycles as well, which the model doesn't show, including lower efficiency and inability to use newer features. Forrester used four years as a common best practice for refresh cycle duration.

- **Storage acquisition cost drives a significant portion of the internal storage model.** The acquisition cost of the storage you buy is inarguably a huge aspect of the total cost. It's not the only factor in how much it costs you to deploy storage internally, but I&O teams probably pay too much attention to it compared with operating cost. It also plays a huge role in the potential economy of scale offered by cloud storage vendors, many of whom claim to use fundamentally cheaper storage systems in their environment — often some version of scale-out software running on commodity server hardware — than individual I&O shops can buy themselves. Forrester sees an average street price for midrange storage with NAS capability between $3 and $5 per usable GB, including all software and hardware. For the model, we've used the middle value of $4 per usable GB. Storage cost is typically benchmarked with usable GB, so this means that the price here includes the RAID protection overhead and system resource overhead that increase capacity amounts from RAW.

- **Number of redundancy copies increases the storage capacity requirement.** This is one of the biggest differences between cloud and traditional storage. When you send 100 TB of data to the cloud, you pay for 100 TB. However, when you store 100 TB of data internally, you have to buy much more than that amount. For protection, firms typically hold three copies of their data: primary, local replica, and remote or backup copy. This number can be higher or lower based on availability and regulatory and other requirements, but three is a reasonably conservative copy count. You could argue that one of the copies might be cheaper as it's a backup copy on tape or deduplicating disk, or that you use space-efficient snapshots that use less capacity, but you also have more versions of the backups for multiple recovery points. For the sake of a clean model, we've stuck with three copies at the single price point of midrange capacity. We encourage you

to increase the number of copies in the associated workbook if you have more data copies or reduce the number if you don't keep three versions or if you feel confident that your copies are cheaper than primary disk.

- **Low utilization means that you have to buy more storage capacity.** Another big difference between cloud and traditional infrastructure is that any inefficiency is on the provider's dime. The burden falls to the provider if it needs to consume twice as much raw capacity to provide high availability or if it needs to buy capacity upfront. Most traditional storage environments struggle with capacity management and have low utilization of raw storage as a result. Typical storage environments see raw capacity utilization (data written over raw capacity on hand) of 20% to 40% due to RAID protection, system resource capacity, growth buffers, and low use of big allocations.[4] Since we're dealing with usable capacity in this model, system resource overhead and RAID protection are already factored in, so we've set the usable capacity utilization for the model at 60%. Yours might be higher or lower, but this is where Forrester sees most I&O environments today.

- **Staff is a big part of internal storage costs.** When you build a storage environment, you have to pay to run it, and hiring capable staff is difficult and costly. The staffing component is included in the cost of cloud storage, which drives a big cost difference. There are two aspects of calculating the cost of staffing: how many employees you need and how much they cost. Use TB per full-time employee (TB/FTE) as the typical metric for storage staffing efficiency. Forrester sees a rough average rate of 150 usable TB per FTE in the industry, although this number can be far higher or lower depending on the environment or technology used. For the salary, Forrester uses $120,000 as a fully loaded admin staff cost. Substitute your firm's standard rate for modeling if you have one. Keep in mind that staffing is not particularly flexible, so when you're considering the impact of cloud, determine whether the move would actually change your staffing levels. If you're only moving a small portion of data, you might not see any impact in your remaining internal environment. If you feel that you wouldn't reduce headcount with a move to the cloud, then zero out the $/FTE column and staff costs won't be considered in the comparison.

- **Facilities and power is a real cost, although often obscured.** If you're like most I&O organizations, you don't have a direct line of sight to how much the facilities or power associated with your infrastructure costs. However, it's a big part of running internal infrastructure, and the more you move to the cloud, the less you have to pay to keep it running. It costs about 1% of the acquisition cost of storage to power and cool it, which isn't a huge number.[5] The data center space itself is harder to quantify. Colocating or building a data center is tremendously expensive, but each TB only consumes a small portion of that and it's not easy to correlate the costs. To be conservative, Forrester assumes that 5% of the original storage acquisition cost captures the total annual cost of power and facilities.

- **Buying storage means paying for maintenance.** Most new storage systems come with three years of support and maintenance included. After that period you usually pay around 15% of the original purchase price per year. When you get storage from the cloud, that cost is borne by the provider, so it's another strike against internal storage. You can modify the years of included maintenance and percentage of original purchase price based on the agreements you have with your vendors. Keep in mind that support and maintenance varies in terms of quality, so if you upgrade to premium support levels, your spend levels here may be higher.

- **Data migration is another hidden cost of storage.** Not only do you have to buy, run, maintain, and power internal storage, but you also have to migrate it from old equipment to new every few years when it's time for refresh. This process is onerous, risk-laden, and the bane of existence for many storage directors. Most firms choose to pay outside consultants to accomplish this, or even modify their permanent staffing levels based on the need to do migration. When you get storage from the cloud, you get perpetual storage that doesn't have the effort or disruption of migration associated with it. Forrester estimates $500 per usable TB for migration planning and execution services for midrange storage systems, but this number can vary widely.

Given all these assumptions and calculations, the all-in annual cost of storing 100 TB of file data internally comes out to $955,500 — close to the $1 million we mentioned above. We encourage you to use the downloadable tool to customize this number for your environment. Also keep in mind that the comparison to cloud storage is a relative one. If there's something that you think would not change or go away if you shifted storage to the cloud — such as staff or facilities costs — then zero it out in the model and it won't be considered in the comparison.

Figure 2 Traditional File Storage Systems Are Expensive To Buy And Run

Internal storage

	Assumptions	Calculations	
TB of actual data	**100**	$400,000	Acquisition cost of base amount of storage
Years expected lifespan of storage	4	300	Usable TB required including data copies
$/usable GB purchase price	$4	420	Total usable TB required for primary, copies, and utilization
Copies of data for redundancy	3	$1,680,000	Acquisition cost of total storage requirement
Typical utilization of storage (excluding RAID and system resource overhead)	60%	$420,000	Annualized storage acquisition cost based on lifespan
Typical TB/FTE	150	2.8	FTE requirement for storage admin of stated TB count
Fully loaded $/FTE	$120,000	$336,000	Annual storage admin cost
Facilities and power charge (of storage acquisition cost)	5%	$84,000	Annual facilities and power charge
Years of included warranty	3	$63,000	Annualized maintenance charge
Percent of original purchase price for additional warranty years	15%	$52,500	Annualized data migration charge
Cost per usable TB data migration	$500	**$955,500**	**Total annual cost of internal storage**

57696   Source: Forrester Research, Inc.

Cloud Pricing Is Generally Simpler, But It's Important To Evaluate What You Get

The pricing for cloud storage is boiled down to a price per GB stored per month in most cases, although there are a few factors that can complicate the pricing. The elements that go into calculating the annual cost of cloud storage for 100 TB of data include the following (see Figure 3):

• **The base rate for cloud storage is the main driver of annual cost.** Each vendor offers different rates, and the rates vary by volume and negotiation, but the standard in the cloud storage industry has emerged as dollars per GB ($/GB) per month. Amazon clearly publishes its rates for its cloud storage, S3, and it changes based on volume, with limited negotiation variability. Other vendors are still ironing out their pricing, but are generally similar with the exception of negotiating better rates as an early adopter. For the purpose of this model, we've started with $0.118 as cloud storage price per GB, based on the rate Amazon would charge for 100 TB stored, but you can enter another number based on the vendor you're considering.

• **You pay a fee to send data to the cloud in most cases.** In addition to paying for the capacity you consume, most cloud providers charge you for moving data into and out of the cloud.

This model assumes that you have all the data on day one, which would be the case if you're migrating existing data from internal storage to the cloud. The 100% number represents moving all 100 TB to the cloud and paying the charge immediately. Once the data is there, you wouldn't pay for importing it again although the model just assumes this as a component of the annual charge; fortunately it's not huge, so it doesn't have a big impact on the end comparison. In the case of ongoing growth, cloud storage provides an advantage over internal storage in that you only pay for it when you generate the data, rather than in advance.

• **You also have to pay to read the data from the cloud.** The fact that you pay to access the data you have sent to the cloud means that active data is more costly than archival data when stored in a cloud repository. It's hard to predict what the actual frequency of access is for most data, but since the charge isn't tremendous, a rough guess will suffice. We recommend that you enter 100% in the data out/month field for extremely active data; 50% for moderately active data; and 10% for minimally active or archival data. We've used 50% in the base model, although it's probably an overstatement. Most employees won't read half of their primary file data every month. If you read every piece of the 100 TB of data every month of the year, then it would cost you $180,000. This drops to $18,000 if you only read 10% of the data each month, so the swing is significant in relation to the total cost.

• **The level of redundancy you get with cloud services varies.** Amazon stores multiple copies of every piece of data you send to its S3 service, to prevent data loss or unavailability. This is comparable to what most firms do with internal storage and should be seen as a requirement for primary data. Our model allows you to input the number of copies included with the cloud service. If the cloud service includes three or more copies, we assume that you wouldn't need to pay for additional ones. But if the included copy count is lower than three, we assume that you would have to pay for extra redundancy for the service to be effective for this use case.

• **Cloud gateways can add functionality and cost to the service.** Cloud gateways are a relatively new category of technology that act as brokers to major cloud service. Using one in your environment means that you don't have to do the API integration with the service provider yourself and you get added functionality such as local data caching, encryption, protocol translation, WAN optimization, snapshots, data migration, improved SLAs, and streamlined contracts and relationships. A gateway can be priced as a hardware/software solution with an upfront cost and an expected lifespan or as a service with a charge added to the cloud provider capacity charge. For example, Nasuni, a gateway provider focused on cloud file storage, has a pricing model that increases the capacity charge from its back-end partners to around $0.58, which includes the gateway and all inbound and outbound transaction fees. To model this, simply increase the cloud capacity charge to $0.58 and then zero out the inbound, outbound, and upfront cloud gateway rates (see Figure 4). Other cloud gateway solutions, such as Ctera, F5, Panzura, Riverbed, StorSimple, and Twin Strata, have different pricing models and functionality, so check with them to see whether they meet your needs and how they would price their offerings.

- **You may need to shore up your WAN links to accommodate cloud data traffic.** The network is a big factor when considering moving from internal storage to cloud-based storage. Moving significant quantities of data across the WAN as opposed to within the walls of your firm could put a strain on your links to the outside world. However, there are many reasons to increase WAN bandwidth in the era of SaaS, IT consolidation, cloud, and remote workers — so a move to cloud storage might not be the straw that breaks the camel's back. Many firms tell us they are increasing the size of these pipes regardless of moving to cloud, so many storage teams may just catch a free ride. Even if you did include a cost here, it should represent only the portion used for data traffic, not the cost of the whole upgrade. We've put $20,000 in as a reasonable annual charge to consider, but there's a great deal of variability here. You may want to zero it out if you think that any network improvement would happen regardless of a move to cloud storage.

The final calculation in the base case for cloud storage comes out to $251,600 to store your 100 TB of data with a public cloud provider, a whopping 74% reduction compared with the cost of running it yourself. Clearly, there's a great deal of promise in cloud storage as an alternative to building and running your own file storage environment. Many I&O execs are hesitant to move their data to the cloud at any price, but doing the proper due diligence on appropriate workloads will help.

Figure 3 Cloud Storage Is More Straightforward And A Lot Cheaper Than Traditional Storage

Cloud storage

	Assumptions	Calculations	
TB of actual data	**100**	$11,800	Monthly cost of cloud storage
$/GB/month cloud charge	$0.118	$141,600	Annual cost of cloud storage
Months/year	12	$10,000	Total data-transfer-in charges
GB/TB	1,000	50	TB out/month
$/GB data-transfer-in rate	$0.1	$90,000	Annual data-transfer-out charges
Initial data in, assumed in annual cost	100%	$0	Additional redundancy capacity charges
$/GB data-transfer-out rate (simplified)	$0.15	$0	Cloud gateway annualized charge
Data out/month	50%	$20,000	Incremental annual network charge
Included copies of data for redundancy	3	**$251,600**	**Total annual cost of cloud storage**
Cloud gateway hardware/ software charge	$0		
Years expected lifespan of gateway	4		
Incremental annual network charge	20,000		

Source: Forrester Research, Inc.

Figure 4 A Gateway Can Make Cloud Storage More Viable, But Also More Expensive

Cloud storage

	Assumptions		Calculations	
TB of actual data	**100**		$58,000	Monthly cost of cloud storage
$/GB/month cloud charge	$0.580		$696,000	Annual cost of cloud storage
Months/year	12		$0	Total data-transfer-in charges
GB/TB	1,000		50	TB out/month
$/GB data-transfer-in rate	$0		$0	Annual data-transfer-out charges
Initial data in, assumed in annual cost	100%		$0	Additional redundancy capacity charges
$/GB data-transfer-out rate (simplified)	$0		$0	Cloud gateway annualized charge
Data out/month	50%		$20,000	Incremental annual network charge
Included copies of data for redundancy	3		**$716,000**	**Total annual cost of cloud storage**
Cloud gateway hardware/ software charge	$0			
Years expected lifespan of gateway	4			
Incremental annual network charge	$20,000			

57696   Source: Forrester Research, Inc.

R E C O M M E N D A T I O N S

CLOUD STORAGE OFFERINGS ARE STILL DEVELOPING: USE CAUTION

While the interest and potential cost savings in cloud storage are tremendous, there is still a limited number of large enterprise I&O organizations that have moved their data to a cloud repository. There are a number of concerns as to whether cloud providers can be an effective replacement to internally managed storage. I&O leaders should consider that:

- **Service-level agreements (SLAs) are limited, and penalties don't cover the risk.** File data represents a key asset for most firms, and access and protection are key elements of good data stewardship. The major cloud storage providers have improved their SLAs and transparency over the past couple of years, but some are still not ready for enterprise prime time. Penalties are especially weak, with little clarity or recompense for lost data or unavailability. For example, you may get a portion of your monthly service fee back as a result of a full-day outage, but it doesn't come close to compensating your firm for the

productivity lost. These contractual issues need to be resolved for the majority of enterprises to consider this move, regardless of cost advantages.

- **Integration with cloud providers can be a major hassle.** Let's face it, as an enterprise I&O organization you don't want to be in the business of coding to APIs and building mechanisms for secure data transfers. Isn't that why you wanted to move to the cloud in the first place? Cloud providers need to go further to make the process easier, and develop a more user-friendly experience. Gateways can go a long way toward this, so consider those products as a way to navigate the integration more effectively.

- **Performance can be an issue when moving data across the WAN.** File data represents a likely use case for the cloud since it's less performance sensitive than block database workloads, but your users won't tolerate significant degradation of experience. Cloud storage might represent a cheaper deal or fewer headaches, but keep a close watch on user satisfaction and data transfer performance. Consider gateways for their local data caching and WAN optimization capabilities to improve performance or other products that might include those features separately.

- **Lower criticality data sets can be a good first to move.** Moving to the cloud represents risk as does any change in your computing strategy. Don't cut your teeth on the most critical files in your environment. Start with archives, or lower importance departmental file shares, and then move on to more performance- and security-sensitive files such as engineering files, CAD/CAM, and executive file shares.

ENDNOTES

[1] Cloud storage has become more viable as offerings become more enterprise focused, but they are very workload specific. Every vendor in the space is painting its products and message with a cloud veneer. But for all this hype, there is little clarity about what cloud storage really means and how it might be used. Currently, only three cloud storage use cases are ready for prime time: 1) whole in-cloud applications with their own storage; 2) backup to the cloud; and 3) file storage in the cloud. As you explore these use cases, you'll need to select vendors that provide the proper underlying facilities, infrastructure, and storage management applications. For more information, see the August 5, 2010, "Cloud Storage Comes Down To Earth" report.

[2] Object storage has gained attention as potentially easier to use, cheaper to deploy, and more highly scalable storage architecture, although the market is still struggling with exactly how and where to use it, and showing some anxiety when it comes to moving away from more widely adopted, standardized solutions like file and block based storage. For more information, see the November 28, 2010, "Prepare For Object Storage In The Enterprise" report.

[3] It's a nice idea to measure and compare all the costs of a given technology process or investment, but it's often not practical. Total cost of operations (TCO)-based financial analysis is held up as the gold standard for technology investment justification, but most firms don't have the rigor to apply the discipline to their

environment. To really implement TCO-based analysis it takes a comprehensive and continuously updated catalog of asset inventory, in-service dates, agreed-upon operating cost rates for activities, and a scheme to divide shared costs among the constituent business processes that use them. A more pragmatic approach is to compare only the costs that change between two scenarios, known as relative cost of operations. For more information, see the August 26, 2008, "TCO Is Overrated" report.

[4] Storage is expensive, and efficiency is critical to reduce costs. Building an effective storage environment is a balancing act between performance, reliability, and efficiency. For years, improving storage efficiency took a back seat to the other goals. Now, with infrastructure budgets being cut and data growth continuing unabated, it's time to take a closer look at storage efficiency to see if there is room for improvement. Most storage teams don't have rigorous processes in place to measure key performance indicators (KPIs) of storage efficiency, and this document is intended to identify the metrics that matter most. For more information on measuring and improving the efficiency of your storage environment, see the October 22, 2009, "How Efficient Is Your Storage Environment?" report.

[5] Storage vendors worldwide have jumped on the "green" bandwagon in their marketing campaigns, but it's hard to determine which technologies move beyond hype and rhetoric to have a real positive impact. In a gloomy economy, initiatives that sound good but have little measurable influence on the bottom line are unlikely to receive funding, so sorting through the claims and identifying benefits that are achievable in the near term is key to a successful green strategy in storage. Power isn't a huge component of overall storage cost, but reducing consumption often aligns with other efficiency objectives. For more information, see the April 17, 2009, "Align Green Storage With Overall Efficiency" report.

FORRESTER®

Making Leaders Successful Every Day

Headquarters

Forrester Research, Inc.
400 Technology Square
Cambridge, MA 02139 USA
Tel: +1 617.613.6000
Fax: +1 617.613.5000
Email: forrester@forrester.com
Nasdaq symbol: FORR
www.forrester.com

Research and Sales Offices

Forrester has research centers and sales offices in more than 27 cities internationally, including Amsterdam; Cambridge, Mass.; Dallas; Dubai; Foster City, Calif.; Frankfurt; London; Madrid; Sydney; Tel Aviv; and Toronto.

For a complete list of worldwide locations visit www.forrester.com/about.

For information on hard-copy or electronic reprints, please contact Client Support at +1 866.367.7378, +1 617.613.5730, or clientsupport@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Research, Inc. (Nasdaq: FORR) is an independent research company that provides pragmatic and forward-thinking advice to global leaders in business and technology. Forrester works with professionals in 19 key roles at major companies providing proprietary research, customer insight, consulting, events, and peer-to-peer executive programs. For more than 27 years, Forrester has been making IT, marketing, and technology industry leaders successful every day. For more information, visit www.forrester.com.

FORRESTER®