PICKING THE RIGHT OPTION FOR HOSTED STORAGE

RACKSPACE® HOSTING | 5000 WALZEM ROAD | SAN ANTONIO, TX 78212 U.S.A

rackspace® HOSTING

Picking the Right Option for Hosted Storage

The demand for storage is growing quickly... but choosing the right technology to support it can be challenging. This white paper helps by describing the four types of data assets and recommending the best hosted storage option(s) for each type.

Demand for Storage is Growing Fast

A recent survey of Oracle User Group members found data growing rapidly at 90% of their firms.[1] Smaller businesses are not immune. One consultant advises SMB managers to plan for 800% storage growth over two to three years.[2]

The demand for storage can grow so quickly it catches IT managers off-guard. For example, a Forrester survey found that 400 IT decision-makers around the world predicted a 19% increase in 2007, but experienced actual growth of 131%; their average data requirements jumped from 38 to 88 Terrabytes in one year.[3]

Picking the Right Technology is Vital

Clearly, every business needs storage... but what type? There are different technologies, and not every IT department knows how to evaluate every choice.

For example, there is a big difference between user-generated content and a mission-critical customer database. Storing both with the same technology can mean costly overkill for one, and underpowered performance for the other.

A mismatch could even slow down your operations, turn off prospective buyers, and stifle your company's growth.

The Four Types of Data

Consider how business information is classified. Data can be divided into structured and unstructured, or split into three types: files, e-mail, and databases. This white paper proposes the four categories of data assets summed up in Table 1.

1. The basics cover all the routine files that any business needs. These files tend to accumulate at about the same pace the business hires new employees.

2. Public assets cover anything the public is supposed to see. There is little need for security for these files, since you want them to be viewed by the outer world. And it's a great option if any of these go viral. Public assets tend to accumulate quickly—especially if the company has an active community of users who generate ongoing blogs, comments, photos, and videos.

RACKSPACE® HOSTING | 5000 WALZEM ROAD | SAN ANTONIO, TX 78212 U.S.A

rackspace® HOSTING

Category	Description	Growth Rate
1. The basics	OS, desktop applications, e-mail, small databases, routine business files	Predictable; grows with company headcount
2. Public assets	Marketing materials, websites, blogs, videos, user-generated content	Very unpredictable, especially with active communities.
3. Private assets	PII, customer files, transactions logs; and all financial, HR, insurance, medical, and school records	Unpredictable; grows with number of clients, employees, and transactions
4. High-performance assets	Any files or databases that require high-availability, high-throughput, and high reliability	Unpredictable; grows with number of clients, employees, and transactions

Table 1: Four Categories of Business Information

3. Private assets cover anything the public is not supposed to see. That includes any personally identified information (PII): name, address, SSN, driver's license, credit cards, and so on. Every business wants to safeguard transaction records and anything covered under FERPA, HIPPA, HITECH, or Sarbanes-Oxley, or SEC rules.

Private assets tend to grow at about the same pace as the business. In other words, as the number of clients, employees, and transactions grows, so does the volume of private files. These are also subject to unpredictable spurts.

4. High-performance assets include heavy-traffic websites and mission-critical databases. This is the only category that can overlap with another.

For example, a large e-commerce site can be considered both a public and a high-performance asset. If the product line grows and sales explode, so will these assets. Even if the site's sales are stable, it could experience a sudden spike at any time.

In Storage, One Size Does Not Fit All

Many IT departments are committed to one or two types of storage technology. The company may have standardized hardware and IT may not have time to investigate other approaches. But this one-size-fits-all design can overlook many of today's choices.

As one commentator says, "single platform shops are becoming as rare as an ice cube in the Sahara."[4]

After all, three out of four asset types can grow unpredictably. If all your assets are stored on a single type of network server with a single architecture, this can severely limit your flexibility. This is especially true when you launch a new app, new site, or new database that falls into a different category of asset. That's the best time to consider a wider set of options.

A Better Solution: Match Technology to the Asset

Many IT departments now use hosted storage to gain better cost, expandability, and performance... and to match the right technology to the right asset.

There are many service providers to choose from, beyond the old-line IT giants. For example, the hosted storage options from Rackspace have been well received by customers and the industry.

"Rackspace has very quietly outflanked giants Microsoft and Google in online storage," declared a recent review on CRN.com.[5] When NetworkWorld compared five storage vendors including Amazon, "Rackspace delivered the best overall performance."[6] And a blogger who ran speed tests found Rackspace much faster: "After 24 hours of monitoring, the average response time from Rackspace was 10 times faster than Amazon."[7]

Rackspace offers a range of storage options, as shown in Table 2. These options provide for every type of data asset and every budget.

Internal or Direct-attached Storage

For any hosted customer, additional hard drives can be plugged into existing servers at the data center until all the slots are filled. This builds more storage capacity in a transparent fashion, at a known cost, and takes care of many SMB needs.

Cloud Storage

Many companies are benefitting from the dynamic expandability and low cost of the cloud. "Cloud storage offers pay-as-you-go, elastic self-provisioning with low prices per unit storage per unit time, making costs easy to predict, control, and map to the workload of an organization's IT assets," said a recent white paper from Oracle.[8]

Asset	Option(s)	Capacity	Ideal For
1. The basics	Direct-attached to remote servers	Physical limit of slots and capacity of drives, up to 50 TB	Any hosted assets, most SMB needs including DB applications and HA Clustering
	Cloud Drive	Unlimited, expands dynamically at flat price per GB	Backups, file sharing and syncing
2. Public assets	Cloud Files	Unlimited, expands dynamically at flat price per GB	Unlimited storage,fast distribution of public assets
3. Private assets and 4. High-performance assets	1: Shared SAN from EMC	in 250 GB, max 5 Terabytes, in 50-GB increments	Entry-level clustering or virtualization with unpredictable growth
	2: Dedicated SAN from EMC	960 TB	High-performance databases
	3: Dedicated NAS from NetApp	1 PB	Rich media websites, virtualization, network shares, cloud shares

Table 2: Hosted Storage Options from Rackspace

Rackspace has two options for cloud storage:

Cloud Files provides unlimited storage with high-speed delivery of public assets. An easy-to-use control panel enables any knowledge worker to upload, store, and share files. The API can be integrated with web apps for even better performance. And the Content Delivery Network can cache files on servers around the world for faster downloads.

Cloud Drive is for businesses that need to store files, back up PCs, and synchronize basic assets. This works like any other drive on a computer, but stores files in the cloud for simple and secure sharing and backup.

Shared Storage Area Network (SAN) from EMC is for businesses that need shared storage or entry-level virtualization or clustering, but aren't sure how much their needs will grow in the future. In this multi-tenant approach, space on the server is shared with other businesses. Some limitations are in effect, such as no support for snapshots or replication.

Dedicated SAN from EMC is for businesses that need fast, highly secure storage for public/private and high-performance assets. This SAN provides block-level storage via fiber channel on a dedicated server that enables snapshots and clones for easy backup and replication.

Dedicated SAN from NetApp is for businesses that need fast, highly secure storage for public, private, and high-performance assets. This SAN provides file-level storage via network-attached storage; ideal for Web 2.0 and rich media that requires fast delivery and high availability. Storage optimization achieved by consolidating multiple versions of the same files (de-duplication).

Conclusions

Most data assets are growing in most businesses today, some rapidly. But storing all assets with the same technology is old-fashioned and inflexible. It's better to match the right storage technology to the appropriate asset.

Hosted storage from a leading provider like Rackspace offers a range of cost-effective options plus a flexible roadmap to handle both steady growth in storage and sudden spikes in demand.

If you'd like to hear from the author, please contact our Enterprise team and we will happily schedule some time to discuss this topic further with you.

Philip Bankey
210.312.1779
philip.bankey@rackspace.com

Katie Schmidt
210.312.1788
katie.schmidt@rackspace.com

Sources

1: Joe McKendrik, "Big Data, Big Issues: The Year Ahead in Information Management," Database Trends and Applications, 30 November 2010
2: Joseph F. Kovar, "SMB Storage A Fast-Growing Need," CRN.com, 23 January 2009 retrieved from www.crn.com/news/storage/212902209/smb-storage-a-fast-growing-need.htm
3: Stephanie Balaouras, "The Continued Shift to Disk-Based Data Protection Solutions," Forrester Consulting, January 2008, page 4
4: Ibid, McKendrick
5: Edward F. Moltzen, "Analysis: Rackspace Beats Google, Microsoft At Hosted Storage," CRN.com, 25 November 2009, retrieved from www.crn.com/news/storage/221901135/analysis-rackspace-beats-google-microsoft-at-hosted-storage.htm
6: Tom Henderson and Brendan Allen, "Cloud storage lives up to the hype," Network World, 23 August 2010
7: Chris Meller, "Amazon CloudFront vs. Rackspace Cloud Files CDN Performance," personal blog, 22 October 2009 retrieved from http://chrismeller.com/2009/10/amazon-cloudfront-vs-rackspace-cloudfiles-cdn-performance#more-642
8: Cris Pedregal, "Cloud-Stored Office Database Backups," Oracle, May 2010, page 3

RACKSPACE® HOSTING | 5000 WALZEM ROAD | SAN ANTONIO, TX 78212 U.S.A | 09272011

rackspace® HOSTING