Enterprise Strategy Group

REPORT

Managing Data at Internet Scale

By Terri McClure
With Steve Duplessie

December, 2008

Table of Contents

New Market Dynamics Drive New Storage Requirements

Within a relatively short time, the majority of capacity under management in the commercial sector will be born as file-based, rich digital content. Just as small, random access file data generated in the Distributed Computing Era dwarfed small random access block data capacity from the Transactional Era, in short order, the large-file, collaborative data of the Internet Era will do the same within organizations. Further, where large orders are still measured in TBs in transactional or distributed environments, they are measured in multiple PBs in Internet computing environments.

The Internet Era of computing is upon us and commercial enterprises are going to get dragged in—whether they like it or not. Web 2.0, cloud computing, and SOA content/data will coexist with transactional and distributed content in commercial enterprises—requiring a mix of price, performance, and functionality that differs from both.

Traditional storage technologies were not designed to scale in the multi-petabyte Internet Era. With traditional storage architectures, new arrays need to be added as capacity requirements increase. As the numbers of arrays under management grows, the storage environment becomes increasingly complex, harder to manage, and more costly to operate. This brings negative consequences to the business: increased time to market, loss of productivity, and decreased flexibility. Traditional storage technologies continue to excel in the areas they were designed to address—namely, transactional and distributed computing—but these solutions fall short for Internet Era requirements.

So how do you build and manage a global system that scales, and do it cost effectively? If it cannot be accomplished with traditional storage architectures, what are the requirements for storing Internet Era data? This paper proposes a new approach, focusing on five key infrastructure requirements:

- **Infinite Scale** – scale in real-time, dynamically, with little to no human intervention
- **No Boundaries** – expand outside the walls of the IT department
- **Operationally Efficient** – leverage commodity components, policy-based automation, and secure multi-tenancy to realize economies of scale
- **Self-Management** – automatically re-balance and optimize without human intervention
- **Self-Healing** – withstand failures and automatically adjust /heal itself

Creating, Accessing, Storing, and Finding Data in the Internet Era

Any device connected to the Internet has the potential to create, access, move, find, manipulate, delete, store, or manage online digital content. In this era, everyone—via most every conceivable device—can be connected to everyone else, eliminating any geographic boundaries that previously existed. As in the Transactional and Distributed Eras, corporations will be responsible for creating the policies and methods that ensure the proper use and protection of digital assets. Unlike the previous eras, any security benefits associated with limited physical connectivity as a defense mechanism have been effectively eliminated.

The Internet Era of commercial computing has been fueled by several factors:

- The Internet itself has reached every corner of the world, effectively creating a flat global network with few, if any, barriers to connectivity.

- Technology advances have enabled the creation of digital content to occur at every point of connection—the core transaction system, the desktop, in handheld devices, cell phones, laptops, smart cards, and countless other devices. Anything and everything that can connect to the network has the potential to create, access, and move online digital data.
- The ability to share digital content with the world now occurs in real-time.

In the Transactional and Distributed Computing Eras, corporate information assets were created and controlled almost exclusively by the business itself. Corporate data may be displayed and accessed via the Internet, but rarely were non-trusted users allowed to create or manipulate corporate information. Conversely, most consider even an attempt by those without permission to be a cyber crime or hack. In the Internet Computing Era, we see a rapid increase in data creation and access points for trusted users, opening up parts of ourselves and our systems to the rest of the world. Customers and other interested parties will change the way business is fundamentally done by providing content, aiding in product development and support, sales, and calling you out on your mistakes.

Corporate computing environments, while lagging behind consumer markets, are slowly but steadily moving into the realm of Web 2.0. Whether as an online community, social networking site, new media, or collaboration, the commercial computing world is going to have to ready itself to participate in the rapidly evolving realities of business. Tools such as SharePoint, blogs, wikis, streaming media, and a host of other digital content creation and management applications are enabling organizations to redefine themselves in almost real-time. New media content is being created for everything from training to marketing and becoming a mandatory component of everyday business. Whether it's blogs or video, content is easier than ever to create—and management will become harder than ever without significant changes.

Characteristics of Internet Era Data

The very nature of data has changed. No longer primarily concerned with block-based transactional data, IT initiatives must deal with a new breed of data: Data that is almost exclusively file- or object-based. Furthermore, the files themselves are changing—becoming larger as their richness increases. Factor in that the ability to create content is becoming ever easier, and it is no wonder that the amount of data we are forced to contend with is growing explosively. This era requires scale-out at the server, network, and data layers in ways that have rarely been seen before. No matter whose data you look at, the growth of capacity is accelerating—and the creation devices, use cases, and makeup of that data will continue to change because when it comes to Internet Era data:

- **There is a lot of it:** Because of the ubiquity of content creation devices and the richness of data, Internet Era information is measured in petabytes rather than terabytes. Even though much of this data is rich data, there remains a mix of large and small files that need to be managed. For example, social media sites often store multiple versions of images—from thumbnails to high resolution.

- **It is global in nature:** Locality of data becomes of increasing importance. Not in terms of where it is stored on a disk drive, but in terms of geographic location. The Internet is a global networking device that reaches the farthest corners of the world. Despite advances in networking and bandwidth, latency—the time it takes information to get from where it is stored to where you are—continues to be an issue. To significantly improve access time, Internet Era data needs to be in geographic proximity to where it will be accessed most frequently. For businesses that have global workforces, faster access time directly translates into greater productivity.

- **It has unpredictable use patterns:** Who would ever guess that a relatively unknown governor from Alaska would be a vice presidential pick in the 2008 US presidential election? Within a fairly short period of time after Saturday Night Live aired a parody featuring impersonations of Sarah Palin and Hillary Clinton, the NBC web site had over six million views of the video clip of the skit, making it NBC.com's most watched skit ever.[1] Palin was still one of the top five internet searches on Lycos a month after the '08 presidential election ended.[2] Nobody could have predicted the web traffic that would be generated

[1] Source: http://downloadmovies101.com/wordpress-1/2008/09/20/feys-palin-video-hits-6-mil-on-nbccom/, 5 Dec. 2008
[2] Source: http://50.lycos.com/, 5 Dec. 2008

on a global basis by the Palin pick. What is the next world event that will generate global interest? Where on the globe will that interest come from? Web 2.0 service providers need to be equipped to respond to these types of surprises by having the right data in the right place at the right time, on a global basis.

Challenges with Taking a Traditional Storage Approach for Internet Scale Data

Difficult to Manage at Internet Scale

Most traditional monolithic NAS systems are designed for "Tier 1" transaction-oriented data, or file data that is relatively small in size and typically associated to a single user. This type of data functions similarly to transactional block data, where scale in users and total I/O operations are paramount. These systems are big iron, using the fastest disk drives and boasting robust high availability features. They do, however, get difficult to manage in large scale deployments because each system is managed as a discrete file server. Managing at Internet scale can get complex, requiring management of hundreds of individual systems with no ability to load balance for efficient controller and storage utilization. And because these systems are designed for transactional data, most can't meet the high bandwidth performance demands of rich Web 2.0 data. These challenges are difficult enough to overcome, but add the challenge of managing storage services such as deduplication, compression, and versioning, as well as low level storage management tasks like load balancing within the arrays, provisioning, and troubleshooting all of the individual systems, and the operating costs skyrocket.

Designed for Operation Within IT's Walls

Today's storage architectures fall short when employed for Internet Era infrastructure requirements if for no other reason than the fact that they are designed to work within the walls of an IT department. Scale-out clustered NAS systems are gaining greater popularity because they overcome the scale limitations of monolithic systems. In scale-out NAS systems, independent nodes can be added to scale processing power, bandwidth, and storage capacity online and independently. Regardless of the number of nodes, the system is managed as a single entity that grows organically. In most, both performance and capacity are automatically load balanced for efficient utilization. Because bandwidth and processing power can be added independently, these systems typically perform well in Web 2.0 applications where rich media makes up the bulk of the data. But most scale-out architectures are still limited to a local installation within the four walls of the IT department. To overcome network latency on a global basis, you'd need to put one in each geographic location. The challenge with this implementation is an inability to share resources, load balance the systems to ensure the right data is in the right place, and set policies across the geographies. These systems are simply not designed to meet geographic distribution requirements.

The challenge of global content distribution is not mitigated with remote mirroring solutions. Most remote mirroring solutions do not allow use as an active/active cluster, so information is not accessible from each location in real time. Considering the scale of Internet deployments, managing and synchronizing remote copies, manual attempts at geographic load balancing, and the extreme, unpredictable performance requirements of Internet data, managing this solution can get ugly, fast. Growth in both the volume of data and the already enormous complexity of enterprise infrastructure can only lead to an inevitable catastrophic breakdown.

Elongated Change Management Cycles

As Web 2.0 migrates from questionable business value to mainstream business deliverables, IT departments are increasingly finding themselves without a plan. How does an IT operation go from taking three months for planning and provisioning infrastructure to support a business application to an infrastructure that can be created and scaled in semi real-time? Internet Era information demands are unpredictable. When every device is a content creation machine, the ability to predict data growth via classic trend analysis or three month planning cycles is out the window. Without any realistic method to determine how fast growth will occur, it is very difficult to plan.

Cloud Storage

A Cost Effective Approach to Handling Internet Era Data Growth

Data storage growth is pushing global enterprises to the limit—many are running out of space, power, and cooling. These enterprises simply cannot *afford* to add one more terabyte of data. Fortunately, we are in the midst of a perfect IT storm, where bandwidth has increased enough to allow enterprises to look at alternative storage models and storage software technology has advanced enough to offer them.

To deal with Internet Era data growth, a massively scalable infrastructure is required—one that offers global data distribution, self-healing, self-management, and multi-tenancy features. The storage infrastructure needs to store files or objects and be accessible via multiple protocols, whether they are NFS and CIFS or web services protocols like SOAP and REST. It must have rich metadata to tag content and use those data tags to apply policies, improve searches, or build custom queries. It must have policies that can leverage metadata to determine where and how the system stores data so it can meet information performance and availability requirements such as how data is protected, for how long, with how many copies, and where it should be stored geographically to best meet requirements. And it should include data services to handle policy, replication, versioning, compression, deduplication, and spin-down. The system should have:

Infinite Scale – The cloud storage system needs to scale performance and capacity infinitely within a unified namespace. A unified namespace enables cloud storage to scale massively while still operating as a single entity. As processor and storage nodes are added, they must be dynamically incorporated into the namespace, system, and load balancing algorithms in real time. There can be no human intervention required; nodes should be self-discovered. Essentially, a cloud removes the limitations of individual devices, thereby removing the boundaries of the boxes and enabling efficient single-level management of multiple locations.

No Boundaries – A storage cloud should have global distribution capabilities that allow for policies to be put in place to automate content delivery to match content popularity, geographic location, and desired retention periods. In other words, the system needs to be able to respond to the demands of Internet Era information—by scaling quickly and having the high bandwidth required to handle large objects with high levels of concurrent data access. For global enterprises and service providers, this means policies can be set to place data in the geographic region from which it will be accessed most, significantly reducing latency.

Operational Efficiency – Cloud storage systems must leverage commodity components, policy-based automation, and secure multi-tenancy to realize economies of scale. Multi-tenancy allows enterprises or service providers to support multiple applications from within the same infrastructure and, in essence, fence them off from one another. This allows service providers to benefit from economies of scale and operational efficiencies by having everything housed within a single system under a unified namespace, rather than separate systems. Support for web services via REST and SOAP, as well as legacy protocols such as CIFS and NFS, is also desirable to remove the need to deploy specialty solutions for each.

Self-Management – A lot of the complexity in managing a global cloud is taken out of the equation with a unified namespace. But that does not go far enough. Management is also simplified with integrated, policy-based data services that handle deduplication, spin-down, compression, and replication. Data services should be native to the cloud—rather than purchased, licensed, managed, and maintained separately—to produce a solution that is fully integrated rather than bolted together. Users need to get out of the business of software license and revision tracking sheets and get back to enabling business advantages through IT.

Self-Healing – Cloud storage needs to protect data, not disks. Data needs to be protected based on policy. We've already seen several storage vendors make the shift from sector mirrors to object mirrors, with good reason. In an object or file mirroring system, all data objects are mirrored for basic protection from drive failures. Unlike RAID 1, where two disks act as mirrors for each other, object-based mirroring ensures the data object is mirrored on another disk. For example, when an object is written to a disk (disk 1), it is mirrored to another disk, disk 2. The next object written to disk 1 is mirrored on a different disk (disk 3) and the next object is mirrored to yet another disk (disk 4). If disk 1 fails, the system does not try to "recreate" disk 1. It recognizes that every object that was on disk 1 is unprotected and it writes new object mirrors distributed across the storage system, based on policy. This allows the storage cloud to leverage the power of massively parallel disk reads and writes to create new mirrors—and ensures the data gets back to compliance with its policy in as short a time as possible. There are vendors doing the same with information-based parity schemes, where parity is spread across nodes based on information objects rather than disk sectors.

Use Cases: When a Cloud Approach is Appropriate

The storage cloud is not for general purpose use. It is not meant for transactional data or high performance file sharing applications. Those applications will continue to need local storage and the more traditional monolithic and scale-out systems will continue to meet those needs. Transactional applications would starve in a cloud environment; network bandwidth and latency would not keep the applications fed with data. Cloud storage is most suited for rich Internet applications, online service providers, and digital media creation and distribution.

- **Rich Internet Applications:** Rich internet applications are anything involving multimedia, generating terabytes of information. Capacity needs cannot be forecasted—the next big thing, like a relative unknown vice presidential pick, could lead to a massive increase in information and interest on a worldwide basis. One well known US social media site has 2–3 terabytes of photographs uploaded every day and serves 15 billion photos a day, which is 300,000 images each second. Typically, new content generates initial interest and intense activity, and then it sits idle for however long the user leaves it posted—potentially forever! With policy-based information management features, popular content can automatically be placed on higher-performing disks in the geography closest to the most intense activity.

- **Online Service Providers:** Multi-tenancy allows service providers, or internal IT departments that act as service providers to the business, to realize economies of scale. Service providers can use the multi-tenant feature and the web service APIs to build an application that offers secure online storage for consumers. And multiple consumers can be served via the same cloud storage infrastructure. Because of the shared infrastructure, managing storage for one hundred or more clients is no more complex than managing for one. For example, a photo sharing and printing site could create a service that allows subscribers to securely store their own digital content in a global vault. The service provider can then choose to monetize the service or give it away for free as a competitive offering.

- **Digital Media Creation and Distribution:** Digital media creation and distribution is another area that could benefit from a cloud storage approach. This is an intensely collaborative process that could involve hundreds of people sharing large files or objects. Creating a global, policy-based network that ensures data is in the right place at the right time would significantly speed the development process. For example, once a film is ready for dubbing into new languages, it can be distributed to the dubbing service provider in real-time, as soon as it is ready. Digital media distribution that leverages a global cloud architectures allows movie studios to get digital films to market anywhere in the world faster and significantly cheaper than through the traditional method of shipping disks.

Summary

Basic data management has been a never-ending problem for IT organizations for many years. New business processes—coupled with the assault of new, large file-based digital content—will likely crush existing norms. There is an absolute need to re-evaluate processes predicated upon the transaction-based computing of the past and focus on finding new methods of dealing with the criteria of the Internet Era. The ways we handle information file storage, backup, archive, search and retrieval, content delivery, and collaboration will inevitably have to change—in terms of both infrastructure and management. In addition, infrastructure files in the form of virtual machine images, reusable web services, file systems, and management databases will continue to proliferate widely. IT managers charged with protecting and managing the dynamic range of files now face an almost untenable situation.

For rich Internet applications, digital media creation and distribution, and online service providers, adopting a globally distributed cloud solution can be the answer that allows them to retain some level of sanity. The cloud, as described here, is an autonomic solution that takes almost all low level storage management tasks and automates them. It allows storage managers to make one decision about the information being stored rather than dozens, reducing the burden of deciding which policy to apply. And that can be automated based on metadata. Taking this approach frees up human resources, space on the IT floor, and capital budget—who could ask for more? But this approach does not absolve managers from practicing common sense and due diligence. The initial setup of an automated system becomes all the more critical to ensure the storage environment meets all of your requirements. Once the environment is set up, tested and audited, the operational burden is significantly reduced. But getting it right up front is critical to the success of any automated system.

ESG Enterprise Strategy Group

20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com