

Ensuring Global Presence in a World of Data

New Standards in Global Data Management and Access Render Old Storage Paradigms Obsolete


Businesses are Engulfed in Data

Large and small enterprises across the world face a rapidly growing problem rooted in pervasive technology: they are engulfed in data. Stored data is expected to grow 800 percent in the next five years; 90 percent of today's data mass is less than 2 years old.¹ Every three years we *double* the amount of data recorded and stored since the beginning of recorded history. The data floodgates have opened.

Proliferation of Low-cost Mobile Devices and Sensors

What's causing this growth? Worldwide, there's an ongoing explosion in the number of handheld mobile devices, digital sensors, and other devices that are "network attached" – and they're all relaying information across the web 24/7. We don't think of it much, but every one of these devices is connected to the web, and with an ever-widening range of apps they're used for everything from banking online to tracking deliveries to signing agreements to reporting the status of mechanical devices. Mobile phones are now central to everyday life, everywhere:

- By the end of 2012, the number of mobile-connected devices will exceed the number of people on earth.
- By 2016, there will be over 10 billion mobile-connected devices, though the world's population will be just 7.3 billion.
- By 2015, mobile internet use will exceed desktop internet use, changing how we use, share, and rely on data.
- In 2000, Nigeria had about 100,000 landline phones run by state-owned telecoms, but today the state telecom is gone, and Nigeria has close to 100 million mobile phone lines, according to Nigerian Communications Commission statistics.



The *quantity* of data is growing...
Beyond the scope of imagination

To function, all of these devices are utterly dependent on instant access to the most reliable data, and, simply by functioning, they all *feed* the stored information mass. They generate an unbelievably large amount of data, and are accessing, changing and updating data constantly.

Every time someone (e.g., 1.34 billion people in China) takes a photo with their smart phone, that photo is added to a database, and then it's replicated, moved and downloaded into other systems as it's shared. If one person's photo of his favorite pet goes out to 50 people, then suddenly those 50 photos are circulated, moved, replicated, edited, deleted, etc. over and over. This happens countless times daily all over the planet.

Highlights

Problem: Overwhelmed with data!

Need a "state-of-the-art" data-management solution that transcends "old" NAS limitations.

Local Access – automatically manage many storage tiers (locations).

Global Access – connect users globally with previously impossible efficiency.

Easy-to-use self-service Cloud access to meet demands of rapid data growth.

Enterprises can't adapt to ever-changing business needs without 21st Century data storage-management-and-access tools.

Access and Management will be Crucial to Sustaining this Explosive Data Growth

This unprecedented growth, like nothing we've ever seen, affects almost all organizations! The *massive* volume of data we generate potentially holds great utility and value, but only if it's managed effectively. Worldwide data proliferation doesn't mean much if the right data isn't getting to the right people at the right time. As available data grows, *demand* increases for easy-to-access, high-utility, very reliable information. Mobile-device users are impatient; they want and expect data to be delivered instantly. Companies with far-flung employees and customers need to share and update data files securely. Providing efficient access to data is an imperative.

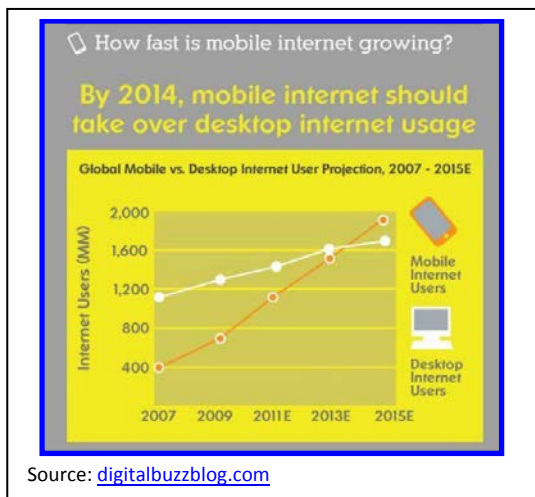
But the very network-attached devices that enable so many users in so many places to access so much data also pose a major data-proliferation problem and a threat to data integrity; they enable files and data to grow and change very rapidly through instant and easy digital duplication and manipulation. Data managers are left with a fundamental and herculean challenge:

Knowing which data file is where, and which file is the *real, current, trusted* original - the one that can be *relied* on – and ensuring that multiple versions are consistent. The opportunity for confusion and errors regarding a data file's status (original, current, updated, reliable, or not) is enormous.


With the entire world able to access data, ensuring the original's integrity, and that everyone *has* the right version, is a hypercritical key to data management. In fact, customer satisfaction today *depends* on a company's ability to deliver *reliable* knowledge instantly to all users everywhere. The related storage and data-control requirements present a serious managerial issue to any organization that must maintain data effectively to succeed.

A Strong Data-Management Foundation is required ... beyond Traditional NAS

So, how do we manage it? How do we keep up with it? This data tsunami must be stored and managed *dynamically* – it's a monster task that requires a very muscular storage infrastructure, one that moves data managers beyond the "old way" of storing, finding, and managing data; one that supports large-scale operations, and anticipates a global reach. It requires a new, very strong foundation, nothing short of "[Watson-like](#)" intelligence.



It's no longer enough to simply store data – today it must be analyzed, applied, leveraged, controlled, accessed, and converted into value by harnessing the actionable insight, new ideas, and new services it holds. This must be done across multiple platforms, in real time, globally. To do this companies are increasingly turning to cloud technology and services, and as they do, many data-intensive organizations are realizing that their own file server systems and public clouds have security, control, and reliability limitations. As data growth explodes, demands on cloud access and management increase dramatically, and a *trusted* "cloud" platform is essential. So, who do you trust?



Limitations of the “Old” NAS – Typical cloud file servers operate with traditional NAS (network attached storage) technology, which organizations with fast-growing, rapidly multiplying file systems are now discovering has serious limits and challenges. Accommodating and managing more data storage as it grows to millions and billions of active files is very complex; with NAS this requires adding more and more separate storage islands or disks, which stand alone (with separate file systems) and, despite being connected together, aren’t designed to “scale out” performance and capacity. It’s a “silo-building” approach to growth. As the number of silos increases they become increasingly hard and costly to manage; they often fragment and imbalance data, underutilize disks, and generate hot spots because they don’t talk to each other very effectively. One silo doesn’t know what data is in the other silos. It’s easy to overload some silos and not others, to the point where performance suffers. Operational complexity grows exponentially.

Old NAS techniques and limitations also complicate the issue of ensuring consistent data integrity, and hamper efforts to attain it. This leads many to go to strange and complex lengths (e.g., multiple and fractured layers of back up in different locations) that end up adding new inefficiencies and cost.

Traditional NAS is no longer able to scale or provide the automatic infrastructure for today’s vast data skyscrapers that must serve and interact with users across the globe. Yet, most enterprises still store and find data in this more cumbersome old way. And these old data-management techniques really slowing enterprises down.

The “New” NAS – Scale-out NAS is a new, emerging data-storage technology gaining traction as enterprises face a mountain of rapidly-growing unstructured data. Its basic concept is to consolidate tens or hundreds of file servers, and store millions to billions of active files in a single virtual container (namespace) and enable easy access. This spreads data over multiple clusters and leverages a distributed file system that can grow as needed. Scale-out NAS is managed from the center, and permits additional file servers to become part of the system on demand. With scale-out NAS it’s no longer necessary to manually move terabytes of data among multiple boxes to balance the load.

Scale-out NAS systems are now available through various vendors (with varied capabilities), and many offer enterprise-level features like parallel access and asynchronous replication.

IBM’s Scale-out NAS Vision stands alone – Moving Beyond Scale-out to a Robust *Global Reach*

IBM’s vision of the Scale-out NAS concept (called “SONAS”) is very robust, capable of handling any job, and highly distinguished from the others. The unique SONAS vision offers a strong foundation that moves data management far beyond the limitations of traditional NAS, enabling users to reliably aggregate vast data-storage facilities globally, through an intelligent, adaptive approach designed to grow with unprecedented scale. This is a giant leap forward.

The future of cloud-based data management and access holds great promise, and IBM’s SONAS vision is a paradigm-shifting foundation for *private* clouds that must offer high security, flexibility, and global reach. IBM is way out in front leading the charge, delivering *necessary* change and redefinition based on an inexorable reality that’s coming fast. IBM’s Scale-out NAS vision has three key elements, all built around the very powerful General Parallel File System (GPFS):



Managing this “Mobile Digital World” is necessary and challenging!

Self-managing Intelligence will be required!

- 1st – The system is very-high-parallel-processing enabled and can support and service billions of files, accommodating today’s high-demand workloads, and expected larger future workloads.
- 2nd – The system offers Multi-cluster capability, allowing multiple Scale-out NAS systems, regardless of location, to function as a “single” NAS system.
- 3rd – The Reach beyond Scale-out NAS element provides *both* local *and* global reach to all data no matter where it is located. IBM has trademarked this element as the IBM Active Cloud Engine™.

Other Scale-out NAS offerings provide **only** the 1st element!

But it is elements 2 and 3 of IBM’s vision that represent a real paradigm-shift, and a revolutionary jump to a “new curve.” It’s a jump very similar in its impact to the quick and sweeping change from Blackberry and Palm smartphones to the revolutionary Apple iPhone technology, which unleashed the *real* power of smartphones, forever changing how we use and interact with data – and we’ve never looked back.

IBM has tested and demonstrated all three key elements over the last 10 or more years, and presented the results in various IBM press releases. The tests have demonstrated GPFS’ very powerful core architecture (see accompanying box), and verified its suitability for deployment at different levels, each appropriate for a particular storage offering and customer requirements.

IBM is implementing *all three* elements of its paradigm-shifting Scale-out NAS vision with a global reach, and offers it in two primary storage solutions: (1) the IBM Scale-Out NAS (SONAS), and (2) the IBM Storwize V7000 Unified. While the GPFS architecture is all there in each case (with all of its capabilities), its full power isn’t needed or cost appropriate in most storage applications, and so isn’t fully activated – though it still packs an unbelievable punch. This permits customer environments the flexibility to grow and support any mix of their ever-growing data requirements. Very unique and very powerful!

With IBM’s unique combination of all three key elements, each solution can act as either an NAS spoke or an NAS hub. **“Truly integrated” global Hub and Spoke designs are now possible!** For example, unlike other NAS vendors, using the IBM Active Cloud Engine (explained below) IBM will offer the ability to integrate all of the following into one global namespace: a.) smaller traditional NAS systems (IBM V7000 Unified), b.) virtualized non-IBM NAS systems (SOD for SONAS to virtualize non-IBM NAS), thus preserving these investments, and c.) super-scaled IBM General Parallel File System (GPFS) installations.

These current IBM NAS storage solutions offer users the paradigm-shifting NAS foundation for IBM’s key-feature rollout over time, allowing them to grow into these fresh, revolutionary capabilities.

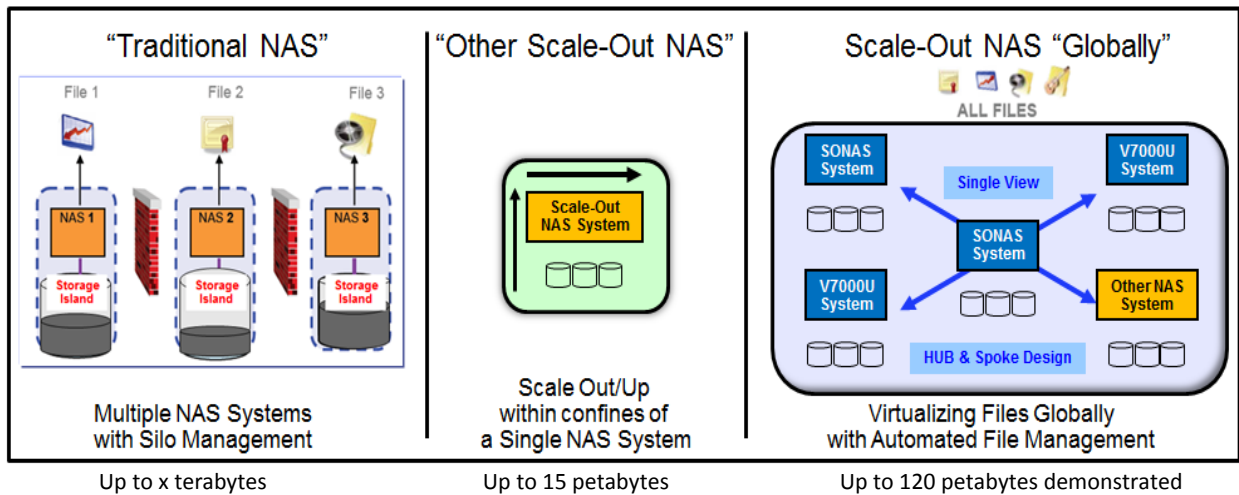
GPFS Lab Demonstrations

Smashing known records in a significant leap, IBM Research Almaden, California, has technologies using GPFS that allow it to combine 200,000 hard drives into a single storage cluster of 120 petabytes — or 120 million gigabytes — that will store one trillion files (or around two billion hours of MP3 music). Today’s largest systems are in the 15 to 25 petabyte range.

GPFS’ “one-of-a-kind scan engine” searches huge amounts of data – and finds what you need – at lightning speeds by working simultaneously on commands, not sequentially. It can read files on a massively parallel basis from multiple disks at read/write speeds of several terabytes per second, and can create more than 30,000 files per second. This is the same core intelligence (scan engine) that’s in IBM’s “Watson.” GPFS represents a true breakthrough in storage performance – it can scan 10 billion files in just 43 minutes (demonstrated in July, 2011), shattering the previous scanning record (set by IBM in 2007) by a factor of 37. It combines multi-system parallelization and fast access, and enables much larger data environments to be unified. Policy-based searches can now be completed in *minutes*, a pace untouchable by competitive scan engines!

Source:

Sebastian Anthony, IBM builds 120 petabyte cluster out of 200,000 hard drives, www.extremetech.com, August 26, 2011, <http://www.extremetech.com/computing/94082-ibm-builds-120-petabyte-cluster-made-out-of-200000-hard-drives>



The key to SONAS’ dramatic and unique capabilities is its seamless integration of three essential components.

First • SONAS includes IBM’s exceptionally high-performance General Parallel File System (GPFS) with its one-of-a-kind “scan engine.”

SONAS’ core base is built on GPFS, IBM’s premier scale-out file system, which is already handling mountains of data daily in over 100,000 businesses worldwide. GPFS grew from the 1990’s Tiger Shark file system research project, originally targeting access to multimedia video servers, and has now evolved to accommodate primetime use in all industries. GPFS enables applications on multiple nodes to share file data with performance and fault tolerance characteristics that surpass those of traditional distributed file systems.

SONAS and its GPFS search capabilities are the foundational piece of today’s data storage and management puzzle. The other piece is managing massive data volume across multiple storage locations and just-in-time access to that data by users all over the world. SONAS’ single-file-system performance is “best-of-breed” while other Scale-Out systems require more resources and more file systems to be competitive. It delivers effective multi-cluster cloud storage in the petabyte age, while controlling escalating data-management costs.

IBM’s SONAS has Three Integrated Components

1. *General Parallel File System (GPFS) and its one-of-a-kind “scan engine.”*
2. *Active Cloud Engine™ with its unique local & global capabilities.*
3. *Cloud Storage Access Self-service Portal Offering*

Each component builds on the others in sequence.

The combination of these components is essential to providing a real data-management solution that tackles the scale of data complexity enterprises face.

Second • IBM's Active Cloud Engine™ (integrated into SONAS) controls access to your data, and creates a high-utility, automated approach to 21st Century data management.

IBM's Active Cloud Engine™ – with Local and Global Access

Today, managing data and managing *access* to data go hand in hand. Ensuring efficient access to data by all potential users requires the integration of *both* a sophisticated data-management system like IBM's SONAS, and highly flexible 21st Century data-access techniques (i.e. Active File Management) that redefine what it means to deliver the right data to the right user/location at the right time. While many in the storage business and many data managers recognize the need for a New NAS, and some are working on Scale-Out-like capabilities, effective data-access techniques have proven elusive.

Fortunately, IBM has solved the issue of providing 21st Century data-access by making its “enterprise class features” available to everyday businesses in high-data-growth areas. To help small and large enterprises cope with exploding data-growth-and-use challenges, IBM has developed *and integrated* into the SONAS base (not merely an add-on) a dramatic leap forward in storage capability with its “Active Cloud Engine™.” Active Cloud Engine is a unique automated data-management tool enabling easy-to-use, and the fastest possible, global access to your data. This *real* 21st Century data-management solution delivers global reach, global presence, and unprecedented growth opportunity through “[Watson-like](#)” intelligence.

Active Cloud Engine™ is available with two distinct capabilities, **Local Access** and **Global Access**, each with impressive features!

Active Cloud Engine™ Local Access (Storage Tiering) – Redefining right data, right place, right time.

Active Cloud Engine™ automatically places data in the right place, at the right time, for the right *reason* – high-use data is placed for *ready* access and infrequently-used data is placed elsewhere. (By analogy, you store the table salt on the table, where it belongs, because you use it every day, but you store holiday decorations in the attic, not on the kitchen table, because you only use them once a year). This simplifies data-management tasks (i.e. storage tiering) and lowers operating costs by improving traffic flow and saving time. Note: the infrastructure required for high-activity data is different than that required for infrequently used data – so the system discerns data location partly on infrastructure capabilities.

The Active Cloud Engine™ system is designed to become *your enterprise's* private “cloud” – that is, your data accessible on your terms to users you choose (inside and outside your organization). You control (define) its access terms, and can set up and use multiple internal and external storage pools transparently; the system scans file metadata at very high speed to determine file placement in appropriate storage tiers and global remote locations.

Your private cloud can be custom configured to define who has access and what the policy rules are; you direct what happens to a file during its lifecycle or in relation to a particular scan, and the system acts automatically on those rules. Its automated access and use capabilities are astounding. IBM Active Cloud Engine makes administrative tasks easy and efficient through an intuitive graphical user interface (GUI) and built-in policy-driven management tools.

IBM's Active Cloud Engine™

*A powerful policy-driven engine, tightly coupled with your file system, designed to manage **massive** amounts of data with ease.*

Local Access

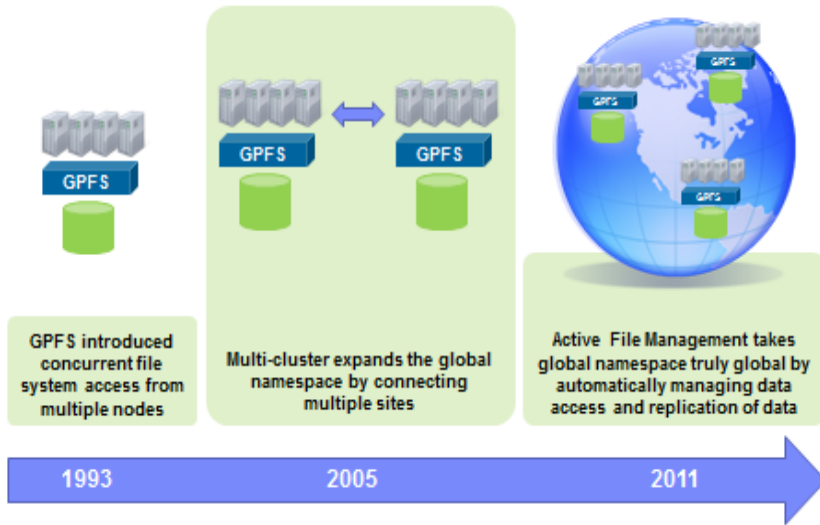
Global Access

It's Ideal Cloud Storage!

Active Cloud Engine™ Global Access (21st Century Access) – wide-ranging capabilities transform what was impossible into game-changing solutions.

What distinguishes IBM Active Cloud Engine is its ability to do things that haven't been possible until now. The Active Cloud Engine does the unique and quite remarkable (and what will soon be essential):

Apply New Techniques to your Old File System – Active Cloud Engine can harness and manage your existing servers and any other NAS (no matter the number or locations) by applying its new management and access techniques and features to your old system. This tight and



SONAS w/Active Cloud Engine™ Worldwide design for Global Presence

Source: IBM

seamless integration with your existing NAS file system(s) allows you to continue running *that* system and eliminates the need for disruptive and costly migration to an entirely new system. Of course, Active Cloud Engine can also accommodate those who do want to migrate their data to new servers.

Single View of Global File System – Also impossible before, and not available anywhere else, Active Cloud Engine enables consolidation by centralizing *all* files under a single file-system view, working with your varied data platforms (regardless of type, media, or physical location). Even better, beyond creating one scale-out system, Active Cloud Engine can unify *multiple* distinct scale-out systems into the single view (regardless of their type), and seamlessly manage all data simultaneously. This is a *very powerful* capability. Users experience the appearance of a single system, and are provided a single view of all files across all systems and locations (all users, no matter where they're located, have the same view). Remote users see and access all the files (the ones your policies give them permission to see) as if they were local, even though they are actually centrally maintained and retained. Without Active Cloud Engine, other scale-out systems are incapable of communicating or sharing outside of their own boundaries.

Global File Access with Worry-free Data Security – Users anywhere on the planet can instantly find and access data they need, and use it, add to it, modify it, or share and distribute it, regardless of location (theirs or the data's). Changes are instantly synchronized centrally. No other system can do this. Multiple remote users can simultaneously access a file in the global namespace, but Active Cloud Engine's single-writer caching capability permits only one party at a time to modify the file, thus ensuring complete integrity and control of master data files, across all data platforms.

Instant File Sharing and Distribution – Effectively store, share, distribute, and delete data – also at very high speed – throughout the globe. Seamlessly deliver the right data (files, images, updates) to those who need it, at the right time, regardless of location. Prepopulate (cache) files to remote locations needing them, before user requests, enabling instant file availability and very fast access at remote sites. The system can also be configured to define participants and automatically share data updates with participants requesting it. For those authorized to access your private cloud, the system, by default, automatically notifies them when data has been changed and resaved. Enterprises can define what files are pushed to remote sites for local viewing and use, and which files on a central site are viewable and accessible by remote sites.

Key New 21 st Century Features	Other NAS	IBM GPFS with Active Cloud Engine™
Scale-Out Design	Yes	Yes
Local Storage Tiering	Limited	Yes
Unique Remote Features	No	Yes, with Active Cloud Engine
Scan and Use Remote Systems	No	Yes
Cache Data Remotely	No	Yes
Distribute Files Remotely	No	Yes
“Pull” Changes on Demand	No	Yes
“Push” Changes Incrementally	No	Yes
Connecting Multiple Sites – (Multi-cluster NAS)	No	Yes
Virtualize Any Other NAS	No	Yes

Incrementally Push and Pull Changes – When data files are modified Active Cloud Engine automatically pushes *just the changes* (not the entire file) back to the central site, thus reducing network costs by minimizing bandwidth use and maximizing speed and efficiency. When users access a file Active Cloud Engine automatically conducts a “staleness” check to verify that the local file being viewed is the latest version; if it’s not, the system instantly pulls the current central version for viewing while updating the older remote file.

Improved File Access for Remote Users Reduces Network Costs – Your centrally stored files can be accessed locally from anywhere in the world, through a continuously updated “local cached copy” of the data that matches the “central copy.” This enables users to continue working and save time without relying on network connectivity. By localizing files where they are needed (i.e., distributing them closer to users) and consolidating files scattered in multiple storage locations the system reduces network costs and improves access and application performance.

Virtualizing (Connecting) Your Data Globally – Your storage system’s various elements (silos) are no longer isolated, and no longer need “hands on” maintenance activity to move data because Active Cloud Engine establishes a “logical connection” between and among the silos, essentially converting them into a single large “virtual” file server with automated storage placement (aka “tiering”).

Scalability for Dynamic Enterprises – Active Cloud Engine, while extremely powerful and capable of handling billions of files, isn’t just for large enterprises. It’s highly scalable, and accommodates rapidly-growing enterprises of any size, meeting their needs, and delivering significant value, at their pace.

With these features and incredibly fast GPFS scanning capabilities, the Active Cloud Engine means unprecedented speed and performance for users administering millions or billions of files and for organizations whose file storage needs are growing rapidly. These automated data-sharing and storage-management features simplify storage protection and deliver an enormous time savings to the cloud operator and its authorized users. Policy-based searches can now be performed in *minutes*, rather than in weeks or months. Dramatically improved data-management efficiency also offers tremendous economic savings, and enhances administrator productivity.

The IBM Active Cloud Engine™ is an essential tool for all data-reliant enterprises, and an exciting opportunity for managerial coherence. For those who need your data, access to this tool helps them do their job better.

Third • IBM’s Cloud Storage Access – enhances your global presence by providing a friendly user interface that accommodates the high demand for real-time access to ever-growing data!

Pressures like increasing workforce mobility and productivity impose great stress on storage systems and IT. As data grows exponentially, management is critical, and “simple” Cloud access will be demanded by all. To meet this unprecedented demand, Cloud access must be omnipresent, self-servicing, and up to date.

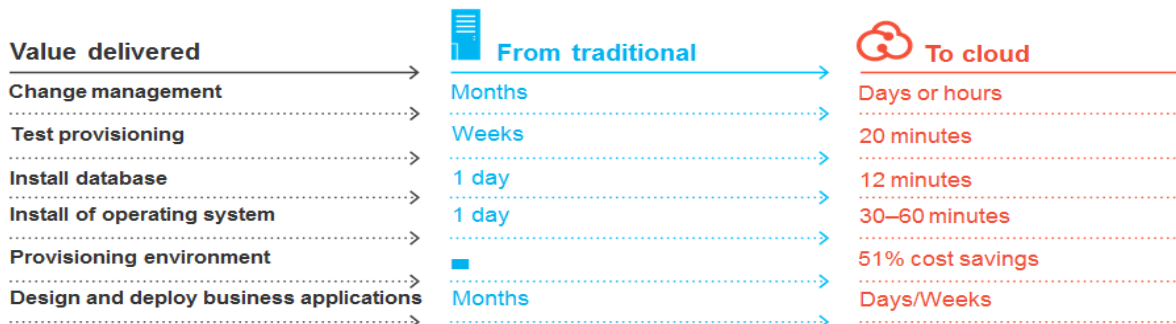


Source: IBM

Private storage clouds are attractive to enterprises because they offer security and value. First, they improve employee productivity by permitting them all to work together simultaneously from any location at all times. Second, they offer scalable on-demand capacity with self-service and policy-based automation. Third, they offer time and money savings through standardized, automated access and administration. The old adage “time is money” is more valid than ever, which renders the Cloud an increasingly appealing way of delivering IT services.

To deliver on its promised value and potential, Cloud storage must provide some key capabilities:


- **Self-Service** – the ability for users to manage their own IT environment in their own way.
- **Quick Access** – getting the right data to the right people at the right time.
- **Pay Per Use** – users pay only for what they consume.
- **Elasticity** – capacity for growth without constraints, and readily releases unused resources on demand.



Source: IBM

IBM Systems & Technology Group Cloud Storage Solutions now offers “**IBM SmartCloud Storage Access,**” software that taps into and fully exploits SONAS’ capabilities - it increases productivity, lowers operation cost and improves storage service quality. It enables enterprises to implement a private cloud storage service that permits users, with just a few clicks, to request and receive storage capacity, share files with other users, and allows administrators to easily monitor and report usage.

IBM SmartCloud Storage Access delivers real value through a simple, intuitive, self-service interface that improves user productivity as it reduces operational complexity.



Offered with IBM SONAS and Storwize V7000 Unified Systems

Source: IBM

Without GPFS’ Scale-Out File System, Advanced Search, Active Cloud Engine™, and Cloud Access capabilities, Enterprises Risk their Competitive Position

We live and work in a shrinking, network-connected world because of technology. The rate of data generation is accelerating. No enterprise can escape these realities. While history has demonstrated that “information is power,” just *having* information is no longer enough in competitive commerce. Today, it’s the ability to control and connect data and *deliver* information that is truly empowering.

To be used successfully and trusted, all the data we generate and store must be readily accessible, easy to find, and *current*. Organizing, accessing, and sharing the data we generate is critical to its utility and its value. Instant access to the most reliable data is no longer optional – it’s a necessity. No one can afford to use old cumbersome search technology any more.

SONAS’ unique Combination of GPFS, Active Cloud Engine™, and Cloud Storage Access Capabilities represent a *foundational* change in capability (similar in scope to the change from Blackberry to iPhone), and as such it’s simply a “must have” data-management tool. All of these capabilities are vital to optimal, global cloud storage services. These technologies will soon be a *necessity*, and it’s fair to say businesses without them will be in trouble, left in the data-storage Dark Ages as culture (especially among businesses) inexorably shifts to SONAS, GPFS and Active Cloud Engine as compelling standards.

Worse, without Active Cloud Engine, you’re wasting resources, will lose thousands of dollars on unnecessary network costs, will be forced to operate in archaic ways, and will continue to suffer through a much greater and unnecessary manual communications work. You’ll be missing the opportunity offered by the “best of breed,” “Watson-level” service when accessing and managing data.

Smart enterprise *needs* these tools. IBM’s SONAS, GPFS, Active Cloud Engine™ and Cloud Access capabilities permit users to take a measured, incremental approach to adapting and scaling their existing data-storage infrastructure to private cloud storage technology, while realizing immediate efficiencies. IBM is a natural choice for this important and inevitable transition.

No other vendor offers the advanced capabilities of IBM’s Active Cloud Engine. The SONAS architecture is sufficiently superior that its competitors are reduced to idle criticisms without competitive merit. Don’t be fooled by IBM’s competitors using fancy words to imply equivalence. They’re not equivalent ... here’s why. Other private cloud systems are essentially one-size-fits-all adaptations of public cloud platforms, which require effort, substantial modification of existing data operations, and disruption to implement. They also don’t integrate effectively with existing data-center file-management operations, and don’t improve on existing capabilities; most require users to buy both cloud computing and cloud storage services.

Why is use of IBM’s personal cloud critical for small and large enterprises right now? Because data volume is exploding exponentially today and it won’t stop anytime soon. Data growth is on a sharp upward track and there’s no end in sight. It’s preparation time.

End

PBW • Probizwriters, LLC

We invite you to contact PBW (www.probizwriters.com), an independent business-writing consultancy, if your organization or team is interested in developing authoritative proposals, business plans, white papers, eBooks, case studies, manuals, reports or other business materials, whether technical or marketing in nature, for internal or external audiences.

Copyright © 2013 Probizwriters, LLC.

All rights reserved.

IBM, the IBM logo, ibm.com, Active Cloud Engine and Storwize are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this document with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM when this document was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml.

Resources and References:

IBM Scale-Out Network Attached Storage:

Most advanced architecture, flexible clustered scale-out solution

<http://www-03.ibm.com/systems/storage/network/sonas/index.html>

Automated File Management with IBM Active Cloud Engine

IBM Document No. TSS03081-USEN-00

Made in IBM Labs: Researchers Demonstrate Breakthrough Storage Performance for Big Data Applications

IBM Press Release, July 22, 2011, [http://www-](http://www-03.ibm.com/press/us/en/pressrelease/35111.wss)

[03.ibm.com/press/us/en/pressrelease/35111.wss](http://www-03.ibm.com/press/us/en/pressrelease/35111.wss)

Endnotes

¹ http://www-03.ibm.com/innovation/us/watson/?csr=agus_watsonopad-20120718&cm=k&cr=google&ct=USBRB301&S_TACT=USBRB301&ck=watson&cmp=USBRB&mkwid=sqNrsoG1w_27280026457_432oop16376