



Transparent Data Movement: The Key to Your Storage Strategy



A RAINFINITY WHITE PAPER

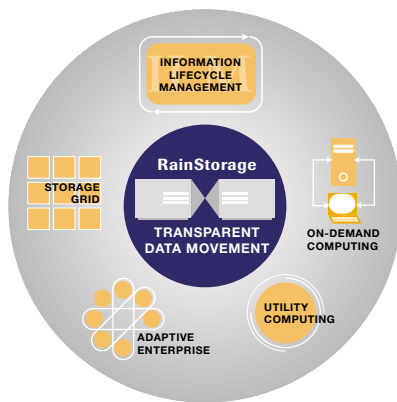
Transparent data movement creates a much more flexible IT infrastructure, enabling companies to better respond to rapidly changing business conditions and priorities.

No matter which solution you choose, the key to your storage strategy is transparent data movement.

TRANSPARENT DATA MOVEMENT: THE KEY TO YOUR STORAGE STRATEGY

Regardless of the storage strategy your company deploys, be it Information Lifecycle Management (ILM), On-Demand Computing, or Storage Grid, you must have the ability to move data efficiently, consistently, and securely across your enterprise. More importantly, you must have the ability to move this data *transparently*. Disruptive data movement can significantly impact end-user and application data access, translating into lower productivity, increased time-to-market and reduced profitability.

Transparent data movement creates a much more flexible IT infrastructure, enabling companies to better respond to rapidly changing business conditions and priorities. No matter which solution you choose, the key to your storage strategy is transparent data movement.




EXPLORING TODAY'S STORAGE STRATEGIES

As you compare storage strategies, you will find that they all share common goals, such as maximizing resources and minimizing costs. These strategies also share certain approaches, such as pooling storage resources, automatic provisioning based on need, and maintaining transparency to users and applications.

Another factor in pursuing storage strategies is customer service levels. Most end-users are unwilling to reduce current service levels as part of a new computing paradigm. Data movement that is not transparent will not improve service level agreements and may, in fact, actually reduce service levels. No matter which strategy you adopt, transparent data movement is the key to accomplishing each of these objectives.


In this paper we provide a snapshot of five storage and computing strategies — Information Lifecycle Management, On-Demand Computing, Adaptive Enterprise, Storage Grid, and Utility Computing — and discuss the critical role of transparent data movement in each.



Information Lifecycle Management

Information Lifecycle Management (ILM) enables companies to proactively manage data throughout the information lifecycle. This approach ensures that an organization's storage infrastructure delivers maximum business value at minimum cost. ILM is achieved by prioritizing data based on business requirements such as accessibility, protection, security and compliance, and aligning that data with the most effective storage medium.

ILM depends on the movement of data to optimize storage resources based on the changing value of the data. Without transparent data movement, the ability for data to flow freely across storage locations, devices and other mediums is severely restrained. The ability to maximize business value and minimize costs is limited to only those time periods or data sets that do not require data access. This is especially problematic with IP-based data. Large files used in collaborative processes such as Electronic Design Automation, Oil & Gas Seismic Analysis, Gene Analysis in Biotechnology, or special effects rendering in the Entertainment industry, all require long periods of access by many users. Many of these processes also involve long-running applications. EDA and seismic analysis applications, for instance, can run for multiple days. Any data access disruption requires these applications to restart. While it is not the only component, transparent data movement is critical for ILM.




ON-DEMAND COMPUTING

On-Demand Computing

On-Demand Computing calls for the automatic provisioning of IT resources based on a company's needs. An On-Demand strategy is based on autonomic computing technologies that enable companies to harness the collective computing power of disparate IT resources and manage them as a single system. IBM's vision of autonomic technologies is to facilitate more intuitive user interaction, providing a self-managing storage infrastructure that can configure, optimize, protect and even heal itself.


On-Demand Computing provides storage as needed, based on changing business requirements. Data must flow freely, without impacting users, to ensure data availability and maximum flexibility. Without transparent data movement, an On-Demand Computing approach could provision additional resources but utilization of the new storage would be gated by what data could move to the new resources. With transparent data movement that concern is eliminated and the degree of flexibility is increased.



Adaptive Enterprise

HP's Adaptive Enterprise is a strategy enabling companies to respond in real time to business changes, and to synchronize business and IT operations to take advantage of those changes. This methodology promotes moving away from a vertical infrastructure that allocates specific resources to individual applications, and adopting a horizontal infrastructure based on agile, modular systems that can dynamically and easily adjust to evolving business needs.


The Adaptive Enterprise aligns IT resources with business needs and requires a flexible IT architecture with rapid provisioning. This approach also requires data to dynamically adjust to continually changing needs and to transparently flow across the environment. As is the case with On-demand computing, transparent data movement is required to ensure flexibility and an agile response.



Storage Grid

Storage Grid refers to an approach that relies on a grid that connects an organization's systems to distribute workloads across storage resources and centralize management of those assets. The Storage Grid "virtualizes" the server layer of storage infrastructure, allowing any application to access data anywhere in the system, regardless of the physical location. Physical storage can be shared by a variety of data types, to maximize storage capacity.

Logical and physical locations of data are separated in the Storage Grid. Previously unconnected storage resources are optimized by linking them in a single "virtual" storage system. For this strategy to succeed, the placement and movement of data must remain transparent to users.



Utility Computing

Utility Computing is founded on the concept of delivering storage on an as-needed basis, similar to traditional utilities such as electricity. The VERITAS model for Utility Computing brings technology and business requirements in line, to provide high data and application availability, optimal system performance and automated adaptability to a changing business environment. This goal is accomplished by grouping hardware into pools of resources that can be rapidly or automatically reconfigured to respond to changing work requirements. Utility Computing relies on management software that determines the workload and controls resource allocation.

Availability of data is the core of the Utility Computing strategy. Utility Computing allows users to tap into any storage resources they need, without awareness of the actual hardware or storage management issues. Transparent data movement is required to insulate the user while ensuring reliable data accessibility.

TRANSPARENT DATA MOVEMENT IS THE KEY


There is an inherent need for transparent data movement with each storage strategy. All the strategies outlined are based on linking storage resources and optimizing them by moving data based on need. In addition, these strategies all strive to provide uninterrupted access to users. Two factors, data mobility and uninterrupted access, drive the need for transparent data movement.

Data Mobility without Disruption

Transparent data movement is important to storage innovation because it enables any storage strategy to overcome the obstacles that have historically hindered efficiency. Without transparent data movement, attempts at optimizing storage — such as data consolidation and migration — result in system downtime. This is at odds with the non-stop data availability required by our 24x7 global economy. Companies today simply cannot afford downtime.

With transparent data movement, organizations can eliminate the tradeoff between efficiency and end-user access. They can maintain 100% uptime and implement a storage system that continuously optimizes their storage resources. RainStorage, a new IP-based storage solution from Rainfinity, makes transparent data movement possible.

RAINSTORAGE: THE TRANSPARENT DATA MOVEMENT SOLUTION



Rainfinity's RainStorage is the industry's first transparent data movement solution that allows data to move between any IP-based storage device without any disruption to applications or end-users. RainStorage allows organizations to perform storage management tasks including data migration, capacity management, directory renaming and storage consolidation on the fly — increasing storage utilization and efficiency.

RainStorage facilitates optimization of storage environments, via all the strategies described above, by eliminating end-user and application downtime. This transparent data movement solution enables storage administrators to perform data management tasks such as data migration, volume expansion, capacity management, load balancing, file server consolidation and other storage management duties without impacting the user.

Providing data location transparency without a single point of failure or data integrity risk, RainStorage operates as a layer-2 switch with file system protocol intelligence and is invisible to clients and IP-based storage devices.

With RainStorage, data mobility becomes a simple process of copying data to the target and updating the global namespace. Without RainStorage, a storage administrator must go through several additional steps, all requiring extended downtime and complicated user coordination.

Meeting the Needs of Today's Storage Strategies

RainStorage is the only solution available for transparently moving data between IP-based storage including Windows and UNIX file servers and NAS devices. RainStorage helps all of today's storage strategies ensure that users can access file-based data at any time.

In ILM, On-Demand, Adaptive Enterprise and Utility Computing strategies, RainStorage enables data to move in response to needs such as demand, business priorities, data lifecycle, storage capacity and performance. RainStorage provides the capability for staging data to newly provisioned storage and managing capacity in real time for storage.

In the same way, RainStorage supports the Storage Grid strategy by providing a non-disruptive method to seamlessly move data throughout the grid. The Storage Grid's virtualization of storage resources into a single system cannot be accomplished without RainStorage. This collection of previously unrelated storage resources must be presented to the user as a single system. The user should not need to be aware of the physical location or the movement of the data, and that is accomplished with transparent data movement.

On Demand Computing, Adaptive Enterprise and Utility Computing all have a similar objective in merging physically separated storage resources into a single, centrally managed system. Even ILM requires a unification of storage resources. Like Storage Grid, these strategies rely on transparent data movement to achieve this goal. With RainStorage's transparent data movement, data can be moved across multi-vendor platforms while still being accessed and updated by end-users.

CONCLUSION

There are a wide variety of high level approaches to drive storage innovation that have the potential to save money, maximize storage investments, and deliver uninterrupted end-user data. The good news is that implementing a transparent data movement solution moves organizations one step closer to their storage strategy while driving tangible benefits today.

RainStorage, from Rainfinity, is a groundbreaking technology that makes transparent data movement a reality. Whether you choose ILM, On-Demand Computing, Adaptive Enterprise, Storage Grid or Utility Computing, the key to your storage strategy is transparent data movement — and the key to transparent data movement is RainStorage.

ABOUT RAINFINITY

Rainfinity ensures fast, efficient, always available solutions for critical networked storage resources and security applications. Rainfinity's solutions are deployed across industries by customers such as Apple, AT&T, Bayer Corporation, Dresdner Bank, Harvard Medical School and Qantas Airways, Ltd. Strategic partners include Check Point Software Technologies, Ltd., Hewlett Packard Company, Microsoft, Network Appliance, Inc., Sun Microsystems, Inc. and Symantec Corporation. The company is headquartered in San Jose, California, with offices throughout the Americas, Europe and Asia Pacific. **More information is available at www.rainfinity.com.**

The Americas: sales@rainfinity.com

Europe (EMEA): sales-europe@rainfinity.com

TECHNOLOGY THAT ACTUALLY IS ROCKET SCIENCE

Rainfinity provides non-stop access and optimized performance of networked resources, from storage to Internet connections and security. Our patented clustering technology, RAIN (Reliable Array of Independent Nodes), was developed by a team of scientists from Caltech (California Institute of Technology) in collaboration with NASA. The same technology is the foundation of Rainfinity's family of products including RainStorage, RainConnect and RainWall.



NON-STOP ACCESS