Quantum

OPTIMIZE ARCHIVING AND CLOUD TIERING WITH INSIGHTS AND AUTOMATION WITH QUANTUM ATFS

COMMON OBSTACLES TO COPYING OR MOVING DATA

There are five main reasons why data is copied or moved from one storage system to another storage system: cost abatement through tiering or archiving, backup and recovery, disaster recovery, collaboration, and compliance. Unfortunately, many organizations use legacy storage architectures defined by old and restrictive laws of data storage. In each case, lack of insights into what data exists, where it is located, and who or what application owns it can affect an organization's productivity and efficiency.

Lack of insights stemming from the old laws of data storage...

- Precludes organizations from understanding their data and knowing how to manage old, orphaned, or rarely accessed files. If you don't know what you have, you can't make informed business decisions about whether to store or delete data that consume storage resources—or with whom it's appropriate to share that data.
- May cause significant delays accessing data and affect productivity when data placement isn't aligned with the needs of the workload or application.
- Will add complexity and only marginally improve efficiency when tiering data, preventing organizations from realizing the full potential and benefits of a tiered architecture.
- Affect the ability to comply with best practices, regulations, and partnership agreements when all relevant data is not easily identified, or proprietary data is shared with the wrong people.

5 KEY REASONS TO COPY OR MOVE DATA

1. **Cloud tiering, tiering, and archiving** are terms commonly used to describe the movement of data from production grade media to lower-cost media in order to reduce costs.
2. **Backup and recovery** creates versions of data and stores them in a secondary repository in order to protect against data loss.
3. **Disaster recovery** is the creation of a copy of data at a different location to ensure the most recent copy of data is available to restart operations in case of failure.
4. **Collaboration** requires multiple organizations to share data in order to conduct research, share content, or deliver end products.
5. **Compliance** refers to storing and safeguarding data that is required by industry regulators, operational best practices, litigation support, or compliance with contracts.

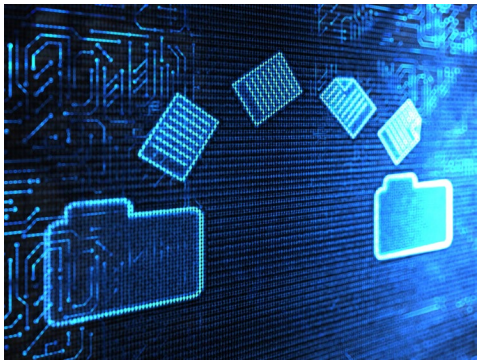

In all these scenarios, data is either copied or moved to another—secondary—system. In order to assure the desired results, it is imperative that only the relevant data is copied or moved to the designated location. This secondary system may be a target on premise, in a different data center, or in the cloud. When using owned secondary systems, the costs are front and center; when using the cloud, there may be hidden costs.

QUANTUM ATFS PLATFORM MINIMIZES HIDDEN CLOUD COSTS

Tiering can be implemented within a system or to an external system or cloud. Organizations have looked to the cloud, both private and public, for ease of collaboration and to address short- and long-term storage needs. The ATFS platform responds to this need and simplifies the movement of data within the system and across clouds while minimizing consumption-based costs as well as the "hidden" costs that can be unpredictable and change with usage of the service.

Examples of some such costs are:

- **Egress Costs:** Putting data in the public cloud is free, but egress can result in unintended costs at multiple dollars per terabyte retrieved. If data is cold and is stored only as an insurance policy, the egress costs may not be noticeable. If data needs to be retrieved periodically, identifying and pulling only data that is needed can minimize egress costs.

Every piece of data retrieved that is not relevant is money wasted. ATFS's ability to categorize and organize data ensures that only relevant data is consuming storage and compute resources.

• **Interface Management Complexity:** Each public cloud has its own interfaces, which makes it hard to move data from cloud to cloud. Interfacing with clouds via ATFS reduces these complexities. Data may be moved across clouds and premise based on metadata and tags without the need to interface with each one separately.

• **Searches:** Once data is in the cloud, private or public, it is hard to track it. What data is being stored and why can be difficult questions to answer without knowing who created the data, who owns the data, and what value it represents to the organization. Every time a search is performed, there is a cost in time and money associated with the action. ATFS has a metadata and tags repository available for real-time searches without having to access the data or put a strain on computing, human, or budget resources.

- **Manual Retrieval:** Traditional copy and move methodologies don't have the ability to manage files and objects per application needs. Once the file is moved to the cloud, it is in the cloud and if it needs to be accessed, it is manually retrieved. Each operation may result in a cost, which can result in a significant financial expense over many operations. ATFS can tag data based on how applications use it. This results in data being automatically copied or moved to where it is most efficient to keep it. In some cases, a copy of the data may remain on premise while another copy is stored in the cloud; data is delivered to the application from where it is most efficient. Machine learning and read-ahead techniques are used to minimize access times and data movement.

OPTIMIZE DATA TIERING AND ARCHIVING THROUGH INSIGHTS AND AUTOMATION

The five main use cases for moving or copying data require visibility into what data exists and to whom it belongs, control over access to data within and outside the organization in a secure manner, and the ability to retain data based on recovery objectives, regulations, or operational needs. ATFS fulfills the requirements of tiering, performance, archiving, and cloud tiering in support of cost abatement, data protection, disaster recovery, collaboration, and compliance. Here's how it works:

- All data entering ATFS is identified, tagged, and classified; data can be searched based on metadata or tags without having to crawl the file system. Information about the data can be accessed in real-time without touching data regardless of where the data is, on premise or in the cloud.
- Tasks may be automated based on a combination of metadata and tags. When an event is triggered, only data with the appropriate tags and metadata are affected, giving you pinpoint accuracy.
- A set of APIs and integrations with schedulers and applications empower workflows, in real time, to determine the placement of data. All movements reflect workflow and application needs, with no manual intervention.
- Automation eliminates human error, ensuring compliance with regulations and best practices. Automation is driven by predetermined business requirements, which are represented by tags; tags are not inherent to the file.
- ATFS supports multiple roles, each with varying degrees of access control. Data from across file systems and directories may be combined to present a virtual file system view. External viewers without access permissions and remote resources may be given read access to a virtual file system's data without letting them access any other data in the system. This facilitates collaboration across users within an organization as well as external partners.

ATFS responded to the new challenges of the cloud era and the need to store and consume data more efficiently by delivering a tiered storage platform that enables insight-driven data and storage management and automates tasks for greater accuracy, consistency, efficiency, and performance.

ST02331A-v01 Nov 2020