

WHITE PAPER

Automated Storage Tiering and Business Transaction Intelligence Delivers Improved Performance and Reduced TCO

Sponsored by: EMC Corporation

Noemi Greyzdorf

Steven Scully

April 2010

IN THIS WHITE PAPER

This white paper focuses on the business benefits and use cases of automated storage tiering. It looks at the role automated storage tiering solutions play in improving performance and reducing TCO as well as the role business transaction intelligence plays in helping IT organizations determine the optimum levels of tiering. Business Transaction intelligence adds an important dimension to tiering by determining the most appropriate allocations of storage resources on an on-going basis.

SITUATION OVERVIEW

The explosion of data being created and the performance requirements of end users put pressure on IT to be more dynamic and deliver optimum service levels without increasing costs. At the same time, data centers are evolving from the traditional physical datacenter to one that is more virtualized and, in some cases, leverage cloud infrastructures where IT is delivered as a service. This has driven organizations to look at a number of technologies that can help them achieve their goals. One of these approaches is to automatically tier applications to comply with service level requirements of the virtual infrastructure.

The concept of using storage tiering is not new; organizations have been aligning storage resources with the demands of business applications for some time. Critical applications have been given faster drives and greater redundancies while secondary and tertiary applications are stored on more cost effective, high capacity drives. This approach has enabled organizations to align business demands with the cost of the underlying infrastructure.

Now that applications have been aligned with storage, organizations have begun to look at ways to further increase the efficiency and effectiveness of their storage investment by placing data within an application on the most appropriate tier of storage. As applications go through a lifecycle so does the data created by an application. The phases of data lifecycle vary by application:

- ☒ Access patterns to data may change over time. When data is initially created, it may be accessed often and require high performance storage. As the data ages, access diminishes and the demands on performance also diminish.
- ☒ Data may experience seasonality. There are times of the year when some data is accessed more frequently than at others. Being able to migrate or tier that data to a more suitable storage resource at any given point in time creates great business benefits and storage efficiencies.
- ☒ The way data is accessed may vary based on the process being applied to it. Data might require highly randomized access initially, but over time, access to the same data becomes more sequential. Locating data on the appropriate storage tier closely aligns the capabilities of storage with how the data residing on it is being used.

Moving static data to higher capacity and more cost effective storage media will reduce an organization's costs while dynamically leveraging flash drives may increase application performance. Combining the two capabilities should deliver higher service levels while reducing storage acquisition costs. To take full advantage of dynamic storage tiering, it is optimal to have management software that monitors application performance over time and provides accurate visibility into performance requirements. Managers can leverage the information gathered to make more informed decisions regarding data placement. However, as the number of applications and their associated data continue to increase, the ability to physically manage the appropriate placement of the data becomes difficult – to the point where it is no longer possible to do manually – and automation becomes a requirement.

In the past, movement of data across storage tiers was very resource intensive, complex and had limited scalability. In order to realize the benefits of storage tiering, it is best for the movement of data across storage tiers to be automated by the storage system and non-disruptive.

- ☒ Data must be moved in a timely manner, in response to how data is being used by the application
- ☒ Data must have the ability to move to faster and slower storage media depending on need.
- ☒ Data movement and its location at any time must be transparent to the application.
- ☒ Data movement must be automatic and driven by policies, without continuous intervention of an administrator.
- ☒ A management framework must exist to track movement of data across storage resources in response to the application's requirements.

As the market seeks ways to reduce costs associated with storage, the ability to move data across storage tiers based on meeting application SLAs is becoming a prominent requirement of solution providers.

EMC FAST AND PRECISE

In response to the demands of the market, EMC has developed software that provides automatic data movement across storage tiers to align the demands of the application with storage resources. The Fully Automated Storage Tiering (FAST) feature allows storage managers to set policies on an application basis that will automatically initiate movement of data across storage tiers. FAST utilizes intelligent algorithms to continuously analyze device I/O activity and optimizes the array to make the best use of Flash, Fibre Channel and SATA Drives. FAST proactively monitors workloads at the LUN level in order to identify those that would benefit from higher-performance drives as well as identify workloads that could be relocated to higher-capacity, more cost-effective storage. Applications' data can benefit from FAST by placing data on the appropriate storage tier for a particular access pattern, protection level, or event. Financial data used for end of quarter analysis and reporting may need to reside on high performance storage during the reporting period, but is content sitting on higher capacity, less expensive storage throughout the rest of the quarter.

FAST is supported across all EMC storage platforms (Symmetrix, CLARiiON, and Celerra). FAST works by defining storage tiers that consist of drive type (SATA, FC, SSD (Flash to EMC)) and RAID levels (1, 10, 5, etc). Once storage tiers have been defined, managers can create policies that dictate how much of each storage tier is used by an application. An application policy may consist of 10% Flash, 30% FC, and 60% of SATA drives. As an application writes and reads data from storage, FAST non-disruptively migrates each volume to the most appropriate storage tier. The benefits of FAST are:

- ☑ Faster application performance (right tier at the right time)
- ☑ Lower storage costs (using higher capacity, lower cost SATA drives)
- ☑ Reduced energy consumption (data spends more time on dense, energy efficient drives and performance needs can be accomplished with fewer flash drives)

To help customers optimize the performance of business applications, Precise software has developed a transaction performance management solution which compliments EMC FAST. Precise monitors application performance and storage usage by connecting the actual business transactions to the data accessed. This includes identifying all database tables and other database objects and associated LUNs wherever they might reside. The Precise Business Storage Optimizer models this information and provides recommendations for tiering this data based on response-time impact for the specific transactions or applications. A storage manager can view this information to determine which workloads are the best candidates for tiering to faster storage like flash to improve transaction or application performance, or to slower less expensive storage to improve the TCO for the specific workload. The intelligence provided by the Precise Solution could be used to set the optimal tiering policies for applications in a FAST environment.

Precise helps managers make intelligent decisions regarding their storage resources by providing "what-if" results of tiering the data. The software can identify specific workloads from previous events and use that information to anticipate future events. Precise can identify what data should be tiered to which storage type in order to have

the most positive impact on the upcoming event. Pro-actively moving the appropriate data to the appropriate storage tier will improve the performance of a pending event like month-end processing or quarter close. The use of Precise software with EMC FAST helps organizations realize the value of their storage resources, reduce waste, and simplify the management of data movement across storage tiers. These combined capabilities provide the ability to tier the right data to the right place at the right time to provide optimal business impact. The Precise software allows storage managers to have the insight needed to determine the best workloads to tier in order to meet transaction SLAs and to reduce cost.

FUTURE OUTLOOK

The market is beginning to recover from recent downturns, but the focus on efficiency is here to stay. Leveraging appropriate storage resources for a given application and use case can significantly reduce storage and operational costs. In order to take advantage of different storage tiers available in a system, there needs to be an automatic way to move data. This movement must be based on policies that are managed by application owners and storage managers.

Users across all industries and organizations will seek ways to leverage faster storage where performance is key and denser storage where capacity is key without incurring unnecessary management costs or complexities. Dynamic, automatic storage tiering, like FAST, offers a way to take advantage of appropriate storage tiers in a simple to use and manageable way when combined with solutions like Precise, that provide the insight into the business workloads.

CHALLENGES/OPPORTUNITIES

By bringing together two solutions designed to help IT organizations make the most of their storage for business applications, EMC and Precise have created the opportunity to make automated storage tiering more intelligent and workload-focused. They also have the opportunity to further integrate their products so that in addition to the Precise software being able to discover the workload requirements, determine policies on a per application basis, and create the storage groups to best support the application, but also be able to then create and assign the policies within FAST thus eliminating an additional configuration step.

There are some challenges that EMC and Precise will have to overcome if they are to maximize the success of their partnership including:

- ☒ Crystallizing the message around the combined value of the solutions for potential customers and the EMC and Precise sales organizations and channel partners. This is a common challenge when vendors meet in the market with solutions.
- ☒ Expanding the current focus from structured data (specifically for Oracle and SQL Server today) to more broad volumes as the challenge of storing and managing the growing amount of unstructured data continues for organizations.

CONCLUSION

Organizations large and small are looking for ways to optimize their IT infrastructure environment without incurring additional costs or creating unnecessary complexity. The first phase in this process was to align application volumes with the right type of storage resource. The next phase is to align a much more granular set of data with storage resources, creating additional cost and operational efficiencies. This process requires the adoption of storage and application management technologies that can accurately report on the usage and help design policies that reflect actual application needs. The execution of these policies must be achieved without adding complexity.

The goal for organizations should be to maximize the business processes, not just improve the efficiency and performance of their storage. This is where the combination of the Precise software and the EMC FAST solution can help. Precise and EMC have developed a joint solution that automates the process of storage tiering to improve performance and reduce total cost of ownership.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2010 IDC. Reproduction without written permission is completely forbidden.