WHITE PAPER

Precise Enables Business-Aware Storage Tiering

Improve Application Performance and Control Storage Costs

By Brian Babineau

July, 2009

Table of Contents

Introduction

Making up an estimated one third of all corporate information,[1] database applications are the lifeblood of many businesses. They are critical to the knowledge workers who depend on them to automate human capital management, financial reporting, supply chain management, and other business processes. But slow application response times and other accessibility issues negatively impact overall productivity and delay the aforementioned business processes, causing immeasurable financial impact.

As database information continues to grow—ESG estimates that it will increase at roughly 25% per year[2]—IT must find ways to maintain application performance. For many, the most common performance 'fix' is buying the fastest storage possible. The resolution is usually overkill as portions of the database application infrastructure do not need such high levels of responsiveness. Recognizing this, organizations often implement tiered storage infrastructures where systems with varying disk drive performance characteristics are used—a more cost effective approach when compared to purchasing a system full of the fastest disk drives. Traditionally, freshly created, more active data is stored on a storage tier supported by the fastest drives and older, less active data is relegated on slower drives. In doing so, IT believes that it is getting the best of both worlds: improving application performance and controlling storage costs.

Current tiering methods based on the age of data or access patterns do not take the value of the information into account, which results in IT spending money to boost performance for less critical information. For example, a table with settled stock trades from a year ago may be moved to slower storage while recently completed new account applications are saved on faster disk. The problem is that the settled stock trades—while older—may still be accessed for critical end of month processing tasks, while the account applications are rarely used after they are approved. This simple example highlights the fact that organizations need to take a more business-aware approach to tiering, in which data that supports more important transactions are stored on the fastest storage while other data is moved to slower, more cost effective tiers.

Achieving a business-aware approach to storage tiering requires organizations to identify individual transactions which hit the storage within a database application environment and locate storage-specific performance bottlenecks that impact these transactions. Precise TPM for Storage is a top-down troubleshooting process that begins with a performance issue impacting business transactions. Precise correlates the transactions, objects, and applications to the physical storage devices which they access, enabling IT to drill down into the storage tier and find the root cause for an excessive IO wait. In addition, the company's Precise for EMC solution also includes a storage tiering simulator which predicts the impact a new storage tier would have on a database application environment. Precise for EMC recommends data placement to any drive—including enterprise flash drives and traditional hard disk drives (Fibre Channel, SATA, etc.)—based on real transaction access and predicts the impact of such placements on transaction performance. This is extremely useful for customers evaluating EMC systems with EFDs because an organization can see what the performance gain would be for set of transactions before actually making an investment.

Database Application Performance Challenges

Dealing with Data Growth and Retention

Despite current economic conditions, a recent ESG research study of nearly 500 organizations revealed that 36% of respondents are planning to increase spending on database applications in 2009. Another 36% said they expected database application spending to remain flat in 2009—a respectable number given that some IT budgets are being cut by 20% or more. These figures underscore the importance of database applications as organizations continue to invest in business process automation.

[1] Source: ESG Research Report, *Medium-Size Business Server & Storage Priorities*, June 2008.
[2] Source: ESG Research Report, *ESG 2008 Enterprise Storage Systems Survey*, November 2008.

Clearly, database applications will remain at the center of many business operations. In turn, these applications will continue to contribute to overall corporate data growth. When asked what IT infrastructure initiatives will impact their storage spending over the next 24 months, 35% of the respondents surveyed by ESG said database implementations/upgrades/migrations—the highest percentage of responses.[3] Also contributing to database application information growth is the amount of data organizations must retain for business reference, compliance, and electronic discovery purposes. The most common average retention period for database information is between six and ten years (see Figure 1).

FIGURE 1. AVERAGE RETENTION PERIODS FOR ARCHIVED DATABASE INFORMATION

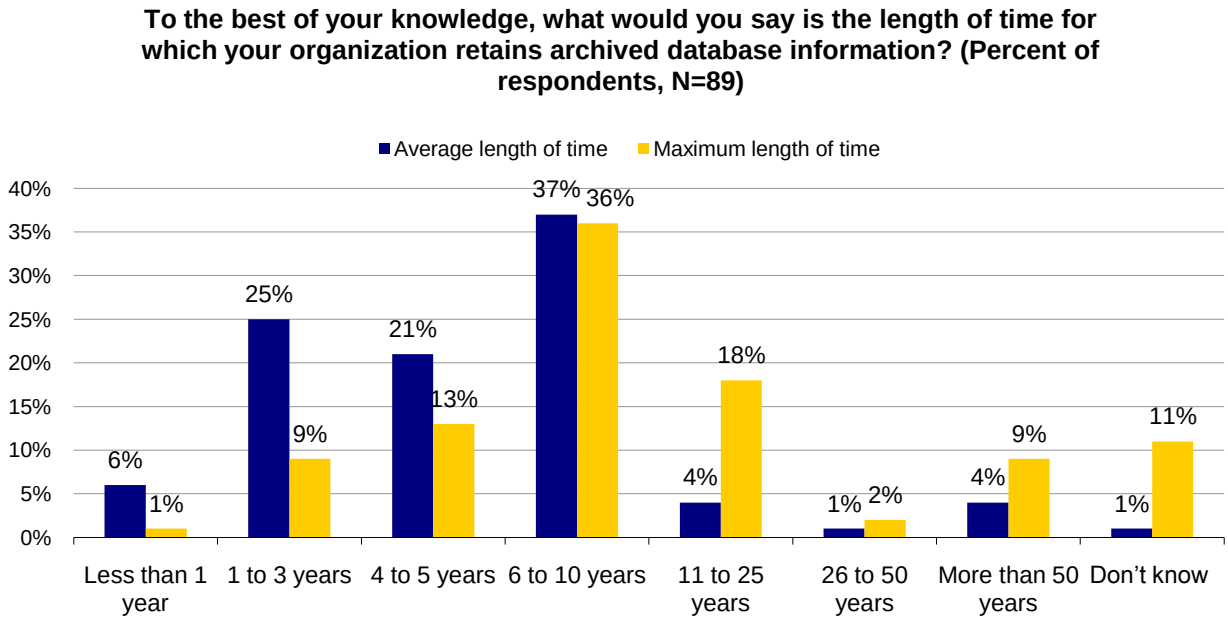

Source: ESG Research Report, 2007 Database Archiving Survey, December 2007

As more information is created and retained, database applications become bloated and response times can slow. Adequate levels of database application performance must be maintained in the face of constant data growth, especially since IT rarely has the luxury of deleting information.

Complex Architectures Add Latency

Database application architectures are usually made up of several tiers, with each layer supported by several hardware and software components. For example, one application may include a Web server layer, followed by an actual application connected to several databases via a middleware tier. Each one of these layers may run across several servers, all of which are connected to a storage system.

When an application is new, the transaction path through the architecture is typically well defined. But as the application matures, new database tables are added, queries are expanded, additional modules may be deployed, and the application may need to connect to an outside data source. For example, an organization may not currently use the Accounts Payable module within Oracle e-Business Suite. If that module is added later, it may require a change to the entire application that would allow it to connect to the inventory management systems as well as a separate data warehouse application. As changes to an application are made, the paths transactions take to complete are elongated, causing IO response times to slow.

When you combine evolving, multi-tier database applications with ongoing data growth, it is not a matter of if response times will slow; it is a matter of when. When database application performance slows, IT usually has to

[3] Source: ESG Research Report, *ESG 2008 Enterprise Storage Systems Survey*, November 2008.

abandon normal daily tasks and focus on what is going wrong. Given the architectural complexity of database applications, diagnosing a performance problem is the equivalent of finding a needle in a haystack.

It is Easy to Blame Storage

If an issue cannot be found or easily fixed, storage is often the scapegoat because application and database administrators rarely have any visibility into storage system performance. Even if there *is* some level of insight into the storage system, it is limited to a logical unit (or volume) within that system and can't be correlated to the business transactions being served.

Storage is easy to blame because usually only one system handles transaction requests from the multitude of Web, application, middleware, and database servers. That storage system may not have been sized appropriately to handle the transaction workload in an attempt to save money when making the upfront investment purchase. In some instances, a storage system is sized appropriately for one application, but then others with different workloads are added to the device. This can cause contention at the individual disk drive level as multiple applications request data from the same spindle at the same time. As a result, the storage system may not have the fastest disk drives or enough memory to keep up with all of the transaction requests.

Luckily, there are many options storage administrators can choose to boost application performance, including:

- Adding more memory (cache) to a system
- Buying newer, faster disk drives to store database data
- Altering the storage system configuration (change logical volume sizes)
- 'Short stroking' the disk drives (only store data on the fastest part of the disk drive)
- Spreading the database data across more disk drives

Given all of these options, IT has several ways to improve overall database application performance without spending time to find the *real* source of the bottleneck. Unfortunately, these fixes may be cost prohibitive and short-lived because none of the approaches take into account the performance requirements of specific transactions within an application. When the transaction profiles change (i.e., a new application module is added, a query is updated, etc.), storage enhancements quickly become obsolete and the next performance issue arises.

Time for a Change

Existing Methods are No Longer Feasible

Due to a lack of visibility when it comes to database application performance issues, storage systems are guilty until proven innocent—and this mindset is unlikely to change anytime soon. However, when the next issue arises, IT departments would greatly benefit from understanding the costs and benefits associated with some traditional storage fixes. For example, spreading database tables across drives or only storing data on the fastest part of the drives (short-stroking) causes storage utilization to plummet. Several terabytes may be required to store just a few hundred gigabytes. ESG recently spoke with the Director of Storage at a large enterprise organization that needed better storage performance for a 500 GB data warehouse application. The organization used a 10 TB system—where the database tables only stored bytes on the outside portion (the fastest part) of the drive—to meet the required response times. The other 9.5 TB were unused.

While this solved database application performance without requiring application tuning, it may no longer be suitable due to shrinking IT budgets. Organizations cannot continue to buy more and more (faster) drives only to use a portion of them—wasted capacity means wasted money. Further, the operational costs—namely, the power needed to run and cool all the underutilized capacity, as well as the data center space to store it—add up fast. What's more, these expenses do not take into account the time and effort involved to constantly move data so that it remains on the fastest drive or on the fastest part of the drive.

With limited budgets, it is unlikely that any IT spending that will increase power, data center real estate, and labor expenses will be approved—even if does mean an increase in database application response times. And once these expense, effort, and time investments have been made, there is no way to measure the realized performance improvement. If the improvement could be quantified, organizations might determine that it justified the cost. Alternatively, they could just as easily realize that the time and expense might have been spent elsewhere for an even greater impact.

Another rarely considered hidden cost of implementing additional hardware is the productivity impact brought on by disruption to the database application environment. Constantly adding, removing, or changing components means system configurations change. At best, some of the configuration changes may require a simple database or application server restart. At worst, a database application may have to be taken offline or rebooted. This worst case scenario may not even be possible for some mission critical applications. ESG research states that 30% of enterprise organizations will experience an adverse business impact—such as revenue loss—if their mission critical applications are unavailable for over an hour. [4]

The Move to Tiering

Buying the fastest available hard disk drives (HDDs) is usually first on the list of solutions to a database application performance issue. However, it may not be economically feasible for these practices to continue and, as such, many companies have started implementing tiered storage infrastructures leveraging a variety of disk drive options including offerings from a range of traditional HDDs (including Fibre Channel and ATA) which have differing performance, reliability, and cost characteristics. IT departments then store information across these tiers based on how old it is or how frequently it is accessed. For example, the most active portions of a database application, such as the index, can be stored on the fastest drives while infrequently accessed parts, such as tables holding old transactions, can be placed on denser, slower, and more importantly, less expensive drives. Similarly, data that has not been accessed in over six months may be saved on the slower drives while newly created database records are kept on the fastest drives.

Tiering to address performance issues has recently gotten more interesting as many storage administrators are salivating over a new tier of storage that can be added to boost application performance: enterprise flash drives (EFDs). EFDs can increase application response times by an estimated 10-20X when compared to the fastest HDDs. The caveat is that they are also 8-10X more expensive than HDDs. These drives are ideal for database applications because they are not hindered by seek times, allowing them to immediately handle IO transactions. When ESG asked over 300 organizations what applications would benefit most by EFDs (sometimes referred to as solid state disks, or SSDs), database/OLTP applications was the top response (see Figure 2).

Storage tiering is generating significant savings and, as a result, over one-third of organizations are going to maintain or increase spending on tiered storage projects in 2009.[5] To maximize these investments in database application environments, companies need to reevaluate the criteria they use to determine how to store data across various tiers. Current methods of data tiering—based on data access patterns or age of information—help control storage costs, but may do so at the expense of the performance of critical application transactions. For example, an old database table storing critical data may become active again due to business intelligence initiatives. With that table being stored on slower disk drives, business intelligence initiatives may grind to a halt as poor application response times impede productivity.

[4] Source: ESG Research Report, *Data Protection Market Trends*, January 2008.
[5] Source: ESG Research Report, *2009 Data Center Spending Intentions Survey*, March 2009.

FIGURE 2. TOP FIVE BUSINESS APPLICATIONS REQUIRING THE PERFORMANCE OF FLASH-BASED DISK DRIVES

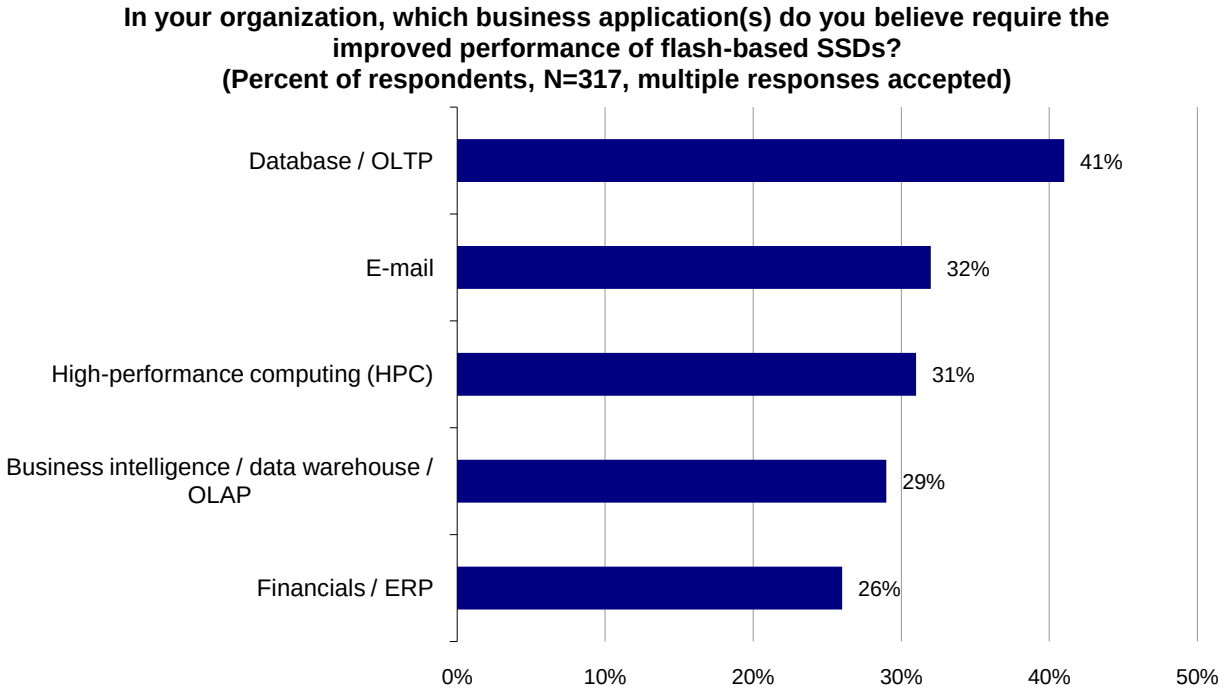

Source: ESG Research Report, ESG 2008 Enterprise Storage Systems Survey, November 2008

Business-Aware Storage Tiering

The Next Phase

Organizations need to tier storage based on the business importance of the information being accessed. Each transaction within a database application environment offers some level of value to the business. By identifying which transactions are the most (or least) important to the business and mapping those transactions to the information in the storage subsystem, organizations gain the insight they need to promote important data to faster storage to optimize for performance. Just as importantly, they gain the insight needed to demote less important data to slower storage in order to optimize for cost.

Questions Arise

Knowing the business importance of transactions serves to heighten awareness around the performance and cost of storage supporting an organization's most critical applications. This heightened awareness brings with it heightened expectations, as well.

In order to live up to these expectations, organizations need a way to predict the application performance impact that a change in storage tiers will have: If IT deploys a faster disk drive, will the performance benefit justify the cost? If information is moved from faster drives to slower ones, will the modification be noticeable? As an example, buying several terabytes of EFDs and receiving only a marginal performance boost is a costly—and inefficient—proposition. Likewise, storage administrators will never hear the end of it (and could put the business at risk) if they put a portion of a database on ATA disk drives and application performance immediately slows. In order to avoid these situations, IT needs some level of insight as to what the performance impact will be if they move data between different storage tiers.

Precise Simplifies Next Gen Storage Tiering

A Transaction-Oriented Approach to Performance

With visibility into pre-packaged database applications (SAP, Oracle, PeopleSoft, Siebel, etc.), databases (Oracle, SQL Server, etc.), middleware development platforms (J2EE, .NET, etc.), and storage systems (EMC, HDS, IBM, etc.), Precise is able to identify and trace a transaction to show a correlated view of a business transaction breaks down into IT tiers (application, middleware, storage, etc). At any point in these tiers, there may be a bottleneck which slows the completion of that transaction down. This insight enables quick pinpointing of application performance problems, allowing customers to take more targeted measures to address them.

When a performance issue arises in the storage portion of the database application infrastructure, Precise can identify the transactions that are impacted and what may be causing the poor response times. It can also determine bottlenecks that result from poor database storage layouts which can lead to 'hot devices,' slow disk drive performance, or inadequate storage system resource (memory) allocations.

Precise also makes recommendations as to how customers should go about resolving application performance issues, including altering the underlying storage infrastructure. Recommendations include utilizing faster storage to support a given transaction or identifying transactions currently executed on high performance disk drives that could be served by less robust infrastructure. The focal point of all these recommendations is the transactions within the application—not tables, indexes, or DBF files. Customers can prioritize storage tiering and data layout based on transaction importance rather than age or frequency of access.

Precise for EMC: Try Before You Buy

EMC storage customers will have access to exclusive capabilities not available in the standard Precise for Storage solution. Precise for EMC includes a storage tiering simulator capability which takes the guesswork out of introducing a new tier of storage. The storage tiering simulator analyzes transaction response times of the current storage infrastructure and accurately predicts what would happen if certain disk drives were replaced with faster (or slower) ones, including EFDs. Precise for EMC provides customers with an estimate of what the performance gain (moving to faster storage) or decrease (moving to slower storage) will be. EMC storage customers can use these percentage estimates to see if an investment in faster disk drives (EFDs) will improve performance of critical business transactions enough to justify the cost.

EMC storage customers can also test the impact of a faster drive and then see how performance scales as more are added to the system. The solution also allows examination of changes in transaction performance at the more granular file or object level as a result of using a faster drive. For example, if a complex SQL query executes against a table on mechanical disk, IT can simulate the query response time if part of the table were moved to an EFD.

Precise for EMC also alleviates concerns when moving certain database application transactions to slower disk drives. When storage administrators tell their database counterparts about such a move, the database administrators may not react favorably because, on the surface, slower drives equals slower performance. Storage administrators can respond by showing the database administrators a report from the simulator that application performance will not decrease if slower drives are used minimizing any internal pushback.

All of the simulations can be executed without disrupting the database application environment. As a result, EMC customers can test a series of configurations and database layout schemes to help determine what storage tiers will deliver the best performance for the investment. Precise for EMC provides the 'benefit' aspect of the cost/benefit analysis customers should run when making an investment in faster drives as well as help them remove transactions that are unnecessarily consuming fast, expensive storage capacity.

The Benefits

Precise correlates real-time performance metrics for all elements of a database application infrastructure. This information helps customers identify and resolve application performance issues from a transaction perspective. When part of the resolution involves a tiered storage infrastructure, customers can leverage Precise to make sure the right data is stored on the most appropriate tier. In turn, this enables organizations to optimize application response times while controlling storage costs.

Precise enables more educated tiered storage decisions as companies store data based on the value of the transaction rather than the age of the information. As EMC storage customers make those decisions, Precise for EMC provides a deeper level of visibility to predict what the performance impact will be by introducing faster storage or moving certain transactions to slower devices. As a result, EMC storage customers can maximize their investments in the latest technology in the marketplace, such as EFDs, while proving that slower disk drives do not always equate to poor application response times.

Conclusion

As long as database applications automate critical business processes, IT will need to constantly watch response times to avoid disrupting business operations. At some point, the combination of data growth, data retention, and complex database application architectures will slow transaction throughput, forcing IT to find a way to boost performance. And, more than likely, many of the tuning exercises will be executed within the storage environment as there are plenty of fixes that usually work within this part of the infrastructure.

IT departments have attempted to address database application performance issues by spreading a small amount of data across a significant amount of storage capacity and by buying the fastest drives available in the market. These alternatives are no longer sustainable as the fixes only last until the next application configuration change. More importantly, they are no longer economically feasible—too much capacity, and therefore too much money, is wasted while the associated operating costs, including labor and power, spiral out of control. To reign in these expenses, many IT departments have turned to tiered storage implementations using fast disk drives for only the newest data and inexpensive, slower drives to handle older information. The problem is that it is difficult to determine the success of today's tiering efforts as they do not take into account the business transaction performance requirements of the data.

With Precise, organizations can base tiering storage based on key application transactions ensuring the most performance sensitive ones use the fastest storage resources available. For EMC environments, if current resources are not enough, Precise enables customers to predict what the impact will be of implementing faster devices before they make a purchase. This particular capability is extremely important right now as organizations evaluate whether or not to make investments in EFDs for their database application environments.

In a recent ESG survey, 62% of organizations stated that a reduction in operating costs will be the best way for IT to justify technology investments over the next 24 months. Using Precise, customers can cut operating expenses by resolving application performance problems more quickly, reducing the impact to the critical business process. By leveraging Precise's detailed storage transaction data, customers can also make more targeted capital investments by implementing the fastest storage resources where the most improvement can be realized. These two benefits should provide more than enough savings for IT organizations to consider a purchase before the next database application performance issue arises.

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

www.enterprisestrategygroup.com