



Noemi Greyzdorf
Research Manager, Storage Software

Dynamic Storage Tiering

October 2009

IDC expects data to grow at a CAGR of more than 50% through 2012. Among the drivers affecting this growth are the increased rate of digitizing records and content; collaboration and communication across organizations within an enterprise and with partners, suppliers, and customers; regulatory compliance, governance, litigation support, and best practices that are driving longer retention of data; and data protection and disaster recovery that are creating copies of data.

As a result, storage managers are trying to match the demands of the application with storage resources. One size fits all no longer is good enough in the majority of datacenters. Business units demand the characteristics that enable growth, productivity, and competitiveness. Additionally, storage managers have been asked to reduce the overall cost of storage and its impact on the datacenter.

To achieve the goal of better aligning storage resources with the demands of the application, storage managers seek a tier-based strategy in which appropriate storage resources can be mapped to the application. There are two situations in which storage tiers can be applied, and they are not necessarily exclusive of each other: aligning the demands of an application with an appropriate storage tier, and taking data that has been created by an application and stored in one location and moving it to another tier of storage. One strategy for accommodating both of these situations is implementing dynamic storage tiering.

The following questions were posed by Avere Systems to Noemi Greyzdorf, research manager of Storage Software at IDC, on behalf of Avere Systems' customers.

Q. Many vendors claim to have dynamic storage tiering solutions. What questions should IT buyers ask vendors regarding their specific dynamic storage tiering solutions to help differentiate them and decide which type is best for their environment?

A. Dynamic storage tiering refers to the automated movement of data across various classes of storage based on performance demands on that data or data life cycle based on predefined criteria. In both cases, the data is moved automatically and seamlessly to the application or end user. Not all data is created equal, and the value of data to an organization may not change over time even if access to that data does decline or even cease.

Dynamic tiering can be performed at the file system layer or at the block layer; this implies that the smallest segment of data that can be moved across storage tiers is a block. In either case, there are clear, predefined parameters that drive automatic migration of data across tiers with no manual intervention required except for at the time of configuration. The migration itself can occur either on a continuous basis or at scheduled intervals. Dynamic

storage tiering is not a one-way movement of data, from faster to slower; rather, it is the dynamic placement of data on storage tiers most appropriate for that data segment at that point in time. This may mean moving data from slower to faster or from slower to slowest.

When you are evaluating dynamic tiering technologies, adhere to the following recommendations:

- Ascertain whether migration is driven by policies or involves manual intervention. You want as little manual intervention as possible.
- Ascertain whether tiering of data is transparent to the application or end user. Users and applications shouldn't have to change how they use storage as a result of tiering.
- Determine whether tiering occurs within a storage system only or whether data can be migrated to another storage system; this may have significant implications on ancillary services such as data protection and archiving.
- Determine whether the administrator must deploy and manage a third-party application to achieve desired storage tiering results or whether this functionality is inherent in the product. Introducing a third-party product may create unnecessary management overhead that results in added costs.

Q. What are the benefits of dynamic storage tiering?

- A. Dynamic storage tiering primarily delivers economic benefits to its users. The ability to automate movement of data to more appropriate storage tiers reduces an organization's investment in expensive storage resources without increasing administrative overhead.

Dynamic tiering also allows administrators to place data on media that is most appropriate to how data is being accessed; if demand for some data increases, the data can be moved to a faster tier, but once the data becomes less active or the demands on the data are not as great, it can be moved to higher-density, lower-performing media. Most organizations recognize that over 50% of all data stored is considered inactive or static. Typically, this data resides on the same storage media as when it was first created, thus consuming expensive storage, datacenter footprint, and power. Moving this data to a more appropriate tier of storage dynamically can reduce the operational expense of power and facilities as well as decrease the capital investment required to host data. This move is typically done without any impact on the application or end user.

Q. When should an organization consider adding dynamic storage tiering to its environment?

- A. Anytime is a good time to add dynamic tiering — if available — to your environment. Dynamic tiering automates the movement of data across storage tiers, aligning application and end-user demands with storage characteristics while reducing operational expenses and complexity. Those environments with significantly varying application demands would benefit from dynamic tiering the fastest. Because not all vendors offer dynamic tiering, it is important to evaluate its impact against the cost of either switching providers or acquiring standalone technology. If an environment is being designed from scratch, then dynamic tiering should be included as a requirement of the architecture.

Claims of dynamic tiering may exist in greater number than actual deployments, so make sure to do your due diligence in understanding how dynamic tiering is deployed by a given vendor and its impact on ancillary operations such as data protection, replication, archiving, and business continuity.

Q. What criteria should IT buyers use to select a dynamic storage tiering vendor?

A. There are two ways to implement dynamic tiering: in block storage or in file-based storage. IT buyers should use the following criteria when selecting a dynamic storage tiering vendor:

- Determine where it makes most sense to deploy dynamic tiering, block or file, and which option will give you the most efficiency and the greatest performance flexibility at the lowest price point.
- Consider how tiering is executed. Dynamic implies intelligent and automated. There could be two ways to tier storage: one based on predefined criteria and one based on real-time performance demands. Depending on your application, one or both of the approaches could be utilized.
- Understand the implications of dynamic tiering on applications and end users. It is important for dynamic tiering to be as transparent as possible to the applications and end users as well as to ancillary services such as data protection, archiving, capacity optimization, snapshots, and replication. Solving a problem in one place by creating a problem in another place is not very efficient or effective.
- Know the storage tiers that are available as options, such as RAID levels, SATA drives, MAID, tape, SSD or RAM, and flash. The various options can help determine whether the tiering will deliver not just greater storage efficiency but also performance levels necessary to power some of the applications adequately.
- Perform a realistic cost analysis that weighs the technology's acquisition and implementation costs, ongoing maintenance costs, and operational costs (such as power and cooling, facilities footprint, and administrative time) against the benefits. You should come out on top, but it is important to know for sure.

ABOUT THIS ANALYST

Noemi Greyzdorf is a research manager with IDC's Storage Software. Ms. Greyzdorf's primary focus is to analyze and track the file systems and storage infrastructure management software markets. The key objective of her research is to provide insights to clients in the areas of market direction, end-user demands, positioning, messaging, and go-to-market strategies.

ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

COPYRIGHT AND RESTRICTIONS

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the GMS information line at 508-988-7610 or gms@idc.com. Translation and/or localization of this document requires an additional license from IDC.

For more information on IDC, visit www.idc.com. For more information on IDC GMS, visit www.idc.com/gms.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com