

StorTrends Information Lifecycle Management

StorTrends Information Lifecycle Management (ILM) offers a fully automated tiered storage with greatly reduced storage costs and administrative time.



© Copyright 1998-2008 American Megatrends, Inc.
All rights reserved.
American Megatrends, Inc.
5555 Oakbrook Parkway, Building 200
Norcross, GA 30093

TRADEMARK AND COPYRIGHT ACKNOWLEDGMENTS

This publication contains proprietary information that is protected by copyright. No part of this publication can be reproduced, transcribed, stored in a retrieval system, translated into any language or computer language, or transmitted in any form whatsoever without the prior written consent of the publisher, American Megatrends, Inc. Trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. American Megatrends, Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

FOR ADDITIONAL INFORMATION

Call American Megatrends at 1-800-246-8600 for additional information. You can also visit us online at ami.com.

LIMITATIONS OF LIABILITY

In no event shall American Megatrends be held liable for any loss, expenses, or damages of any kind whatsoever, whether direct, indirect, incidental, or consequential, arising from the design or use of this product or the support materials provided with the product.

LIMITED WARRANTY

No warranties are made, either express or implied, with regard to the contents of this work, its merchantability, or fitness for a particular use. American Megatrends assumes no responsibility for errors and omissions or for the uses made of the material contained herein or reader decisions based on such use.

DISCLAIMER:

Although efforts have been made to assure the accuracy of the information contained here, American Megatrends expressly disclaims liability for any error in this information, and for damages, whether direct, indirect, special, exemplary, consequential or otherwise, that may result from such error, including but not limited to the loss of profits resulting from the use or misuse of the information contained herein (even if American Megatrends has been advised of the possibility of such damages). Any questions or comments regarding this document or its contents should be addressed to American Megatrends at the address shown on the back cover of this document. American Megatrends provides this publication "as is" without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability or fitness for a specific purpose. Some states do not allow disclaimer of express or implied warranties or the limitation or exclusion of liability for indirect, special, exemplary, incidental or consequential damages in certain transactions; therefore, this statement may not apply to you. Also, you may have other rights that vary from jurisdiction to jurisdiction. This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. American Megatrends may make improvements and/or revisions in the product(s) and/or the program(s) described in this publication at any time.

Executive Summary

Today, data centers are facing an indomitable challenge of provisioning colossal capacity of storage space at an affordable price yet meeting ever-increasing performance demands. The biggest threats that data centers and the storage industry are facing today are power consumption, housing space and environmental concerns.

According to latest research, power consumption in data centers constitutes a large share of consumed power the world over. Therefore, some technology innovations are required to ensure that these challenges are met in good bearing. The noteworthy leap that the storage industry is forced to take in this regard is Storage Tiering. Here, the capacity to be provisioned is divided into separate pools of storage space with various cost/performance attributes. At the top resides the Tier 1 pool, which is most expensive and highest performing. The bottom tier is occupied by more cost-effective storage arrays.

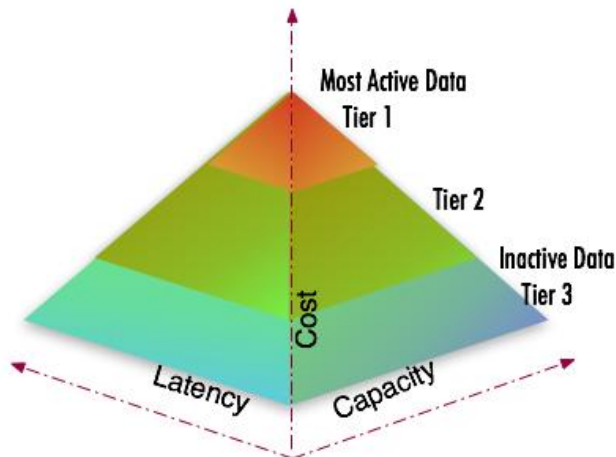


Figure 1: Data organized into different storage tiers

The next challenge is to devise a sophisticated software layer that intelligently places data into the different tiers according to their value. This concept is variously known as data classification or ILM. This ensures that more valuable data is automatically placed in the higher tiers, and relatively less used data is relegated to lower tiers. This, therefore, when properly implemented, provides the most optimal cost performance architecture.

How is this done?

In StorTrends, storage is not merely seen as a container of data. Another important dimension of intelligence is appended to every block, transitioning blocks of data into blocks of information.

Data + Intelligence = Information

This intelligence associated with every block of data, forms very vital metadata, which automatically tracks the access patterns to these blocks. Therefore, data is first classified; then moved at the block-level from tier to tier, based on frequency of access. It is important to note that data no longer needs to be manually classified and transported. The distinctive highlight of StorTrends ILM is that it is neither time-consuming nor administrator-intensive, especially when compared to its more conformist counterparts. Here, the data classification is done intelligently in the background. The entire transport operation is transparent.



Figure 2: ILM done at the block level adds intelligence to the storage system

It has been statistically seen that up to 80% of data may become inactive over a period of time (usually within months). Consequently, ILM will “demote” these inactive data blocks to the lower tier SATA drives, thus optimizing space and performance. Conversely, during certain bursts of activity, when stale data gets accessed, StorTrends ILM automatically “promotes” these blocks to Tier 1.

Tiered Storage and ILM go hand in hand. Tiering provides a mechanism of pooling storage space into various cost/performance categories. Once this is done, there has to be a mechanism in place that would determine which block of data resides where, according to the value of the data. In the absence of ILM, tedious and expensive data classification strategies have to be implemented, that may also require painstaking administrative intervention. StorTrends, with its elegant ILM implementation, does this job transparently and behind the scenes. The major tasks performed by ILM are promotion and demotion of blocks between the tiers.

In StorTrends, extensive statistical information is maintained for every block of data and thereby it can intelligently distribute large volumes across multiple tiers, ensuring that active portions of data reside in Tier 1 while un-accessed data is migrated to lower tiers. It may be worthwhile to mention here that some of the competitive storage servers that implement tiered storage do so only at volume levels and not at block levels, which is less efficient.

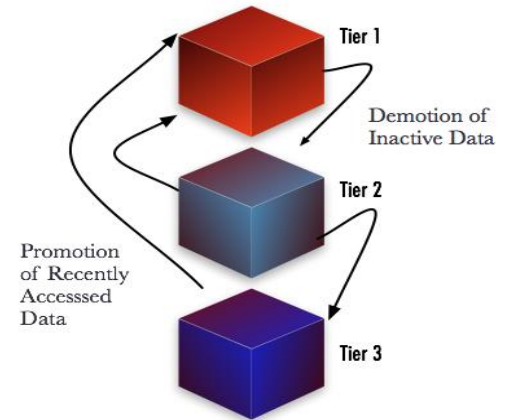


Figure 3: Promotion and demotion of data between different tiers

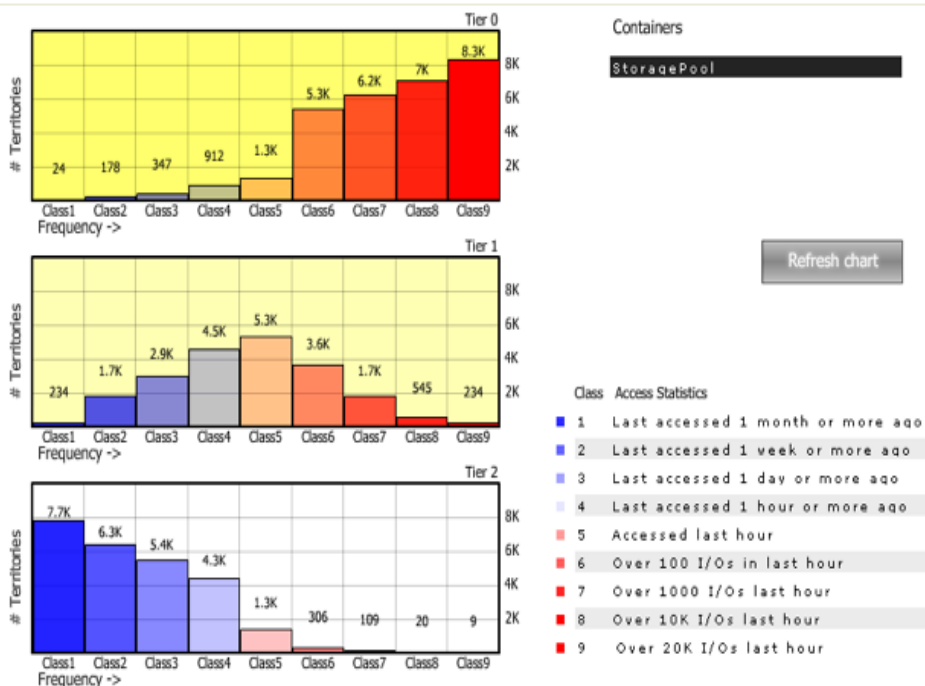


Figure 4: The StorTrends servers automatically rearrange data to match their value with the cost of the storage that they are stored in

The built in workflow management engine optimally schedules these background operations so as not to hurt the high priority client I/Os. Using the built-in workflow management engine, which is a part of its SRM offering, StorTrends optimally schedules the background data placement activities. Here, using carefully designed task prediction algorithms, the Background ILM tasks are scheduled to kick in at times when the storage server is reasonably lightly loaded.

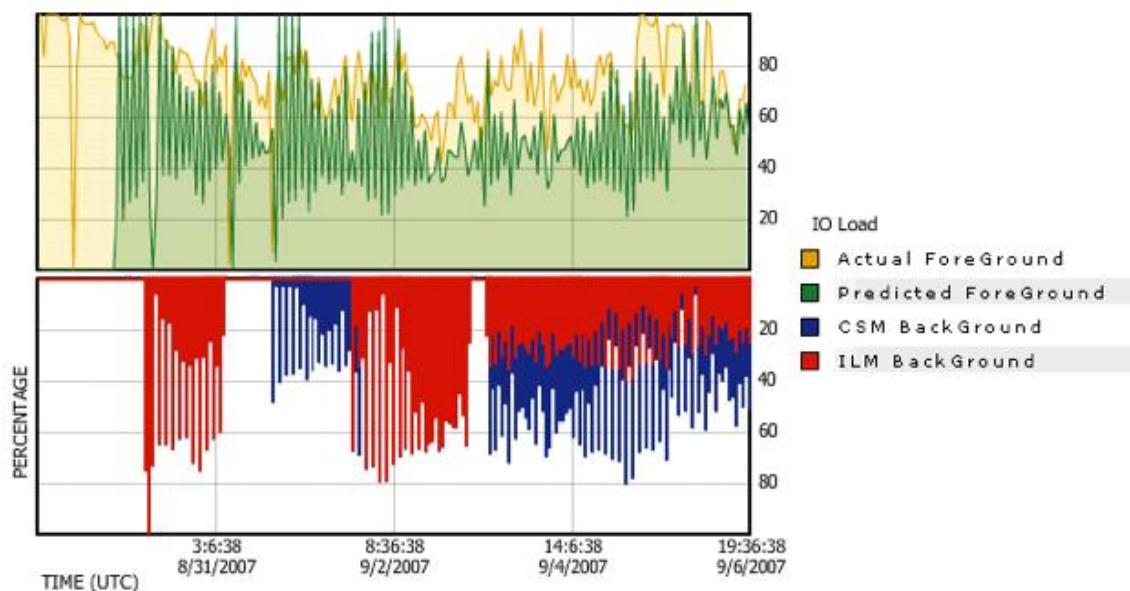


Figure 5: Workflow management predicts periods of relative inactivity and schedules data migration tasks during those periods

StorTrends Tiering and ILM: A deeper look

Today, in the storage world, the appetite for capacity is a train that has no stops. Obviously, a storage server has to be designed with capacity in mind and an eye towards scaling up and scaling out. The challenge here is to simultaneously contain the costs so that the TCO (total cost of ownership) does not go through the roof.

The other primary task of a storage server is to cater this capacity at an acceptable performance criterion. Normally, when consolidated storage servers are deployed for applications like Microsoft Exchange or SQL databases, the performance demand is very high. Latency becomes a key problem that has to be contained. Therefore, in conventional storage servers under such demanding situations, the only way is to use high performance/high cost SAS or FC drives to furnish the entire storage capacity.

Let us consider the design of a consolidated storage server for a heavy duty MS Exchange environment requiring 20 TB of capacity. As mentioned above, in conventional storage servers, the performance requirements will force the deployment of SAS or FC drives. Employing the largest SAS drive (300 GB) available in the market today, this will require approximately 67 drives, if configured in RAID 0. In order to provide resilience against single drive failure, some sort of RAID with redundancy has to be utilized. In transactional servers, there is a prevalent traffic of small block random I/Os, and RAID 5 (and its derivatives) are ill-known for their small random-write performance. Therefore, in a performance-tuned environment, a RAID 10 may be called for. This would require 134 drives. Now it can be clearly seen that the cost of the drives, power requirement, and space footprint will become prohibitive.

Power Requirement: Assuming every drive requires 15W of power, this will require approximately 2kW of electricity to keep the drives running. A rule of thumb for the data centers is that it requires an equal amount of power for cooling. So the net power requirement is 4kW.

Energy Requirement: This needs about 35,000 kWh per year.

Annual Utility Bill: Assuming a moderate \$ 0.10 / kWh, the total annual electrical expenditure just for this storage server would be around \$3500.00.

According to a recent quote from Gartner, 50 Percent of Data Centers Will Have Insufficient Power and Cooling Capacity by 2008.

The StorTrends Advantage – A Green Data Center

From above, we see that the initial cost of hardware acquisition, as well as the recurring maintenance expenditure, becomes an almost insurmountable barrier.

Statistically, it has been observed that Storage Servers that employ thin-provisioning cut down the capacity requirement to almost 25% compared to legacy servers that lack this feature. So, a thin-provisioned StorTrends server may do well in the above scenario with just 5 TB of provisioned capacity. Further, this capacity can be segregated into 2 tiers, one using RAID 10 SAS drives (Tier 1), and the other using RAID 5 SATA drives (Tier 2). The logistics used here would be to have one-fourth of this capacity designated as fast (Tier 1 Storage) and the balance in a more cost-efficient Tier 2 category. The partitioning has been chosen in concordance with the fact that about 75% of dated data usually becomes inactive or dormant. Assuming RAID 10 SAS drives (146 GB), this would require about 14 drives in Tier 1. For Tier 2, we can safely use a RAID 5 array, as performance becomes a moot point for this tier, which holds inactive data. Assuming 1 TB SATA drives, this will require 5 drives. Let us use 1 extra drive as a hot-spare for this array, for a total of 6 SATA drives. It may be worthwhile to mention here that SATA drives have now achieved the enterprise class tag (ES2 series from Seagate), where they have comparable MTBF (1.2 million hours) to SAS or FC drives (1.4 million hours). In StorTrends, we may also choose to implement this secondary tier using RAID 6, which will impart dual disk redundancy by just adding one more drive.

Power Requirement: Assuming every drive requires 15 W of power, approximately 600 W of power will be needed (inclusive of cooling).

Energy Requirement: This would require about 5,000 kWh per year resulting in an annual Utility Bill of \$500.00 (which is much cheaper than the corresponding figure of \$3500.00 for conventional storage arrays).

An Eco-friendly Data Center

One rather intangible, yet vital, concern in such a data center is environmental awareness. Today, global warming and pollution are major hazards that cannot be ignored. There are both regulatory as well as financial incentives to reducing carbon dioxide emissions, which often result in a direct cost saving due to increased carbon credits.




Figure 6: Carbon dioxide emissions of traditional storage servers versus StorTrends servers

At a stipulated 1.55 lbs. of CO₂ emitted/ kWh of electricity consumed, the CO₂ produced here would be approximately 3.8 tons/annum. The corresponding figure in a storage server scenario that does not implement tiering and ILM, would be 27 tons/annum.

The Salient Features of StorTrends ILM

Vast enhancement in performance: *Inactive data is automatically moved from SAS drives to cheaper, high capacity SATA drives. Alternatively, data can also be migrated from RAID 10 tiers to RAID 5 tiers.*

Transparent, Bi-directional data flow: *Data can be promoted or demoted from tier to tier according to access patterns. This is done at fine-grain block levels rather than at volume levels. This is an automatic background operation, and therefore does not necessitate any additional data classification software or manual administrative intervention.*

Considerable Cost Advantage: *Inactive data is automatically stored on cheaper, high capacity drives, thereby increasing available storage space for a reduced overall cost.*

Reduced Time Consumption: *Manual data classification and data movement are archaic technologies. In contrast, StorTrends ILM presents a fully automated tiered storage processing.*

Reduced Business Downtime (with snapshots) and lower Disk Expenditures: *Frequent snapshots are taken and stored on lower cost, high capacity SATA drives, thereby increasing recovery points. Consequently, operational efficiency is amplified. Note that since the snapshots are only used during recovery or backup, they automatically get relegated to cheaper tiers.*

Efficient Partitioning of Tiers: *Tiers are defined using 3 elements – disk drive type, RAID level, and rotational speed.*

Mixing of Tiers per Enclosure: *A single JBOD can hold multiple tiers. These tiers can be realized using different RAID levels, since it is not advisable to mix SAS and SATA drives in the same enclosure due to Rotational Vibrational Interference (RVI).*

Efficient Streamlining of background data migrations: *Powered by an elegant workflow management engine, StorTrends schedules the background demotions to happen at a time when the storage server is least busy – thus ensuring minimal hindrance to the high priority client I/Os.*

Green Data Centers: *With Tiered Storage and ILM, StorTrends enables data centers to have storage with much fewer disks causing low footprint, reduced electricity costs, and reduced CO₂ emissions-leading to an eco-friendly data center.*

Conclusion

Storage tiering in enterprise-class storage is becoming a highly desirable feature today. It is only a matter of time before the environmental and performance benchmarks of a tiered system become critical parameters on which decisions of storage system procurement will be based. In an environmentally conscious society, and in a storage market where the range of capacities, performance and costs offered by drive vendors is bewildering in its variety, tiering is no more a luxury but a necessity.

StorTrends servers offer a greater cost advantage and performance when compared to competitors. What sets StorTrends ILM apart, is its automated tiered storage processing, which leads to highly reduced administrative overhead. StorTrends' distinctive workflow management engine effortlessly demotes and promotes data as a transparent, background activity. On a final note, it is important to realize that with its Tiered Storage and ILM, StorTrends enables data centers to reduce footprint, electricity costs, and CO₂ emissions, for the creation of a greener and more eco-friendly data center.

Why AMI?

AMI offers a wide array of disaster recovery and high availability solutions for your business needs. We provide services that range from storage needs analysis to the design and implementation of a custom disaster recovery solution. We can help your business plan for when things are at their worst, while reducing the cost and complexity of your storage environment. For more information on AMI StorTrends solutions, visit www.StorTrends.com, email to sales@ami.com, or call (800) U.Buy.AMI.