

SNIA Data Management Forum



ILM and Tiered Storage

January, 2006

By: Michael Peterson,
Program Director, SNIA Data Management Forum



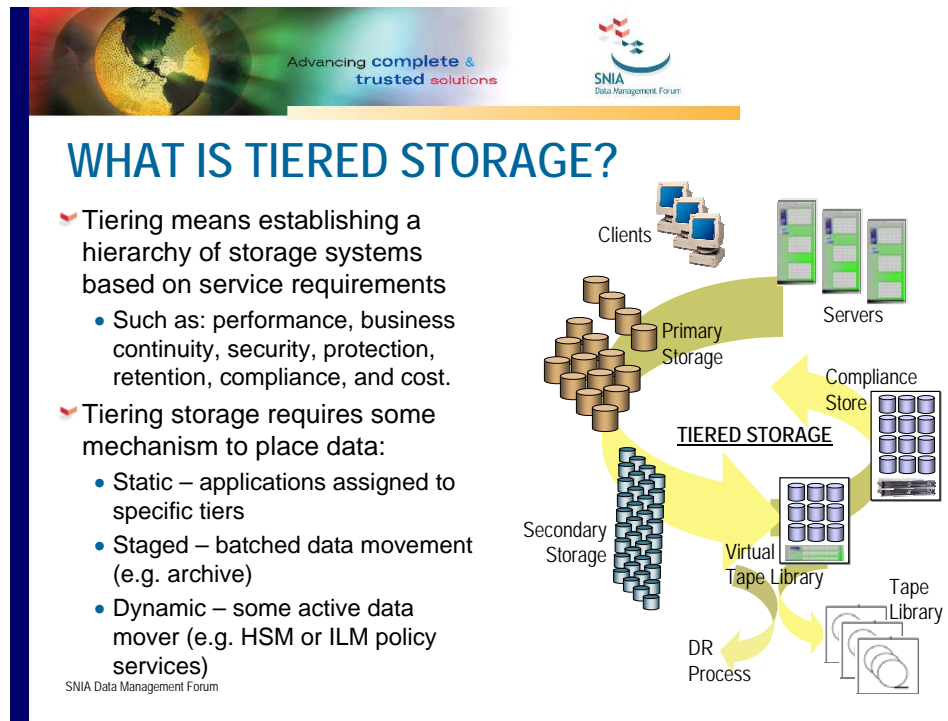
ILM and Tiered Storage

Tiering of storage is an important new strategy for Information Technology (I.T.) services. Yet, amid the many discussions of the value of tiered storage from a business perspective or its usefulness for I.T. as a preferred cost reduction method, we constantly hear vendors selling tiering as Information Lifecycle Management (ILM) as though this is all there is to ILM. Not so! Tiering storage is an important component of an ILM practice, but not all.

Tiered Storage, a Primer

Tiered storage means to establish a hierarchy of storage systems based on service requirements and cost. Why leave inactive or expired data on your expensive primary enterprise arrays? What if there were a lower cost alternative? Well, there is, tiering. The cost advantage tiered storage offers to I.T. is immense and proven. Most enterprise implementations see a return on investment in under 6 months with reductions in purchase and operating costs in the millions-of-dollars (USD) per year¹.

Adoption of tiered storage as a datacenter architecture is underway. This transformation is fueled by the explosion of low-cost, high capacity Serial Advanced Technology Attachment (SATA) disk array systems. These capacity oriented systems have successfully found a home in the datacenter within a tiered storage architecture for secondary storage, backup, .online archive, and compliance stores.



Now, add the intelligence of ILM-based practices. First identify the value and requirements (information classification), then set policies and service level requirements that map to your storage resources (data classification), and then implement tiering practices. By applying policies

¹ Source: Strategic Research Corp. 2005

defining information, data, and security service requirements such as retention, placement, protection, business continuity, archive, and compliance policies to information, based on its value or business requirements, the practices can be automated and operations costs reduced. This is also the area where the SNIA is working on standards for ILM by extending the current SMI-S (Storage Management Initiative Specification) standard. Once complete, SMI-S for data, information, and security services will instrument the automation of ILM practices.

Definitions (according to SNIA-DMF)

Tiered Storage: - Tiering means establishing a hierarchy of storage systems based on service requirements such as performance, business continuity, security, protection, retention, compliance, and cost. A company may implement a number of tiers such as primary disk, secondary disk, backup disk, archive disk, compliance disk, tape backup, tape archive, optical archive, etc.

Information Lifecycle Management, ILM: - ILM is described as a '*standards-based, business driven management practice*' or more simply as '*information-based management.*' ILM is not a specific service such as archive or tiering; rather it is about establishing business requirements for information and then via classification methods, operating to standard practices and service levels. Using this approach, we talk about a service as 'operating to ILM-based practices' not about the service "as ILM".

The key differences between ILM and Tiered Storage

Once again, tiering is not ILM. It is common to see the two confused. Tiered storage is how we arrange the repositories. Calling tiering 'ILM', is analogous to saying backup is ILM or archive is ILM. In a tiering context, ILM is how we manage the repositories based on the requirements over time. ILM aligns the business requirements and business processes with the service level requirements in an automated fashion, from data creation to data deletion. Whereas, tiered storage should be viewed as one of many components in the implementation of an ILM-based practice.

Applications Support Tiered Storage

Many applications have the ability to operate to ILM-based practices, including tiering. You can start small and easy. Begin with a practice like backup by implementing a backup array or virtual tape system connected to your tape automation and let the backup media manager and virtual tape system automate the tiers. Or, implement an integrated email archive system which integrates multiple tiers for primary storage and archive or compliance stores. Then, when you get more ambitious, free up some of that space on primary storage by deleting the expired data or allow the database to archive it to secondary storage. You'll find that the leading databases and some key add-on tools instrument and support tiering very nicely. Similarly, many ECM, Enterprise Content Management, products do the same.

Want to do more? Start using storage resource management, SRM, tools to profile your unstructured information. Find the active, inactive, reference, and expired data and then set up classes and operate to service requirements based on the value of the information to the business. Move that inactive or expired information off primary storage.

Yes, tiering storage requires some mechanism to place data and there are three basic methods:

Static – applications assign information to specific tiers

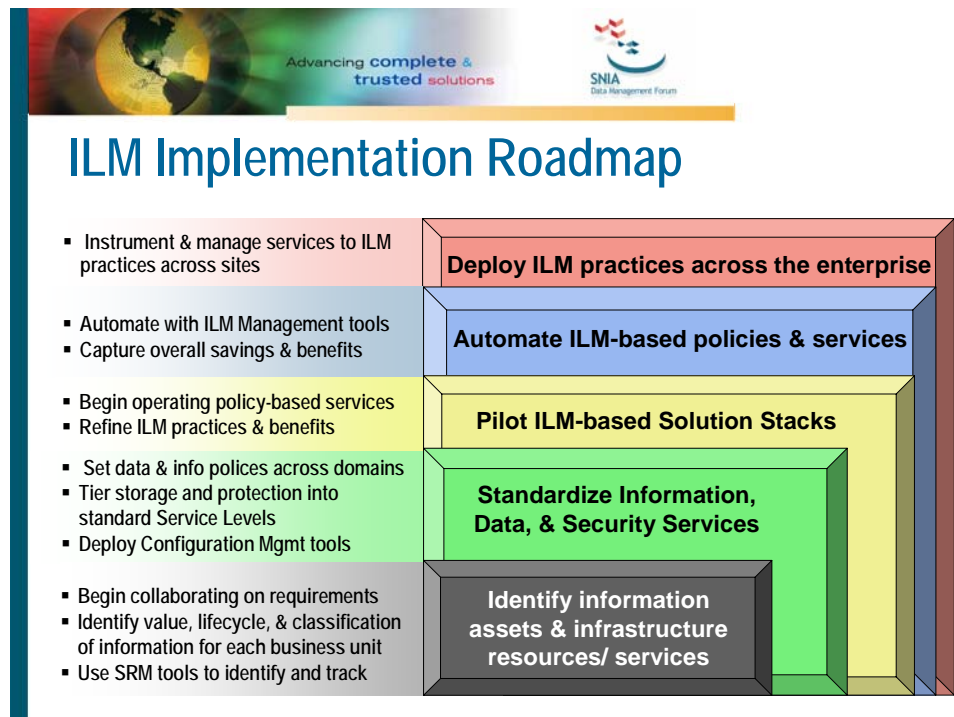
Staged – batched data movement (e.g. archive)

Dynamic – some active data mover (e.g. hierarchical storage management, HSM, or ILM policy services)

The biggest challenge we hear from customers as they begin implementing ILM-based tiered storage is in getting agreement on the information and data classification requirements. This is the crux of establishing successful ILM practices. A successful deployment begins with collaboration. The objective is to have the business participate in setting the requirements. It takes collaboration with the business groups, the records & information managers, the security managers, and sometimes legal and finance to get this done, but it is worth the effort.

Tiering is a “Phase-II” activity in the course of implementing an ILM program. It is an important early step because so much benefit can be had which demonstrates the power of ILM as a management practice.

So, our recommendation is to add intelligence to your tiering efforts by integrating it into a broader ILM-based practice. You can get help with understanding this further from the SNIA’s End-User



Council and the Data Management Forum at www.snia.org/dmf and at our upcoming conferences, Storage Networking World and Enterprise Information World.

By: Michael Peterson
 Program Director, SNIA Data Management Forum
 January 2006

www.snia.org/dmf www.enterpriseinformationworld.com