



Taking on the Classification Challenge

Lessons learned from the EMC ILM Showcase

Table of Contents

Preface	3
The First Challenge: Getting Business to “Buy In”	4
Educating business owners	4
Bridging the communication gaps	4
Scoping the Classification Effort	5
Gathering Data: The Interview Process	5
Avoiding talk about technology, cost	6
Shifting mindsets	6
After the Interviews: Aligning Applications.....	7
Classification Applied: Migrating Applications.....	7
Is Your IT Organization Ready?	8
Conclusion	9

Preface

EMC's IT management in 2003 shared many of the same challenges as our customers. Environments were dispersed—with over five data centers in operation—and large amounts of information, over 900 TB, predominantly used expensive high-end arrays. Fragmented storage across physical assets created operational challenges and low asset utilization (around 45 percent). No formal service levels or service-level management processes were in place; instead, point solutions were the default choice for storage deployment. Delayed procurement cycles were prevalent since EMC's customers took priority over internal requirements for all available storage shipments. Capacity management was reactive, manually intensive, and unreliable, resulting in low utilization/allocation levels—and with information growing at exponential rates, EMC was facing skyrocketing storage operations costs and increasing inefficiencies over the next several years.

An executive decision was made: implement information lifecycle management (ILM) at EMC using the same knowledge, tools, and resources we use at our customers' sites. The fundamental initial objective was twofold—cut costs and improve operational efficiencies and effectiveness, primarily by:

- Consolidating storage devices within each environment
- Linking service requirements to storage architecture to deliver a range of service capabilities
- Aligning applications to the infrastructure based on business requirements
- Avoiding future storage purchases by utilizing storage assets more efficiently—specifically, taking captive pools of storage, combining them in a new environment, sharing across applications, and removing old equipment

Thus began a comprehensive, seven-workstream ILM effort at EMC that is expected to yield a five-year storage avoidance of 1.2 PB and cost avoidance of millions of dollars that were required to support that storage growth. Furthermore, it is creating a flexible and scalable storage environment that positions IT to manage growth more effectively—our core business driver. Aside from achieving our own business goals, this effort has allowed us to experience the ILM journey from the customer's point of view, giving us a unique understanding of the fears, concerns, challenges, and roadblocks inherent in such an undertaking. It allows us to speak to customers not just as an ILM provider, but as a fellow organization that has “lived through” the experience.

In this paper, we—specifically, EMC® Consulting—share some of the key challenges faced, steps taken, lessons learned, and successes achieved during the first step in our ILM journey at EMC: Classification.

ILM is a strategy to align IT infrastructure with the business based upon the changing value of information over its lifecycle. *Typically, the critical first step in implementing ILM is classification*—a process of categorizing information by service level which helps determine how to build an optimized IT infrastructure to match. The classification process includes identifying business service-level requirements, developing a service catalog model to capture the full range of IT services required, developing supporting tiered storage, aligning applications to new tiers for cost optimization and improved utilization, and migrating the applications accordingly. It may also include the development of policies for the management and automated movement of data to meet business requirements.

Classification is foundational to truly understanding your applications' IT requirements and getting the most value by putting those applications into a new, tiered, service-oriented environment—and to implementing ILM. The classification effort will help you understand what you need to do, understand the boundaries, take action to put your applications into a new environment, and ultimately, enable dynamic movement of those applications for maximum, measurable value.

The First Challenge: Getting Business to “Buy In”

A prerequisite to a successful classification initiative is that business owners understand the value of the process and be fully willing and able to engage as needed in terms of time, resources, input, championship, and support. Many classification teams hit their first (and perhaps biggest) challenge right out of the gate when the business owners they must engage are reticent or unwilling to participate.

We found that many business owners perceive/fear that the complexity of the classification process and the time, physical resources, effort, and cost they will have to commit to it will not generate adequate payback and they cannot justify the undertaking to their management. Others fear change and the potential negative impact of classification to the business. Still others buy in to the classification concept, but lack the required resources and funding they think they need to get started. We also found that many business owners have difficulty balancing the perceived value of a classification effort with the perceived impact the classification will have on applications. Especially if they have been utilizing tier 1 storage for all their applications, they equate moving their applications to a lower tier of storage with a lesser service level—and are thus reluctant to engage.

Educating business owners

It is essential to address these fears and misconceptions and gain buy-in by educating business owners at the outset. The classification team must be able to make a compelling business case to owners and successfully communicate the benefits of ILM and classification to the user community. Discussions should be held with core business owners about both the qualitative and quantitative benefits of tiering on the business. Subsequent data and performance analysis can be performed to justify the effort and to reassure the business that the expected results will be confirmed during migration and testing.

We were able to successfully present a message that aligned with corporate direction—showing how the classification effort would provide value back to EMC internally via storage avoidance and cost reduction through the resulting tiering and consolidation. We discussed balancing the inherent complexity that a new tiered environment would bring—managing a blended environment, moving data, spreading applications, operational complexity—with the value that ILM would bring, demonstrating tradeoffs and benefits. By finding a message that resonated with business owners, we were able to gain buy-in at all levels, from users up to the executive level.

Bridging the communication gaps

Two critical communication gaps must be addressed in the process. The first, not surprisingly, is between business and IT. Creating a trusted advisor relationship is key to bridging this known disconnect. Business/application owners are not used to having conversations with IT about what their applications really need; instead, they are used to handing off tactical requests—and expecting that IT will deliver. Changing the conversation to focus on the positive impact of classification on applications will help eliminate questions in business owners’ minds. And as you build a successful track record with demonstrated classification “wins,” you will build credibility with new businesses and establish IT as a trusted advisor that can listen to the organization’s business needs and execute against them.

We discovered a second gap between the executives who fund the projects and the business owners who are responsible for interfacing with IT during the project. We found that even after executive VPs saw the business case and “bought into” the classification/alignment initiative, their word and sponsorship never filtered down into the business units it would affect.

It is critical that everybody in the organization be on board with this effort. Business owners need to be educated and socialized to have productive conversations with IT. This can be accomplished fairly easily over a relatively short period of time through short (60 to 90 minute) training sessions with each of the business owners to be engaged to explain the migration process, what their piece will be, and when their resources will be needed. This effort will help them better prepare for your requests when they come, your program will run more smoothly, and your requests will be more realistic.

There is a very fine line between “application alignment” and “data classification.” Application alignment is more of a static structuring of where an application will sit in a tiered environment, as opposed to examining specifics of the data—e.g., age of data within tables, where and how it should be archived, incorporating automated data movement, etc. We chose to look at applications and where the individual parts of the applications should go—since our goal was to tier to what was needed to support our business case. As opposed to focusing on archiving of stale data, we focused on putting applications where they would increase utilization, lower-cost storage, and enable us to make these savings. We did capture information during our conversations about data usage—what data was required and when it could be archived—so that if incorporating policy became a requirement at some point in the future we could take advantage of it. We are, in fact, now leveraging our content-addressed storage technology to archive data based on policies in selected environments.

Scoping the Classification Effort

We chose to approach our classification effort at the application level (and the associated environments) as it would allow us to achieve our objectives of consolidation and cost avoidance and deliver the most return on our investment. Keep in mind that you can take a phased approach to classification and start at the application level but drill deeper later into the data within a single application if it becomes necessary. We did do this opportunistically—i.e., where we saw benefit without creating unnecessary risk.

It may not be necessary or feasible to ultimately migrate all of your applications as dictated by the classification results due to the amount of effort that would be required and the potential risk that would not be offset by large benefits. Therefore, it is important to consider which applications will deliver the most impact toward your business case. Out of our initial pool of 200 applications, for example, we determined that 90 percent of our business case—core savings focused internally on storage avoidance—would be delivered by classifying approximately 100 applications (the “big-hitters”) and a small group (30-40) of associated “tag-along” applications.

Identifying and focusing on the “big hitters” for classification helps bring the potentially daunting scope of the project into a far more manageable perspective. This makes it much easier to understand where and how to get started. It also allows you to engage your own internal resources more quickly to continue the classification effort on “lesser” applications; you can utilize outside experts to get you through the initial complexities of the process—to help you understand the business, show you how to get the classification planned and executed—and then transfer the process into your organization once the key target applications have been classified.

Gathering Data: The Interview Process

Without classification, implementing a tiered architecture is little more than a technology refresh. The classification process determines what those tiers will look like and where your applications/data will reside on those tiers. In our experience, this information is best obtained through in-depth discussions with business stakeholders about the value of applications/data to the business. For each application, questions should be designed to gather service requirements—e.g., availability, performance, local/remote system recovery, local/remote data recovery, archive retention period, and data usage.

Our goal was to matrix both the criticality of the applications as well as performance requirements and usage at different points in time—information that would help us to subsequently determine where applications should go in the environment both for the near- and long-term. To facilitate the information-gathering conversations and ensure that the information captured was complete and consistent across applications, we leveraged the EMC Consulting organization’s standard approach and toolset, which included a template/worksheet to guide the interviewer (we customized it to be client-specific) and an online template in which to capture responses.

After interviewing a handful of business owners using these tools, we discovered that the interview template made it easy for the interviewer to quickly click through and not capture enough depth of information to truly understand the application and its requirements. Furthermore, the interviewer was not pushed to probe further to gather more detail. We quickly shifted to a more conversational interview approach, focusing on truly understanding each application: all moving parts, down to the database, instance, environment, type of level, backup and refresh levels, development, test, production, how it works, plus interfaces between other applications that supported the core applications. We also realized that it was important to have all the players in the room—the application owners (i.e., the team that determined business requirements) and the technical team (the developers or the DBAs supporting the application).

This shift in approach not only yielded a valuable best practice, but it clearly demonstrated the need for the team to be flexible and adaptable throughout the classification process.

Avoiding talk about technology, cost

We have found that an important strategy in conducting interviews with business owners about their applications is to take technology and cost out of the discussion and focus instead on business process. Our conversations focused on the business processes that each specific application automates and the business impact of those processes being interrupted or slowed. There are several benefits to this strategy. It puts business users in a more comfortable space; they understand how their business operates but may be less knowledgeable about technology. It forces them to be more objective about the steps in a process and what happens when change is introduced into the process. And with cost out of the conversation, they don't have to filter answers against what money they have to spend versus which technology they favor.

The interviewer benefits by getting a set of pure business requirements/business impact statements to which technology can later be applied. Also, because reintroducing technology and cost into the discussion typically reignites user concerns, the interviewer is in a position to have objective conversations and negotiations with business owners about requirements versus cost. This gives both sides in the discussion an opportunity to hone in on what's really important. (As an added benefit, the individuals who conduct these interviews develop a knowledge of the business that most IT infrastructure personnel never develop. By the end of the project, they know more about how business units operate at the business level than most IT shops in the industry.)

Shifting mindsets

Every business owner thinks his or her applications are the most important. They tend to react emotionally when discussing their applications. They know that if there's a problem with an application—if there's downtime or inaccessibility, for example—they are going to “hear about it” from the staff. It is important, therefore, to shift the mindset of business owners during the interview process away from an emotional response and force them to think instead about what problems with their applications really mean to the company in terms of impact. Will there be dollars lost? A significant impact to customer service? Focus the discussion on what the application truly needs as opposed to allowing business users to demand putting applications on the highest priced storage simply to avoid being paged at inconvenient times because of downtime. (Note: A business impact analysis can be performed, if necessary, to help business owners assess revenues lost to applications going down.)

After the Interviews: Aligning Applications

Classification is what typically determines the tier/platform on which applications and data are stored. The value of classification to the business is not storage-centric, however, but rather about determining the service requirements of applications and data to the corporation and then aligning the different layers supporting that application (database, connectivity, storage) with the infrastructure. A cross-functional design team should be brought in to review the documented tiering recommendations and validate that they can be deployed from a technical perspective.

We took a “tiered service approach” to the alignment process, using a service catalog to align applications with tiers based on a level of service. Service levels were divided into four primary areas: production, archive, operational, and disaster recovery. Business units were then presented with a preliminary storage design and a strategy for aligning, deploying/integrating their application against numbered tiers (tier 1, tier 2, tier 3, etc.) By aligning applications with the right tier of service first—not a product—it becomes easy to subsequently assign the right product to each tier.

We found that tiering every component of an application may not be warranted because of the size of the environment, the complexities of refreshing and of splitting applications across multiple tiers, and the additional hardware required. This is why it is important, as discussed earlier, to focus the classification interviews on those applications perceived to have the “biggest bang” in support of the business case. By clearly understanding the requirements (e.g., availability, recovery, performance) for this select group of applications, you can establish how many discrete components within each application on which you need to focus. This allows you to create a strong business case that shows how you will align your applications to different pieces of the new environment and demonstrate how the alignment supports your business-migration strategy. It also allows you to unemotionally ensure that your business users are on board with what you want to do before you actually do anything. (Putting an escalation path in place to a high-level sponsor—i.e., the CIO—will help in the event that you are unable to gain consensus, or agreement from the business.)

It takes time and practice to achieve a balance between the effort expended and the cost savings realized in tearing an application apart and trying to figure out where to put it. When we took into consideration the complexity and risk involved in doing this we ended up challenging a lot of decisions we made early on. From experience, we are now able to take an application, look at it, and quickly create an environment that achieves what is needed from a business case and cost standpoint, while avoiding complexity.

Classification Applied: Migrating Applications

Classification is an enabler that must be leveraged further to provide value by moving applications onto new tiers. As the potential impact to the business is greatest during these events, the migration approach and methodology must include robust governance and quality assurance checks to mitigate risks. Several valuable lessons were learned during this effort:

Map out the entire migration process at the outset. The challenge in taking action against the classification results and moving applications onto new tiers is that the business will not stand still while these processes are in progress. Timelines must be adjusted accordingly at the outset to accommodate the likely loss of resources to other activities over the duration of the project. Gaining commitment for resources or to actively engage in the migration activities will be lengthier and more difficult. Ongoing business initiatives such as application development, commitments to release new code, and technical refreshes also add a layer of complexity into the migration process; the team must plan the migration around these activities as well as take into consideration the impact the migration will have on them (remember that a lot of data affecting a lot of business applications will be moved in the process).

Mapping out the entire classification and migration process at the very beginning, socializing the plan with all parties to be involved at any point, gaining their buy-in, and only then starting execution (instead of notifying people along the execution path) minimizes the risk having to make adjustments while the migration is in process or having to repeat the migration within the same quarter or year. The organization will be in a position to help with the alignment as they see customer demands materializing instead of having to squeeze them around ongoing activities on short notice. Doing this will also bring changes to more of an exception state than a regular occurrence.

Document and educate. There are a surprisingly large number of things that technical staffs do not know about their environments; lack of adequate documentation is a key factor. A concerted data collection and discovery effort is important, therefore, to ensure that the migration team fully understands the application. It is also important to educate the IT organization up front about the migration process. ILM represents a significant cultural shift and can entail some difficult conversations to change how people frame their requirements. These conversations will be easier if you educate IT before—not during—the alignment interview.

Address all elements of the infrastructure in the plan. When planning for migration, you must take into consideration the impact the migration will have on all elements of the infrastructure—for example, server remediation and backup and recovery—and not just storage. Workstreams must be synchronized accordingly. Otherwise, improperly aligning the storage consolidation and tiering with the new backup, recovery, and archiving architecture can introduce problems where none existed before—for example, longer backup windows.

Validate your environment. It is very important to do exhaustive data collection about the environment up front—even if this adds time to the effort and requires augmenting the storage or server teams or data center operations with additional staff to do the work. Otherwise, you are likely to uncover unknowns about your environment during the migration that will delay or stop the process—such as mislabeled or misrouted cables that must be rerouted or removed before the migration can continue. If you fully understand your environment at the outset, you can staff your project more appropriately—perhaps with more people as well as with more people in lower-skill areas (e.g., hardware support staff or PC technicians to do physical tracing in the data center instead of storage administrators who must be pulled off of other activities to do the job). Also, much of this work can be done during normal business hours and other times that will not affect operations.

Tactically, business looks at IT as a cost center. Demonstrating success in the classification process transforms this perception, enabling business to look at IT as providing a critical service that delivers not only what they want, but what they truly need, while not over- or under-delivering. This changes IT into a business service provider—a more strategic position.

Is Your IT Organization Ready?

It is essential to understand the limitations of your IT department before engaging them in a classification effort. Having an adequate number of resources with the appropriate best practices, tools, and knowledge can spell the difference between success and failure, especially in areas such as operational alignment, capacity planning, and service catalog definition. Resources must understand your operational environment, be able to help your organization understand the structure of what your service catalog should look like, and be able to execute a detailed design of a to-be architecture.

In our experience, most IT shops are very thin in terms of project management capability (their core mission is to develop, build, and support solutions), but are usually solid on technical staff. A successful classification engagement needs a solid project management team that's experienced in dealing with large-scale projects and that brings best practices and knowledge to the table. Our IT department looked to us to help them with governance, project manager methodology, planning, scheduling, and tracking to keep the program on scope, within budget, and on schedule.

Classification, consolidation, and tiering have created a scalable, manageable foundation for future ILM growth at EMC.

EMC Consulting brings a unique mix of information management and technology infrastructure experience, strategic business focus, and IT best practices to help companies implement ILM strategies. EMC addresses emerging business and technology challenges arising from the exponential growth of information, increased regulatory guidelines, and the need to cost-effectively manage and protect information assets. More information about EMC Consulting services can be found at www.EMC.com/services.

Conclusion

Classification is foundational to truly understanding your applications and getting the most value from IT assets by putting them into a new tiered, dynamic environment—and to implementing ILM. The classification exercise will help you understand the real business requirements from IT, understand the boundaries, take action to put your applications into a new environment, and ultimately, enable dynamic movement of those applications for maximum, measurable value. This helps IT understand the business process, the value of the data to the functional unit, and how to align with the established service levels.

Inarguably, the application/data classification exercise is a significant undertaking with many known challenges and complexities that raise legitimate concerns in both business and IT about risk, impact, and value. Our classification experience validates undertaking this journey. Clearly, the move alone from tier 1 storage to lower-tiered storage has the potential to avoid millions of dollars in infrastructure purchases over time. But classification is not just about dollars and a technical refresh. It is perhaps more importantly about organizational, cultural, and operational change that impacts the way IT service is provided to the business. The customer organization must not only understand this but also embrace and internalize this concept up front in order to realize maximum impact.

Classification is an ongoing effort. Business is dynamic, and business requirements for information change over time. Furthermore, most organizations choose to break the task down into manageable increments based on their business schedule and the speed of change. Value management is an important component of the classification effort. Users want to see tangible benefits that justify your moving their applications. It is important to measure those benefits—e.g., cost savings or new service levels (“earned service value”)—and then come back to the customer to demonstrate quantifiable results against the original business case.

Our business case was based on 118 TB of real data at 30 percent growth. Over the two years subsequent to embarking on this project, stored data has exploded to close to 1,500 TB as a result of robust business growth, acquisition, and more. Nonetheless, our measurements show that over the full life of our business case (five years), our effort will net within 90 percent of what we had targeted originally. We expect to achieve a reduction in the number of storage arrays from 205 to 86, SAN fabrics from 63 to 14, and switches from 70 to 30. We expect to avoid more than \$15 million in costs. And we will have created a flexible and scalable storage environment that will allow our IT organization to manage growth more effectively.



where information lives®

EMC Corporation

Hopkinton
Massachusetts
01748-9103
1-508-435-1000
In North America 1-866-464-7381

EMC², EMC, EMC ControlCenter, AlphaStor, ApplicationXtender, Captiva, Catalog Solution, Celerra, CentraStar, CLARAlert, CLARiiON, ClientPak, Connectrix, Co-StandbyServer, Dantz, Direct Matrix Architecture, DiskXtender, DiskXtender 2000, Documentum, EmailXaminer, EmailXtender, EmailXtract, eRoom, FLARE, HighRoad, Navisphere, OpenScale, PowerPath, Rainfinity, RepliStor, ResourcePak, Retrospect, Smarts, SnapShotServer, SnapView/IP, SRDF, Symmetrix, TimeFinder, VisualSAN, VSAM-Assist, WebXtender, where information lives, Xtender, and Xtender Solutions are registered trademarks and EMC Developers Program, EMC OnCourse, EMC Proven, EMC Snap, EMC Storage Administrator, Acartus, Access Logix, ArchiveXtender, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, C-Clip, Celerra Replicator, Centera, CLARevent, Codebook Correlation Technology, Common Information Model, CopyCross, CopyPoint, DatabaseXtender, Direct Matrix, EDM, E-Lab, Enginuity, FarPoint, Global File Virtualization, Graphic Visualization, InfoMover, Invista, MirrorView, NetWin, NetWorker, OnAlert, Powerlink, PowerSnap, RecoverPoint, RepliCare, SafeLine, SAN Advisor, SAN Copy, SAN Manager, SDMS, SnapImage, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, UltraPoint, UltraScale, Viewlets, and VisualSRM are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.
© Copyright 2006 EMC Corporation. All rights reserved. Published in the USA. 4/06

White Paper
H2145