

Right Focus, Right Time — EMC Documentum CSS Brings Content Awareness To Data Migration Policies

Analyst: Anne MacFarland

Management Summary

Progress is more than just a matter of direction and velocity. Often there are barriers to taking the next step that must be struck down or sidestepped. Achieving progress is then a matter of having the right tools at the right time to strike down those barriers. In such situations, capabilities not needed for immediate forward progress are needless overhead.

When an enterprise tries to align the quality of service given by different tiers of storage with the business value of the information being stored, it enters a situation where IT application administrators and storage administrators must collaborate with line-of-business folks in something which deeply involves each of their areas of expertise but is also on the fringe of what all of them expected to be doing. The situation is more likely to breed an unnatural calm of conflict avoidance - or a perfect storm of discussions rather than a reasonable velocity of progress - unless there is a vehicle that all three parties can understand and work with to set the rules governing the use, placement and migration of enterprise information.

EMC has recently released a tightly-focused software offering it calls *EMC Documentum Content Storage Services*. CSS links the business value of the information, as classified through Documentum's Content Management application, with data migration capabilities needed to implement business-centric Information Lifecycle Management (ILM). This product gives enterprises large and small a way to get the business folks and the storage folks collaborating productively to set policies governing how enterprise information should be treated over its lifecycle. It enables the charge-back to make these negotiations productive and a tiered storage environment a more obvious boon to the enterprise. Most importantly, it gives them a way to get started.

Read on for more details.

IN THIS ISSUE

› The Tiered Storage Imperative	2
› The CSS Solution	2
› How It Works.....	2
› Conclusion.....	3

The Tiered Storage Imperative

With the development of storage system utility programs to migrate information non-disruptively to the application and end users, IT systems have a way to deftly and unobtrusively lower the costs of maintaining access to enterprise data. It is a matter of orchestrating the movement, and there are many traditional products that do a piece of what is needed, but not the whole job.

Storage management software sees basic file characteristics and produces file usage metrics, and so can migrate data files when they reach a certain age, or when their usage falls beneath a threshold. This is called Hierarchical Storage Management (HSM), and it has been around for years, from back in the days when migrating data was a real pain (which limited HSM's popularity). Furthermore, **HSM typically does not recognize the more granular characteristics of the content of the files, nor their business value.**

Enterprise Content Management (ECM) software applications develop profiles for the business value of different kinds of content over their lifetime. Files, once created or captured, can be versioned, forwarded, amended, approved, audited, and finalized and still have residual value for analysis and documentation. Each of these stages may involve a different set of access rights or a different write authorization. The value of some information wanes quickly, but in other cases the value of information is increased in the aggregate, like research data, or in hindsight, if it turns out to have been important. ECM relates the information to the workflow and business process of which it is a part, and automates the changes to the directories and other structures governing and promoting its use.

EMC Documentum Content Storage Services puts together both capabilities into a unified interface, focusing on the unstructured information that comprises, by some estimates, up to 80% of enterprise data. The rich characterization that Documentum provides, now extended to enable data movement, fulfills enterprise demands for flexible data use and enterprise desires for cost control.

The CSS Solution

EMC Documentum CSS manages a hierarchy of information objects with a console by which these rules requirements may be specified and changed as needed, a repository that stores these rules and documents their change over time, and a rules engine that enforces the rules. Documentum Content Management characterizes these objects with tags for a variety of attributes specific to their context and use. Documentum can link information to related information to form composite objects, or an aggregated virtual file, which can then be subject to a single policy. These composite objects allow rules to be more comprehensively written in a way that is congruent with how business is done.

EMC Documentum CSS works with storage from many vendors including IBM, HP, and HDS, as well as with all EMC products. It works with tape as well as disk storage.

The CSS API exposes the Content Management policies to relevant enterprise applications, like SAP, which can call CSS to invoke a policy-driven event. This lets application management contribute its expertise to the CSS system.

CSS lets enterprises more easily build a tiered storage system to reduce costs, but it is also useful to enterprises yet to embrace tiered storage. For them, content management lets them manage their information better. And, of course, CSS can be used to regulate the treatment of data subject to governmental retention regulations. It is designed to work well with EMC *Centera*.

How it Works

The rules engine makes it easy to create and manage policies using three dimensions of *where, when, and what*.

Where

Retention policies are built on a CSS policy engine that covers the three critical aspects of data retention:

- Where the data is originally put.

- Where it is moved, under what circumstances.
- Documentation of the movement of the data, giving a sort of *curriculum vitae* of the data.¹

When

Sensible migration of data curbs the costs of the data archive, and also expedites activities like data shredding (for in proper records management every information life-cycle ends in destruction). How CSS determines when data should be moved can be determined in three ways:

- **The migration can be *event triggered*.** This would be appropriate where the books are closed or a process is finished and the data turns from a workflow element to documentation, which perhaps should be locked into inalterability.
- **The migration can be *ad hoc*.** This capability is needed for the kind of unstructured information that just doesn't get used as often and would fall between the cracks if there weren't something like usage metrics to invoke a move to a lower tier of storage
- **The migration can also be *batch scheduled*.** This is a good way to get the migration done during known times of inactivity, particularly for information networked on a LAN.

What

Because of the rich content management components, the characteristics by which CSS classifies information are detailed enough to allow security and access rights and quality of service to be implemented at a granular level – so the data gets the individual treatment it deserves.

Because with CSS you can have an auditable log of policy changes, you can see how policies are working, and can reap expertise further optimize the system and extend its reach in the enterprise should you choose to do so. It is all a matter of having the tools for the job at hand.

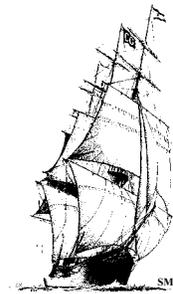
¹ While this does not give the fine detail of information use, Documentum has partners, such as Authentica, who can add in that capability.

Services

If an enterprise is still befuddled by the nature and extent of its unstructured data, EMC has professional services to help with the task². Like many such services, these are an accelerant of a task that many enterprises do not have the time or personnel to do.³

Conclusion

The launch of EMC Documentum Content Storage Services comes as enterprises are recognizing the *need* to consider both the business value-lifecycle and the usage-lifecycle of the data, and the *difficulty* of translating these considerations into policies. For existing Documentum customers, the new storage management capabilities will pay for themselves quickly. For customers thinking about using content management software for the first time, the combination of ECM and tiered storage gives them a chance to save money on storage costs while getting a better handle on their enterprise content.



² These professional services are not to be confused with the *Software Services*, as EMC call this Documentum Content product.

³ See **The Clipper Group Explorer** dated February 4, 2004, entitled *Services - A Vital Tool for Accelerating Enterprise Responsiveness*, available at <http://www.clipper.com/research/TCG2004007.pdf>.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Anne MacFarland is Director of Enterprise Architectures and Infrastructure Solutions for The Clipper Group. Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 28. (Please dial “1-28” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.