



Research Note

Cover Your Assets: How to Protect Your ECM Investment

This Data Mobility Group report pulls the cover off enterprise content management (ECM) to expose a weakness that can (and has) cost customers hundreds of thousands of dollars per incident in lost productivity alone. Then it takes a brief look at CYA Technologies, a company that earned its reputation by helping customers protect their ECM investments.¹

In a recent report about data protection, DMG sought to illustrate how even the tiniest impact on operational efficiency and resiliency can have an enormous impact on a company's bottom line. Using payroll data from a 2002 U.S. Census Bureau Economic Census, DMG calculated the average cost of one minute of productivity across the top 50 companies in four healthcare sectors. The result was an astonishing \$13,415 dollars per company per minute—more than \$3.35 million dollars per year.²

Over the past 15-plus years, many tens of thousands of corporations built business cases for content management predicated on anticipated operational savings and risk management. Executives and managers alike envisioned frictionless workflows, greater collaboration and effortless repurposing of digital assets—largely based on a common desire to streamline operations. Content management was (and still is), in the eyes of many, a fast track straight to the bottom line.

While the long-term operational benefits of content management seemed straightforward, managers soon realized the journey toward intelligently

Copyright © 2002-2008 Data Mobility Group, LLC. All Rights Reserved. Reproduction of this publication without prior written permission is forbidden. Data Mobility Group believes the statements contained herein are based on accurate and reliable information. However, because information is provided to Data Mobility Group from various sources, we cannot warrant that this publication is complete and error-free. Data Mobility Group disclaims all implied warranties, including warranties of merchantability or fitness for a particular purpose. Data Mobility Group shall have no liability for any direct, incidental, special, or consequential damages or lost profits. The opinions expressed herein are subject to change without notice.

managed content was anything but simple. Schedules slipped, costs rose and a number of unanticipated obstacles threatened to undermine substantial investments. For the organizations that persevered during the early stages of ECM implementation, the business process insight alone was worth the effort.

Practitioners began to understand that content management is not a finite project, but a process of continual improvement—an ongoing commitment to operational efficiency and intelligent risk management. However, they failed to foresee that their investments would be periodically and inadvertently ambushed by the very systems designed to protect them.

The Hidden Weakness: Partial Data Loss

Costly incidents of partial data loss attributed to misconduct, incompetence, application failures or simple human error are the product of outmoded data protection solutions developed at a time when technologists were tasked with device- and facility-level disaster recovery (DR), not the protection of individual information assets and metadata.

“Adequate data protection, like wearing a safety belt in an automobile, is something rarely thought about until the worst happens, and then we find ourselves grateful we had the protection, or wishing we had it.”

In business, as in life, the worst problems are never immediately obvious. Nobody, not even the companies that designed these complex environments of databases, file systems and multiple interdependent applications, fully understood the challenges of protecting them from incidents of partial data loss.



Corporate users and content management engineers defined business assets as objectives, workflows, tasks, projects, forms, documents and files. In contrast, storage system engineers and IT departments perceived business assets in terms of disks, tapes, bytes, blocks, and volumes. These two dramatically different worldviews continue to be a source of frustration and misunderstanding between business users and IT, information managers and storage administrators.

For the longest time, existing backup practices and applications were thought to be perfectly adequate for the job—a notion that would later prove to be untrue. Content management systems continued to be deployed by the thousands worldwide, each with its own faulty data protection scheme hidden beneath a façade of structure and control.³

It was not until the mid nineties, after many casualties of inadequate data protection, that a handful of companies began to question if traditional data protection methods were sufficient for systems that track, store and manage complex relationships and dependencies—the very heart of content management.⁴ They intuitively understood the limitations of existing disk-, volume- and block-level DR schemes—all of which lacked relational awareness and diagnostics—and embarked on a mission to find a more complete solution.

Supplementing Traditional Backups With Partial Data Recovery Capabilities

Traditional disaster recovery solutions are indisputably unable to cost-effectively and efficiently protect content management systems from partial data loss scenarios. Such systems are designed for recovery from disasters that incapacitate or destroy facilities, hardware or applications. It is true that ten years ago companies had no choice but to work with the only options available at the time, and any DR protection was better than none at all. But, today companies have a choice and they choose not to be economic hostages to inadequate data protection practices.

Thankfully, a handful of time-tested partial data recovery solutions are available to address the needs of ECM customers. Toward the end of this report, DMG will highlight one such solution designed from the ground up to protect ECMs from partial data loss.

The following is a set of guidelines against which every data protection solution must be measured to pass muster for modern-day content management:

- **Say no to full restores**—Partial data loss should never, under any circumstance, require a full system restore/rollback. Full restores of the type used for DR are widely disruptive and costly, and inevitably result in unpredictable additional data loss tradeoffs. The objective is to quarantine data loss and minimize the impact on unaffected personnel, systems, and content.
- **Say no to offline restores**—It follows that a content management application should never need to be taken offline in order to perform partial data recovery. Offline means unavailable to everyone. In today's 24/7 economy, the impact on global productivity alone—from partners and suppliers to employees and customers - can be devastating.




- **Say no to resource-intensive recovery**—The impact of information recovery on IT resources should be minimized. Partial data recovery should rarely, if ever, require the effort of more than one person.
- **Say yes to minimum recovery times**—The ability to recover specific affected content without a full system restoration dramatically reduces recovery time objectives (RTOs). Affected content should be fully restored and accessible as quickly as possible.
- **Say yes to proactive integrity checks**—A backup is only as good as the information it contains. Garbage in, garbage out applies here. Few companies regularly test their backup processes, and fewer still test the integrity of the backup data. Periodic integrity checks ensure data consistency and reliability.
- **Say yes to consistent automated retention policy enforcement**—A solution designed to protect against partial data loss should always destroy protected content in accordance with established retention policies to avoid any possibility that the content would be inadvertently restored to the ECM at a later date.
- **Say yes to relational awareness and preservation**—Content management's complex web of workflows, access controls, metadata, and information assets must be thoroughly captured and preserved without loss or corruption. This forms the foundation upon which all partial data recovery activities depend.

Conclusion

After hundreds of thousands (and often many millions) of dollars and countless hours of investment into reengineering and streamlining operations, companies must not gamble the hard-won operational and economic gains of their ECM implementations on incomplete and inadequate data protection.

Companies can choose to maintain the status quo while unrelenting partial data loss continues to erode profitability, or they can choose to fix their existing data protection strategies to minimize risk and costly lost productivity.

Data Mobility Group believes IT managers should implement partial data recovery solutions alongside existing DR to preserve the operational integrity and value of their ECM investments. Failure to do so is the equivalent of throwing money and caution to the wind. 



CYA Technologies: Protecting ECM Investments Since 1998

Since 1998, CYA has earned its reputation as a recognized expert in solutions ensuring the integrity, granular recoverability and asynchronous replication of information within ECM systems. CYA is protecting hundreds of companies against the risks of partial information loss, which accounts for 80 percent of all ECM data loss incidents according to research from AIIM. Partial information loss incidents are common, and can be caused by anything from simple typing mistakes to intentional malfeasance or system failures.

Traditional backup and recovery solutions simply don't understand the information architecture of ECM systems and ignore metadata—the audit trails, digital signatures, renditions and workflow associated with each file. Recovering a single important file stored in an ECM system without CYA can impact the ECM system availability and ECM user productivity—not to mention take a lot of IT time.

CYA Technologies is applying its ECM granular recovery solutions to life sciences, financial services, aerospace, and several other industries with increasingly complex regulatory demands. CYA can help companies ensure compliance with record retention, auditing and document workflow management demands by protecting the critical metadata associated with each document—metadata put at risk with each partial information loss incident.

DMG is aware that CYA Technologies is not the only solution available to protect against partial data loss in an ECM environment, but we feel it is one of the best, time-tested and mature options available. But do not take our word for it. If you have an ECM, talk to the folks at CYA and add their solution to your short list.

For more information on CYA Technologies, visit www.cya.com.

Endnotes

¹ For simplicity, DMG will use the acronym ECM and the phrase “content management system” interchangeably throughout this report.

² The value of one minute of productivity gained or lost per day, per company, summed over a period of 250 work days.

³ This is not to imply that vendors misrepresented their products. In fact, they were no more aware of the problem than the IT departments of customers. It simply was not considered to be a “problem” at the time.

⁴ Even today, some of the most recent advancements in state-of-the-art data protection, such as continuous data protection (CDP) fail to address the needs of content management.