

Gilbane Beacon
Guidance on Content Strategies, Technologies, and Practices

Content Management Interoperability Services (CMIS)

Addressing Contemporary Requirements for Content Integration

By Dale Waldt, Senior Analyst
December 2009

Sponsored by members of the OASIS CMIS Technical Committee including:

OASIS 



nuxeo

Microsoft



Contents

The Opportunity to Embrace Integration	3
Business and Technical Background.....	4
How CMIS Works	5
CMIS Benefit Scenarios	7
A Contemporary Solution to Integration.....	10
To Learn More About CMIS	11

The Opportunity to Embrace Integration

The deployment of multiple content management (CM) systems within a single organization is no longer a business trend, but a business fact. AIIM reports in its 2009 industry overview that “a single [enterprise content management] system concept” is in place at just 35% of the 568 companies surveyed this year.¹ Mergers and acquisitions, a steady stream of new CM products, and the rapid evolution of innovative approaches like cloud computing virtually ensure that multiple disparate CM systems will define the enterprise IT landscape going forward. As a result, a growing number of companies are wrestling with the significant business challenges that result from the proliferation of CM systems within a single enterprise: content and data in isolated silos, operational inefficiencies due to disjointed processes, expensive and costly application development, and compromised business agility because of integration complexity.

Executives want to make decisions based on strategic goals and objectives, not on whether systems can be integrated with reasonable effort. Business buyers of CM want to be able to embrace integration so they can implement best-of-breed, purpose-built systems for specialized applications like digital marketing and electronic discovery. They need systems that connect externally as well as internally, supporting relationships throughout their value chain. Developers want cleaner interfaces that make integration seamless to system users. They want to remove cost obstacles to building composite applications and preparing their applications for cloud computing.

IT and business managers need new approaches to making disparate content management systems work together in useful ways. They can no longer afford to be saddled with significant integration challenges caused by systems with proprietary features and methods for communicating and accessing the data stored within. Historically these systems utilize specific connectors or adaptors to interoperate. The time and money required to integrate just two systems is burdensome. Cobbled-together integrations are not reusable or easily maintained. Worse, tremendous amounts of business value lie dormant simply because content assets cannot be leveraged across the enterprise.

The Content Management Interoperability Services (CMIS) specification is a market-driven response to CM integration challenges. CMIS is being developed under the auspices of OASIS (Organization for the Advancement of Structured Information Standards) by an OASIS Technical Committee. CMIS specifies a standardized content management repository

¹ AIIM Industry Watch: State of the ECM Industry 2009

access method, providing a powerful, consistent way of connecting CM systems. It is designed to reduce integration complexity and the development time and costs required to provide interoperability between disparate systems. With CMIS, creating integrated content ecosystems with two or more repositories will be easier, cheaper, and faster.

Gilbane believes that CMIS has the potential to be a game-changing standard, not only through its promise to facilitate affordable CM, but also as an enabler of whole new classes of high-value, information-rich applications that have not been possible to develop in the past. This paper describes the CMIS specification itself, the business and technical drivers behind its creation, and the benefits organizations can expect to see with its deployment.

“Multiple content management systems are becoming the norm for our customers. With point-to-point connections, the seams always show.” -- Ethan Gur-esh, Microsoft

Business and Technical Background

Content management tools have evolved into a diverse mix of applications designed to support different aspects or processes within the enterprise. Connecting these disparate repositories allows the content to be accessed as needed, even blended together in new and interesting ways. Some come with out-of-the-box connectors, but interoperability often requires custom integration development. Add the mix of operating systems, applications and database platforms, and diversity of data models, and the ability of systems to interoperate well is even more challenging. These connections are sometimes fragile and do not achieve all the desired functionality.

Increasingly, content providers find they need to present their information in a more easily consumed format, perhaps organized into “views” for specific audiences. Often this means that data from multiple repositories needs to be combined into an integrated, organized format, in spite of data model and operating environment differences. Limited interoperability of these systems can prevent an organization from fully achieving its business goals. New and interesting content “products” can be achieved when data and system accessibility are improved.

Recognizing these challenges, several vendors collaborated to define a standardized approach to repository interoperability. Originally, EMC, IBM, and Microsoft worked to create a draft specification that would provide a common method for repository interoperability. They were joined by Alfresco and several other providers of enterprise content management platforms. The initial work of this group was submitted to the OASIS Content Management Interoperability Services (CMIS) Technical Committee (TC) where it

has been further developed by a broader group of companies. Participants in the OASIS CMIS TC currently include Adobe, Alfresco, ASG, Booz Allen Hamilton, dotCMS, EMC, fme AG, IBM, Magnolia, Microsoft, Nuxeo, Open Text, Oracle, Quark, SAP, Saperion, Sun Microsystems and others.

How CMIS Works

The developers of the CMIS specification took into consideration the most common and widely used features and properties of repository connectors and integration found in a range of systems. The goal is to create an interoperability standard that can easily be adopted and supported by the widest range of vendors and user organizations.

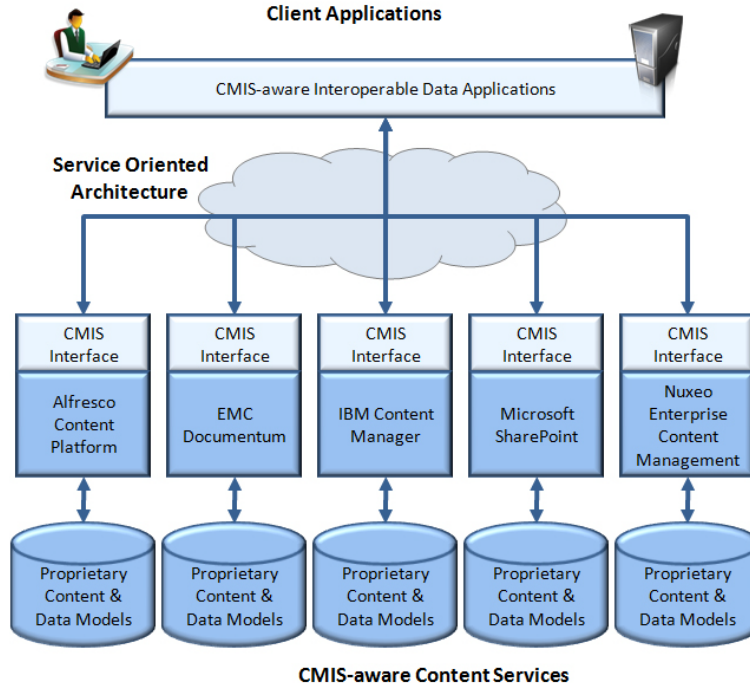
“CMIS helps break down the walls of information silos and allows data to be presented in ways not constrained by the technology.” – Florent Guillaume, Nuxeo

The CMIS Object Model provides a standard, unambiguous method for accessing repositories regardless of its computing platform or environment. The CMIS specification currently defines four common types of objects that can be addressed and managed in a CMIS-ready repository:

- Document objects;
- Folder objects;
- Relationship objects; and
- Policy objects.

Document and *folder* objects are defined, as their names indicate, to classify objects stored as documents and folders respectively, similar to the way many file systems and repositories define and manage content. The *relationships* object is used to define relationships between document objects (including links), but also supporting other means of grouping and relating objects in much the same way relationships are defined in relational databases. *Policy* objects are used to manage policies for accessing and exposing the content.

Two access methods are defined in the current CMIS specification: a *Web Services interface* and an alternative *Atom Publishing Protocol (ATOM Pub) method*. The Web Services method provides a WSDL (Web Services Definition Language) description of the interface to allow developers to connect systems to it. SOAP/WSDL Web Services have been widely adopted for enterprise applications. The ATOM Publishing method relies on the predefined ATOM specifications and processes. AtomPub is an increasingly popular protocol for lightweight applications. CMIS is positioned to support additional protocols in the future.



Example of diverse repositories interoperating through CMIS interface to provide client applications access to wide range of content

Within a Service Oriented Architecture (SOA), a Web Service can provide platform-independent access. This access method and content structures would be defined in a WSDL (Web Services Definition Language) and would use either a SOAP for messaging, or communicating, with the repository. A simpler interface can be provided via CMIS based on REST and the ATOM Publishing protocol. Either method enables programming-language and environment-agnostic repository interfaces to be built and deployed quickly. Both are easily understood by developers seeking to establish a new connection to another's content resources. And since the repository is based on a standardized generic access model, developers will be less dependent on proprietary connectors.

*"Imagine the possibilities if developers could 'build once, run many' instead of hand-crafting specific proprietary integrations."
– Ken Bisconti, IBM*

CMIS and Alternatives: A More Complete Solution

Why should companies consider CMIS over other approaches to facilitating integration? Connecting CM systems in better, streamlined ways has been the goal of several efforts, but none have approached the core challenges in an open and universally applicable way. Aside from proprietary connectors, precursors to CMIS tended to be constrained by the environments in which they can be deployed, or by limited adoption.

Not much is directly comparable to CMIS in the market.

- JSR 170, the Java Content Repository (JCR) API specification coming out of the Java community, and WebDAV (Web-based Distributed Authoring and Versioning) offer similar functionality, but they are not competing as much as complimenting each other and CMIS. An existing JCR, WebDAV, or proprietary system could also add a CMIS layer and, possibly, extend the life of that system.
- Several vendors have proprietary product offerings that provide integration to a variety of other systems. CMIS is an open version of a portion of the functionality that they provide. Over time, CMIS will reduce the need for proprietary tools and eventually may replace them completely. Meanwhile, CMIS and proprietary systems can work together.
- Web portlets, such as those based on JSRs 170 and 283 Content Repository for Java Technology API provide some similar features, but are not the same as CMIS and have limited adoption. Developers can plug in a presentation layer that can pull in content from other repositories, but portlets do not help make the content consistent with other content (portlets are agnostic to content). CMIS goes much further towards unifying the look and feel of the content. CMIS has a defined Domain Model (combination of data model and defined services) that enables consistent, robust access and updating of the content.

Compared to these approaches, CMIS is platform independent and will work in any operating environment. It defines a more robust set of content objects and actions to enable a common access method to a wide range of tools and systems. Lastly, CMIS offers the benefit of confidence that comes with the adoption of an international standard. It is based on well-defined and stable concepts that system developers can rely upon for a long time to come.

*"The ECM market needs CMIS to achieve its potential. CMIS adoption will drive the marketplace to achieving new applications and products, and not just from vendors, but also from user organizations."
– Razmik Abnous, EMC*

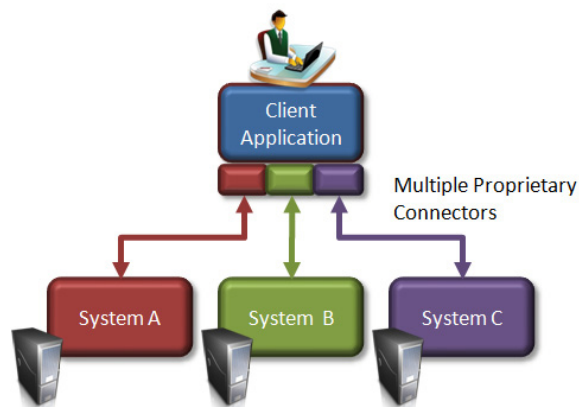
CMIS Benefit Scenarios

CMIS was conceived to address business issues. It offers tantalizing opportunities for CM buyers and developers because it addresses both top- and bottom-line issues. It is fairly easy to see how development costs and schedules can be reduced with a standard repository interface. Another class of truly exciting benefits lies in the new and innovative applications

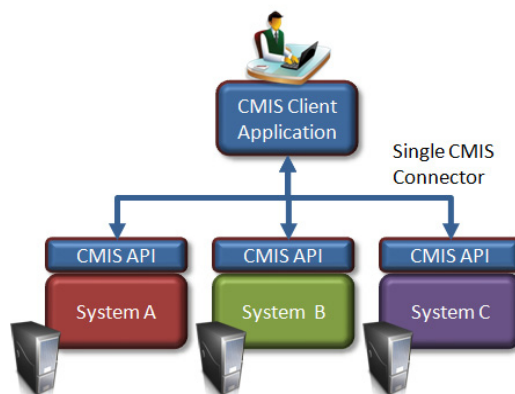
that this capability will enable. Here we lay out four benefit scenarios that business and IT managers can expect from CM systems that are CMIS-enabled.

Reduce Development Costs and Schedules

A common interface into repositories will eliminate much of the integration work currently required to create a unified view into disparate systems. Consider the simple architecture below. Without CMIS, the Client Application will require the development or configuration of three proprietary connectors. Developers will need three skill sets – one for each disparate proprietary system to be accessed.



With a CMIS-compliant interface for each system, as illustrated below, the developers of the Client Application really only need to focus on one connector (shown in blue) in order to talk to the CM systems, assuming that they all speak CMIS.



Also consider the impact of replacing any one of the four systems in either diagram. The amount of development needed to accommodate that change is significantly less with CMIS since only the changed systems are likely to require modification; the Client Application will most likely continue to operate without requiring any changes to be synchronized between

client and server applications. In the non-CMIS architecture, several applications and connectors may be affected by a change in only one of the systems.

Leverage Existing Investments

CMIS is envisioned as an integral part or feature of content management systems. Even so, some vendors intend to create CMIS-based middleware to ease integration with non-CMIS compliant systems and applications. With this approach, existing legacy applications can be configured to communicate with the middleware using CMIS, thus extending the life of the legacy system until it is migrated to a CMIS-compliant repository. The cost of migrating an entire system to a new repository, even to gain CMIS benefits, may be more than adding a middleware layer and building the custom interface and connectors needed to communicate with the middleware.

Enable New Applications

CMIS eases the effort required to achieve seamless integration, freeing business applications from the constraints of content locked in information silos or propriety data formats. By enabling access to content stored in repositories that were previously cumbersome to integrate, information becomes usable in ways not possible before.

End-users will see optimized processes that combine content from multiple sources in a unified view. CMIS-enabled business applications will eliminate the need to interrupt a business interaction with a customer, for example, just to switch content sources – from a ticket number in a support tracking system to a knowledge base of technical documentation to the ERP system with information about the customer’s specific product configuration. With seamless integration, processes that require information from multiple repositories can follow a logical flow rather than a flow dictated by the types of connections or systems. In this case, the benefits of CMIS drive all the way to better experience for the end-customer.

Information from one repository can be combined with content from others and be presented in new ways. For instance, a mashup of company and partner data might combine product catalog data from an internal source with external distributor addresses plus map information. The end-result is an applications that assists potential buyers in locating in-stock products available for purchase. Regulatory information from many jurisdictions can be integrated and presented in “views” by topic or for specific audiences, such as buyers of a

“CMIS will make it possible for business units to deploy systems independently and focus on application needs rather than on infrastructure considerations. This is how it should be.” -- David Choy, CMIS TC Chair

product that carries different warnings and disclaimers depending upon their geographic location (e.g., prescription drug labeling in North America and the European Union).

Enhance Content Discovery

In some environments, content applications can be static with little or no change to the configuration and available data over time. Other content applications are much more dynamic, requiring more frequent “scanning” of existing resources and rapid integration of new content systems. CMIS will shine in these dynamic environments by reducing complexity, cost, and time to add new CMS repositories on the fly and by making that content immediately available in a unified view.

CMIS enables flexible discovery processes where data can be quickly configured and accessed without lengthy delays waiting for a new connector to be developed or installed. A lawyer researching a case may need to establish a new connection to a content resource as part of the discovery process during which evidence data is gathered and reviewed in support of a case. A researcher in the intelligence community may need to monitor a range of systems with content in multiple formats, but receive results in a unified, consistent format that is easy to import into the target intelligence application. In these two examples, individual repositories can be quickly added or excluded from a user’s view as business rules dictate if they use CMIS as a common access interface.

A Contemporary Solution to Integration

In the course of Gilbane’s work as analysts and consultants, we see many organizations throughout the world and across a variety of industry segments struggling to develop and deploy content-centric business applications that deliver real business value. Integration of systems is fundamental to the success of these applications in nearly every case – integration of content management systems with databases, content creation tools with content management systems, translation memories with content management systems, content management systems with ERP systems, and on and on. The inability to rapidly integrate a variety of systems at an affordable cost with readily available developer skills remains one of the biggest challenges facing companies today. Traditional integration approaches also impact growth of the content management market, as companies are forced to make do with outdated systems that cannot be replaced because the cost to do so is prohibitive. The burgeoning growth of social media platforms and the compelling business

“CMIS will ultimately become the foundation for developing next generation content collaboration and social computing applications.” – John Newton, Alfresco

benefits of cloud computing are stimulating a demand for more contemporary approaches to building integrated applications that bring true innovation to their organizations.

Earlier in this paper, we described CMIS as a “market-driven response to CM integration challenges.” It is clear that OASIS and the TC members are working to deliver capabilities that CM buyers and developers need to advance their CM practices and applications. We also noted that CMIS has the potential to be a game-changing standard because it opens new possibilities for vendors and internal IT organizations alike to deliver operational efficiencies, cost savings, improved customer experience, increased revenues, and content-driven innovation that has not been feasible to date. At the end of the day, companies simply need better approaches to integrating systems in order to succeed in the very unpredictable business climate of the 21st century. Business agility increasingly separates the winners from the losers, and agility is perhaps the biggest single benefit that CMIS offers.

We are encouraged by what we see in terms of CMIS as market reality – the support of leading vendors, coverage in the business and trade press, and attention from the analyst firms who educate IT buyers. Although support of the standard in software products is essential to adoption, we remind users that they play an equally essential role in ensuring widespread adoption by requiring CMIS support in the products they buy.

To Learn More About CMIS

Ask About CMIS Support

Organizations that are kicking the tires of new CMS systems and technology are very likely to have integration requirements incorporated into their technical designs. We strongly suggest making specific inquiries about CMIS support when investigating potential suppliers. Consider asking questions such as:

1. Do you intend to fully support CMIS? When can we expect support to be available?
2. Are you participating in the development and review of CMIS? Are you an active member of the OASIS Technical Committee developing it?
3. Where can I find out more about your support for CMIS? Do you have demonstration applications or materials available for downloading?
4. Do you have reference customers or applications that you can share?
5. What metrics and benefits can be found with a CMIS-aware system?

Participate in the Standards Process

Members of OASIS can participate in the further development and review of the CMIS specification. The specification itself and supporting documents can be found at the CMIS Technical Committee Web page at <http://www.oasis-open.org/committees/cmis>.

Toolkits, Sample Applications, White Papers and More

Vendors participating in the CMIS TC have put demonstration applications and documentation on their Web sites to allow interested users and developers to better understand CMIS, its capabilities, and their implementations based on CMIS. We encourage readers to visit the Web sites of the sponsors of this paper – [Alfresco](#), [EMC](#), [IBM](#), [Microsoft](#), and [Nuxeo](#) – and of all members of the OASIS CMIS TC.