TANEJA GROUP
TECHNOLOGY ANALYSTS

OPINION

Building the Cloud - It takes more than an API

October 2009

Charged with marketing buzz like we've never seen before, the potential of cloud technologies seems enormous, and undoubtedly poised to alter the IT landscape in a lasting way. Yet with an apparent cloud version of nearly every technology imaginable, the cloud often leaves the business customer wondering how to assemble the pieces. Across this space, the enterprise customer is slowly but surely realizing that at least part of the promise of cloud is in fact born upon *the ability to assemble the pieces*. Fundamental to this are APIs that are accessible in nearly every cloud technology, and in turn let customers wrap these technologies in their own applications, or interface them with nearly any other technology imaginable. But woe to the customer who thinks APIs are the only key to unlocking the potential of the cloud. Effective use of the cloud will not be determined by solely the API, but also the capabilities of the provider.

Iron Mountain has long been a key vendor to businesses around the globe – starting from a foundation of storing critical records and backup data, but built over time into a comprehensive set of digital information management services that range from data conversion to eDiscovery. In early 2009, on top of an aggressive services initiative, and alongside a cloud infrastructure built for the enterprise, Iron Mountain announced a cloud storage API. But more recently and more importantly, Iron Mountain has brought forward a foundation for a cloud "ecosystem" that reflects their realization that cloud is about more than just hosted storage, and more than just an API. This ecosystem, launched from the intersection of an Iron Mountain API, Iron Mountain information-centric digital services, and Iron Mountain's comprehensive Developers Program is bound to attract the interest of anyone looking to build cloud storage-enabled services with the business in mind.

What's in an API?

In the world of cloud, it seems at first glance that an API is everything. The ubiquity of the internet has in fact been determined by a psuedo API – in the form of how the basic HTTP protocol interacts with systems dispersed across a world-wide web of connectivity. While it took a while for the industry to come to this realization, interactions with web pages or any other system over an HTTP connection is made up of a set of single command transactions that can be very easily thought of as the most fundamental data management actions. For decades, the practice of information technology has revolved around creating, altering, and accessing digital information, and these core actions are well served by the core set of HTTP commands.

Today, encapsulating APIs inside of HTTP or similar protocols, and simultaneously making transactions "stateless" so that they can be easily distributed anywhere, is

common practice. In developer parlance, most such APIs make up either SOAP/WSDL – an XML and protocol pairing – or RESTful APIs that commonly make use of the HTTP protocol. More recently, we're seeing a new generation of internet-distributed *storage* services take flight, based around just such architectures. One such service is Iron Mountain's cloud storage Archive Services Platform, and it offers both approaches – a RESTful HTTP-based API as well as SOAP/WSDL.

Without a doubt, such APIs and protocols can make data storage in the cloud flexible and extensible. Actions are taken on a single individual object, but the service itself may be able to manipulate that object in powerful ways, and return different representations of it – the object itself, metadata, a subset of data, or more. Moreover, these core actions can be used to drill down into underlying layers of possibilities – e.g. an HTTP request can be made to identify possible actions that could be taken upon a file object. With such commands, data can be combined and manipulated by different systems and applications strewn across the Internet.

It takes more than an API

But in reality, what good is an API if it is just remote access to a bucket of storage that does little more than a fair effort writing and reading of data? In our view, not much, and that is why Iron Mountain's story is much deeper than an API.

In our view, the promise of the cloud rests in new capabilities, beyond the reach of traditional approaches, that can be delivered on top of data or applications in the cloud. But it is about doing so while *maintaining* or *superseding* key characteristics of traditional technologies.

Unfortunately, more often than not today's cloud services require a compromise in the key characteristics businesses expect their storage to have. Moreover, the business may not realize the compromises, because the conversation often turns to costs while glossing over a more important conversation about capabilities. One choice in cloud services may optimize cost, while sacrificing availability, security, or reliability. Another may optimize cost while stranding data in a bucket of cloud storage that is limited in capabilities when compared to other offerings on the market.

In contrast to how cloud is often delivered with a partial set of capabilities or a compromise in security, availability, reliability or other characteristics, Iron Mountain's Archive Services Platform steps up to the plate and follows through. Specifically, Iron Mountain has leveraged both their capabilities, and an "ecosystem" approach into a cloud offering that is enterprise-class, full featured and ready for the diverse business use cases bound to be found across the internet.

Harnessing Iron Mountain Capabilities

While the landscape is still evolving, it is clear to us that Iron Mountain has been building an enterprise cloud with a specific vision in mind. In our opinion, that vision is carefully matched to the requirements of the

business and leverages Iron Mountain's long standing expertise in delivering data management services with reliability, availability, and security. That experience is born of well over 50 years of enterprise information storage in deep record, tape, and digital repositories. The associated transportation, physical storage, indexing and retrieval processes may make the complexity of most best-in-class enterprise data center operations pale in comparison.

In turning to the cloud, Iron Mountain has carefully exercised and extended those capabilities. The result is an infrastructure that likely supersedes the highest expectations for security, isolation, data integrity, availability, and reliability. Iron Mountain today can boast a track record of 99.99% uptime across 10 years of operations and technology evolution, largely because of their sophisticated data centers built inside of "mountain" facilities that span *acres*. Moreover, when it comes to cloud storage Iron Mountain has been attentive to delivering capabilities with the business customer in mind – such as certificate-based authentication versus user ID SSL sessions. With a focus on business capabilities, and an infrastructure built to match, Iron Mountain may have a real leg up on the competition.

Building the Iron Mountain *Ecosystem*

But what Iron Mountain has really recognized, is that the cloud isn't about a storage service, nor about an API. Rather, the cloud is about the *ecosystem*. In the cloud, an ecosystem is the set of capabilities, services, access, and partners that transform a remote pool of storage into a *platform* for innovation that can deliver the deep reuse and extension of information. In turn, Iron Mountain has augmented both their capabilities and their API with a full-featured Developers Program. The resulting ecosystem is in fact at the heart of Iron Mountain's latest, and in our view, biggest announcement.

In the Spotlight: The Iron Mountain Developers Program

In September of 2009, Iron Mountain launched a formal developer program they had been building in stealth for several months. As a key component of building an ecosystem, Iron Mountain has created a program with the deep technical support required by ISVs seeking to develop services around this new information storage model, including free access to Iron Mountain cloud storage, extensive API and technical documentation, and personalized support. Partners and developers can sign up today by visiting Developer.IronMountain.com and completing an application. Upon acceptance, the developer will find they have access to a comprehensive support system that will help them understand how to leverage the *information-centric* capabilities of the Iron Mountain cloud, including rich metadata services, content indexing, search, retention and protection, and more. With the ISV and key business developers in mind, Iron Mountain has prepared the organization to deliver personalized, white glove service. Developers receive free access to resources, and customized support for their efforts that will undoubtedly help jumpstart the Iron Mountain digital cloud storage ecosystem.

O P I N I O N

Taneja Group Opinion

Often missed when it comes to cloud discussions, the cloud-enabled enterprise of tomorrow will require "depth" in their cloud services. An effective ecosystem is *required* before any storage service will have depth. Without depth, cloud storage easily becomes about little else than cheap storage on the far side of a potentially inconvenient connection. But with depth, the cloud can become a gateway to information services that supersede the capabilities of any single enterprise on the planet.

While Iron Mountain has been attentive to the capabilities of their API – serving up both SOAP protocols and RESTful HTTP-based APIs so that both ease of use and sophisticated capabilities are accessible – we are confident Iron Mountain hasn't stopped there. Our confidence rests on their continued announcement of ecosystem supporting programs, as well as their sophisticated portfolio of core capabilities that Iron Mountain possesses – *information-centric* capabilities that cross the gamut from information classification to media format transcoding. Iron Mountain's enterprise information management familiarity already shows in the strength of the fundamentals underlying their cloud storage – few offerings on the market today offer the automated ability to index content, WORM-lock data, slice data at multiple depths and create extensive metadata that can be associated with objects. As the Iron Mountain ecosystem matures, we're confident that richness and depth will only increase.

While the market for services in the cloud is rapidly becoming crowded, for these services to survive by serving business customers the foundation must be strong, with demonstrated expertise in servicing their information needs. But simultaneously, the service itself must grow to be deep and rich through the right programs that enable both the vendor and ISVs interested in delivering specialized services to excel. On both counts, it looks like Iron Mountain is getting it right.

Email Taneja Group:
What are the key factors driving your decisions around cloud storage?