Archive Database Instances - It's All About Options

Recently I wrote an article talking about how I think the Cloud is very murky - full of all sorts of abstract ideas and messages. My main point in the article is my assertion that the Cloud is ultimately about choices and flexibility.

As we began to work with a number of our customers on Cloud solutions of varying degrees, we started to ponder the big question of *"How do we get from here to there?"* in our own thought processes. Our thoughts instantly moved to the question - how can we produce maximum ROI with minimal impact?

We then examined common issues in the WMS space that we should tackle. At the very top of our list? ***Archive Servers***. Archive Servers are a necessary evil in the WMS world - where you need to have your production servers very lean. However, standing up a whole server (virtual or physical) along with the database necessary to virtually duplicate a production instance, but on a massive scale - is just very expensive. Think about the enormous cost of creating, maintaining, synchronizing schemas, and caring for these huge database servers that are, for the most part, idle 90% of the day. Let's look at the characteristics of an Archive Server in the WMS world:

- It is a Write Once/Read Many (more likely Occasionally) Database
- No transactions or concurrency issues
- Massive amounts of data over time
- Necessarily long lifetime of data
- Long-term data protection is critical - as contracts, customer compliance, or government compliance drives retention

Talk about a great candidate for the Cloud!

It did not take us long to make the decision to invest in this idea. We call it Simple Cloud Archiving (or SCA for short). With SCA, we also created solutions for some of the gaps in many archive processes. Features Include:

- A configurable archive log that allows you to understand exactly when something got archived - and where it went.
- NoSQL databases mean no schemas to synchronize. Database changes - no problem. Upgrades - no problem.
- A number of Local and Public Cloud options for storing the data.
- A standard Deep Storage/Near storage structure that allows you to keep data in very inexpensive Deep storage perpetually for, literally, pennies per month per Gb.

- A technology agnostic file format for Deep Storage. So if your requirements change in the future, or wish to re-create your archive elsewhere, you can.
- Use of Amazon AWS technology where appropriate:
    - You maintain control over the S3 (File Storage) and DB platforms
    - You can move the data to a different DB or platform
    - You maintain the security keys and Access Control Lists
    - Costs are very low for typical archive scenarios

Simple Cloud Archiving Supports:

- Full Cloud support: Deep Storage in Amazon S3, near storage a choice of NoSQL Databases
- Partial Cloud support: S3 for deep storage, Near storage by local NoSQL or less expensive SQL alternatives
- Local Deep/Near storage: Use file system for deep storage, NoSQL or less expensive SQL alternative for near storage
- Local Direct storage: You can send archive data directly to a less expensive MySQL database or use your own file system for deep storage.

The Scenarios looks like this:

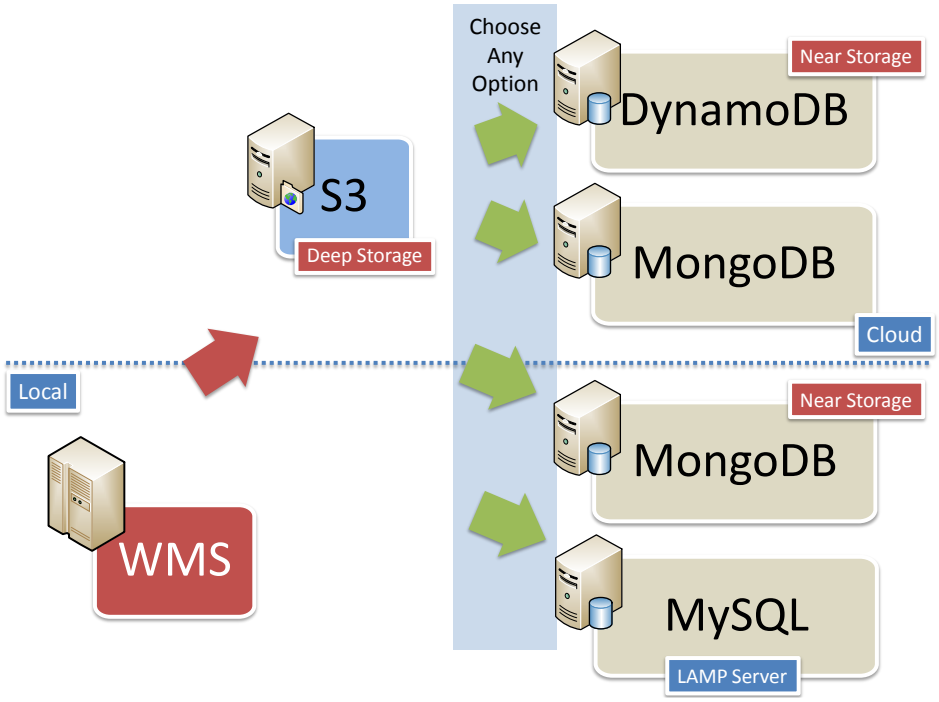

How does it all come together?

Customers of GetUsROI may use several new innovations created by our System Architects:

S3 For MOCA - MOCA components that can write and read data to/from S3.

NoSQL Access for MOCA - MOCA components that understand how to interact with some of the newer NoSQL databases - currently MongoDB (for Open Stack) and DynamoDB (for Amazon AWS)

Simple Cloud Archiving for MOCA - SCA pulls it all together in a set of capabilities that allow you to utilize the cloud for archiving as if it was part of the native system.

S3 and Amazon DynamoDB example:

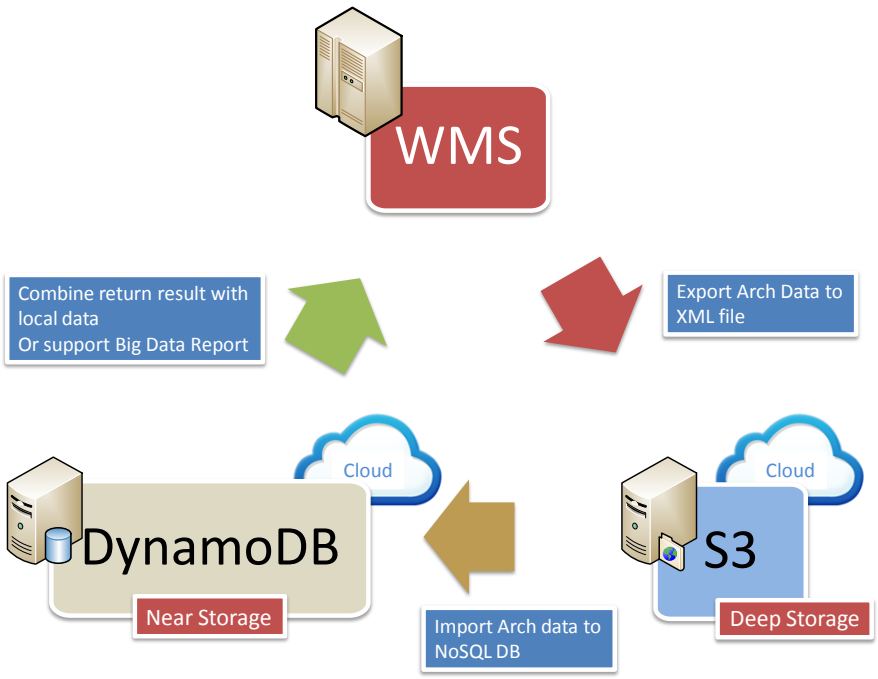

Our technologies are not for sale. However, they are available at no cost to our customers.

Visit Service Oriented Scripting or GetUsROI.com for more information, or call (262) 264-0477 to get started.