How does the 99.999% uptime design work?

Mercury will have an SSIS package that goes out and pulls data from each engine source – including the attribute engine. When it goes out to a system like UDM, it will look to see if UDM has completed a new run since the last update. If it has, it will then load the new data incrementally into a copy over table. This does not affect end users since they hit the production table in their views. Once the data is copied over and validated, the controller table will be updated and then one more important step will take place: there will be a synonym change that the user views point to that will change the production table over to the copy over table. This will be transparent to the end user, but because there will be a metadata view that defines what time updates occurred to each table, the standard business practice will be to include metadata information with each report run. Once the transition to the new copy-over table has occurred, the former production table will become the new copy-over table where data will be a 1 day lag table, which will serve as a failover backup in case something goes entirely haywire. While there should be no reason to suspect anything would ever cause the 1-day lag table to be used, it does serve as a model for disaster recovery in adhering to the six-sigma data availability model.

What is the attribute engine that is mentioned?

Essentially, any business category buckets that are not inherently placed on the files MSI receives, will be processed in stored procedures which will build crosswalk tables. These crosswalk tables will be imported into an attribute table for members, claims, revenue types, and providers. Examples would be Aid Category, Programs, PAK, etc. These can also include MSI derived fields. This keeps the ETL process separate from the attribute logic thus increasing efficiency and accuracy in load times, while building an object-oriented attribute engine.

Why are replicated tables needed?

Essentially, this creates six-sigma availability to business data. It also allows for a uniform way to track historical changes without having to store everything in separate tables. This is a true transactional model, but enhanced to support analytics through the attribute engine. Replicated tables also allow for more creative disaster recovery models as well as removing potential bottlenecks in disk usage and processing power by increasing flexibility for DBAs and database storage experts.

The best part?

It's ready to go April 30th, 2014 and can retrofit into the existing MSI model.