



Dynamic Services or an Easy Way to Change Your Model

*By: Sergey Babkin, Platform Developer
December 2008*

In the database world it's an issue familiar to every DBA: the application needs to update the database model, how do we do it? The business analysts come up with the new table layout, and now it's up to the DBA to implement it.

You would think that since SQL is full of ALTER statements, we would just figure out what to change, and do them. And then add the new tables and delete the ones no longer used. It's not so easy in practice. First, the data needs to be converted too. Say, if some table needs to contain denormalized information which was formerly split into two tables, it's not good enough to just create or alter this table. The information needs to be taken from the old tables and converted into the new format. Second, it's always convenient to keep the database schema on hand, for installation on the new machines. This means that the changes to the schema need to be done twice: once when editing the schema, and second time as a sequence of the ALTER statements. Add the issues with stopping the application for the duration of conversion, and the backup consideration on top of it, and things become complicated.

The streaming model meets similar challenges. But we at Aleri went beyond the simple SQL syntax to make the changes easy.

The foundation of everything is the Aleri architecture based on data consistency. The idea is that the data is not thrown away quickly. Instead it's preserved, and the contents of the derived streams may be recalculated afresh based on their input streams.

The Aleri Streaming Platform uses this property to automate the changes to the model. You don't need to pick out the differences manually and write the ALTER statements. Instead you just develop the new model as usual. After that you simply apply the new model to the running Platform. It will do the rest. It will identify the differences between the old and new model and modify the running instance accordingly, as illustrated on fig.1. The unchanged parts of the model are left unchanged. The data for the new or modified parts of the model will be automatically regenerated from their inputs.

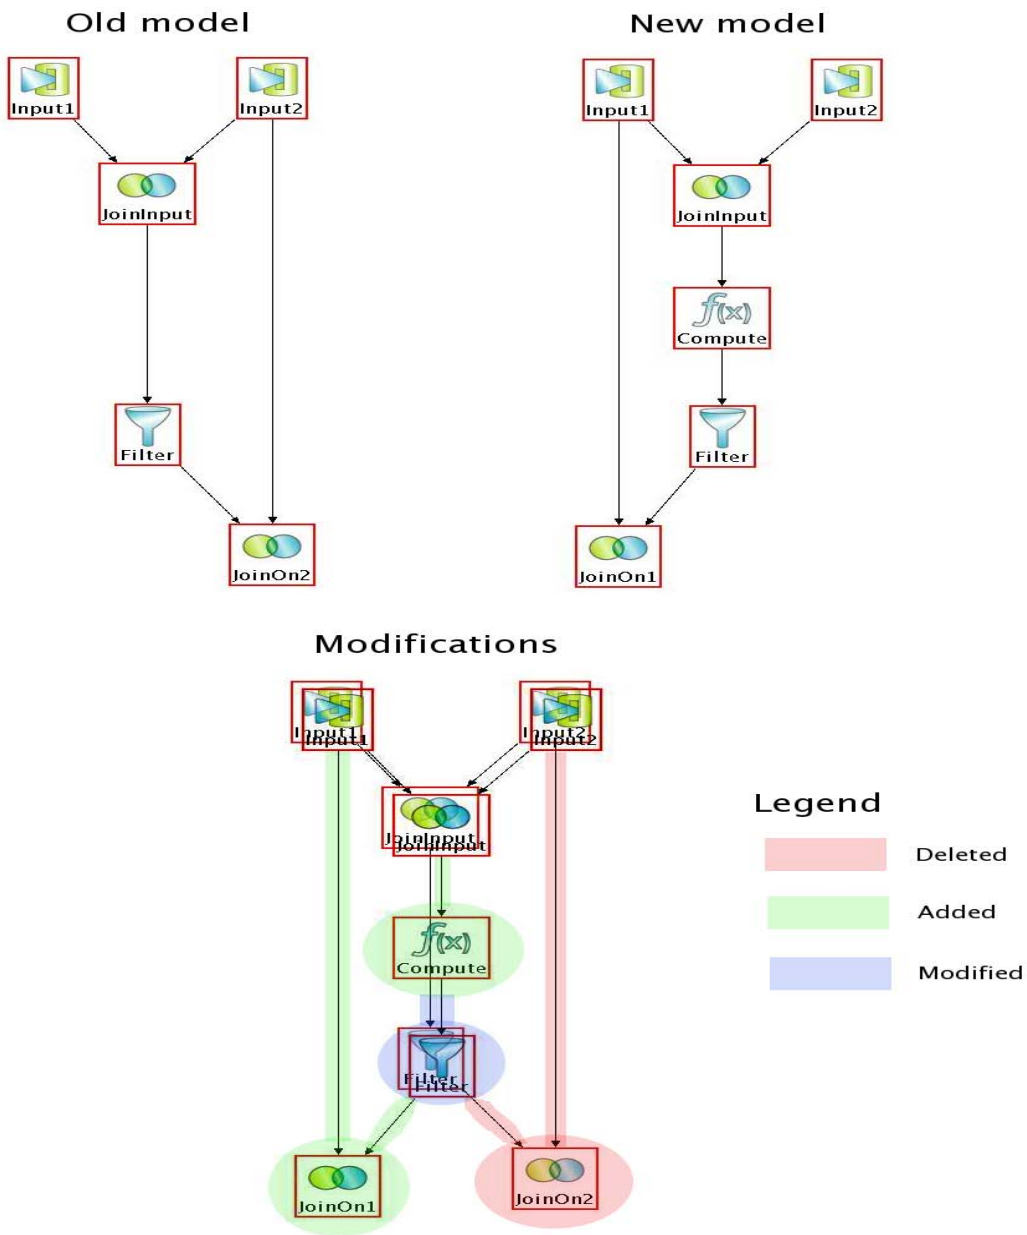


Figure 1: Identifying the modifications

Persistence of the data on disk is one of the distinguishing Aleri features. The dynamic modifications support it as well: the Platform automatically makes a backup of disk files before applying the new model. The backup files are written with compacted data, so they won't take as much space on disk as the full Log Store size.

And the best part is: the dynamic modifications are transparent for the applications supplying data to and consuming data from the Aleri Streaming Platform. The applications see only a temporary delay as the running model is being transformed.



What if a source stream gets changed? The source streams receive data directly from outside of the Platform, that can't be easily regenerated. Or in many practical cases they can, with a little effort. If a source stream is completely new then of course the only way to populate it is to load the new data from outside after applying the new model. This will cause the Platform to update all the derived streams accordingly, in a fully consistent way, as if the data had been there all along. But if a source stream is only modified or if it combines data from two old source streams then the contents can be converted directly when the new model is applied. This conversion is more difficult than for the derived streams. It has to be programmed manually, as for a common database schema change. But the Platform provides the tools: What is a streaming model? Exactly such a program to convert data from input format to the output format. So the conversion can naturally be described as a small temporary model, written in the same Aleri languages, using the same Aleri tools. This "conversion model" would get instantiated on the fly, do its work and then would be automatically discarded, leaving the newly formatted data for the main model.

To summarize, the answer to the question "What can be changed dynamically on a running Platform?" is: "Anything to anything".

UNITED STATES

Two Prudential Plaza, 41st Floor
Chicago, IL 60601
t +1 312 540 0100

430 Mountain Ave.
Murray Hill, NJ 07974
t +1 908 673 2400

UNITED KINGDOM

New Baltic House
65 Fenchurch Street
London EC3M 4BE
United Kingdom
t +44 (0) 20 7821 1110