

Build an Effective RFID Architecture

Companies that are deploying RFID systems will be quickly overwhelmed by a flood of data, unless they follow three core principles for building an operational data management architecture.

Feb. 2, 2004—You could water your lawn pretty quickly if you could use the fire hydrant on your street. But if you hooked up your garden hose directly to the hydrant, you'd wind up with a hose blown apart, a big pool of water and a very angry fire chief. If you deploy an RFID system by directly hooking RFID readers to your backend systems, the results will be similarly disastrous. What's needed is the right architecture.

RFID readers will deliver a gush of data. Our estimates are that pallet-, case- and item-level tracking, combined with data generated by RFID readers as items move within the enterprise, will increase the volume of data by 100 to 1,000 times today's levels. Even early RFID pilots have generated data volumes that would burst a few hoses. In its field trial, the Auto-ID Center used short-cuts—for example, storing only the last scan for complex analysis and decision making—to process the 15,000 events a second generated by an RFID reader.

Effective RFID implementations must borrow from the architectural principles developed for financial trading systems, process control and large-scale network management. Like RFID systems, these systems process terabytes of data, correct errors in real time, correlate events, detect patterns, re-organize and cleanse data and recover from faults—all in real time. RFID architectures should embrace three central principles of these systems.

1. Develop an operational data management architecture. An operational data management architecture is a software system that captures events at the “edge” of the enterprise, where operational activity occurs, rather than in the center, where business-oriented transaction processing occurs. These architectures act as a cache for transactional, enterprise systems, and then communicate via middleware once actionable information is identified.

Operational data management architectures are common in the financial industry—for example, trading systems collect market data in real time, execute trading strategies heuristically, identify trading opportunities, and forward the results to traders for action. These systems eventually generate the trades tracked by back-office systems, but billions of dollars are invested every year on the operational stages needed that precede the actual trade. This kind of data management and applied intelligence will become more prevalent as RFID adoption encourages companies to deploy intelligence at the point of operational activity.

2. Employ Savant “concentrators.” You could water your lawn using water from a hydrant if you could step down the pressure and flow. One way to deal with the gush of RFID data is to develop Savant “concentrators” that help control the flow and provide filtering and aggregation of EPC event streams (see figure below).

Savants are distributed software systems developed by the Auto-ID Center to act as the central nervous system of the EPC Network. A Savant takes data from an RFID reader, does some filtering, handles product lookups and sends the information on to enterprise applications or databases. Current conventional wisdom says to simply deploy a Savant and you're off; but Savants are just one piece of the puzzle.

Today's Savants provide interfaces to register a user EPC event receiver. A Savant concentrator is a software component that would run on a server and apply intelligence to control the volume of data. It begins to address the problem described by 7-Eleven CIO Keith Morrow, who said that RFID will require gulps "of bandwidth and access and storage such as we've never contemplated."

The concentrator is a type of operational architecture that starts with the Savant, which grabs EPC data from readers, stores the data in an operational database, checks for errors and operational anomalies, and applies business intelligence to identify critical business-level events, then propagates the data via the enterprise's middleware to another database, to a user or to both. This architecture has multiple layers of middleware, including Savants, message-oriented middleware, and enterprise application integration, as well as multiple data stores (an operational data store and enterprise data store).

3. Create pipelines for distribution. Developers often start small projects that are successful. But as the project scope expands, data volume grows far more rapidly than expected. The system can't scale and the project winds up becoming a disaster (think exploded garden hoses and pools of water). RFID deployments are particularly susceptible to this "success disaster" syndrome. Increasing numbers of RFID readers will lead to more streams of data, which can quickly overwhelm the system. The way to handle the load is through pipeline processing.

Once you have a Savant concentrator architecture in place, set up pipelines to handle the streams of data. Pipelines separate streams of EPC data to handle load and coordinate or process the data streams after they have been captured. To achieve additional scalability, you may need more than one concentrator distributed on a set of distributed machines, each one controlling a domain of RFID activity, distributed on multiple machines. For instance, let's say you have deployed EPC readers at a distribution center with 30 dock doors. Half the readers are tied into one Savant concentrator and half into another. You should create several pipelines from each concentrator. One will feed data into a local inventory database, one into a regional database, and one into financial applications.

The Auto-ID Center's Savants offer a way to manage data in very small pilots. But as rollouts begin later this year, companies will quickly find they are going to be overwhelmed by data, unless they follow these three basic principles and build an RFID architecture that will scale, not fail.

Mark Palmer has 15 years of experience in the operational data management and middleware industry as a chief technology officer at YouthStream Media Networks and director of product strategy at IONA Technologies. He is currently RFID technical evangelist with ObjectStore, a division of Progress Software. To comment on this story, click on the link below.

Copyright ©2005 RFID Journal, Inc. All Rights Reserved