

Search for:

- [Subscribe](#)
- [Search](#)

- [Subscribe](#)
- [Search](#)


- [News](#)
- [Insights](#)
 - [Editor's Notes](#)
 - [Expert View](#)
 - [Trends](#)
 - [White Papers](#)
 - [Ask The Experts](#)
- [Industries/Topics](#)
- [Events & Resources](#)
 - [Events](#)
 - [Event Recordings & Videos](#)
 - [Get Started](#)
 - [RFID Journal Glossary](#)
 - [RFID Journal Awards](#)
 - [Magazine Archive](#)
 - [FAQs](#)

Select Page

Cache Is King

I was talking recently to David Habib and Rob Rosales of Swisslog, the warehouse solutions company, about the impact of RFID data on their warehouse management system (WMS). I wanted to understand better how their WMS would use RFID data—in particular, how fast they needed to get data from an Electronic Product Code tag to the programmable logic controllers (PLCs) that control packaging systems.

“Oh, any response time faster than 1/100 second will do,” said Habib. “For example, our machines need to retrieve data on the height of the product for layering on a pallet. The PLCs need that information in real time to adjust the sizing. You can’t have the line stopping for four or five seconds while you get the data from the Web.”




The EPCglobal Network will allow companies to track goods in the global supply chain and get information about specific products—but it needs to be able to get the data related to a given EPC to participants in the supply chain in a timely manner. The problem is to architect a system that can scale and meet the response times required by various participants (see Leveraging the Internet of Things).

At MIT’s Auto-ID Lab, we are building an EPCglobal Network simulator called the Realm to investigate how we can meet these requirements. The simulator has been requested by EPCglobal’s Architecture Review Committee to provide a test bed for various architectures that might be adopted as standards for the network. It has a workflow engine that can script and generate many different “instances” of various supply chains that will involve different uses of the network.

In the case of warehouses, we need to be able to get data associated with EPC codes much quicker than, say, the response time of any message that goes out to the Web. We know from building Web-based information portals that the only way to get that response time is by caching the data—pre-positioning

the data into machine memory from secondary storage, such as disks or the network.



The caching of Web pages and database tables is one of the techniques we commonly use in building Web-based information systems. But the trick in caching is to know when you need to “shoot down the cache” and regenerate the data. Typically, a caching system spans several “layers” of the software architecture, and it can’t be retrofitted later.

We are implementing a caching system with the first release of the simulator, in first quarter 2006. It is based on the ideas of abstract data sources that allow a programmer to deal with databases, in-memory cache and XML BLOBs (collections of data stored as a single entity in a database) in a transparent manner.

The simulator should enable us to better understand the trade-off between disk-based databases, which are fairly slow, and in-memory caching, which is much faster. But database technology is advancing to include in-memory databases, so both technologies are converging. Implementing the latest technologies in our simulator will allow us to make informed decisions about the network and its response times.


John Williams is the director of MIT's Auto-ID Lab in Cambridge, Mass., and an associate professor and director of the Intelligent Engineering Systems Laboratory at MIT.



- ABOUT
- ADVERTISE
- CONTACT

FOLLOW US ON

- Follow
- Follow
- Follow
- Follow



© 2024 Emerald X, LLC. All Rights Reserved

ABOUT CAREERS AUTHORIZED SERVICE PROVIDERS Your Privacy Choices TERMS OF USE PRIVACY POLICY