

# Towards an Enterprise Location Service

Christof Roduner  
Institute for Pervasive Computing  
ETH Zurich  
8092 Zurich, Switzerland  
roduner@inf.ethz.ch

Christian Floerkemeier  
Institute for Pervasive Computing  
ETH Zurich  
8092 Zurich, Switzerland  
floerkem@inf.ethz.ch

## Abstract

*Radio Frequency Identification (RFID) technology has a lot of potential to automate business processes and improve visibility across the supply chain. To realize the full potential of RFID, an IT infrastructure is required that manages readers, filters and aggregates raw RFID data, but also facilitates data exchange among the supply chain partners. In this paper we first provide an overview of the EPC Network, an architecture developed by the Auto-ID Center and its successor organisation EPCglobal, that addresses the above requirements. We then propose the concept of an Enterprise Location Service that captures all location information about products and assets within the enterprise to enable what we call business process debugging.*

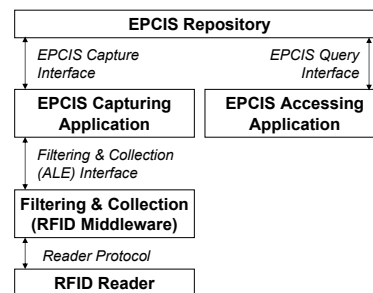
## 1 Introduction

Radio Frequency Identification (RFID) technology has recently seen growing interest not just from the research community, but also from a wide range of industries such as retail and pharmaceutical. In these domains RFID technology holds the promise to eliminate many existing business problems by bridging the economically costly gap between the virtual world of IT systems and the real world of products and logistical units [5]. Common benefits include more efficient material handling processes, elimination of manual inventory counts, and the automatic detection of empty shelves and expired products in retail stores [1, 7]. Many benefits commonly associated with RFID require however data sharing across the supply chain [1]. In order to realize those benefits, an infrastructure for the inter-organizational sharing of data collected by means of RFID is needed in addition to an RFID middleware that manages readers and aggregates raw RFID data. The *EPC Network*, originally proposed by the Auto-ID Center [8] and further developed by EPCglobal Inc., is currently the predominant standardization effort that addresses all of those issues.

In this paper, we will provide an overview of the state of the art of the EPC Network. We will highlight how the EPC Network enables users to track and trace objects in a business context. The main contribution of this paper is the concept of an additional component in the EPC Network – termed *Enterprise Location Service*. It aggregates all location information about the whereabouts of products and assets captured within the enterprise and makes these location data available to services that interpret the location information in the appropriate business context.

## 2 The EPC Network

At the heart of the EPC Infrastructure lies the Electronic Product Code (EPC) [3]. Unlike the EAN.UCC number encoded in barcodes, an EPC can be used to identify products not just at a class-level, but at an item-level (i.e., two identical bottles of a softdrink are assigned different EPCs). EPCs are stored on EPC-compliant RFID tags using the air protocols also defined by EPCglobal [2].



**Figure 1. Roles and interfaces of the EPC infrastructure.**

Building upon the Tag Data Standard [3] and the air interface protocols, the EPC infrastructure consists of a number of *roles* and *interfaces* that need to be deployed within a

company in order to process EPC tags in an EPC-compliant way [4]. The EPC Network specifications do not define the individual components, but rather roles and interfaces that must be implemented. These roles (in bold type) and interfaces (in italic type) are illustrated in figure 1 and explained in the following sections.

## 2.1 RFID Reader

The RFID reader is responsible for interacting with EPC tags within range of its antenna on the one hand and a host application on the other hand. Communication with the latter is done through the Reader Protocol interface, which provides commands to inventory tags (i.e. to read its EPC), to read data other than the EPC from tags, to manipulate the EPC on a tag, to write data other than the EPC, and to lock or kill a tag.

## 2.2 Filtering and Collection

The Filtering and Collection role is responsible for controlling the RFID readers and is therefore often referred to as RFID middleware [6]. In particular, this role is responsible for the following tasks:

*Reader Coordination:* As soon as several readers cover the same physical area, radio-frequency interference can occur and decrease system performance. Readers must thus be coordinated in order to minimize these effects. *Access Coordination:* In an environment with multiple applications, the Filtering and Collection component mediates between the different clients in case of concurrent access. *Data Filtering and Aggregation:* The raw reads output of readers is generally not suitable for client applications. It must be transformed into event streams that are easier to process by clients. In many cases, applications aren't interested in all EPCs detected by a reader, but rather in the EPCs of a specific object class or simply in the number of EPCs of a certain type, which requires filtering and aggregation capabilities. Another example are reports summarizing all EPCs seen over a certain period of time with all duplicate reads eliminated. *Logical Readers:* A single reader is often not sufficient to reliably cover the entire physical area relevant to a business process. A loading dock will for example have to be equipped with several readers that should be exposed to client applications as a single logical reader. This grouping is done by the Filtering and Collection component.

This functionality is exposed through the Application Level Events (ALE) interface, which can provide data both in a synchronous and in an asynchronous way (standing queries).

## 2.3 EPCIS Capturing Application

The EPC Information Service is the main element in an EPC-enabled information system. It allows EPC-related data to be shared within and across organizations, thus increasing the visibility of tagged items along the supply chain or in some other business context.

The EPCIS Capturing Application's main task is to transform the events obtained through the Filtering and Collection interface into EPC-related business events by being aware of the business process in which the EPCs were collected. While the former events do not contain much more than the EPCs of the observed objects, the latter are enriched with business context. An event obtained through the ALE interface might read for example: "At location L, in the time interval T1 - T2, the following case-level EPCs and one pallet-level EPC were read." – Since the EPCIS Capturing Application is aware that the readers are located in a place where pallet commissioning takes place, it generates an EPCIS level event that reads as follows: "At location L, at time T2, it was confirmed that the following cases were aggregated onto the following pallet."

If necessary, it is also the EPCIS Capturing Application's responsibility to keep track of multiple Filtering and Collection events and transform them into a single EPCIS event. Moreover, the Capturing Application will acquire any additional information such as interactive user input or data from business information systems that are needed to infer the precise business context.

## 2.4 EPCIS Repository

While the events that are available from the Filtering and Collection interface are of transient nature, EPCIS events are normally sent to a database for long-term storage since they contain historical information that relates to a business process. An organization adopting the EPC Infrastructure will thus have to deploy a component acting as the EPCIS Repository. It is expected that this role will be incorporated into standard enterprise software packages. An EPCIS Repository implements two interfaces: First, the EPCIS Capture Interface that accepts incoming events from Capturing Applications. Second, the EPCIS Query Interface, from which historical data can be retrieved using a query language similar to SQL, though far less expressive.

## 2.5 EPCIS Accessing Application

When an application makes use of EPCIS events in a repository, it acts in the role of an EPCIS Accessing Application. Typical examples of such applications are warehouse management systems or other systems carrying out business processes relying on EPC data. Even though

the three EPCIS roles (repository, capturing application, accessing application) are introduced as separate building blocks in EPCglobal's specifications, it is assumed to be common for a single application such as a warehouse management system to include all three roles.

### 3 Enterprise Location Service

As we saw in the previous section, the primary objective of current EPCIS specifications is to facilitate the exchange of RFID triggered business events. This is important because many benefits associated with RFID result from the data exchange among supply chain partners. In this section, we outline the need for an additional component within the EPC Network that does not target the benefits of RFID resulting from the exchange of RFID among the supply chain, but from the use of RFID within the enterprise itself.

#### 3.1 The need for a repository of all location data captured within the enterprise

The current approach of the EPC Network is to store location data within the appropriate business context. More specifically, there is no repository for simple RFID readings indicating no more than that a certain EPC was detected at location  $L$  at time  $T$ . Instead, RFID readings in the form of ALE level events are always interpreted in the context of a business process before they're sent to an EPCIS repository.

While this approach is useful in some domains, we believe it has its shortcomings in others. Since the RFID data that are generated but not handled by any capturing applications are effectively lost, there is no way to trace back the cause of an exceptional event.

For example, in the case of a safety incident, an incomplete shipment or a product quality problem, the data available from EPCIS repositories might not be sufficient to efficiently find and deal with the cause. Because these cases are not part of everyday business operation, they're not mirrored in an EPCIS repository. In order to take advantage of RFID readings that might be useful when examining such situations, a dedicated capturing application would be required. However, this is not a viable approach for most users: First, it is too costly to provide an application when the data it captures is needed only rarely. Second, it is impossible to foresee many of the exceptional situations in which RFID readings are a valuable basis for problem analysis and solution.

Apart from the main issue of storing only business events rather than location data, there are also other reasons why, in certain scenarios, the EPCIS alone is not sufficient to take full advantage of an enterprise's RFID infrastructure:

- In the current architecture framework, there can be an arbitrary number of EPCIS repositories potentially containing records on a certain EPC. It is therefore impossible for an application to know from where it can retrieve the complete location trace of a tagged object. EPCglobal is planning to propose a so called Discovery Service that would allow applications to obtain a list of all repositories holding information on a given EPC. However, detailed discussion of such a service has not begun yet.
- Assuming the availability of a service that could point an application to all relevant EPCIS repositories, the burden of contacting them is still placed on the client. This generic functionality should clearly be offered as a service.
- Long-term storage of EPCIS events will be handled very differently by the various repositories. Some might retain data only for a very short time, whereas others might keep them for years. It is thus difficult or even impossible for an application to reassemble the location trace of an EPC because some repositories might already have pruned the data that are needed.
- The focus of a typical system incorporating the EPCIS Repository role is generally on the business aspects. It is not optimized to deal with pure location information in an efficient way.

We therefore propose a single repository that captures all location data generated within the enterprise, provides consistency and makes these location data available to various services within the IT enterprise architecture that interpret the data in a given business context. Rather than keeping a location model individually up-to-date within each application, the applications would access the location service to retrieve location information.

#### 3.2 Example Scenarios

In this section, we will present three application scenarios in which a company would benefit from such an enterprise-wide location service.

##### 3.2.1 Quality Problems

When a product is returned to the manufacturer because of poor quality, it is often desirable to track down the problem in the production process and the number of other products in the same batch. A history of all objects that were in physical proximity of an item when it was produced can be of great value. That way, a pharmaceutical company could for example detect what exactly caused the impurity in a certain production batch.

### 3.2.2 Audits

Many companies must make sure that their production processes adhere to regulations imposed by law or internal standards. It might for example not be allowed to keep one chemical substance in close physical proximity of another, highly reactive substance. In order to ensure compliance with regulations, it is common practice to conduct audits. In these cases, it might be useful for an auditor to have historical location data at hand. By analyzing samples from the location history, dangerous conditions could be detected for example in a safety audit or environmental audit.

### 3.2.3 Business Process Reengineering

A location service could also assist enterprises in analyzing and improving their existing business processes by providing the data needed to understand precisely how existing practices perform. In this case, data available from EPCIS repositories will be of little value since it only reflects what is already possible with the existing setup. However, in order to find novel ways of carrying out a business process, uninterpreted location data is needed.

In summary, an enterprise-wide location service allows for what we call "business process debugging": It provides a means to conduct ex-post evaluation of processes. Since the data available from such a location service are neutral with respect to business context, new applications can easily be introduced and get full access to historical data.

## 3.3 Requirements

As a first step towards an enterprise-wide location service, we will discuss some of the key requirements in this section.

*Integration of enterprise-wide location infrastructure.* A location service must integrate the data available from the company's location infrastructure. This includes RFID readers, but also dedicated positioning modules equipped with a GPS receiver for example. Clients must be provided with an interface for easy access to aggregated location information.

*Persistency.* Historical data can prove useful for applications that have not been foreseen. Incoming location data should therefore be stored in a persistency service.

*Efficiency.* The amount of location data we envision to collect can become difficult to deal with. A location service must be efficient with respect to throughput (i.e. volume of incoming data), query processing and disk usage.

*Location model.* In order to support a broad range of application needs, it might be useful to expose location data both through a symbolic and a geographic location model.

*Querying and real-time notification.* Some applications might need instant notification when a certain condition

holds. A location service should offer both one-time queries and standing queries.

*Access control.* The data collected by such a location service reveal many details about a company's business processes. It is thus highly critical and must be protected from unauthorized access.

## 4 Conclusion

This paper provides an overview of the EPC Network, a networking infrastructure, developed by the Auto-ID Center and its successor organization EPCglobal, to manage RFID readers, to filter raw RFID reads, and to convert RFID reads into appropriate business events. In addition to the various components of the EPC Network, we see a need for an enterprise location service that captures all location data including RFID data captured within a company. The persistently stored location data are then made available to applications throughout the IT enterprise architecture. We believe that such a service is essential to facilitate what we termed business process debugging.

## References

- [1] K. Alexander, T. Gilliam, K. Gramling, M. Kindy, D. Moogimane, M. Schultz, and M. Woods. Focus on the Supply Chain: Applying Auto-ID within the Distribution Center. Technical Report IBM-AUTOID-BC-002, Auto-ID Center, 2002.
- [2] EPCglobal. Class 1 Generation 2 UHF Air Interface Protocol Standard Version 1.0.9, 2005.
- [3] EPCglobal. EPC Generation 1 Tag Data Standards Version 1.1 Rev.1.27, 2005.
- [4] EPCglobal. EPCglobal Architecture Framework Version 1.0, 2005.
- [5] E. Fleisch and M. Dierkes. Ubiquitous Computing: Why Auto-ID is the Logical Next Step in Enterprise Automation. Technical Report STG-AUTOID-WH-004, Auto-ID Center, 2003.
- [6] C. Floerkemeier and M. Lampe. RFID Middleware Design – Addressing Application Requirements and RFID. In *Proceedings of sOc-EUSAI 2005 (Smart Objects Conference)*, Grenoble, Oct 2005.
- [7] A. Krohn, T. Zimmer, M. Beigl, and C. Decker. Collaborative Sensing in a Retail Store Using Synchronous Distributed Jam Signalling. In H. W. Gellersen, R. Want, and A. Schmidt, editors, *3rd International Conference on Pervasive Computing*, volume 3468 of *Lecture Notes in Computer Science*, pages 237–254, Munich, May 2002. Springer-Verlag.
- [8] S. Sarma, D. L. Brock, and K. Ashton. The Networked Physical World – Proposals for Engineering The Next Generation of Computing, Commerce & Automatic Identification. Technical Report MIT-AUTOID-WH-001, MIT Auto-ID Center, 2000.