Whitepaper:

RFID Infrastructure – A Technical Overview

Abstract

Geared to general interest readers and entry-level practitioners, this paper takes an in-depth look at the elements of an RFID infrastructure, including its tags and readers and protocol layers, and the roles they play. History has taught us that a technology cannot and will not be deployed pervasively and globally without a robust set of standard protocols specified between these entities.

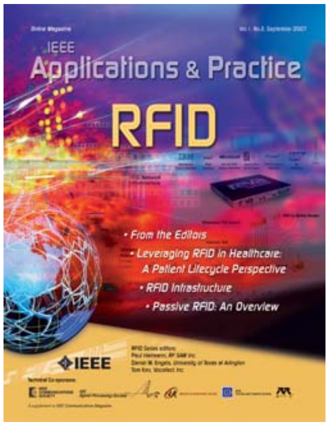
IEEE
Applications & Practice
RFID
• From the Editors
• Leveraging RFID in Healthcare: A Patient Lifecycle Perspective
• RFID Infrastructure
• Passive RFID: An Overview
IEEE

Originally published in IEEE Applications & Practice, September 2007, Vol. 1, No. 2

Introduction

Geared to general interest readers and entry-level practitioners, this paper takes an in-depth look at the elements of an RFID infrastructure, including its tags and readers and protocol layers, and the roles they play. History has taught us that a technology cannot and will not be deployed pervasively and globally without a robust set of standard protocols specified between these entities.

First generation RFID systems were deployed at a single site usually with a handful of readers communicating over dedicated links to one or a few application servers. Such architecture (see Figure 1) works fine for pilot and proof-of-concept projects, but does not scale up readily to enterprise implementations with more readers, more sites and more applications.

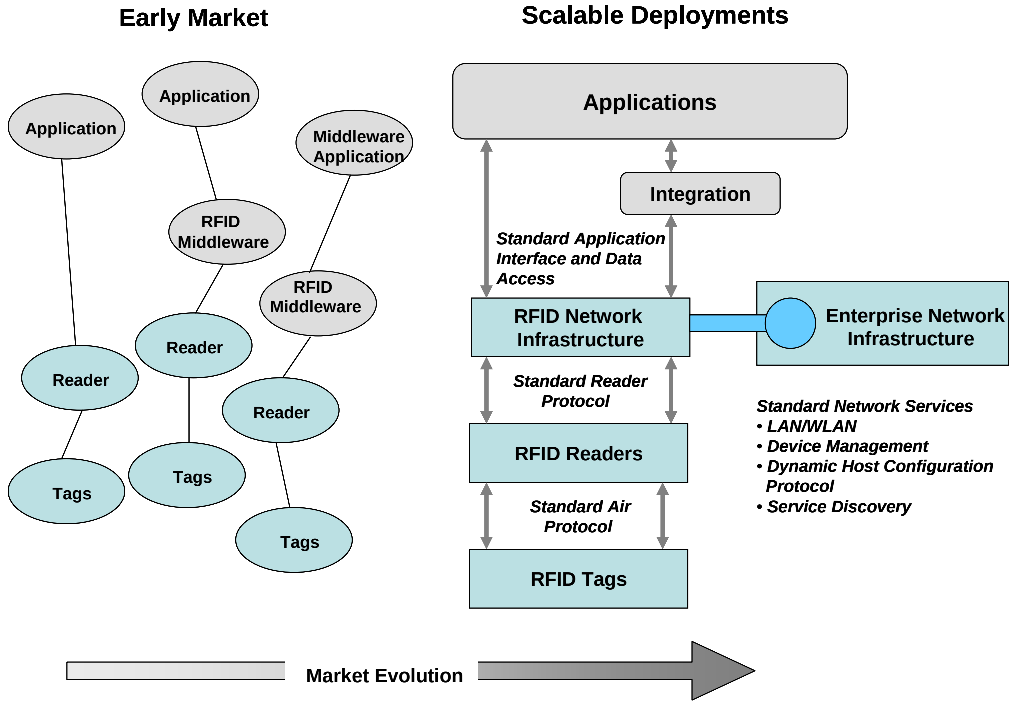

Figure 1: Evolution towards an RFID infrastructure

On a global scale, RFID readers could easily become one of the most densely deployed and numerous network devices in the world, with many analysts predicting over 100 million RFID readers connected globally within 10 years. But, before even contemplating a vision of ubiquitous RFID deployment, today's enterprises are experiencing scaling challenges for even very modest deployments. In fact, the issues associated with rolling out, managing and operating 5 or more readers at more than a couple of facilities are significant and most IT departments are not equipped to support field trials involving server-based RFID middleware for extended periods of time. Connectivity at such scale challenges a network's ability to absorb compounding demands. From radio frequency (RF) contention and bandwidth management to data management, back office integration and operational support—

the proliferation of RFID technology is an escalating RF, network and data management challenge.

Similar situations have unfolded many times before with technologies such as switched LANs, WiFi LANs and storage-area networks. Initial deployments of these technologies were standalone and relied on middleware, but as they took hold in corporate networks they rapidly evolved into standards-based infrastructures operating at a much greater scale.

The operating frequency of the reader—from 10 kHz to 5.8 GHz—as well as the method of coupling the signal to the tag and the range are the key differentiating criteria for RFID systems. Coupling can be via electric, magnetic or electromagnetic fields, with the range varying from a few millimeters to hundreds of meters. Our article only discusses passive UHF RFID systems.

Passive UHF RFID technology offers the best read range, read rate performance, readability through a wider range of materials, and all this at a very cost-effective manner. Due to these benefits, it is the most prevalent RFID technology being deployed across all industries (e.g., retail supply chain, manufacturing, food, pharmaceuticals, consumer electronics, etc.). In addition, most of the current industry standardization efforts are focussed on this technology. These systems operate from 860 MHz to 960 MHz, their tags employ electromagnetic coupling and backscatter, and the tag read range is about 5 meters.

Infrastructure Elements

The RFID infrastructure consists of the elements that manage the devices and tag data. Consumers of the data are the client network elements (typically end-user applications). The network elements between the tag and the clients form the conduit that transports tag data to the applications, and convey tag operational commands to the RFID devices. At a minimum, the RFID infrastructure (see Figure 1, again) comprises tags, readers, RNCs (Reader Network Controllers) and applications running for example, on enterprise servers. In addition, other devices could also be in the network such as RFID/bar code readers, I/O devices (such as electric eyes, light stacks and actuators), bar code/smart label printers and applicators.

Typically, a reader transmits an RF signal in the direction of a tag, which responds to the signal with another RF signal containing information identifying the item to which the tag is attached, and possibly other data. The tag may also include additional field-writable memory store, and integrated transducers or environmental sensors for providing data such as the temperature or humidity of the environment. The reader receives the information and provides the tag data to the RNC which may do further processing before sending the data on to the applications.

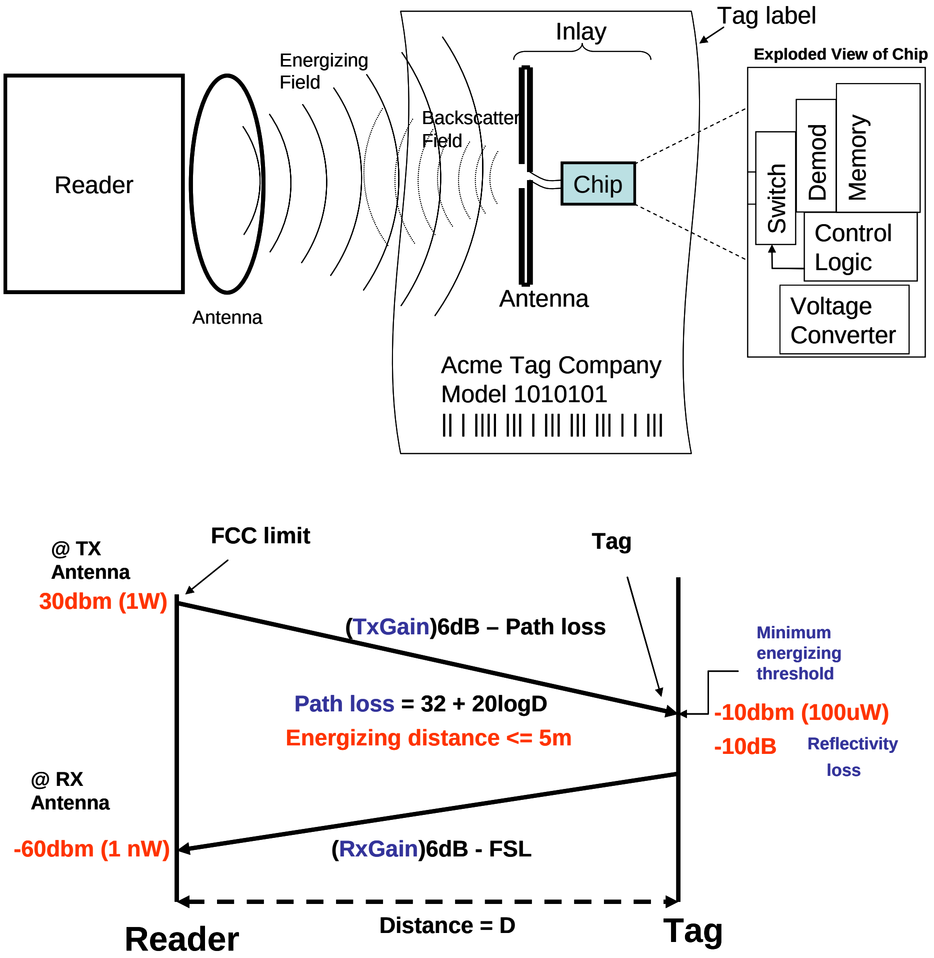

Figure 2: Reader Tag Interaction

There are three types of tags. At its simplest, a tag includes a small antenna connected to a microchip. This tag is *passive* in the sense that it has no integrated power source, such as a battery. Instead, the tag harvests the electromagnetic energy emitted by the reader, converting that energy to the DC power for operating the microchip. The tag then transmits information to the reader by backscattering part of the energy it receives. Figure 2 illustrates the interaction between a reader and a passive tag. As shown, a tag label consists of an adhesive label that is embedded with a tag inlay (the tag chip plus printed antenna).

A *semi-passive* tag has a battery to operate the microchip, and also uses backscatter to communicate with the reader. Its range is no longer limited by the need to power the tag from the RF field, but by the sensitivity of the reader's receiver (which may be able to receive very weak signals at -90dBm or lower). Such tags have a considerably longer read range than passive tags.

The third type of tag is *active*, powered by a more potent internal battery so it can actually transmit an RF signal in response to the reader, rather than backscattering the reader's signal. This enables a broader range of functions, such as tag-to-tag communications and security. The internal battery may also power built-in environmental sensors and maintain data and state information dynamically in an embedded memory in the tag.

Figure 2 also illustrates the path loss on the forward path (reader-to-tag) and the reverse path (tag-to-reader) for a passive tag. The maximum range over which the reader can communicate with the tag depends on the transmit power of the reader, the environment through which the RF signal travels, the presence of interference, the minimum energizing threshold of the tag, and the receive sensitivity of the reader. Losses due to signal attenuation and multipath interference reduce the range. Attenuation is low or negligible for gases in the atmosphere, such as nitrogen and oxygen, and also for paper, cardboard and certain plastics. Materials like metal and liquids have a stronger attenuating effect depending on their thickness.

The *Air-Protocol* defines the signaling layer of the communication link, the reader and tag operating procedures and commands and the collision arbitration scheme for identifying a single tag in a multiple-tag environment. This last process is known as *singulation*.

An RFID reader typically has an RF front end that serves one or more antennas, an RF signal processor, an air-protocol processing engine that implements the air-protocol message decoder/encoder and state machine and algorithms, and the network interface processor to communicate with the upstream network elements. Some readers have support for digital I/O ports. These ports are used for connecting serially to sensors, triggers or other controllers.

Reader designs cover a wide spectrum based on different factors:

- Number of antennas: multiple (typically 4 to 8) antennas or a single integrated antenna
- Processing complexity: data processing includes business intelligence (sometimes referred to as "smart" readers), or just RF intelligence (sometimes referred to as "thin" readers)
- Tag access functions: some perform all air-protocol operations (read, write, lock and kill tags), others just inventory the tags
- Connectivity: Ethernet, serial, or wireless
- Number of digital I/O ports: none or multiple (typically 1-4)

The Reader Network Controller (RNC) plays the role of the RFID infrastructure layer. It resides logically above the reader layer as an extension of the enterprise network. It transforms a collection of autonomous readers and devices into a reliable and scalable network. RNC functionality includes real-time adaptive control and management of readers and devices, location-aware tag and sensor data processing, and standards-based data services for the applications using the RFID data. This functionality could be implemented in standalone hardware, as standalone software running in an enterprise server, as software integrated with enterprise middleware or directly with RFID-enabled applications. The choice for deployment would primarily depend on the complexity of device management operations and control, the data load and processing requirements and the application services requirements.

Infrastructure Functions

The infrastructure comprises three interlocked communications paths: data processing (the Data path), device management (Management path) and device control and coordination (Control path), as shown in Figure 3).

The Data path refers to the tag and sensor information collected by the readers and forwarded to the Reader Network Controller and applications. With the advent of sophisticated air-protocols like the second-generation UHF Gen2, and deployments of large number of readers, the need for reader control and coordination in the architecture becomes important. Likewise, with diverse types of devices deployed at a facility or in an enterprise, device management and monitoring (the Management path) becomes very important too.

For a bit of background, "UHF Gen2" refers to the communications protocol that defines the physical and logic requirements for passive UHF readers. It is shorthand for EPCglobal Class-1 Generation-2 protocol.

Broadly speaking, an RFID infrastructure must take care of:

- Reader operations
- Tag data processing
- Device management and monitoring
- Inter-enterprise and intra-enterprise tag data and event dissemination

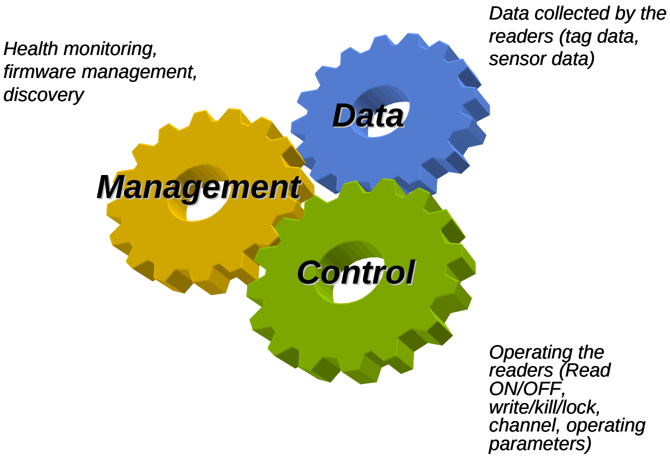
Health monitoring, firmware management, discovery
Data collected by the readers (tag data, sensor data)
Data
Management
Control
Operating the readers (Read ON/OFF, write/kill/lock, channel, operating parameters)

Figure 3: RFID infrastructure network communications paths

Reader Operations

A reader typically performs either inventory or access operations on a tag population. *Inventory*, as the name suggests, identifies a population of tags using a sequence of air-protocol commands. Using a singulation algorithm, the reader isolates a single tag reply and reads the EPC memory contents from the tag.

Access is used to describe the further operation of communicating with (reading from and/or writing to) other memory regions on a tag. Similar to the inventory operation, access comprises multiple air commands.

Reader operation deals with controlling and coordinating the readers to maximize system-wide RFID performance. Although the communication between a reader and tag is local, the interference impact due to that local communication is global. This means that a single reader's operating parameters that maximize the performance of the local communication between the reader and its set of tags may not translate to a "global" (that is, a system) performance maximum.

Maximizing a single reader's performance does not necessarily maximize system performance.

The system-wide tag inventory, access rate and latency are key performance parameters for an RFID network infrastructure. With multiple readers, system performance can be affected by reader-to-tag and reader-to-reader interference.

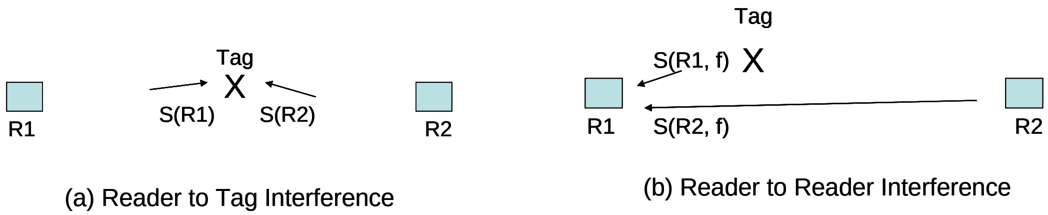

(a) Reader to Tag Interference

(b) Reader to Reader Interference

Figure 4: Types of interference

Reader-to-tag interference occurs when multiple readers simultaneously energize the same tags, which confuses them and prevents them from being read. In Figure 4(a), when the difference between the signal strengths received from Reader 1 and Reader 2 |S(R1) – S(R2)| is less than the tag's tolerance margin, the tag gets confused. A filter in the tag can reject some interference but, currently, a 6 to 15 dB difference in power level between two colliding reader signals, known as the tolerance margin, is required for the tag to respond to just one reader. Filtering and threshold technology improvements in the tag, not the air protocol, can improve the tolerance margin, but this has other drawbacks:

- Vendor dependence—reliance on proprietary tag technology is increased in what should be an open system
- Decreased reader-to-tag interference margin allows stray signals to corrupt the reader-tag interaction
- Silicon costs may be associated with the required circuits

Reader-to-reader interference occurs when a reader picks up another reader's transmit signal at, or near, the same frequency. In figure 4(b), during the interaction with the tag, Reader 1's receive filter has its band-pass set to accept signals at a particular frequency, f. If a signal at the same frequency, f, is received from another reader, R2, and, that signal S(R2, f) is much stronger than the tag response S(R1, f), R1 may not be able to decode the tag's reply. R2 (the interferer) has an advantage over victim tag because its signal decreases with distance by $1/d^2$ versus the $1/d^4$ attenuation of the passive tag's response.

The Reader operation control strategy involves three subcomponents:

- Physical: the forward and reverse link parameters between reader and tag
- Tag inventory: the singulation strategy used
- Data access: the air protocol operation sequence

A successful strategy involves dynamic control of all three of these control components in response to real-time events (such as tag movements at the readers, for example); both external RF and self-interference; and regulatory constraints.

Tag Data Processing

Since the introduction of the Gen2 Air Protocol in 2005 individual reader performance has improved considerably (read rates are near 100%). As the price of readers goes down, end users can economically deploy larger numbers of readers. In return, they expect a much richer set of information from their systems. This includes fine granularity and information on the precise location of the tags, an indication of the direction of motion for tags in transitional locations, and accurate tag group associations (that is, which items comprise the packing case, which cases comprise the pallet, and so on).

However, as end users begin to scale their Gen2 deployments, they experience one of the challenges of second-generation RFID implementations. This phenomenon, which does not yet have a consistent name (it's been called "unwanted reads," "unintended reads," "cross reads," and "wrong positive reads"), relates to the most fundamental difference between RFID and optical-bar-code technology that it is replacing: the inability to know precisely which among the tags that a reader reads are the intended tags, and which are not. (This problem doesn't exist with barcode scanning, which reads only a single item at a time. The operator also knows exactly what item is read by virtue of having aimed the optical scanner at it.)

The problem of unwanted reads can be illustrated by a simple example. In Figure 5 two adjacent dock doors are each set up with an RFID reader. Two antennas face the same direction, but cover two different doors. Antenna A1 is at Door 1, and A2 at Door 2.

Confusion begins when a pallet of tagged goods approaches Door 2. The antenna facing right on Door 1 has a very wide field-of-view in the area approaching Door 2. And with the increased sensitivity of Gen 2 tags, it is not uncommon for the antennas on Door 1 to see some of the tags that move through Door 2.

There are four interesting points of time during the transit of the pallet: time periods a to b and c to d, when only Antenna A1 is able to read some tags from the pallet; time periods b to c, when both antennas A1 and A2 are able to read the tags from the pallet

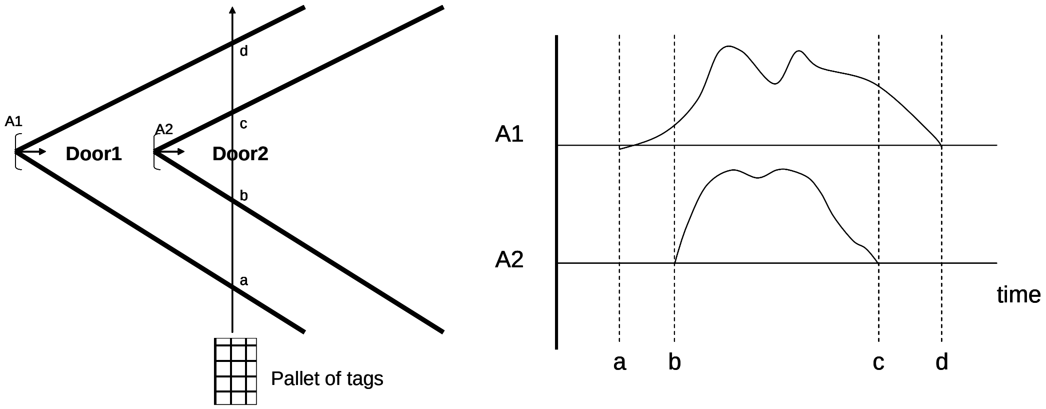

Figure 5: Fields of view of antennas A1 and A2 on dock doors overlap [left]. Tag read count by antennas on the moving pallet at different times [right].

The reads made by A1 are the unwanted reads. Their consequence is that the application using these readers will observe the same tags at both doors and will not be able to infer the actual door through which a tag passed. Nor can the readers make the tag group association—the pallet, case, or item the tag is on. The addition of more dock doors and readers, which may be necessary for the operations, just compounds the problem.

Reads of tag by an antenna does not necessarily determine the actual tag location.

Various means have been employed to alleviate this problem:

- Use of motion sensors on the door to trigger the reads. Thus, when tags pass through Door 2, the reader on Door 1 would not be turned ON. However, nothing prevents the simultaneous use of both doors, in which case unwanted reads will be seen.
- Use of narrow beam antennas and shielding. Such antennas reduce the time windows (a, b) and (c, d), and the shielding reduces the reads by A1 in the time window (b, c). Unfortunately, unwanted reads are not eliminated completely, and it is not a general-purpose solution—it is not logistically, economically, and, in many cases, physically possible to add shielding between dock doors and passageways.

The optimal approach to eliminating unwanted reads takes a different tack. It leverages the information residing in a full reader system view of a facility. It relies on information such as the spatial relationships between the antennas and their locations, the read rates of the antennas, and the tags observed by the antennas.

Device Management and Monitoring

There may be little or no onsite IT support for networked RFID readers deployed in environments that range from shop floors and retail stores to hospital rooms and warehouses. In such situations, automated configuration and discovery of the RFID readers and devices as they are set up is essential.

Once they're set up, RFID system administrators will need device management and monitoring tools, health and performance monitoring, and firmware management of the readers and other devices in the infrastructure, will be critical. The RFID infrastructure may span multiple sites, which drives the need for remote configuration and monitoring tools. In addition, for a robust infrastructure, redundancy needs to be built at each layer, with failover capabilities.

Tag Data and Events

The tag data collected and the business events generated by the infrastructure need to be disseminated for closed-loop, open-loop and cross-enterprise data collection. This data exchange can involve many processes and potentially many companies. More than likely, business events are generated based on a combination of RFID and non-RFID events.

The exchange of data could move two ways. Product suppliers could notify retailers that their goods are in transit, and retailers, in turn, could provide suppliers with visibility to goods as they flow through their supply chains, through to the selling floor or point-of-sale. A rich cross-enterprise visibility of RFID tagged objects and their associated business context makes possible useful response to EPC data and events. For instance, in a retail application, the supply chain partners could collaborate to closely control the retailer's inventory and maximize sales.

Standards in the Infrastructure

Currently, worldwide RFID standardization is driven primarily by EPCglobal and the International Standardization Organization (ISO). Both offer existing and emerging standards which cover most aspects of RFID, starting at the air interface and spanning the enterprise data exchange. In addition, regulatory standard bodies in each country are responsible for defining the RF spectrum used by RFID devices.

Figure 6 provides a macro view of the infrastructure with the different standard protocols.

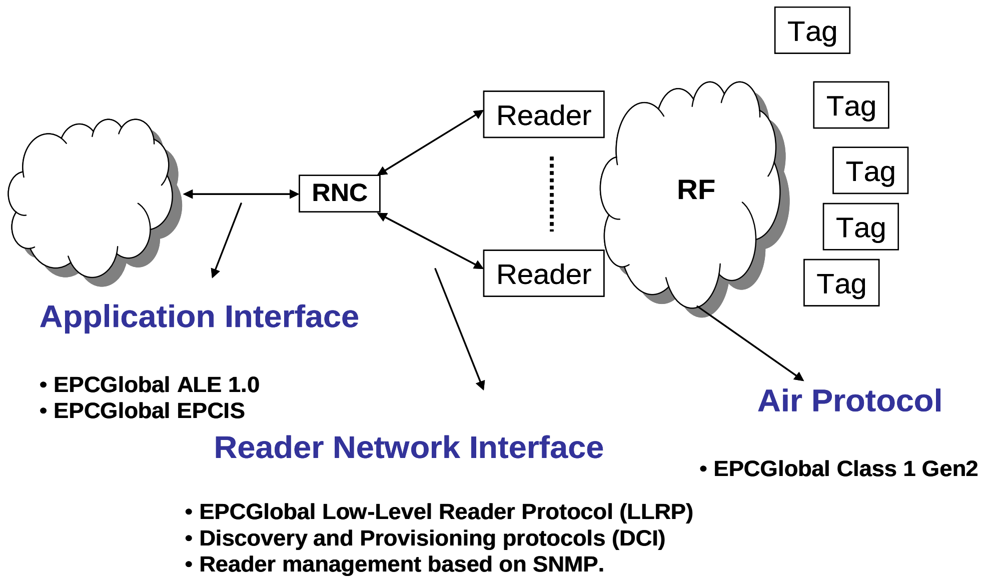

Figure 6: Standards in RFID infrastructure

Regulatory Domains

Table 1 lists a number of regulatory domains, with the frequency band, operating power and spectrum sharing technique specified for each.

Table 1: Examples of regulatory domains

Region	Frequency	Power	Spectrum Sharing Technique
USA	902-928 MHz	4W EIRP*	Frequency hopping
EU	In transition. See text.	2W ERP	In transition. See text.
Japan	952-954 MHz	4W EIRP	Listen before talk (LBT)
China	840.5-844.5 MHz 920.5-924.5 MHz	2W ERP	Frequency hopping

*EIRP = Effective Isotropic Radiated Power

26 MHz of spectrum is available for RFID in the United States, from 902 MHz to 928 MHz. This is divided into a maximum of 52 non-overlapping channels, each 500 KHz wide, with a maximum power of 4W EIRP. Since this spectrum is part of the unlicensed band, the Federal Communications Commission requires that radio transmitters in this spectrum pseudo-randomly change channels every 400 ms to prevent a given radio from monopolizing the spectrum (called the frequency-hopping spread spectrum (FHSS) technique).

In Europe, only 3 Mhz is available for RFID, from 865 MHz to 868 MHz. This is divided into 15 channels each only 200 KHz wide. Only 10 channels allow high-power operation at 2W EIRP, with the others reserved for lower power transmissions. The cooperative use of the spectrum is achieved by applying spectral mask constraints, frequency agility in the transmitting radios and a listen-before-talk (LBT) protocol.

In the LBT protocol, the radio devices triggered to transmit on a given channel must first “listen” to make sure the channel is clear (no signal above -96 dBm), to avoid collisions. If the channel is not available, the radio device may switch to another channel and try again. Because of the power of an RFID transmission (2W EIRP), the propagation distance for a single reader to detect another reader at the LBT level (-96 dBm) could be hundreds or even thousands of meters. This severely limits the number of simultaneous RFID reader transmissions within a facility, and would seriously limit the prospects for effective use of large-scale RFID deployments (which might comprise tens to hundreds of readers in a facility, with different facilities in close proximity).

Several approaches were adopted to address these limitations:

- The RFID community voluntarily agreed to use only channels 4, 7,10 and 13 (see Figure 7) for RFID reader transmissions, and the short range radio device (SRD) community which relies on non-RFID radio devices also in the 865–868 MHz band) voluntarily agreed to use only channels 1, 2, 3, 5, 6, 8, 9, 11, 12, 14, 15.
- As an interim solution, an ETSI Technical Specification [ETSI is a standards organization that develops communications standards for Europe], ratified only this past March, recommended means for LBT synchronization (networked control plus a wireless signaling mechanism). Synchronized LBT would apply the same LBT rules, but to a group (a “system”) of RFID readers operating simultaneously on the same channel.
- Work is currently under way to remove the LBT requirement from channels 4, 7, 10 and 13. This recognizes that with the adoption of the four-channel plan described above, dense-mode reader operation, coupled with the means of dense-mode reader synchronization, LBT in the four channels reserved for reader transmission would no longer have practical value. This phase is expected to be completed in 2008.

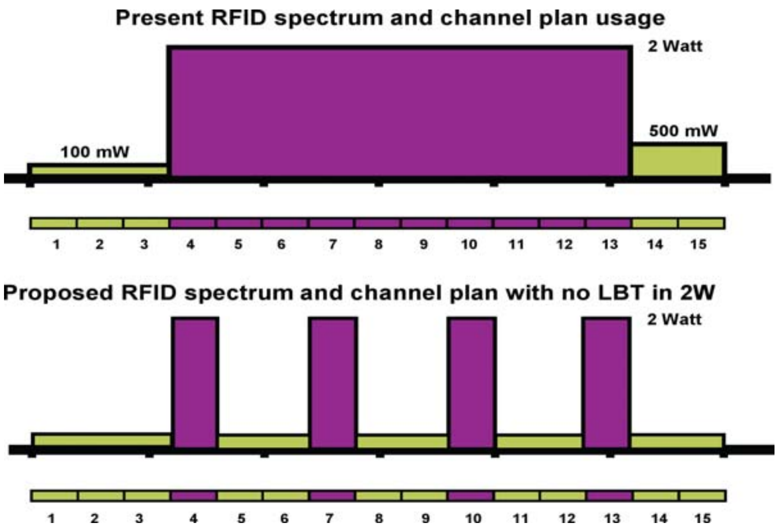

Figure 7: Channel Plan in Europe

The Air Protocol

The ISO 18000-6C/EPC Global Class 1 Gen2 (Gen2) specifications define all aspects of the air protocol for communications between readers and UHF passive tags. The Gen2 air protocol provides several benefits in terms of better performance and richness of tag operations:

- Increased data and singulation rates than the earlier Gen1 protocols
- Interference rejection from RFID and non-RFID users in the unlicensed band, and also signal-dependent backscatter.
- Simultaneous inventory of the same population of tags by multiple readers
- Selection of the subset of tags that participate in the inventory
- Password protection with secure locations in tag memory
- More operations are allowed: read/write/lock/kill

A Gen2 tag memory is logically separated into four distinct banks. The reserved memory contains kill and access passwords, if they are configured. The EPC memory contains the EPC identifier and some control bits. The Tag Identifier (TID) memory has information about the tag manufacturer so a reader can identify the optional and custom features supported by the tag. The user memory stores user-specified data.

Reader Operations

The EPCglobal Low-Level Reader Protocol (LLRP, see Figure 6, again)) is a flexible interface protocol between the RNC and the RFID Reader. It provides control of RFID air protocol operation, timing, and access to air protocol command parameters. It supports a wide range of underlying reader hardware/firmware and provides full access to the underlying air protocol capabilities. LLRP also provides support for RF monitoring if the reader can perform that function. And it supports reader operations in all regulatory jurisdictions.

A variety of applications require operations on the RFID tag data. They may range from reading EPC IDs to performing other tag access operations exposed by the air-protocol like read, write, kill, lock, and so on. The LLRP interface provides a scalable mechanism to manage such access operations at the readers. This is all helped by the LLRP's very rich set of tag data, event and error-reporting abilities.

A standard data and-control protocol for the readers allows for a uniform software infrastructure. This has resulted in important benefits such as predictable and consistent system-level performance, common support and installation expertise, a common set of performance monitoring tools, best-in-breed reader device selection, and lower operating costs.

Reader Management

The EPCglobal reader management interfaces (known by their initials of RM and DCI) provide the ability to provision and configure readers, and manage and monitor the health of the readers in a deployed infrastructure. In essence, RM and DCI allow RFID to meld into an existing IT infrastructure by making use of standard network protocols.

The Simple Network Management Protocol [SNMP] is an established standard used in today's networks that specifies the messaging protocol and transport layer for getting and setting device information, event notification, and security facilities. The EPCglobal Reader Management (RM) protocol specifies an SNMP-accessible Management Information Base (MIB) for monitoring the health of a reader. The MIB is a structured representation of Reader Object Model elements that conforms to the SNMP specification. The RM standard enables reader monitoring to be performed by the existing monitoring facilities of enterprise networks.

The EPCglobal Discovery, Configuration and Initialization (DCI) protocol specifies the means by which readers and RNCs enable network connectivity to other devices and application servers, exchange configuration information, and initialize their operation. This process would typically occur in advance of a data and control protocol (such as LLRP), which will then be used to control the operation of the readers to provide tag and other information to the RNC. Specifically, DCI provides a standardized means to allow a reader to discover one or more RNCs, the RNC to discover one or more readers, and for the reader to obtain configuration information, download firmware, and initialize operations.

Enterprise Data Interfaces

The EPCglobal Filtering and Collection Application Level Events (ALE) standard provides an interface to obtain consolidated EPC data from a variety of sources, decoupling the application consumers of the EPC data from the physical capturing devices.

The EPCglobal EPC Information Services (EPC-IS) family of interfaces provide standard event capture and query capabilities for obtaining and sharing data about RFID tagged objects both within and among cooperating enterprises. It supplements enterprise resource planning and operating systems and enables functions like track-and-trace, product authentication and diversion detection across supply chain partners in multiple vertical markets.

Benefits of Standards

From the end users' perspective, standards-based products create a rich competitive environment, which in turn breeds technological improvements and price reductions. Successful standards like the Class 1 Gen2 air protocol and now software protocols like LLRP, ALE and EPCIS, through universal acceptance, lower system design costs for everyone—broadening niche markets into mass markets.

System integrators and end users benefit from: devices that fully support standard interfaces and thus provide guaranteed interoperable deployment; configuration and data management capabilities offered by standard products allow for fine tuning of the infrastructure to optimize for widely varying application environments; investing in a standards-based infrastructure ensures the long-term value of the investment.

Conclusions

RFID was pioneered with an unstructured architecture of autonomous readers connected to business applications and underlying infrastructure through custom middleware. A long list of other technologies also started out as unstructured solutions, but the successful ones evolved into well-defined infrastructure layers, and melded themselves into the already existing enterprise network infrastructure. RFID is no different—it would never move beyond an interesting niche technology if it remains unstructured and autonomously anchored in middleware.

Strong technical and economic drivers compel the change from self-defined to infrastructure-managed RFID solutions. As the complexity of standalone reader deployment runs up against real-world scales of production, RFID is gravitating towards the standardized adaptations that have shaped earlier networking successes.

References

[ARC] EPCglobal Architecture Framework,
http://www.EPCglobalinc.org/standards/Final-EPCglobal-arch-20050701.pdf

[LLRP] Low Level Reader Protocol 1.0,
http://www.EPCglobalinc.org/standards/EPCglobal_LLRP_Ratified_Standard_20April_20042007_V1.0.pdf

[ALE] Application Level Events Standard 1.0,
http://www.EPCglobalinc.org/standards/Application_Level_Event_ALE_Standard_Version_1.0.pdf

[EPCIS] EPC Information Services,
http://www.EPCglobalinc.org/standards/EPCglobal_EPCIS_Ratified_Standard_12April_2007_V1.0.pdf

[TDS] EPC Tag Data Standard,
http://www.EPCglobalinc.org/standards/EPCglobal_Tag_Data_Standard_TDS_Version_1.3.pdf

[Gen2] Class 1 Generation 2 UHF Air Interface Protocol Standard 1.0.9: “Gen2”,
http://www.EPCglobalinc.org/standards/Class_1_Generation_2_UHF_Air_Interface_Protocol_Standard_Version_1.0.9.pdf

[RM] Reader Management Standard,
http://www.EPCglobalinc.org/standards/RM_Ratified_Standard_Dec_5_2006.pdf

[SNMP] IETF RFC 2578, Structure of Management Information Version 2 (SMIv2)

[ETSI] ETSI TS 102 562, Electromagnetic compatibility and Radio Matters (ERM); Improved spectrum efficiency for RFID in the UHF band.

About the Authors

Pattabhiraman Krishna [IEEE Senior Member] is a founding engineer and chief systems architect at Reva Systems, in Chelmsford, Mass. He is editor of the Low Level Reader Protocol (LLRP), a global standard for interfacing with RFID readers. Krishna was selected as GS1 EPCglobal's Software Action Group Person of the Year in 2006 for his significant contributions to RFID standards development.

He brings over 10 years experience in the networking field having led architecture and design teams at both emerging and established companies including Coriolis Networks and Digital Equipment Corp. Krishna holds five patents, with several more pending. He has published articles in numerous peer-reviewed journals and conferences. He is a member of the IEEE Communications Society, serves on the editorial board of *IEEE Applications and Practice* magazine, and is a member of the Association for Computing Machinery's SIGCOMM. He received a Ph.D. in computer science from Texas A&M University. Krishna can be reached at pkrishna@revasystems.com.

David Husak [IEEE Senior Member] founded Reva Systems in August 2003 and as Reva's chief technical officer focused the company from the start on standards leadership and strategic architectural issues of the RFID market. As a result, Reva products were installed in Fortune 500 companies within months of their introduction, and Reva has won industry awards for the clarity of its RFID effort.

Husak was selected as "CTO of The Year" in the 2006 Technology Leadership Awards presented by the Massachusetts Technology Leadership Council for contributions to the development of innovative business technology. He had previously co-founded and been CTO of C-Port Corp., a fabless communications semiconductor company, where he was the principal architect of the category-creating C-5 Network Processor. C-Port was sold to Motorola in May 2000.

Husak was the founding engineer and system architect at Synernetics Inc., the pioneering Ethernet and FDDI LAN switching company, which was sold to 3Com. Prior to that, he developed LAN interface hardware at Apollo Computer. Throughout his career, Husak has contributed extensively to network industry standardization efforts. He holds seven patents, with several more pending. He received a SBEE in Bioelectrical Engineering from the Massachusetts Institute of Technology and has completed graduate work there in communications systems. He can be reached at dhusak@revasystems.com.

Reva Systems
100 Apollo Drive
Chelmsford, MA 01824 USA
Tel: 978-244-0010
Fax: 978-244-0055
Web:www.revasystems.com

This white paper is for informational purposes only. Specifications subject to change without notice. Reva and Reva Systems are registered trademarks of Reva Systems Corporation. All other trademarks or registered trademarks are the property of their respective owners.

Part #1209-001