Proof of Possession: Using RFID for large-scale Authorization Management

Eberhard Grummt[1,2] and Ralf Ackermann[1]

[1] SAP Research, CEC Dresden
{eberhard.oliver.grummt|ralf.ackermann}@sap.com
[2] Technische Universität Dresden

Abstract. In inter-organizational supply chains, sharing of distributed, item-related information gathered using RFID can enable novel applications. Access control (AC) is needed to selectively disclose information to authorized participants. Given the large amount of data and the number of participants, common AC approaches would require extensive manual efforts. These efforts can be reduced significantly by the ability to prove physical possession of items to other companies. We examine how such Proofs of Possession can be designed. Based on two promising approaches, we introduce the concept of a Possession Service that may become a key factor in addressing the AC challenges in future supply chains.

1 Introduction

Using physical items to perform access control is a proven approach to enhance security. With the proliferation of the "Internet of Things" [3], billions of physical items with the ability to digitally carry identifying information will be linked to information stored in networked databases. In commercial scenarios such as RFID-enabled supply chains, both the data on the items and in the databases is potentially confidential, as it can contain mission-critical information like vendor-buyer-relationships or quantities. It seems self-evident to leverage the physical items themselves to manage access to their digital representations. Since the flow of an RFID-tagged item through a supply chain implies business relationships, periods of physical possession and their chronological order can be leveraged to infer access rights to item-related information. We name this approach *Possession-centric Access Management (PCAM)*. In contrast to tokens traditionally used in access control scenarios, basic RFID tags employed in supply chain management (SCM) are not permanently bound to a specific person or organization, but travel from company to company. They often do not contain secrets and are seldom tamper-proof. These properties pose hurdles to the realization of the PCAM approach. Still, it is attractive considering the amount of information that is automatically acquired in RFID-enabled supply chains and needs to be securely shared between companies to enable novel applications [2, 1, 13]. We argue that the ability to *prove* current or past physical possession of an item can form a sound basis for authorization decisions.

Based on a problem statement in Section 2, we systematize how such *Proofs of Possession* (PoPs) can be constructed and managed, discussing several approaches' differences regarding reliability and infrastructural requirements (Section 3). In Section 4,

we introduce the novel concept of a *Possession Service* that can enable organizations to evidence whether other organizations were in possession of an item of interest or not. We present related work in Section 5 and conclude with a summary in Section 6.

2 Problem Statement

We investigate how physical access to RFID-tagged items at chosen points in time can be proved to remote parties while and after a company possessed them.

2.1 Definitions

Let $C = \{c_1, c_2, \ldots, c_n\}$ be a set of companies that successively handle an Item $i \in I$. With $T = \{t_1, t_2, \ldots\}$ being the ordered set of all values in a discrete time system, the relation *Possession* is defined as $Poss \subseteq C \times I \times T$, where a tuple (c, i, t) is in $Poss$ when a company c possessed item i at time t, i.e. $(c, i, t) \in Poss \Leftrightarrow c\, possessed\, i\, at\, t$. We define the relation *Possession Period* as $PossPer \subseteq C \times I \times T \times T$ with $(c, i, t_{start}, t_{end}) \in PossPer \Leftrightarrow \forall t \in T \exists (c, i, t) \in Poss, t_{start} \leq t \leq t_{end}$.

A *Proof of Possession (PoP)* is defined as information that enables others to verify, with a certain probability, if a company c really is or was in possession of an item of interest i at a chosen point in time t or during an interval (t_{start}, t_{end}).

A *Claim of Possession (CoP)* is a statement made by a company that it possesses or possessed an item i at t or during (t_{start}, t_{end}). We denote claimed possessions as elements of the relations $Poss'$ and $PossPer'$.

We distinguish between *Possession Time*, comprising time values at which actual possessions exists, and *Validation Time*, referring to the time when PoPs are evaluated by participants. A *Verification Service* (VS) is an abstract concept for a networked service operated by a party that is trusted by all $c \in C$, e.g. to calculate and store PoPs, manage secret keys, certificates, or transactional information. The concrete meaning and tasks of a VS is explained in the context of the several approaches and may differ slightly.

2.2 Use Case

As an item i travels through a supply chain, each $c \in C$ gathers and stores information related to i, the aggregation of which represents the item's history. The involved companies (i.e. the route of the item) are not completely known beforehand. All companies agreed on a generic policy that permits any company that possessed a specific item to access the related information stored at any of the other companies. Thus, the access control decision includes determining if the requester has ever possessed the item or not. Fig. 1 illustrates a scenario where i has passed through n companies, each of which operates a database containing information about i. c_n wants to access information related to i stored at c_1. c_1 does not know the route i took, in particular it does not know c_n. However, c_1 can rely on a PoP to grant c_n access. Note that c_n "found" c_1 in the first place by using a *discovery service*, which again might have performed access control based on a PoP. These initial steps are not depicted.

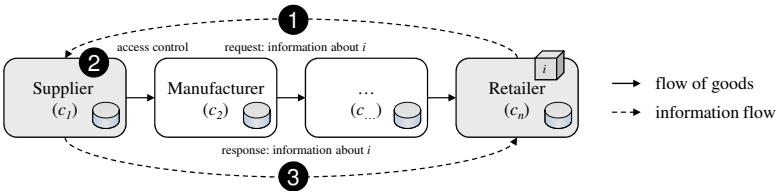

Fig. 1. The Access Decision Problem: how can c_1 decide if the request from c_n is legitimate?

2.3 Challenges and Evaluation Criteria

The main challenge is to maximize the probabilities with which a PoP is correct and with which a wrong CoP is identified. We consider that participants might issue wrong CoPs. A malicious company c could manipulate values for i or t. As we assume that every company $c \in C$ can be identified reliably, spoofing of identity is not addressed in this paper. When designing a system facilitating PoPs, the following high-level criteria need to be considered:

C1: it should be difficult to "prove" non-existent possession. We further distinguish
C1a: it should be difficult *before* a company had actual possession and
C1b: it should be difficult *after* a company has had actual possession
C2: it should be hard to repudiate possession

The difficulty of the operations needed to construct an invalid CoP that is accepted as a PoP can be measured in computational and monetary resources, in the amount of colluders needed, and in the amount of a priori information an attacker needs. Additional requirements and criteria strongly depend on the application domain and can therefore not be stated in detail here. Tag costs, infrastructure dependencies, and trust in participants and third parties are a few examples.

3 Methods for Constructing Proofs of Possession

3.1 Method Alternatives

In the simplest case, one could treat a CoP as a PoP, meaning one would need to trust the issuing company not to make any false statements. If such a PoP was used in access control, access would be granted to anyone who claims "I am authorized", which is obviously not a satisfactory solution.

Manual signature A similar approach is to have each company $c \in C$ digitally sign their CoPs using software under their control, i.e. $PoP = CoP + Sig_c(CoP)$. This way, assuming a reliable public key infrastructure (PKI), c could at least be sanctioned for issuing wrong PoPs. Yet, none of the criteria C1 and C2 are met.

Signature by trusted reader The signing process can be done by a "trusted" RFID reader r whose signature key $k_s(r)$ is not known to c and is physically protected against extraction. The public signature verification key $k_v(r)$ needs to be logically

bound to c, for example using a directory and certificates issued by a trusted third party (TTP). The integrity of the reader and possibly its location may be certified by a TTP, too. If the tags can be cloned or emulated, false PoPs can still be generated without tampering the reader hardware.

Linking and reasoning on supply chain transactions This approach is based on the fact that an item can only be at one location at any point in time and is successively handled by companies, each of which usually knows its predecessor and its successor. A company c_m can issue a CoP at Possession Time of the form $possper'(c_m, i, t_{received}, t_{sent})$ along with c_{m-1} and c_{m+1} to identify its predecessor and its successor. If the majority of participants tells the truth, wrong CoPs can be identified as logic discrepancies. The PoP can be calculated by an instance that has access to all the CoPs submitted by the participants. Ideally, this instance could reconstruct the *Possession Chain* defined as $PossChain(i) = \bigcup_{c \in C}((c, i, t_{start}, t_{end}) \in PossPer)$ for every item $i \in I$.

Tags with static secrets If a tag contains a static secret that can only be determined by companies possessing the item, a proof of knowledge of this secret can be used as PoP meeting C1a, but not C1b. At creation time of the tag, the secret could be shared with a VS that can later on perform the verification of CoPs by comparing the stored and the submitted secrets. Alternatives are using *zero-knowledge proofs* to prevent the VS from learning anything about the secret, and using the secret as a signature key. The latter approach would need a PKI for individual tags.

Tags with secret generator Instead of using a static secret, a tag could continuously generate unpredictable "secrets" for example by signing or encrypting the tag's id together with the current timestamp using the tag's secret key $k_s(i)$. This approach requires actively powered, physically shielded tags with clocks. At Possession Time, a company c would send a signed CoP $poss'(c, i, t)$ along with $check = \{(i, t)\}_{k_s(i)}$ to a VS. The VS would then verify c's signature and if $check$ was really calculated by i. The latter can be realized using both symmetric and asymmetric cryptography.

Tag constellations Simple RFID tags expected to be dominant in SCM contain an Electronic Product Code (EPC) which is not secret and therefore not suited as a PoP. But a *set* of EPCs (called a *constellation*) can be considered a secret if its distribution of EPCs is sufficiently chaotic. This secret can be used as described in *Using tags with static secret numbers* and even be changed by repackaging items.

Chosen, temporarily valid secrets To avoid the need of expensive secret generators, a tag's secret could be altered *manually* at every transition of i from c_m to c_{m+1}. c_m could choose a secret s_m and share it with the VS and with c_{m+1}. c_{m+1} acknowledges s_m to the VS and chooses a new secret s_{m+1} which it again shares with the VS and with c_{m+2}, and so on. The VS keeps a list of secrets and their validity intervals. Instead of overwriting the secret on the tag, each company could slightly alter it, for example by inserting some random bits at random positions. So some kind of fingerprint of every company remains on the tag.

3.2 Storage Alternatives

Proof-on-Tag means that the PoPs are stored directly in the memory of the tag attached to i. Data protection against outsiders becomes mainly a concern of physical protection of the tags. However, a hostile participant in the supply chain has the whole range of attack possibilities such as probing, cloning, destruction, and removal. Especially where no permanent network connectivity is available, PoT has advantages. PoT requires more complex tags, providing an increased amount of (re-)writable memory and possibly means to perform access control. In order to access the PoPs, physical access to the respective tag is needed.

Proof-on-Network means that PoPs are stored on a remote networked server. We discussed several methods in Section 3.1 that rely on such a server. Most importantly, using this approach the PoPs can be accessed from any location with a network connection. The tag hardware requirements are lower. However, the server needs to be trusted, forms a single point of failure, and thus an attractive target for attackers. Companies will depend on its availability and the network infrastructure.

Hybrid Approaches can be used to combine the benefits from both of the aforementioned approaches. As we pointed out, the ability to write onto tags is a viable means to prove possession and to transfer small amounts of data directly bound to objects between and only between the companies that exchange goods. An auxiliary networked service trusted by all supply chain participants can increase security and availability of the possession information.

3.3 Evaluation

As a first step towards evaluating the different approaches, we reviewed them using the criteria introduced in 2.3 and the complexity of the required tags. The results are shown in Tab. 1.

Table 1. Evaluation of the approaches

Approach	**Tag reqs.**	**C1a**	**C1b**	**C2**
Manual signature	low	no	no	no
Signature by trusted reader	low	no	no	no
Linking and reasoning on SC Trans.	low	yes	yes	yes
Tags with static secrets	medium	yes	no	no
Tags with secret generator	high	yes	yes	no
Tag constellations	low	yes	yes*	no
Chosen, temporarily valid secrets	low or medium	yes	yes	no

"Low" tag requirements mean that simple tags like EPC Gen 2 Class 1 tags can be used. "Medium" means that rewritable memory or other features not available in the aforementioned class of tags are required, while "high" requirements imply complex computational functionality or active power supply. "low or medium" in the last row

refers to the fact that the secrets can either be stored on the tags (leading to medium requirements) or be transferred between the companies using a network connection.
(*) Note that to meet C1b, repackaging is neccessary.

4 Towards a Secure Possession Service

In this section, we describe early work on a networked *Possession Service* (PS) and associated protocols enabling to reliably answer questions of the general form $poss'(c, i, t) \in Poss \rightarrow \{true, false\}$. Based on our evaluation, we identified *linking and reasoning on supply chain transactions* and *chosen, temporarily valid secrets* as the two most promising approaches we strive to combine. They both do not need sophisticated tags and meet the criteria C1a and C1b. To meet C2, i.e. making it difficult for c_m to repudiate possession, using information submitted by other companies such as c_{m-1} and c_{m+1} is about the only viable approach because an illicitly behaving c_m could simply not submit any CoP at all. Fig. 2 depicts the two main ideas: companies choose a new secret when an item arrives (s_1 and s_2) while the PS maintains an internal representation of the Possession Chain together with the secrets.

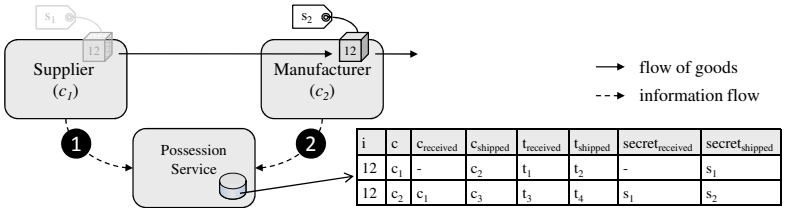

i	c	$c_{received}$	$c_{shipped}$	$t_{received}$	$t_{shipped}$	$secret_{received}$	$secret_{shipped}$
12	c_1	-	c_2	t_1	t_2	-	s_1
12	c_2	c_1	c_3	t_3	t_4	s_1	s_2

Fig. 2. Interaction between Possession Service and companies

The two main functional requirements are (i) calculating and storing PoPs based on CoPs issued by companies handling an item at Possession Time, and (ii) evaluating and answering CoPs at Validation Time. In the following overview, we present first requirements and insights regarding the different components of the envisioned Possession Service.

Query Interface (QI) Using this interface, legitimate requesters can answer the following questions: (a) did c possess i at all? (b) is c currently possessing i? (c) when did c possess i?

Capture Interface (CI) This interfaces's main purpose is to receive supply chain events such as CoPs and different security tokens submitted by individual companies. They are then passed to the *Possession Chain Reconstruction Component*.

Access Control (AC) Access to both the QI and the CI needs to be controlled as mission-critical information can be obtained using the QI and integrity can be compromised using the CI. AC can either itself be based on the PCAM, or employ explicit assignment of rights using Access Control Lists (ACLs) or Role-based Access Control (RBAC). The query interface might also be limited so companies can

only request certificates attesting *own* possessions. These certificate can in turn be used to gain access to other companies' information systems.

Possession Chain Reconstruction Component This component uses information received via the CI to decide whether it is correct or not and to integrate it into an internal representation of an item's Possession Chain. Depending on the information submitted by the individual companies, knowledge of secrets and spacial and temporal reasoning can be employed.

We want to decouple the Possession Service from any actual access control decisions to information systems. What kind of rights (read, write, delete etc.) a company c_l grants to c_m under which kind of possession relation (past or present) shall be up to the individual companies.

5 Related Work

Using physical items, and especially RFID for controlling access to physical resources is not new [10]. In the context of RFID and security, current literature focuses on reader/tag authentication and secure communication [9, 7].

RSA's *SecurID* and related technologies employ small tokens generating one time passwords that can be verified by a remote server. Combining this technique with RFID could yield a very secure mechanism for proving possession. The *Authenticated RFID model* [8] uses signatures of the manufacturer permanently stored on the tag. Each company handling an item submits "supply chain event information" to a networked service. It is signed by PKI-enabled "Authenticated RFID readers" that also support verification of the manufacturer's signature. Optionally, timestamps related to the events can be written onto the tag to enable further plausibility checks.

The approach of linking and reasoning on supply chain transactions is currently discussed in the area of anti-counterfeiting [11], especially in the pharmaceutical industry. It however relies on the assumption that all involved parties submit correct information to a central service. In the same context, Ilic et al. [6] propose that companies establish temporary 1:1-"ownership"-links between physical items and their networked electronic pedigree records. These links can be used to infer and delegate access rights. The key idea is that only one company at a time is allowed to establish such a link, which is supposed to correspond to physical possession. The authors however do not discuss how actual possession can be proved to the pedigree record service provider and how the "links" can be maintained in detail.

The data structures in a PS are similar to those expected to be found in discovery services. While in the EPCglobal Architecture Framework [12], discovery services are not yet specified, early approaches can be found in [5]. In contrast to discovery services in other application areas, access control is crucial because of the sensitive nature of data, both regarding read and write access [4]. To our best knowledge, this problem has not been studied in detail before. A reliable proof of possession would be a suitable criteria to decide if a party is granted write access to a discovery service or not.

6 Summary and Future Work

If reliable proofs for the past and present possession of items can be generated and shared between supply chain participants, access control to item-related data can be simplified because less manual setup of access rights would be needed. In reality, policies would not be as simple as "temporary physical access ⇒ unlimited data access", so temporal constraints as well as exceptions, for example to hide suppliers from wholesalers, would be needed.

PoPs with RFID tags can be constructed either *using item-specific information from a single source* or *combining item-specific information from different sources*. *Item-specific information* can be secrets on the tag which may be static, dynamically generated, or manually chosen. Leaving digital "traces" on the item, maybe using watermarking techniques, is a promising approach made possible by (re-)writable RFID-tags.
Information from different sources can be combined to increase resilience against minorities of misbehaving companies. Regarding our evaluation criteria, C1a is generally easier to achieve than C1b. If a tag contains a secret s that every company can only read out when possessing the tagged item, and companies do not disclose s to others, than knowledge of s proves possession.

PoPs can be generated by individual companies to prove possession to other companies. Alternatively, an external service can be used by a group of participants to assist in constructing, storing and querying PoPs. The main problem with such a service is that it has to be trusted not to disclose any of the information it gathers about business relationships between the companies that use the service. The same issue applies in the context of discovery services for individual items.

In our future work, we will examine how these limitations can be addressed by distributing the Possession Service and protecting its data so only minimal trust needs to be put into single instances. Case studies need to be examined in order to evaluate to what extent PCAM can be used in practical settings and what its advantages are against manually setting and delegating access rights. We are also concerned with the question how access control agreements that leverage PoPs can be negotiated, verified and enforced in dynamic groups of supply chain partners.

References

1. Indranil Bose and Raktim Pal. Auto-ID: managing anything, anywhere, anytime in the supply chain. *Commun. ACM*, 48(8):100–106, 2005.
2. Dursun Delen, Bill C. Hardgrave, and Ramesh Sharda. RFID for Better Supply-Chain Management through Enhanced Information Visibility. *Productions and Operations Management Journal*, January 2007.
3. Elgar Fleisch and Friedemann Mattern, editors. *Das Internet der Dinge: Ubiquitous Computing und RFID in der Praxis*. Springer-Verlag Berlin Heidelberg, June 2005.
4. Eberhard Grummt, Markus Müller, and Ralf Ackermann. Access Control: Challenges and Approaches in the Internet of Things. In *Proceedings of the IADIS International Conference WWW/Internet 2007*, volume 2, pages 89–93, Vila Real, Portugal, October 2007.
5. Mark Harrison, Humberto Moran, James Brusey, and Duncan McFarlane. PML Server Developments. White paper, Auto-ID Centre, University of Cambridge, Mill Lane, Cambridge, CB2 1RX, United Kingdom, June 2003.

6. A. Ilic, F. Michahelles, and E. Fleisch. The Dual Ownership Model: Using Organizational Relationships for Access Control in Safety Supply Chains. In *Advanced Information Networking and Applications Workshops (AINAW '07)*, volume 2, pages 459–466, Niagara Falls, Ontario, Canada, May 2007.
7. Ari Juels. RFID Security and Privacy: A Research Survey. *IEEE Journal on Selected Areas in Communication*, 24(2):381–394, February 2006.
8. Joseph Pearson. Securing the Pharmaceutical Supply Chain with RFID and Public-key infrastructure (PKI) Technologies. White Paper RFIDPH01, Texas Instruments Radio Frequency Identification Systems, June 2005.
9. Pedro Peris-Lopez, Julio Cesar Hernandez-Castro, Juan Estevez-Tapiador, and Arturo Ribagorda. RFID Systems: A Survey on Security Threats and Proposed Solutions. In *11th IFIP International Conference on Personal Wireless Communications – PWC'06*, volume 4217 of *Lecture Notes in Computer Science*, pages 159–170. Springer-Verlag, September 2006.
10. Melanie R. Rieback, Bruno Crispo, and Andrew S. Tanenbaum. The Evolution of RFID Security. *IEEE Pervasive Computing*, 05(1):62–69, 2006.
11. Thorsten Staake, Frédéric Thiesse, and Elgar Fleisch. Extending the EPC Network – The Potential of RFID in Anti-Counterfeiting. In *SAC '05: Proceedings of the 2005 ACM symposium on Applied computing*, pages 1607–1612, New York, NY, USA, 2005. ACM Press.
12. Ken Traub, Greg Allgair, Henri Barthel, Leo Burstein, John Garrett, Bernie Hogan, Bryan Rodrigues, Sanjay Sarma, Johannes Schmidt, Chuck Schramek, Roger Stewart, and KK Suen. The EPCglobal Architecture Framework – EPCglobal Final Version of 1 July 2005. http://www.epcglobalinc.org/standards/Final-epcglobal-arch-20050701.pdf, July 2005.
13. Samuel Fosso Wamba, Louis A. Lefebvre, and Elisabeth Lefebvre. Enabling Intelligent B-to-B eCommerce Supply Chain Management Using RFID and the EPC Network: A Case Study in the Retail Industry. In *ICEC '06: Proceedings of the 8th international conference on Electronic commerce*, pages 281–288, New York, NY, USA, 2006. ACM Press.