Trusted Information Sharing

Mark **Scardina**

Dmitry **Lenkov**

Abstract

Trusted and secure information exchange has become a prime concern and a problem to solve for government agencies and multiple enterprises. In this paper we propose a technological approach for building solutions for the problem. This approach is based on trust agreements represented as XML REL licenses and trust propagation using license distribution models.

Table of Contents

1. INTRODUCTION

Information interchange whether intentional or inadvertent, authorized or unauthorized, malicious or benign, has become a prime concern for executives and government decision-makers in today's environment of digital information, instant global distribution, and cyber-crime. Security alone is not a sufficient condition to control or motivate information sharing in today's distributed systems environment. The 9/11 Commission concluded that it was not the security of information but the inability for government agencies to easily share information that hampered prevention and directed the president to implement a "trusted information network". This need for trust equally applies in today's regulated world, where governments, medical centers, universities, and businesses alike have a real need to share data with principals who are trusted and appropriately authorized.

When working to prevent, prepare for, or respond to crises and disasters, federal intelligence and military agencies as well as federal, state, and local government departments, such as police departments and fire departments, must share highly sensitive information. Given concerns about immediate terrorist threats and critical infrastructure vulnerabilities, they must quickly find ways to securely share information while providing a level of assurance and control that both protects sensitive information and encourages organizations to share that information with others outside their domain. Exchanging sensitive information across organizational boundaries requires that data be controlled, managed, and limited in its access and distribution.

Businesses also demand faster, better, more secure communication processes without sacrificing trust. Over the years, this has led to the development of more sophisticated and robust processes in the physical world, driving the creation and success of overnight delivery services and the acceptance of the facsimile for business use. In today's environment, enterprises recognize the value of electronic communication as a faster and more convenient means of doing business, but security and trust remain elusive.

The challenge for organizations is how to exchange sensitive or valuable data without losing control of it or having it end up in the wrong hands. It is very important that the mutual trust between agencies and government departments is preserved, as is the confidentiality of the digital content being exchanged. To preserve the trust and allow this trust to be propagated across multiple agencies and/or enterprises requires a new technology that supports exchange of content (digital assets) based on trusted information agreements between organizations. This technology should be simple enough to encourage broad use and acceptance, yet robust, scalable, and flexible enough to be used in mission-critical and crisis situations.

The first section describes basic principals of Trusted Information Exchange (TIE) systems and introduces basic elements necessary to build trusted information exchange systems. The next section discusses how ISO-REL (MPEG 21000-5), an international rights expression language standard, is used to represent information agreements as licenses and how a license-based trusted information exchange model can be implemented. The use of XML Signatures to verify the authenticity of the licenses, and XML Encryption to protect the privacy and confidentiality of the license is also discussed.

In conclusion, this paper will address how industry-specific XML Schema-based extensions (profiles) to the ISO-REL standard can meet specific information sharing requirements. A specific example of such extensions as applied to the government sector will be presented and an actual trusted information flow demonstrated.

2. TRUSTED INFORMATION EXCHANGE BASICS

2.1. Basic Principals

The following key principals emphasize the guidelines for the development of the technology for the trusted information exchange:

2.1.1. Principal of Trust

If an information exchange system is not trusted by the relevant agencies, then its effectiveness will remain limited no matter how widely dispersed the network is. Government agencies and other entities must trust that information will be handled properly, in a manner that does not jeopardize national or corporate security. The use of licenses, based on REL profiles, as trusted information agreements establishes the mechanism to enforce a trust model when combined with a secure infrastructure and thus controlling the exchange based on these licenses.

2.1.2. Distribution and Coordination

The most basic principle for a successful networked environment is that information must flow in a distributed, yet coordinated way. Thus information must flow through the environment in a decentralized, non-hierarchical (peer-to-peer) manner. However, the coordination and control of its flow is critically important. In the approach, previewed in this paper, the coordination and control are supported by the issuance of licenses as trusted information agreements. Trust propagation is implemented through authorized re-issuance of licenses for further distribution of information

2.1.3. Balance

Balance between the need to share and the need to preserve information secrecy (Confidentiality) is also critically important. Different information tagging such as classification levels can be used to manage different information classes.

2.2. Basic elements

The basic elements necessary to build trusted information exchange systems are represented in Figure 1.

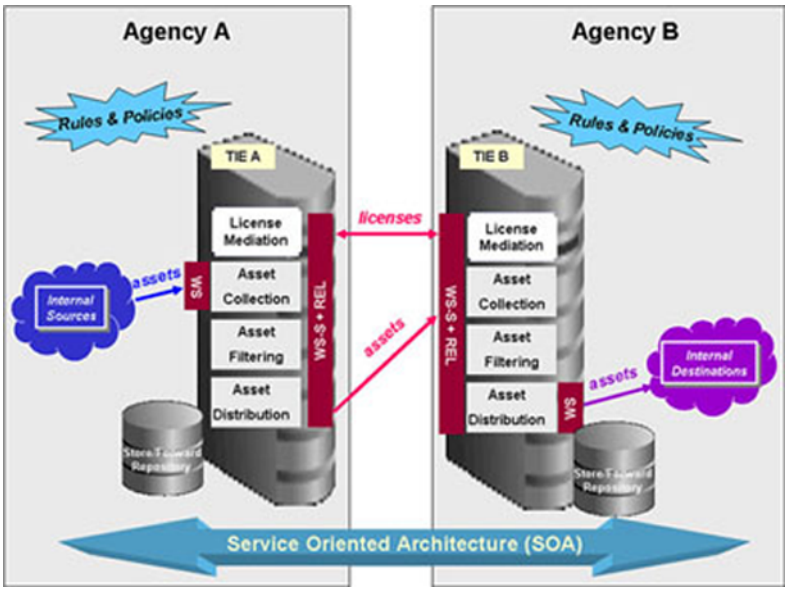

Figure 1. Trusted Information Exchange System

Trusted information exchange (TIE) systems are a solution that allows mutually trusted entities to share content over the Internet based on granted rights and apriori agreed conditions. It has a Service Oriented Architecture (SOA) and uses a content repository to store and forward content to trusted entities on the Internet. It could also be implemented to support content exchange with mobile or intermittently connected devices.

2.2.1. SOA (Service Oriented Architecture)

SOA separates a particular service's implementation ("how") from its interface ("what"). As long as the consumers interact with the Service under the terms (protocol, usage, and interface) of the "service" contract, they can use the service without actually knowing the details of the implementation. This architecture allows for a cross-domain solution.

2.2.2. Internal Sources

These are applications, often legacy applications encapsulated with web services, which are used by the content owning agencies (departments) to submit content, or assets, (information, digital content) to a local TIE for its distribution to other agencies. Such an application is often a vertical solution that has been built to share information within an organization and does not have security/infrastructure capabilities for content exchange or distribution (to trusted entities) over the Internet.

2.2.3. Asset Collection

This is a process of accepting and processing assets from internal producing applications via a web service, filtering, encrypting and temporarily storing them and their associated metadata into corresponding repositories.

2.2.4. Internal Destinations

These are often applications that are used by agencies or departments, to process assets delivered via their local TIE from other agencies or organizations. Supporting the SOA design they are often legacy applications encapsulated with web services.

2.2.5. Asset Disposition

This is a process of distributing assets acquired from of other agencies or organizations, for internal processing.

2.2.6. Asset Acquisition

This is a process of accepting and processing assets from external agencies or departments' TIE's, filtering them, decrypting them, and temporarily storing them and associated metadata into corresponding repositories for further distribution.

2.2.7. Asset Distribution

This is a process of distributing assets collected from internal applications to remote TIEs of external agencies, or organizations. This process is preceded by the license mediation.

2.2.8. License Mediation

This is a process of distributing licenses to or receiving licenses from remote TIEs of external agencies, or organizations. This process involves license generation and license validation on the other side. The license validation involves license metadata validation and approval workflows. All licenses are encoded in REL.

2.2.9. Rights Expression Language: (REL)

The ISO standard (ISA-MPEG-21/5) language allows content owning agencies and their authorized distributors to express the rights and conditions under which an entity can use/distribute digital content. REL uses standards based encryption (XML-Enc), Digital Signatures (XML-DSig), and PKI within its XML syntax to provide for both the authentication of the license itself and the entities involved in the transaction.

2.2.10. Asset Filtering

This includes several filters to check assets in order for them to be safely distributed to other TIEs or to internal consuming applications. These filters can include cipher processing, virus and malicious code scanning, “dirty” word searching, and file type checking among others. They could also include a review by an authorized person.

2.2.11. Rules and Policies

Policies are constraints that guide multiple processes, such as filtering processes and license distribution, within a particular agency. Policies are expressed in a written language and are based on guidelines specific to this agency. Rules are algorithmic representations of policies. They are used by this agency’s TIE to apply policies to a remote TIE’s processes.

2.3. Trust Propagation Models

There are several propagation models of how licenses and assets can be distributed between several TIE-enabled agencies or organizations. These extend from two basic models.

2.3.1. One-Step Distribution Model

This distribution model is static. No new TIE (TIE-enabled agency) can be added to the distribution flow without the source TIE knowing it. The important advantages of this model are:

- The distribution of licenses and assets is directly accountable. The original agency directly negotiates license grants with all other agencies. And it knows at every moment which assets and where they are distributed.
- The identification of assets is local to the source TIE. This is sufficient for referencing assets in licenses and for auditing and monitoring purposes

A significant disadvantage of this model is that the source agency cannot delegate to another agency further distribution, for example, to local offices of this other agency.

2.3.2. Multi-Step Distribution Model

This model is dynamic in that each TIE can be both a consumption point and a distribution point. It has the following significant advantages:

- The load of license and asset distribution on the source TIE can be minimized.
- The source TIE does not have to keep information about all secondary downstream TIEs involved and about all transactions with them.

However, this model introduces some complications. They include:

- Downstream trusted information agreements (licenses) are decoupled from the source TIE’s direct control
- This distribution model requires a more sophisticated auditing model.

3. REL LICENSES FOR THE TRUSTED INFORMATION EXCHANGE

The MPEG Rights Expression Language (REL) [I] is an ISO standard for a general purpose authorization language. It defines a machine readable, XML-based language for expressing rights agreements over digital resources bound to parties under specific conditions. The REL framework for creating precise, reliable, and secure licenses for content is an essential component of any effective and scalable peer-to-peer trusted information exchange management (TIE) system.

ISO-REL can be used to create licenses that address a wide variety of business models. The ISO-REL is designed to be domain agnostic and comprehensive. It is also very flexible, and adaptive to the specific scope and requirements of applications within a particular domain. Specifically, the language provides extension and profile mechanisms for these purposes.

3.1. REL License Basics

REL licenses consist of the basic elements described in Figure2 that abstracts the expression of rights and their bindings and conditions. The detailed description of these elements and additional features such as variables and delegation control is contained in [I].

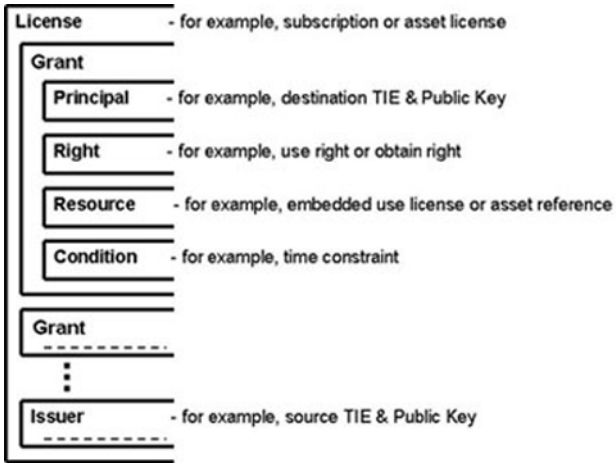

Figure 2. REL License Elements

3.1.1. License

It is the basic construct of REL. It conveys all the necessary authorizations or grants for the use and / or distribution of the content or asset.

3.1.2. Grant

It is the element within a License and grants to a Principal, to whom the Grant is issued, a Right for a certain action in regard to a Resource under certain Conditions that must be met before the Right can exercised

3.1.3. Principal

It presents the identity of a party to whom rights are granted subject to certain conditions by an issuer.

3.1.4. Resource

It is the object to which a Principal can be granted a Right. It can also be an embedded grant.

3.1.5. Right

It is a verb, or action, that a Principal can be granted to exercise on a Resource under certain conditions.

3.1.6. Condition

It specifies the terms, conditions and obligations under which Rights can be exercised. Condition can also modify the Resource or the Principal.

3.1.7. Issuer

It identifies a Principal who issues the rights to a certain party subject to certain conditions.

3.2. REL License Profiles and Extensions

The ISO-REL is designed to be domain-agnostic and comprehensive. It is also very flexible and adaptive to the specific scope and requirements of applications within a particular domain. Specifically, the language provides extension and profile mechanisms for these purposes.

From the extensibility point of view ISO-REL standard has two architectural parts plus an initial vertical extension. The core, standard extension, and multimedia extension, as well as their XML Schemas, are normative parts of the overall ISO- REL specification. Other parties may, if they wish, define their own (possibly domain-specific) extensions to the REL. This is accomplished using the existing, standard XML Schema and XML Namespace mechanisms.

Just as an application domain may need to develop its own ISO-REL extension to capture its own domain-specific elements (e.g. rights, resources and conditions), it may also wish to identify a conforming subset of the ISO-REL to meet its specific needs in order to facilitate compliance and interoperability among its applications and systems. This is where profiles come into the picture.

The profile specifies a particular use of the base REL core standard in order to meet specific needs in an optimized yet interoperable way of a particular community (Trusted Information Exchange community in our case). Thus, the profile provides a specification to build interoperable special-purpose systems, as well as providing a standard reference for profile compliant systems. Clearly, more than one profile may be developed to achieve different objectives.

3.3. Trusted Information Exchange Licenses

Trusted Information Exchange (TIE) functionality can require several types of licenses. These licenses are defined as per REL schemas with certain domain specific extensions. For example, government related domains might include constructs such as topic, classification, relevancy, validity internal, etc.

These domains may also give rise to different license types. The TIE makes use of several including a confirmation license, asset usage license, asset distribution license, subscription license, and subscription termination license.

3.3.1. Confirmation License

The confirmation license is used to confirm identities and operational status of communicating TIE systems. It is assumed that actual identification of principals and the exchange of certificates, such as x.509's, happen through "out of band" processes, outside of TIE framework. The source TIE sends to the destination TIE a confirmation license, which contains destination TIE public key and identification information. This communication acts as an electronic confirmation of

the destination TIE setup to communicate with the source TIE. Validating, signing and returning the confirmation license by the destination TIE completes this communication.

3.3.2. Asset Usage Licenses

The asset usage license communicates a grant that conveys the destination TIE's digital identity, the asset to be retrieved, the right that may be exercised and conditions to be satisfied. Conditions can also include metadata associated with the asset. Finally the digital identity of the issuer is included for authentication. This license type looks like this:

```
<license>
 <grant>
  <keyHolder LicensePartId="Destination TIE Id">
    <info>
     <dsig:KeyValue> ....</dsig:KeyValue>
    </info>
  </keyHolder>
  <use/>
  <digitalResource>
            ...Asset Id
  </digitalResource>
  <expiration>
            ...Expiration date and time
  </expiration>
  <assetMetadata>
            ...Asset metadata
  </assetMetadata>
 </grant>
 <issuer>
  <keyHolder LicensePartId="Source TIE Id">
    <info>
     <dsig:KeyValue> ....</dsig:KeyValue>
    </info>
  </keyHolder>
 </issuer>
</license>
```

Figure 3. Sample Asset Usage License

3.3.3. Asset Distribution License

The asset distribution license makes use of the embedding feature of the ISO-REL to allow for multi-step distribution. It includes the Asset Usage grant as a resource and "issue" as the right. It may also communicate metadata associated with the right or principal as illustrated below.

```
<license>
 <grant>
  <keyHolder LicensePartId="Destination TIE Id">
    ...
  </keyHolder>
  <issue/>
  <digitalResource>
            ...Asset usage grant – (see above)
  </digitalResource>
  <expiration>
            ...Expiration date and time
  </expiration>
 </grant>
 <issuer>...</issuer>
</license>
```

Figure 4. Sample Asset Distribution License

3.3.4. Subscription License

As TIE can employ a “topic” model under which assets can be distributed, it requires a subscription license as a way of registerng interest in a set of assets. The subscription topic defines a set of conditions over the subscription vocabulary. Sets of conditions can be provisioned separately by a TIE internal application or during destination or source TIE provisioning.

The subscription license communicates to a particular destination TIE the set of constraints (conditions) that will be applied to asset or its metadata distributed to this TIE. Each condition can be as simple as a name-value pair, for example, topic = emergency. Or it can be more complex, for example, specifying a time limit for asset requests after the license for this asset has been issued. The set of conditions can be specific for a particular destination TIE or a group of destination TIEs, or it can be defined for all destination TIEs

Similar to the distribution license, subscription licenses include the asset usage grant or asset distribution grant as their resource but use “obtain” as their Right as illustrated below:

```
<license>
 <grant>
  <keyHolder LicensePartId="Destination TIE Id">
   ...
  </keyHolder>
  <obtain/>
  <digitalResource>
          ...Asset usage or distribution grant with variables and constraints.
  </digitalResource>
  <expiration>
          ...Expiration date and time
  </expiration>
 </grant>
 <issuer>...</issuer>
</license>
```

Figure 5. Sample Subscription License

3.3.5. Subscription Termination License

The subscription termination license communicates to the destination TIE system that this source TIE system will send no more assets qualifying for the subscription and is a simple variation of the subscription license above.

3.4. Sample License Mediation

While beyond the scope of this paper, licenses must be both generated and consumed. This process involves validating both the structure of the license as well as its types and the operational semantics that they may represent.

A key element of this mediation is the authentication capability afforded by the inclusion of both the Issuer’s and Principal’s digital identity. This allows for PKI-based bi-lateral authentication without a central root of trust or trust authority. This trust model can also be “chained” through the use of the distribution licenses previously described.

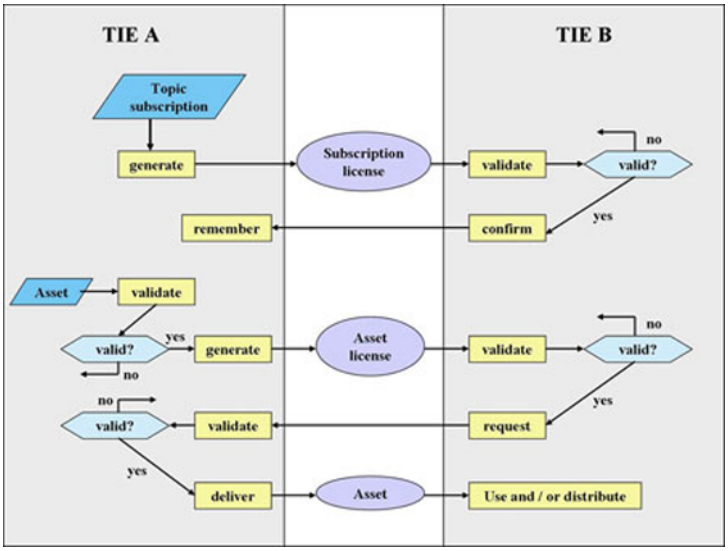

Figure 6. Sample license mediation scenario

Figure 6 represents a sample flow starting from confirmation license generation and ending with an asset being distributed. This is only a sample. Real flows can be significantly more complex.

4. CONCLUSION

Data and information management systems can no longer be viewed as “silos” of dedicated information. The “need to share” has begun to redefine the “need to know” as companies and agencies can no longer afford to extend access to there internal systems externally. Information is an asset that needs to be fully exploited; however, it can only be done if systems are in place to securely protect and track them. Whether it is a terrorist watch list, an MP3 song, a location-based service, or a medical record, it can no longer remain in one place. Thus data management systems must be able to work without boundaries as in a “trusted information network where TIEs act as trusted information routers securely linking these information systems together.

5. REFERENCES

I. ISO/IEC 21000-5, Information Technology – Multimedia Framework (MPEG-21), Part 5, Rights Expression Language

Biography

Mark **Scardina**
Oracle Corporation [http://www.oracle.com]
Redwood Shores
California
United States of America
Mark.Scardina@oracle.com

Mark Scardina is Oracle's XML Evangelist for Server products and is the Group Product Manager for the CORE and XML Development Group tasked with providing the XML infrastructure components used throughout the Oracle product stack including the XML Developer's Kits. Mark represents Oracle on the W3C XSL Working Group and chair of Oracle's XML Standards committee. He is a frequent speaker at industry trade shows and conferences and is co-author of The Oracle9i XML-Handbook.

Dmitry **Lenkov**
Oracle Corporation [http://www.oracle.com]
Redwood Shores
California
United States of America
Dmitry.Lenkov@oracle.com

Biography not received.