High Assurance Information Exchange based on Publish-Subscribe and ABAC methods

Anders Fongen and Federico Mancini
Norwegian Defence Research Establishment (FFI)
Emails: {anders.fongen,federico.mancini}@ffi.no

***Abstract*—The presented effort employs a combination of publish-subscribe distribution and ABAC (Attribute Based Access Control) methods to control the information exchange between security domains. It follows strictly the "separation of duty" principle so a message router only has infrastructure duties while the identity management entity deals with management of authorizations and security policies. The presented work also implements a novel model for message protection and subject authorization. One characteristic of the resulting transfer protocol is that an external bump-on-the-wire device can verify the integrity of the messages and that the security policies are observed. This device can be carefully constructed for the purpose of high assurance and offer fail-safe mechanism in case the message router is malfunctioning or compromised.**

I. Introduction

Military computer nodes handle information of different classification levels, and it is common practice to collect nodes of similar classification levels into *domains*, inside which the nodes can exchange information with fewer restrictions. Between domains, the information exchange must be closely monitored and controlled.

In order to make our analysis applicable to more than a military classification hierarchy we propose the term *Community of Trust* (CoT) indicating that the focus is on the varying security requirements, rather than the actual classification level of the information. The term CoT includes what elsewhere is called *security domains*.[8]

Between the CoTs, exchanged information may be subject to inspection and control by *Policy Enforcement Points* (PEPs), which are non-bypassable units with the duty to stop information which is not approved for exchange. In everyone's home router there is a firewall with this duty, which makes its policy decisions based on stateful packet inspection of *implicit* data in protocol headers and payloads. Although efficient for intrusion protection, a firewall is not well capable of stopping an information leak from high to low side.

A PEP would rather make its decisions based on *explicit* data, often termed *metadata*, which is applied into the data structures by a trusted source and validated by the PEP before the decision whether to allow transfer is made.

A PEP which bases its decisions on trusted metadata is sometimes known as a *guard*.[10] A guard inspects and validates metadata in the form of *security labels* which indicate the security classification of the information, on which it makes its decision if it can be "released" to the opposite CoT.

The concept of a guard is simple and easy to implement as long as the structure of the information objects is well known. It suffers from a number of weaknesses though:

- The validation process may require revocation information from a PKI.
- The security labels may not indicate the authorization of the security label creator.
- The exchange policy is not directly given by the metadata, but must be derived based on configuration data managed by the guard administrator.
- The policy configuration is asynchronous with regard to the information flow, which hinders the information from being *bound* to a specific policy.

The contribution of this paper is a model for information exchange where the exchange policy and the authorization of the sender are included in the information messages. This model allows a PEP to be operated as an infrastructure device without regard to policy or metadata management. It also offers a more general framework for information exchange, inside which the exchange between military security levels is a special case.

The presented model is implemented over an experimental system for Publish-Subscribe distribution which employs the ABAC model for authorization control of senders and receivers. A *message router* replaces the guard as a PEP and bases its policy decisions on the message metadata called *subscriber requirements*.

The resulting message data structure allows an external bump-in-the-wire unit to ensure that the policy is being observed and to stop information that violates that policy. This unit can be built for high assurance in order to serve its function between highly classified networks.

The remainder of the paper is organized as follows: Section II will introduce the exchange model on which we build our analysis. Section III and IV will briefly describe the prototype software used for demonstration of our policy enforcement principles, and Section V presents some implementation details. Section VI gives and in-depth analysis of our chosen policy principles, and Section VII introduces the high-assurance publication inspector (HAPI). Section VIII reports from the prototype evaluation and Section IX relates our work to related research efforts. Section X gives concluding remarks and suggests further research on the topic.

Dette er en postprint-versjon / This is a postprint version.
DOI til publisert versjon / DOI to published version: 10.1109/MILCOM.2014.45

II. The Exchange Model

The exchanged information and the policy under which it is managed should be strongly bound. This requirement is not observed in other policy enforcement systems like XACML [1], where the flow of information and the flow of policy updates are independent and asynchronous. In such systems, it is impossible to know which policy is employed for a given message.

The *stakeholders* should decide the access policy for a message. The PEP is only an infrastructure device and does not risk anything during information exchange. The sender and the receiver are the ones concerned about the confidentiality, authenticity and integrity of the information and should set the security requirements. The receiver should set the requirements to the sender and vice versa.

A. Access Rules

The exchange model presented in this paper is based on the ABAC model for access control. The sender and the receiver are assigned sets of *subject attributes* by the identity provider (IdP), and they both set up boolean expressions called *access rules* which are evaluated over the other parts' attribute set to decide if "access" should be granted or not. E.g., an access rule formulated as

```
$clearance="secret" and $country="Italy"
```

will evaluate to `true` for attribute sets with these values for the attributes named `clearance` and `country`. The conditions on attributes are built using boolean operators like `EQ`, `LT`, `GT`, `INRANGE`, `HASTOKEN`, `STARTSWITH`, `CONTAINS` and the attributes can then be combined into access rules with standard logic operators like `AND`, `OR`, `NOT`.

The term "access control" is a slight misnomer, since we are evaluating the authorization of the parties for sending and receiving messages, not to gain access to a service or resource.

B. Subscriber Requirements

Through an access rule the sender (also called the *publisher*) selects the subset of authorized receivers. The authorized receivers all have attribute sets which evalute to `true` with regard to this access rule. An access rule used for this purpose is called a *subscriber requirement*.

Formally, let AS denote the set of all possible attribute sets. Any subject (sender or receiver) will be assigned an attribute set $as \in AS$ by their identity provider (IdP). The attribute set is embedded in their *identity statement* (cf. Section IV-A) and sealed by the signature of the IdP. The attribute set is therefore well suited for conveying information about the authorizations of the subject.

The total set of access rules is denoted AR. Any access rule $ar \in AR$ can be evaluated over an as using the function $match(ar, as) = true|false$. Consequently, the selection can be expressed as a function:

$$select(ar, AS) \equiv \{as \in AS | match(ar, as)\} \qquad (1)$$

C. Publisher Requirements

Likewise, access rules can be formulated by the receiver (also called *subscriber*) to set requirements to the attribute set of the publisher. This access rule is called a *publisher requirement*. The subscriber should never receive data from publishers which do not meet the publisher requirement. In the Publish-Subscribe system of this paper the publisher requirement is included in the *subscription*.

The subscriber requirement is used to protect the *confidentiality* of the information, while the publisher requirement is used to protect the *integrity*, in the sense that the information is assured to be generated by competent and approved publishers.

D. Policy Authorization

A publisher could be restricted with regard to subscriber selections. E.g., within a Bell-LaPadula authorization framework [3], a subject with access to highly classified information should not be allowed to send information to receivers with access only to lower classifications. Where two CoTs (Communities of Trust) are connected through a PEP, a special authorization could be required in order to "release" information to the other CoT.

If that authorization should be represented as subject attributes it would create a strong and undesired coupling to the actual vocabulary of attributes used in subscriber and publisher requirements. Our choice has been to represent that authorization as an access rule called *policy rule*, $pr \in AR$.

Informally, the policy rule creates a boundary for the possible set of receivers of a message, in the same manner as the subscription requirement does. But while the subscriber requirement is assigned by the publisher, the policy rule is assigned and sealed by the identity provider.

Formally, the subscriber requirement sr of a message must relate to the publisher's policy rule pr in the following way:

$$match(sr, as) \Rightarrow match(pr, as) \qquad (2)$$

which also leads to

$$select(sr, AS) \subseteq select(pr, AS) \qquad (3)$$

When this condition is met, we say that pr is *Wider-Than* sr, and will use the notation $pr \geq sr$ to indicate this for the remainder of the paper. The Wider-Than property of access rules is shown in the Venn diagram in Figure 1 and formally analyzed in Section VI.

The policy rule is stored in the form of a subject attribute, which binds the subject to the resulting policy authorization under the authority of the identity provider.

E. Flexible authorization arrangement

In a multi-CoT environment, we assume that subject attributes are assigned consistent with the subject's CoT. E.g., that all subjects in an Italian CoT have the attribute `nation=IT` and those in the Norwegian CoT have the attribute `nation=NO`. Given such arrangements, any publicaton originating in the Norwegian CoT with the subscriber requirement `$nation=NO` will never enter the Italian CoT.

Dette er en postprint-versjon / This is a postprint version.
DOI til publisert versjon / DOI to published version: 10.1109/MILCOM.2014.45

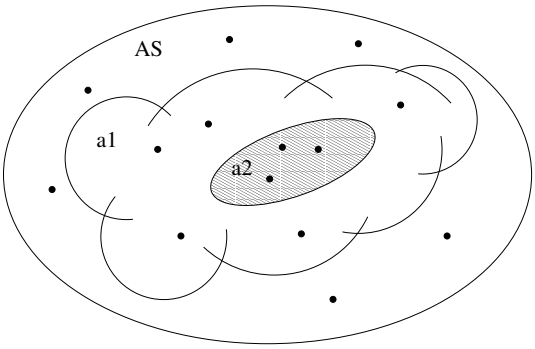

Fig. 1. The Wider-Than relation between access rules a_1 and a_2 shown in a Venn diagram. In this figure, $a_1 \geq a_2$ since $select(a_2, AS) \subseteq select(a_1, AS)$. The dots indicate instances of $as \in AS$.

Likewise, a sender with the policy rule `$nation=NO` will never be allowed to assign a subscription requirement like `$nation=IT` to a message. The ability to send messages across the PEP may therefore be subject to authorization, governed by the identity provider.

Extra arrangements may be necessary to ensure that the attributes are assigned in a consistent manner. Section V presents such an arrangement.

Even authorization patterns like the Bell-LaPadula model [3] can be obtained through a combination of policy rules, subject attributes and subscriber requirements. A message with subscriber requirement like `$clearance=SECRET` will require the attribute `clearance=SECRET` in the receiver's attributes. If the same receiver is given the policy rule `$clearance=SECRET` it will never be allowed to create a subscription requirement like `$clearance=RESTRICTED` in order to send to receivers with a lower security clearance. Consequently, "write-down" is impossible.

III. The Gismo PubSub system

For the purpose of experimentation with identity management, access control systems, cross domain authentication and publish-subscribe distribution a software prototype called *Gismo PubSub* has been built. A part of the prototype is also called *Gismo IdM* which contains the identity provider, certificate authority and software classes for authentication, service invocation, access control, service discovery and TPM protection. Section IV will present Gismo IdM in more detail.

The Gismo PubSub software consists of Message Router (MR) code and client API classes which set up publication listeners, create publication and manage the exchange of identity credentials between clients and MRs.

In a publish-subscribe (pubsub) environment, the message flow is mediated by *topics* and *subscriptions*. Receivers express interest in messages annotated with certain topics through subscriptions, which is why pubsub receivers often are called *subscribers*. In Gismo PubSub, the flow of messages between MR instances is indeed mediated this way, but that property is left out of the following discussions for the sake of focus on security properties. A full description of Gismo PubSub including the message flow mediation is given in [7].

Subject Distinguished Name
Subject Public Key
Subject Attributes
Valid from–to
Issuer Distinguished Name
Issuer's Signature

Fig. 3. The structure of the Identity Statement

IV. Gismo IdM

The text in this section is previously published in [6], and included here as background information.

The presence of an identity management system is essential to the management of subject keys and attributes. The identity provider (IdP) will serve as a trusted third party (TTP) and issue attestation of both keys and attributes.

Gismo IdM was developed in order to study the necessary properties for an IdM used in a multi-domain wireless mobile network used by a coalition tactical force.[5]

In Gismo IdM, existing PKIs are kept for reasons of investment protection, but encapsulated by a number of Identity Providers (IdP), each serving a Community of Trust (CoT). The members of a CoT share the IdP's public key as their trust anchor. The IdP issues *Identity Statements* (IS) to attest the public key and attributes of a subject. The IS is given a short lifetime and sealed with the signature of the IdP. Due to the short lifetime, no revocation arrangement is necessary.

The architectural overview of Gismo IdM is shown in Figure 2. Observe that the CoT members are never exposed to PKIX protocols or data objects (X.509 certificates or revocation lists). The key properties are explained in the following paragraphs:

A. Authentication support

The identity provider (IdP) issues *Identity Statements* (IS) which bind the public key of a subject to its identity, analogous to X.509 certificates. Identity statements are issued to local subjects registered in the IdP, as well as to subjects who can display an IS issued by a different IdP to which this IdP has a *trust relationship*. The structure of an IS is shown in Figure 3.

The subjects (either client or server) authenticate themselves during service invocation by the use of their identity statements and their private keys. Different authentication protocols have been designed with the purpose of generating as little network traffic and as few protocol round trips as possible. [5]

B. Integrated access control

Included in the identity statement is a set of attributes which describes properties of the subject in the form of name-value pairs. The attributes can describe *roles* of the subject and enter into access control decisions based on the Role Based Access Control (RBAC) or the Attribute Based Access Control

Dette er en postprint-versjon / This is a postprint version.
DOI til publisert versjon / DOI to published version: 10.1109/MILCOM.2014.45

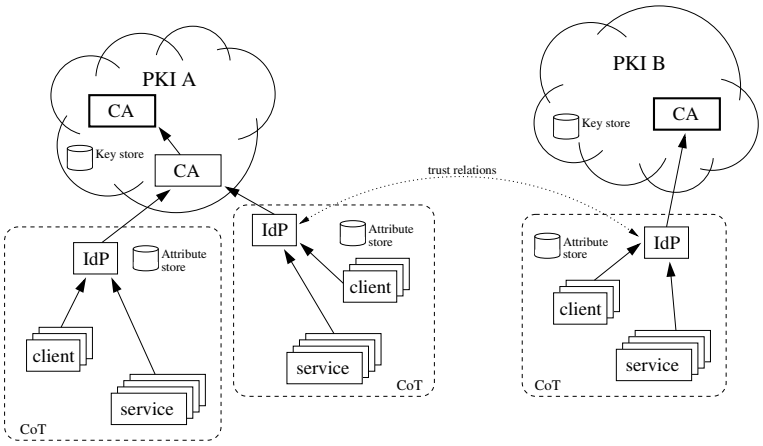

Fig. 2. The functional components of Gismo IdM. Observe that the IdP serves one single CoT. Key management is handled by the PKI whereas the attribute management is done by the IdPs on the CoT level

.

(ABAC) model. They can also describe other properties of the subject, e.g., preferred language, proficiency level etc.

Attributes may be evaluated by an access rule during service invocation as discussed in Section II-A in order to obtain ABAC type control.

Just as authentication is a symmetric process in Gismo IdM, the access control decisions are made both in the service and in the client. The client expresses an access rule which is evaluated over the service's attribute set, and accepts or rejects the service response accordingly.

C. Cross-CoT operations

Clients can authenticate themselves to a different CoT as indicated in Figure 2, provided that there exists a trust relationships between the two CoTs. A client obtains an identity statement from its IdP, then passes on that IS to the IdP of a foreign CoT. The foreign IdP can issue a *guest IS* containing the same information, but signed by the foreign IdP. Since the guest IS's signature will be trusted by servers in the foreign CoT, it can be used to authenticate to these servers. Server authentication requires a *cross domain IS* issued from one IdP to the other, so a signature chain back to the client's trust anchor can be constructed. The middle part of Figure 4 shows the protocol that takes care of this. The IdP of CoT A, termed IdP_a, issues a "native" identity statement to the client, which is given to IdP_b, which in turn issues a guest identity statement.

V. Message router as a PEP

The properties of the Gismo PubSub Message Router (from now on only called MR) make it well suited to work as a PEP. The MR keeps a list of neighbors (clients and other MRs) and their aggregate subscriptions, and only sends publications out through an interface if authorized subscribers are found along that path. To be precise, only connected clients are checked for authorization, not neighbor MRs. This is done for reasons of implementation efficiency. It is always the "last MR" along the route that checks the subscriber requirements.

The message router can serve the PEP function given the attribute arrangement which was discussed in Section II-E.

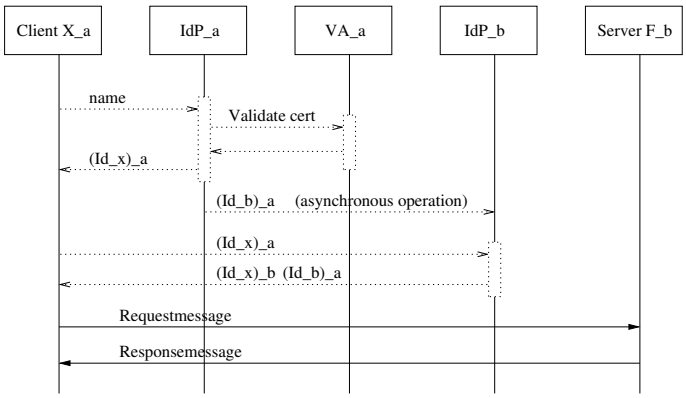

Fig. 4. The authentication protocol for a stateless service. The symbol $(Id_x)_a$ indicates the identity statement for Subject x issued by the IdP for CoT a. $(Id_b)_a$ indicates the cross-CoT for the IdP in CoT b, issued by the IdP in CoT a. S_x indicates signed by subject x, E_x encrypted to subject x.

We see, however, that the attributes need to be assigned in a consistent manner from a group of identity providers who are governed by different CoTs and therefore may not be sufficiently trusted.

For this reasons, a more realistic arrangement of MRs as PEPs is to put two MRs, owned by each CoT and given *guest identity statements* for their interconnection (cf. Section IV-C). The assigned attributes are interpreted as the *aggregated authorization for all subscribers reached through that MR*. By using guest IS, the authorizations of the MR is given by the CoT it receives publications from, not by its own CoT.

For an MR pair that serve as a PEP, publications will not be passed between them unless the subject attributes of the receiving MR satisfy the subscription requirement attached to the publication. This arrangement is a safeguard against fraudulent issuing of attributes in untrusted CoTs. The originating CoT has much stronger control of what is being released for other CoTs based on the attribute management of its own CoT. Figure 5 illustrates this arrangement.

Dette er en postprint-versjon / This is a postprint version.
DOI til publisert versjon / DOI to published version: 10.1109/MILCOM.2014.45

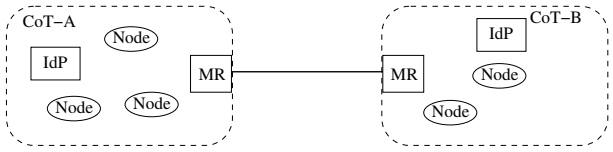

Fig. 5. Two Message Routers (MR) connected to form a PEP, each MR guards its own Community of Trust. Message flow between them are determined by the subject attributes issued by the opposite IdP, in addition to the Subscriber Requirements of the messages.

VI. Analysis of the Wider Than-property

If we take the easier case where every attribute can only assume two values, true or false, then an access rule ar becomes a boolean formula in propositional logic, where only the NOT, AND and OR operators are used. In this case the problem of asking whether given two access rules $ar_1, ar_2 \in AR$ then $ar_1 \geq ar_2$, becomes equivalent to deciding whether $ar_2 \implies ar_1$ is a tautology. The *TAUTOLOGY PROBLEM* asks whether a propositional formula f is true for every possible truth assignment of its boolean variables, and it follows that it is coNP-complete with a simple reduction from the complement of the classical *SAT PROBLEM* which is NP-complete [4]. Since any instance of the *TAUTOLOGY PROBLEM* can be reduced in polynomial time to a special case of the Wider-Than property, then also deciding whether one access rule is a subset of another is coNP-Complete. The reduction consists in taking $f = ar_1$ and then setting $ar_2 = TRUE$ so that asking whether f is a tautology is reduced to asking whether $ar_1 \implies ar_2 \equiv TRUE \implies f$ is a tautology. It is easy to see that this is true if and only if f itself is a tautology. This particular case corresponds also to a real possible instance, namely when the subscriber requirements are empty, so that any subscriber would be a valid one. Then we would have to check whether the policy is indeed a tautology because it has to evaluate to true for any possible combination of attributes, i.e., any possible subscriber. Since even in this restricted case the Wider-Than property is hard to verify, we can expect that with more complex attributes the problem becomes even harder. A hint is given by the fact that if attributes can be expressed as general predicates, then access rules become first-order logic formulas. Verifying whether such formulas are valid (a generalization of tautology) is an undecidable problem.

Therefore, in its general form the Wider-Than problem does not seem to admit an efficient algorithm under the $coNP \neq P$ hypothesis, but only some more or less optimized form of iteration over all $as \in AS$ on Equation 2. Given a limited number of attributes a brute force approach might still be efficient enough, but in general an approximation is necessary for the practical use of the Wider-Than property of access rules. Two approximations have been investigated:

1) Both the subscription requirement and the policy authorization must match the respective attribute set for every receiver of a message (clients and PEP MRs).
2) The expressiveness of the access rules may be restricted so that the Wider Than-property can be determined in finite time.

Approximation 1 guarantees also that publications are delivered only to the subscriber group defined by the attribute set $select(pr, AS) \cap select(sr, AS)$, but only as long as the verification is performed on trusted nodes. Some publications where $select(pr, AS) \setminus select(sr, AS) \neq \emptyset$, i.e., where the subscriber's requirement define an attribute set that in some cases violates the policy rule, might still get out of the sending CoT. This can happen if the MR at the perimeter of the receiving CoT has an aggregated as which does not violate either sr or pr. However other nodes longer down in the distribution path may not enforce the verification on pr and sr correctly, resulting in the delivery of publications to subscribers that should have not been in the receiving pool at all. This can happen anyway, but in this particular case we delivered publications that the publisher should not have been authorized to publish in the first place. If the Wider-Than property could be properly verified in the CoT where the message originates, such messages would be blocked before reaching untrusted nodes.

Approximation 2 will reduce the expressiveness to a level which is still expected to be useful in practice, but avoids the operations that creates the NP-complete properties. The chosen restrictions of the expression are:

1) Every attribute can only appear in the expression once
2) Only AND-operations are allowed to construct the access rules from the attributes

We effectively create a vector representation of the access rule and give each attribute name a certain index in the vector before doing pairwise operations on each vector index. AS is given the representation of a high dimensional space, and the access rules create subspaces which are guaranteed to be contiguous. Given this representation the Wider-Than property means that one space is embedded in another, which can be determined by inspection of the respective dimensions.

The computation based on this approximation can be done anywhere where the two access rules are known, and would preferably happen as the publication is sent to the first MR.

The rest of this section will present the chosen algorithm for the Wider-Than property for access rules with restricted expressiveness according to approximation no 2. The set of restricted access rules is denoted AR', where $AR' \subseteq AR$. The access rule is composed of a series of boolean variables with AND-operations in between.

The boolean variables for $ar \in AR'$ can contain the operation $=$ (EQ), $<$ (LT), $>$ (GT), and $= ..$ (IN RANGE). Given the variable $\$x$ and the numeric values a, b, c and d, the Wider-Than relation between the boolean variables are as follows:

$$
\begin{array}{llll}
(\$x = a) & \geq & (\$x = b) & if \quad a = b \\
(\$x > a) & \geq & (\$x > b) & if \quad b \geq a \\
(\$x > a) & \geq & (\$x = b) & if \quad b > a \\
(\$x < a) & \geq & (\$x < b) & if \quad b \leq a \\
(\$x < a) & \geq & (\$x = b) & if \quad b < a \\
(\$x = a..b) & \geq & (\$x = c..d) & if \quad c \geq a \wedge d \leq b \\
(\$x = a..b) & \geq & (\$x = c) & if \quad a \leq c \leq b \\
(\$x > a) & \geq & (\$x = b..c) & if \quad b \geq a \\
(\$x < a) & \geq & (\$x = b..c) & if \quad c \leq a
\end{array}
$$

For all other combinations of the four operators, the Wider-Than property does not apply.

Dette er en postprint-versjon / This is a postprint version.
DOI til publisert versjon / DOI to published version: 10.1109/MILCOM.2014.45

For evaluation of the Wider-Than property of the access rules, boolean variables on the same vector index (refers to the same variable) are evaluated according to the rules above and combined in an "AND"-operation. Nonexistent variables (null) in a vector are regarded as Wider-Than any variable or null value. Likewise, if a variable is not null in pr it should also be not null in sr otherwise the Wider-Than property should evaluate to FALSE. This to avoid the hard case mentioned earlier where one should verify whether pr matches all instances of $as \in AS$. E.g., the two access rules:

$$a_1 = (\$a = 1) \wedge (\$b > 3) \wedge (\$d = 4..10)$$
$$a_2 = (\$a = 1) \wedge (\$b > 5) \wedge (\$c = \text{“}opx'') \wedge (\$d = 6..10)$$

are given the representation

$$a_1 = [(= 1), (> 3), null, (= 4..10)]$$
$$a_2 = [(= 1), (> 5), (= \text{“}opx''), (= 6..10)]$$

Since all pairwise elements satisfy the Wider-Than property the relation $a_1 \geq a_2$ is true.

These rules are easily applied to the expression tree that represents access rules, in order to establish the Wider-Than property and to assess the required relation between the policy rule and the subscription requirement.

VII. High Assurance Publication Inspection

The publication data structure as used in the Gismo PubSub includes the following information items:

- Publisher's identity statement, including policy rule
- Publisher's signature
- Subscriber requirement
- Information object
- etc.

Anyone who sees a publication can do several checks regarding its integrity, provided that they share the same trust anchor (the IdP's public key).

1) The identity statement can be validated, including the IdP signature and the validity period.
2) The signature can be verified, so ensure that the originator is identified and that the content is untainted.
3) The Policy Rule can be extracted from the attribute set in the identity statement, and the Wider-Than relation between the policy rule and the subscriber requirement can be calculated.

If the publication is on its way to a message router, it is also possible to verify that the MR is authorized to receive that publication, provided that its identity statement is available for inspection. The authorization is decided by applying the `match` function (Equation 1) to the attribute set of this identity statement and the subscriber requirement of the publication.

Followingly, an "High Assurance Publication Inspector" (HAPI) located between to MRs as a "bump-on-the-wire" as shown as in Figure 6 can verify the

1) Integrity of the publication content
2) Authorization of the sender

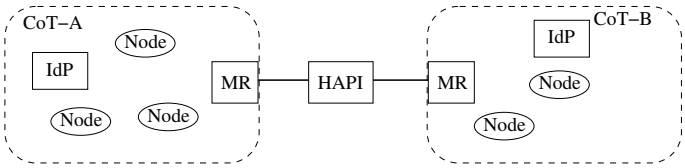

Fig. 6. One configuration of an High Assurance Publication Inspector (HAPI) positioned between two message routers (MR) belonging to separate CoTs, each having their own identity provider (IdP).

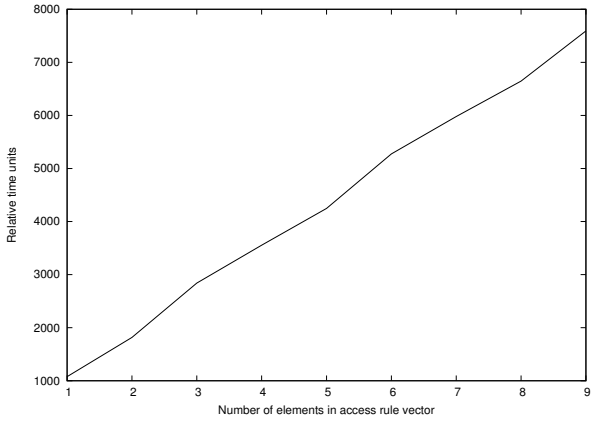

Fig. 7. Run-time properties of the algorithm to determine the Wider-Than property of access rules. The horizontal axis represents the number of elements in the vector representing the access rule.

3) Authorization of the receiving MR

provided that the trust anchor and identity statement of the message router is available.

The HAPI will not add any functionality, since all checks are already done in the message router. Its advantage is that it may be built to higher assurance since it is simpler and contains less logic than the message router, and because it does not need any trust relationship to auxiliary servers for revocation information or validation assistance. It will serve as a safety net for a malfunctioning message router.

Since the HAPI is intended to sit non-bypassable between two MRs it will have to inspect all transactions between the two MRs, including the authentication and exchange of subscriptions. It can verify the structural and cryptographic integrity of these transactions, but not guard against steganographic information leaks caused by an authorized sender infected by malware. Nor will it guard against a malfunctioning or compromised identity provider.

VIII. Prototype evaluation

The presented model has been implemented and added to the existing code for Gismo PubSub, and has been verified on a functional level. The evaluation confirm the correctness of the model and incorrectly constructed publications are shown to be rejected.

Quantitative evaluation is not likely to detect any performance or scalability problems in the suggested model, since operations like serialization, encryption and signature generation and -validation are computationally far more expensive

Dette er en postprint-versjon / This is a postprint version.
DOI til publisert versjon / DOI to published version: 10.1109/MILCOM.2014.45

than the calculation of the Wider-Than property of access rules. The subject attributes are stored in a hashtable (with retrieval cost $O(1)$ so the calculation involves $O(n)$ number of operations, where n represents the number of elements in the two access rules being compared. The results from an experimental study of the run-time properties of the Wider-Than algorithm is shown in Figure 7 and confirms this assumption.

IX. Related Work

The "Content based Information Protection and Release" (CPR) initiative within NATO employs a related principle for adding metadata to information which is inspected during transfer.[2] The CPR initiative proposes builds on the XACML model but makes some modifications for terminal identification. The separation between metadata generation and policy administration in CPR is intended and justified by the strict security regime found in military organization. The use of *policy rules* proposed in this paper is likely to provide similar control and more flexibility at the same time.

The CPR model spends quite a lot of efforts on filtering of compound information object to allow a "permitted view" to pass through an authorization control. Besides adding considerable complexity, the utility of such selection is doubtful. The CPR authors seem to regard a compound object as a tuple, while it in reality is a *directed graph*, and a node selection is very likely to produce an inconsistent and disconnected object graph except under very restricted circumstances.

Finally, the CPR proposes authentication and access control of the client only. Our model does this both ways, which we strongly believe reflects better the security threats of fraudulent servers and phishing attacks. A related matter is that the underlying XACML mechanism performs the access control at the time of the request, not the response, which makes an important difference for asynchronous invocation methods and in messaging/pubsub systems.

Regarding the different approaches for verifying the consistency of the policy rule with the subscribers requirements, we can refer to the work on trusted labelling in [9]. If we think of the generation of the subscriber requirements as a labelling of the publication, we quickly end up with the same challenges. How can a third party know whether a label was correctly applied and that it actually contained what the user who created it intended? In our scenario we have an advantage, namely a trusted policy rule that restricts the possible attributes, but we still need a trusted entity that can perform this verification. In [9] an incremental approach is suggested: first a central trusted server can be used by clients to generate the labels by enforcing possible restrictions, then the trusted service can gradually be integrated in the client themselves as the technology becomes available and the infrastructure is adapted to the new approach. We suggest a possible way to integrated the trusted service in the clients, while still maintaining a trusted component in the infrastructure as safety-net. Our advantage is the presence of a certified policy rule, which mitigates the problem of having to trust the whole client rather than just the labelling application. Without a policy rule one does not have any limitation on the attributes each client can use to define the subscriber's requirements. This means that in order to prevent information leakage we would have to scan all publications against a database with publishers authorizations and policies as in [2], losing all the advantages of the proposed approach.

X. Conclusion and future research

This paper proposes a method for information exchange between security domains based on a novel mechanism for access control. The proposed access control model is quite different from the XACML model, since the policy is attached to the information flow rather than being distributed independently. Our approach creates a loose coupling between the policy authority and the infrastructure service, and allows the stakeholder to present their security requirements rather than receiving them from a separate policy authority. Our approach allows the message routers to be stateless and zero-configured devices which lend themselves well to load balancing and fail-over arrangements.

Future research will include investigation of design details of the HAPI device and the ARG application, the development of a prototype and the evaluation of its performance and utility.

References

[1] OASIS eXtensible Access Control Markup Language. https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml. Online, Accessed March 2014.

[2] Alessandro Armando, Matteo Grasso, Sander Oudkerk, Silvio Ranise, and Konrad Wrona. Content-based information protection and release in NATO operations. In *SACMAT'13*, Amsterdam, The Netherlands, June 2013.

[3] Matt Bishop. *Computer Security, Art and Science*, pages 124–142. Addison-Wsley, 2003.

[4] Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, STOC '71, pages 151–158, New York, NY, USA, 1971. ACM.

[5] Anders Fongen. Federated identity management in a tactical multi-domain network. *Int. Journal on Advances in Systems and Measurements*, Vol.4, no 3&4, 2011.

[6] Anders Fongen and Trude Hafsøe Bloebaum. Trusted service discovery through identity management. In *IEEE MILCOM*, San Diego, USA, 2013.

[7] Anders Fongen and Federico Mancini. Identity management and integrity protection in publish-subscribe systems. In *IFIP IdMan 2013*, London, UK, 2013.

[8] Anders Fongen and Mazda Salmanian. Communities of Trust in Tactical Coalition Networks. In *IEEE MILCOM*, Baltimore, MD, USA, 2014.

[9] S. Oudkerk and G. Lunt. An Incremental Approach to Trusted Labelling In Support Of Cross-Domain Information Sharing. NC3A, The Hague, Netherlands, 2011.

[10] Konrad Wrona and Geir Hallingstad. Development of high assurance guards for NATO. In *Military Communication and Information Systems Conference (MCC)*, Gdansk, Poland, 2012.

Dette er en postprint-versjon / This is a postprint version.
DOI til publisert versjon / DOI to published version: 10.1109/MILCOM.2014.45