

Communities of Trust in Tactical Coalition Networks

Anders Fongen* and Mazda Salmanian†

*Norwegian Defence Research Establishment (FFI), anders.fongen@ffi.no

†Defence R&D Canada, mazda.salmanian@drdc-rddc.gc.ca

Abstract—The need for information exchange between security domains has traditionally been approached through the use of guards and security labels. Although these technologies are thoroughly researched and exist in mature implementations, they offer simplistic approaches with several shortcomings. In this paper, we build on the “guard” model and present a framework for trusted information exchange which accommodates a wider range of use cases, network topologies, and authorization models. Our approach can be used on a range of practical levels, down to dismantled soldiers and sensor networks. Central to this framework are the concepts of *Communities of Trust and Policy Enforcement Points*.

I. INTRODUCTION

In a connected tactical network, groups of nodes would have common interests in sharing information; though they may not trust one another. Therefore, the information exchange needs to be restricted and carefully controlled. The groups may restrict, even hide, network details, identity credentials, and information meta data from each other. While nodes of different groups (nations) may share interests in a common information domain as one group (forming a *Community of Interest*), the trust that they share inside their respective nations is not extended to the nodes of other nations. They therefore constitute different *Communities of Trust* (CoT). An example of a Community of Interest (CoI) could be the medics of the two nations military in an operation that requires them to communicate across their two CoTs.

In this paper we present three problems of information exchange: between CoTs with different non-hierarchical security classification systems, through transit CoTs, and the control of proper source authorization. We argue that the simplistic exchange principle of a “guard” is challenged in these situations. We submit that there is value in establishing a framework for coalition information exchange, such as the one we are proposing here.

A. Non-hierarchical classification

One would expect that a military network is configured in a “System High” security mode of operation. This means that a user must have the required clearance level to access all the information in a CoT; the user may only be restricted by the required need-to-know (NTK) authorization principles, where applicable. The System High mode facilitates the exchange of information in the sense of allowing information to be sent from lower to higher security levels or between peer classification levels, implying a hierarchy of security levels. One solution to such domain access regime is to attach *security labels* to the information, which is inspected by a *guard*. A guard is a non-bypassable unit in the connection point between the two domains. Guards and security labels are well

researched in the literature; a Common Criteria Protection Profile allows guards to be evaluated for High Assurance [11].

It is plausible, however, that information exchange may take place between domains (CoTs) that do not align with one single classification hierarchy. For example, governmental and medical information is governed by legislation and policies unrelated to military policies. Exchange mechanisms must mediate between these access and protection regimes. We will show with our proposal that information exchange between CoTs can pass through *policy enforcement points* (PEPs) which ensures that the information is passed according to a policy that the source (generator of the information) has designated and tagged to the information. The concept of a PEP builds on the guard concept and improves its shortcomings in many ways.

B. Transit CoTs

Ordinary information sharing among CoTs may cause information sent from CoT-1 to CoT-2 to be passed on to CoT-3. This situation requires a decision to be made whether or not the exchange policies from CoT-1 to CoT-2 should be applied to the exchange from CoT-2 to CoT-3. This reasonable proposition requires the exchange point between CoT-2 and CoT-3 to be able to interpret the corresponding security related meta data and to make transfer decisions accordingly. In a *guard centric exchange policy* the policies are attached to guards, and CoT-2 will not be able to observe CoT-1’s reservations regarding the transfer of its information to CoT-3. In a *data centric exchange policy* the exchange policies are attached to the information itself and verified by the PEPs along the transfer path, and the information can be passed to CoT-3 while the policy of the source CoT is observed.

C. Authorization to initiate an exchange

Only authorized subjects should be allowed to initiate exchange of information between CoTs. The authorization should be checked prior to the transfer process and (optionally) validated by the receiver to ensure that the information came from an authorized entity. The intention of this arrangement is to prevent mis-labeling and intentional compromise of sensitive information.

D. Shortcomings of the guard model

Exchange guards, as presented in [10] and [8] support a guard centric exchange policy, not a data centric exchange policy as we propose. Their reliance on centralized PKIs make them unfit for tactical environments. They lack a model for sender authorization and do not present any arrangement for passing information through a path of guards. Guards do not authenticate themselves to the communicating parties.

Other shortcomings of the guard model include lack of arrangement for guard discovery, load balancing, or fail-over operation. Our proposed framework includes PEP discovery and supports stateful and stateless operations where load balancing and fail-over mechanisms require so. In summary, **our contributions** include a framework for the exchange of information between CoTs, one that preserves the originator’s release policy, accommodates a non-hierarchical access regime, supports sender authorization, and which can be applied to dynamic networks, e.g. dismantled soldier Mobile Ad hoc Networks (MANETs) or sensor networks.

The remainder of the paper is organized as follows: we provide a detailed descriptions of the CoT and PEP concepts in section II and III, respectively. In section IV, we discuss practical uses and benefits of CoTs and PEPs. Section V includes a summary of our work and suggests future research in this area.

II. CoT DESCRIPTION

The term Community of Trust describes a group of network nodes that have established trust relationships with one another such that they allow a less restricted flow of information between themselves. A CoT may be a group of soldier radios belonging to the same platoon or a group of nodes with the same security classification in a strategic network. The CoT concept is useful as it allows the individual nodes to implement relaxed protection mechanisms with adequate “fenced in” border control.

We present the concept of CoT as a graph of link connections. In a network of computer nodes there may be subsets of nodes that could meet the following criteria:

- Traffic travels between the subsets (intra-CoT traffic).
- They are governed by the same authority and policy (for policy enforcement and authorization).
- They have high trust in each other.
- They form a connected graph through links (or routes).

The nodes that meet the aforementioned criteria form a CoT. A CoT node can be a member of more than one CoT; therefore, the CoT graphs may be overlapped or disjoint. In the following sections, we will explore these criteria in further detail.

A. Intra CoT communication

Communication between nodes that belong to the same CoT is relatively unrestricted, and the following policy may be observed:

- Nodes authenticate themselves prior to communication.
- Access control decisions are made prior to information exchange.
- Communication is encrypted if the link is “black”.
- Data objects are not inspected during transfer.

Communication within a CoT (over intra CoT links) is similar to ordinary IP traffic, and the nodes are likely to be configured as belonging to the same IP network.

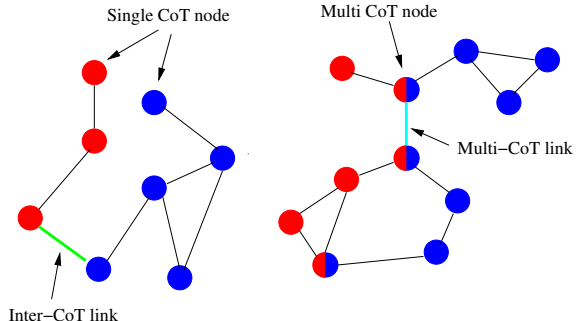


Fig. 1. Disjoint and overlapping CoT graphs, showing inter-CoT links (in green) and multi-CoT links (in blue)

B. Inter CoT communication

There is relatively little trust between communities, so communication between CoTs must be subject to strict policy enforcement, which should not be relaxed by the individual user or service. Information between two nodes from different CoTs needs to pass through *Policy Enforcement Points* (PEP) which may enforce exchange policies upon inspection of the data. Nodes from different CoTs may communicate through a dedicated link, e.g. a VPN tunnel through a black network or through an inter-process communication (IPC) channel. Figure 1 shows how nodes belonging to different CoTs may be organized in a network. The nodes of the network may be *single-CoT* or *multi-CoT*, which are concepts that will be described below.

C. Single CoT node

A node whose processed information belongs to the same CoT is called a single-CoT node. Links to other nodes in the same CoT are called *intra-CoT links*, which follow the established policy for intra CoT communication. The relaxed separation of information between user groups of the CoT is provided by the operating system of the nodes; therefore a cryptographic separation on the communication channel is not necessary. Figure 2 illustrates the structure of a Single CoT node.

Links to nodes from other CoTs are called *inter-CoT links*, which differ from intra-CoT links by the presence of a Policy Enforcement Point (PEP). The PEP will inspect traffic to allow or deny the transfer of a data object based on its observations. The functional perspective of the PEP is instrumental to the presented model and will be discussed in Section III.

D. Multi CoT node

A node that processes information from several CoTs, i.e., it is a member of more than one CoT, must employ a robust separation architecture to meet the confidentiality requirements of the CoTs. The notion of inter-CoT and intra-CoT links still hold for a multi-CoT node depending on the other endpoint node’s CoT partitions.

There are PEPs between the partitions of a multi CoT node to control the exchange of information, as shown in Figure 3. These internal connections are effectively inter-CoT links, but can be implemented by inter-process communication (IPC)

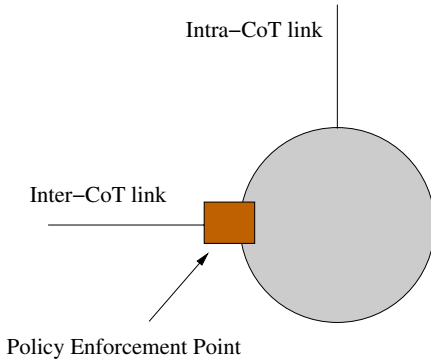


Fig. 2. The structure of a Single CoT node

channels as well as network connections. Multi-CoT nodes are specialized resources that are useful constructs for nodes and services common to several CoTs, e.g., an aircraft that supports more than one nation of a coalition force.

Between multi CoT nodes there exists a concept of a *multi-CoT link*. This is a type of link that carries traffic related to several CoTs duly separated and protected. It may be implemented with a bundle of VPN tunnels, each connecting node partitions in the same CoT. Each tunnel is functionally equivalent to an intra-CoT link and it does not require an associated PEP pair.




Fig. 3. Multi-CoT links connecting Multi CoT nodes

III. POLICY ENFORCEMENT POINTS

Information flow between two CoTs needs to pass through a Policy Enforcement Point (PEP). A PEP is directional and belongs to one CoT. Its main responsibilities are:

- Block incoming information suspected to be fraudulent, unauthorized, harmful, deceiving etc.
- Block outgoing information to protect confidentiality
- Ensure the authenticity of the peer PEP

On an inter-CoT link there are two PEPs guarding information flows in opposite directions. They belong to different CoTs and therefore they enforce independent security policies, using different technologies and algorithms.

A PEP is a “bump-in-the-wire” [9] and does not have any Application Program Interface (API) for the purpose of binding

information objects to security policies. Somewhat similar to a firewall, a PEP binds information to policies through *inspection* of content or traffic characteristics at different levels of transparency. To the nodes on each side of the PEP it may look like:

- A router, which forwards IP packets without modification,
- A firewall, which replaces IP addresses (e.g. network address translator - NAT) and may restrict UDP traffic,
- An HTTP proxy,
- A chat server or any other application level gateway.

If the PEP looks like an IP router, any application using the IP protocol is able to send data through it. If the PEP looks like an e-mail server, it could extract meaningful information about the structure of the transported information objects. Thus, a PEP can have conflicting requirements for application compatibility (a low abstraction level) versus inspection capability (a high abstraction level). In the following paragraphs, the operating principles of a PEP will be discussed.

A. Binding of information content to security policy

The handling, exchange, and use of information may be subject to security requirements, expressed by a security policy. There must be a binding between the information (whether it is represented as a stream or an object) and the actual security policy. The binding can be represented in several ways:

- 1) Implicit from CoT: the security policy is mandatory for all information stored, handled or used within the CoT,
- 2) Implicit from content: the security policy is chosen based on protocol information and inspection of the information content,
- 3) Explicit from meta data: the security policy is chosen expressed by meta data of the information.

It should be clear that option 1 or 2 will exclude the use of a data centric exchange policy (cf. Section I-B) since the data centric policy needs to be explicitly formulated in the form of meta data. For the rest of the paper, it is assumed that the binding is based on meta data in those cases where data centric exchange policy is discussed.

B. Non-forgability

It is important that the mechanisms which bind the information to its security policies are also protected from forgery or tampering. Anyone with access to modification of the information should not be able to modify the security policy binding. This requirement excludes the use of most meta data arrangements, like ID3, EXIF, document information found in Microsoft Word, or meta data elements in XML, unless they are protected from modification. The protection mechanism must not only protect the meta data from modification, but also protect the content from modification, e.g., prevent adding more sensitive content after the meta data has been assigned. The explicit binding mechanism also requires the PEP to validate the authenticity and authorization of the instance that applies the meta data.

C. Protected meta data

There are well understood mechanisms for protecting data from modification. *Message authentication codes* (MAC) provide simple mechanisms for sealing information [3, p.298]. A MAC arrangement requires the subject making the seal and the subject which verifies it to share a secret key. The scalability of a MAC arrangement is limited, where every pair and group of correspondents maintain shared secrets [2, p.436].

A different arrangement may use *digital signatures*. Digital signatures (signatures for short) are results from cryptographic computations which seal the information through the use of a *private key*. The corresponding *public key* can be used for verification of the binding. Since private keys used for signing are usually individually issued, a digital signature also proves the originator of the meta data. The calculation of a digital signature requires that the entire information content be known and accessible to computations; this makes it impossible to apply digital signatures to information streams. An introduction to the theory of *public key cryptography* may be found in [2, ch.8].

D. Statefulness of PEPs

Depending on the construction and the abstraction of the PEP it may operate without memories from earlier operations. Operation in a stateful manner requires memory from previous transactions. If the PEP operates as a firewall with NAT, then the PEP will have to remember the IP/Port mapping for open TCP connections and consequently operate in a stateful manner. A PEP operating as an HTTP proxy may operate in a stateless manner where only one transaction is monitored at a time. The advantage of a stateless PEP is transparency during crash recovery, load balancing and fail-over mechanisms, since restarting other nodes is not required.

The stateless property also refers to the management state of the PEP, and it is a requirement that a stateless PEP is not configured with a guard centric exchange policy (cf. Section I-B). A guard centric exchange policy is a dynamic property and considered to be a part of the PEP's state space.

We note that two PEPs may need to coordinate their activities, guarding traffic in opposite directions. A service invocation will consist of a request message in one direction and a response in the other. Both messages need to be approved for the transfer by both PEPs even if they employ different inspection and action strategies. Policies for incoming traffic must be reasonable and allow normal requests to pass through for processing.

For stateful PEPs, this need for coordination is more complicated, since *state synchronization* requires communication between the PEPs with a carefully defined message ordering. Restart of one PEP may require restart to the other in order to preserve a coordinated state in both.

E. PEP discovery

A sender may not know if the intended recipient is a member of different CoT, and may not know the network endpoint address of the exchange service. The transport service middleware must be able to solve both problems and shall identify *if* this is an inter-CoT operation and assist the sender




Fig. 4. The structure of an Inter-CoT graph

in finding the PEP's network endpoint address, which protocol to use and which credentials are required for the operation.

These tasks are called *PEP discovery* and will not only identify which PEP to use and how to use it, but also authenticate the PEP to make sure that it is not a fake node.

Discovery information may be associated with the recipient's IP address, e-mail address, or target URL. Depending on the protocol layer of the PEP abstraction, the recipient will be indicated in the PDU (Protocol Data Unit) and can be used in a lookup operation. The result of the lookup operation will be the necessary information to complete the transport operation. For the lowest abstraction layer (the PEP looks like an IP router), the client's routing table may suffice.

F. Routing across the inter-CoT graph

A route across one or more CoTs to a destination may be represented by a graph where the CoTs are the nodes and the inter-CoT links (including the PEP pair) are the vertices. We call this graph the *inter-CoT graph*. A prerequisite for such an arrangement is that at least one CoT has several PEPs and inter-CoT links. Figure 4 shows the structure of one inter-CoT graph.

Inter-CoT links may be established statically or dynamically, like MANET links. Similarly, the PEP discovery mechanisms must recognize the present graph in order to indicate the next routing hop, and the transport protocols must solve reliability problems related to loops, stale routes, and lost PDUs.

An inter-CoT graph also raises interesting challenges related to content inspection and policy decision. A PEP using guard based exchange policy only applies policies enforced in its own domain, even if data objects originate in a different domain. However, it is the policy of the originating domain that should be enforced, knowing that the data object could be exposed to transit-CoTs. This means that the content should be tagged with the policy of its originating CoT, i.e. a data centric policy.

A complete and general solution to this arrangement requires more research. A simplification of the problem would be to require that all inter-CoT traffic only traverse one inter-CoT link (PEPs form a complete graph), where only one PEP pair would be responsible for policy enforcement. In a mobile

tactical network, use of long range radio connections (e.g. a vehicle mounted radio system) may eliminate the need for more hops and limit a connection to one inter-CoT link.

G. Security labels

The term *security label* describes a category of protected meta data (cf. Section III-C). Security labels are elements added to documents to bind security related designations to different parts of the document. A security label may be trusted to bind unmodified information, created by a digital signature holder that can be identified, to a security policy. A label is protected from modifications by the digital signature of the originator. A label is valid for a designated duration of time and its content translate directly into the required security policy. Label handling requires proper infrastructure services for the generation and validation of signatures, as well as adequate protection for the processes themselves.

H. Validation of security labels

A PEP must validate a security label's digital signature before it processes the label content. Validation of a digital signature is a well-known challenge. The validation steps outlined in RFC 5280 include (but are not limited to) the following:

- Verifying the correctness of the hash value
- Verifying the correctness of the signature value
- Validating the public key certificate corresponding to the signature, using cross certificates and revocation information sources as required.

The validation procedure may securely identify the originator and that the originator used an approved key. Though, validation does not reveal that the signature belongs to an *authorized* originator.

I. Authorization control

In order to assess the authorization of the signing originator, several approaches can be taken, including:

- 1) Assuming that all certificates issued (by this particular authority) belong to authorized originators
- 2) Noting certain elements of the certificate that indicate the originator's authorities
- 3) Employing a separate service for verifying authorization.

Approaches 1 and 2 may be sufficient in small scale contexts, where there are not multiple label values requiring different authorizations, and where one issuing authority is used for one application only.

Approach 3 will either require a distinct service infrastructure, e.g., as outlined in XACML [1], or require that the public key certificate include necessary information to assess the authorization. The information could be a collection of name-value attributes for *attribute based access control* (ABAC). The subject attributes (included in the public key certificate as private extensions) must also be issued by a trusted source and be protected against tampering. This is seldom performed in practice because:

- The lifetime of subject attributes are shorter than the lifetime of the key itself. Therefore, private extensions may increase the frequency of issuing certificates which, in turn, increases the size of revocation list as previous certificates must be revoked.
- The subject attributes may be issued by several authorities or a different authority than that of the public key. A practical public key certificate with several signatures does not exist.

A possible approach is to let a certificate authority issue key certificates, and let a subordinate authority add subject attributes into the same certificate object. This approach allows attributes and keys to be issued by separate authorities, and subject attributes may have shorter lifetimes than the keys. The Gismo IdM [4] implements this hierarchical principle and demonstrates protocols and data structures that work well in a tactical coalition network.

J. Trust relations and system evaluation

Trust relations inside a system indicate the extent of a trust domain that must be inspected in its entirety during security evaluation of the policy decision process. This is similar to a TCB (Trusted Computing Base) concept, where the entire domain, including the hardware, the OS, and the application software is subject to security evaluation. When a PEP employs a policy (based on an explicit information or meta data) it establishes a trust relation with the originator that assigned the meta data labels. The validation of the security label and the subsequent access control relies on the integrity of these trust relations. When implicit methods are used (i.e., when policies are assigned based on content inspection in a guard-centric manner) there are no external trust relations formed.

IV. APPLICATIONS AND BENEFITS

Although the restrictions on inter-CoT communication enforced by PEPs appears to limit the collaboration between coalition partners, they actually have the opposite effect. The lack of trust, which is described in the introduction of the paper, would otherwise inhibit the partners to connect their networks and to collaborate. Even though they share interest within an information domain they will not accept the risk represented by unification of their IP networks. Section IV-F will discuss the technical challenges by letting the coalition partners share IP routes in order to improve the connectivity of the MANET, yet deploying PEPs on every node for policy enforcement duties.

Given the flexibility and functions discussed in Section III, one notes that PEPs adapt better than guards to the dynamic environment of tactical operations because their data centric policies provides a better security harmonization across the network.

The remainder of this section will address a range of practical applications and benefits related to the CoT and PEP concepts.

A. Identity Management framework

For the purposes of authentication and access control during link establishment and validation of security labels,

an infrastructure is needed to verify the public keys and the originator's credentials, because it is not feasible to rely on manual deployment of keys and credentials. The service that is provided by such an infrastructure is called *Identity Management*, which is a more comprehensive concept than *Public Key Infrastructure*.

One of the responsibilities of an identity management system is cross domain trust relations. In a *cross-domain* operation, where the credentials of each side of the inter-CoT links are issued by independent authorities, a *trust relation* is needed between the authorities for successful validation to take place.

In a tactical operation, an identity management system must limit its requirements for network capacity and connectivity. Identity management services and related protocols (like signature creation and verification) tend to be expensive in terms of network capacity. In addition, keys and signatures are long (e.g., 2048 bits), revocation lists are large and frequently distributed, and validation servers require uninterrupted connectivity.

FFI (Norwegian Defence Research Establishment) has established a significant body of research on identity management in tactical coalition networks. Beside the integration of authentication and access control, the *Gismo IdM* development effort has resulted in efficient protocols for service invocation where authentication and access control are offered without extra protocol round trips. *Gismo IdM* [4] is an experimental prototype whose architecture and protocol design should be considered for future research and development in securing wireless tactical coalition networks.

Although a PKI is normally the nearest choice for management of keys and identities, it is flawed by a number of poor design decisions. The idea of certificate revocation is one of them, which is a costly and confusing mechanism that takes up much network capacity, and leaves the validating parties with possible dilemmas, and exacerbates the problem related to cross-domain operation.[5] The other problem is its lack of support for access control decisions, since a public key certificate is not able to convey attributes or roles. Access control therefore requires separate infrastructures which should not be necessary.

B. Policy and authorization representation

This paper has not discussed in detail possible representations of exchange policies or sender authorization. Many representation forms are possible and their investigation remains to be done. One possible form is to apply an ABAC (Attribute Based Access Control) where policies are expressed in the form of *access rules* which are boolean expressions evaluated over the *subject attributes* of the recipient or the next PEP. Likewise, the sender's authorization is expressed as an access rule (not as subject attributes which would be the normal way to express authorizations) in order to restrict the possible set of recipients for exchanged data.[6]

Pivotal to this approach is the presence of an identity management system (cf. Section IV-A) with the authority to issue trusted attestations of subject public key and attributes.

C. CoT graph of non-adjacent nodes

Under certain circumstances the nodes of a CoT are not connected with links, but through routes from an underlying infrastructure. One example could be that the intelligence community of coalition partners wish to share information which is still protected from other users in their respective networks. The architectural framework of a link-connected CoT also applies to this route-connected CoT. On the conceptual level, a CoT "route graph" will work similarly to an inter-CoT graph as shown in Figure 4, but on the implementation level there are certain differences that must be observed:

- The routes that connect nodes of the same CoT (akin to intra-CoT links) will need to be protected by a tunnelling arrangement.
- The establishment of a tunnel (e.g., by IPSec) may involve separate authentication and key management, in addition to the authentication which takes place between PEPs.
- While link level peer discovery may be performed automatically through broadcast messages on the communication media, route discovery will require either manual configuration or employing external discovery services.
- Identity management services (cf. Section IV-A) may be obtained through network connections outside the tunnels, where a larger pool of reliable network resources may be more accessible.

PEP pairs will intercept and monitor the traffic for control purposes on either side of routes in inter-CoT connections. Care must be taken to deploy the PEPs so that they cannot be bypassed.

D. Relevance to tactical and sensor networks

Security in tactical and sensor networks is a fundamental concern. In addition to the inherited network security concerns of wired networks, these networks endure unique risks and vulnerabilities that are associated with their open medium, flexible topology, dynamic membership rules, simple network-formation algorithms, and the routing capabilities afforded to each node. While these MANETs offer the benefit of reducing the required deployment and management resources compared to those of fixed networks, they do, however, exist with risks that may be mitigated by the CoT concept.

Due to the dynamic nature of MANETs, the use of one particular node as an exchange point to other networks may not be adequate. In the event of *partitioning* of the network into "islands" there should be fail-over mechanisms where other nodes can resume the responsibilities of a PEP. A fail-over mechanism sets the following requirements to the PEP architecture:

- The PEP must be stateless, so that it can start its operation in the middle of application sessions.
- Exchange policies should be data centric (attached to the exchanged information object) to avoid configuration inconsistency between PEPs.

- There must be a PEP discovery arrangement which allows MANET nodes to locate the present PEP and to decide its access protocol.
- The PEP should be able to authenticate itself to MANET nodes and to prove its authorization to offer a PEP service.

Except for the first bullet point, existing guards do not meet these criteria and are therefore unlikely to work well in a tactical environment like a MANET. The model proposed in this paper, however, meet these criteria, since the CoT assumes data centric policies, stateful PEPs and discovery arrangements.

A MANET will need communication to higher echelons through a “reach back” link. The general rules for separating concern and responsibilities for the information exchange presented in this paper will facilitate the interoperability between the security mechanisms of the different CoTs.

E. Communities of Interest vs. Communities of Trust

The term Communities of Trust is established as a contrast to the term Communities of Interest (CoI). The distinction is necessary in order to understand why groups of computer nodes need their activities protected from the surroundings, and share information with the same surrounding at the same time. The two terms indicate that a willingness to share specific information is distinct from the trust which exists between nodes in an unprotected network.

The separation between nodes in a CoI can be stronger than between nodes in a CoT. It is therefore possible to regard CoI as a construct that overarches the CoT in the sense that one CoI may consist of several CoTs. The opposite is possible although the separation between members and non-members of a CoI will be weaker.

F. CoTs in a fully unified network

A coalition MANET where nodes from different nations share routing information and routing services may improve its connectivity and transport capacity.[7]

In theory, it is possible to build CoTs within a MANET. Intra-CoT communication would in this case need to be encrypted since the network path could go through foreign nodes. Any pair (or group, in case of multicast communication) of nodes would need to set up VPN tunnels prior to communication. A VPN tunnel takes time to set up, and a large number of potential tunnels create a demand for large scale key management.

The PEPs used to control information exchange would be built into each node, which requires a hardened operating system (or a separation kernel) that is able to offer the necessary separation and protection. Furthermore, it is nearly impossible to avoid information leaks between CoTs that are caused by traffic flow analysis, routing information and software exploits.

Conclusively, it is possible to obtain reasonable CoT separation even in a grid where all routes are shared. However this is costly in terms of architecture (hardened OS), system administration (key management) and performance (VPN tunnel creation).

V. CONCLUSION AND FUTURE RESEARCH

In this paper, we have outlined a framework for the organization of a coalition network, or other partitioned networks with need for separation and security. The model allows for a fine grained and secure information exchange through exchange points (PEPs) which enforce an exchange policy and controls the authorizations of the sender and the receiver.

The PEP concept can be built with different abstraction levels, depending on the need for inspection and the required application transparency.

An important property of the proposed framework is the emphasis on data-centric policies, stateless PEPs, PEP discovery and identity management, which allows the exchange mechanisms to be applied to tactical/sensor networks as well as strategic networks.

This proposal raises many research questions, including:

- How can inspection strategies safely apply security policies?
- How can an architecture be designed for non-bypassable PEPs?
- How can identity management services be deployed in tactical networks?
- How can written security requirements be mapped to executable handling rules?

The research questions will be investigated in the future.

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] Andrew S. Tanenbaum and Marten van Steen. *Distributed Systems. Principles and Paradigms*. Prentice Hall, 2002.
- [3] Colouris G. and Dollimore J. and Kindberg T. *Distributed Systems. Concepts and Design*. Addison Wesley, 2005.
- [4] 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.
- [5] Anders Fongen. Optimization of a public key infrastructure. In *IEEE MILCOM*, Baltimore, MD, USA, Nov 2011.
- [6] Anders Fongen and Federico Mancini. High-Assurance Information Exchange with PubSub and ABAC. In *IEEE MILCOM*, Baltimore, MD, USA, 2014.
- [7] Marian Hauge, Margrete Brose, Jostein Sander, and Jon Andersson. Multi-Topology Routing for QoS Support in the CoNSIS Convoy MANET. In *Military Communication and Information Systems Conference (MCC)*, Gdansk, Poland, 2012.
- [8] Øyvind Hvinden, Alan Murdock, Michael Rudack, Martin Booth, Martin Diepstraten, Alberto Domingo, Sven Kuehne, and Leon Schenkels. Information Exchange Gateway Roadmap. Technical Report Reference document 2666 Draft version 1.01, NATO C3 Agency, March 2010.
- [9] R. Shirey. Internet Security Glossary, Version 2. RFC 4949 (Informational), August 2007.
- [10] Konrad Wrona and Geir Hallingstad. Development of high assurance guards for NATO. In *Military Communication and Information Systems Conference (MCC)*, Gdansk, Poland, 2012.
- [11] Konrad Wrona and Nadja Menz. Protection Profile for the NATO High Assurance ABAC Guard (HAAG). Technical Report TR-2012-SPW008418-13-4, NATO C3 Agency, March 2013.