Coalition Networks for Secure Information Sharing

CoNSIS Task 3
Final Report

Version 1.0

3 December 2012

EXECUTIVE SUMMARY

The Coalition Network for Secure Information Sharing (CoNSIS) is a multinational project focusing on technologies and methods that will facilitate the partners' abilities to share information and services securely in ad-hoc coalitions, and between military and civil communication systems, within the communications constraints of mobile tactical forces.

The work done within the CoNSIS project has been divided into a number of tasks, each focusing on a different aspect of interoperability issues. This **Task 3** report covers the security mechanisms and methods employed for integration and interoperability of heterogeneous, coalition networks. The likely next expansion of military networks will be into highly mobile platforms on the tactical edge. The communications constraints of such mobile tactical networks call for efficient security mechanisms.

The CoNSIS project explores the concept of a common core network that transports data between user domains called coloured enclaves. The transport network concept provides several benefits for users and operators due to the ability to share communication infrastructure. One benefit is that coloured enclaves belonging to different security domains[1], both national and coalition security domains, can use the same transport network. Thus, separate communication infrastructure for each security domain is not required. However, associating national networks together to form a global coalition TN poses certain threats to the national networks. Mechanisms and procedures that thwart these threats and protect the authenticity, integrity and availability of the networks are addressed by Task 3.

Key management is crucial in mobile tactical networks. Coalition partners that move into an area need the proper cryptographic key in order to start communicating using a crypto device (e.g. an IPsec device). Task 3 addresses key management solutions that use digital signatures, which require a Public Key Infrastructure (PKI). However, the PKI services are not designed for the communications constraints of tactical networks. Traditional PKI services require a high amount of system resources, in terms of networking capacity and networking connectivity. Task 3 investigates scalability and optimization opportunities of the PKI protocols.

Task 3 addresses confidentiality, authenticity and integrity protection of the user traffic between coloured enclaves. We also address traffic flow confidentiality (TFC) solutions that may be used in scenarios where the data traffic traverses third party networks, i.e. public networks or military coalition networks not supporting traffic flow confidentiality.

Further, cross-domain solutions that support information exchange between different security domains are investigated and demonstrated. Interconnection between two different security domains has traditionally been achieved using solutions such as diodes, only providing a one-way information flow. However, a guard based approached may be used to provide a more flexible two-way solution. Task 3 also investigates cross-domain solutions for exchanging management data between a classified domain (coloured enclave) and an unclassified domain in the transport network.

[1] A *Security domain* is defined as a collection of entities to which applies a single security policy enforced by a single authority.

Lastly, protected and controlled communication between civilian and military networks is addressed. The presence of non-governmental organizations (NGOs) in a conflict zone is frequently seen, and their operations may be safer and more efficient through communication with military forces. The NGO requirements for a Civilian-Military communication arrangement are expected to be different from the military requirements for such an arrangement. Task 3 investigates and demonstrates security technologies that can meet both the NGO and military requirements.

Table of Contents

1 INTRODUCTION 1

1.1 CoNSIS 1
1.2 Motivation 2
1.3 Report structure 3

2 TRANSPORT NETWORK CONCEPT 4

3 TRANSPORT NETWORK PROTECTION 6

3.1 Core Network Protection 6
3.1.1 Introduction 6
3.1.2 Category-1 flows 7
3.1.3 Category-2 flows 7
3.1.4 Category-3 flows 9
3.1.5 Category-4 flows 11
3.1.6 Future work 11
3.2 Network Authentication Header (NetAH) 11
3.2.1 Introduction 11
3.2.2 NetAH Implementation 12
3.2.3 Experiment 1 - NetAH protected QoS signaling 12
3.2.4 Experiment 2 - NetAH to build an authenticated C-TNS 13
3.2.5 Conclusions/future work 13
3.3 MLS routing 14
3.3.1 Introduction 14
3.3.2 Experiment setup 14
3.3.3 Classification of router links 15
3.3.4 MLS privileges and routing of packets 16
3.3.5 Conclusions 16

4 KEY MANAGEMENT ISSUES 17

4.1 Scalability of a tactical PKI 17
4.1.1 Introduction 17
4.1.2 Overhead evaluations 17
4.1.3 Overhead reduction 19
4.1.4 Recommendations 21
4.2 Optimization of PKI protocols 21
4.2.1 Introduction 21
4.2.2 PKIX standards and the use of COTS software 21
4.2.3 Certificate revocation 22
4.2.4 The functional components and operations of a PKI 22
4.2.5 PKI operations - size of data units 23
4.2.6 PKI optimization opportunities 23
4.2.7 Cross domain operation of a PKI 24
4.2.8 Conclusion 24
4.3 IPsec discovery protocol (IDP) 25
4.3.1 Introduction 25
4.3.2 Functionality 25
4.3.3 Usage with MIKE key establishment protocol 25
4.3.4 Experiment 26
4.3.5 Lessons learned/future work 26
4.4 MIKE study 27
4.4.1 Introduction 27
4.4.2 Scenario 27

4.4.3 Outline of MIKE 28
4.4.4 Assessment 28
4.4.5 Proposed enhancements 29
4.4.6 Conclusions 29

5 PROTECTION OF USER TRAFFIC 30

5.1 CROSS-DOMAIN INFORMATION EXCHANGE 30
5.1.1 Introduction 30
5.1.2 Cross-domain information exchange experiment 31
5.1.3 Cross-domain access control study 32
5.1.4 Conclusion 34
5.2 PROTECTED COMMUNICATION BETWEEN MILITARY AND CIVILIAN NETWORKS 34
5.2.1 Introduction 34
5.2.2 Technical Requirements 35
5.2.3 The Prototype Configuration 35
5.2.4 The GISMO IdM architecture 36
5.2.5 Service Invocation 37
5.2.6 Messaging Protocols 37
5.2.7 SOAP guard and confidentiality labeling 37
5.2.8 Conclusion 37
5.3 MULTILEVEL SECURITY AND CONSIS NETWORK MANAGEMENT 38
5.3.1 Introduction 38
5.3.2 Current architecture 38
5.3.3 Concept 38
5.3.4 Cross Domain Guard 40
5.3.5 Management proxy 41
5.3.6 Conclusion 41
5.4 TRAFFIC FLOW CONFIDENTIALITY 42
5.4.1 Introduction 42
5.4.2 TFC architecture 42
5.4.3 Experiments 43
5.4.4 Conclusions 46

6 CONCLUSIONS 46

REFERENCES 48

GLOSSARY 50

1 INTRODUCTION

Coalition Networks for Secure Information Sharing (CoNSIS) is a multinational project consisting of members from Germany, France, USA, and Norway, with participants from both research institutions and industry. The objectives of the CoNSIS project are to develop, implement, test, and demonstrate technologies and methods that will facilitate the partners' abilities to share information and services securely in ad-hoc coalitions, and between military and civil communication systems, within the communications constraints of mobile tactical forces.

The project has focused on practical application of information infrastructure technologies in a network of networks, consisting of a variety of low capability network technologies. The work done within the CoNSIS group has been divided into a number of tasks, each focusing on a different aspect of interoperability issues. This **Task 3** report covers the security mechanisms and methods employed for integration and interoperability of heterogeneous, coalition networks.

During June 2012 CoNSIS conducted a large-scale joint distributed experiment in which all the different aspects of technical interoperability were tested; integrating the work of all the task groups of CoNSIS. The joint experiment was based on a scenario involving both military and non-governmental organizations. The scenario takes place in a country torn by civil war. An international coalition is involved in this conflict to protect civilians and initiate the peace process. The scenario is described in the CoNSIS System and Experimentation Architectures document [1].

1.1 CoNSIS

The CoNSIS areas of work are broken down into five major tasks. Task 3 is responsible for the security area. The four other tasks are as follows:

- Task 1 - Communication Services
- Task 2 - Information and Integration Services (SOA)
- Task 4 - Management
- Task 5 - Architecture, Test & Demonstration Coordination

Task 1 provides activities that are undertaken to support a general goal of an overall NII infrastructure based on IP technology. The focus of this task is on demonstrating solutions that will work within the communications constraints and dynamic topology imposed by highly mobile tactical networks. Communication services within tactical systems are analyzed towards their ability to support SOA core services (e.g. discovery), real time services (e.g. VoIP and VTCoIP) and streaming services (e.g. TADIL).

Task 2 demonstrates the applicability of the SOA approach in a multinational military environment, federating the SOAs of each nation. The task is taking into account the constraints of security and the constraints applicable to highly mobile tactical forces (including limited bandwidth, intermittent communications, high rate of change of network topology, and the need to make decisions quickly).

Task 4 explores, specifies, and demonstrates mechanisms for automatic management services and service levels in coalition networks. The main challenge is to automate the end-to-end management across multiple security domains during changing operational and network situations. This requires mechanisms to detect changes and operational policies that define the actions to be taken. A second area of interest is the detection of a jammer attack autonomously per vehicle or on a cooperative basis.

Task 5 develops an overall Experimentation Architecture for CoNSIS. This architecture will define the way in which the deliveries of tasks 1 to 4 are to be integrated. The task will also carry out the overall co-ordination and planning of the CoNSIS project. It will provide reporting and dissemination of the results of CoNSIS during and upon completion of the project. The intention is to demonstrate technical results that can transition to an operational demonstration/scenario (outside this MoU).

1.2 Motivation

Task 3 investigates, develops, tests, and demonstrates security mechanisms for use for integration and interoperability of heterogeneous coalition networks. The likely next expansion of military networks will be into highly mobile platforms on the tactical edge. The communications constraints of such mobile tactical networks call for efficient security mechanisms.

The CoNSIS project explores the concept of a common core network that transports data between user domains called coloured enclaves. The core network is called a transport network (TN) and consists of national and coalition networks. The transport network concept provides several benefits for users and operators due to the ability to share communication infrastructure. One benefit is that coloured enclaves belonging to different security domains[2], both national and coalition security domains, can use the same transport network. Thus, separate communication infrastructure for each security domain is not required. However, associating national networks together to form a global coalition TN poses certain threats to the national networks. Examples of threats are denial of service attacks, installation of false routes in the routing protocol and modification of the QoS signaling in the IP header. Mechanisms and procedures that thwart these threats and protect the authenticity, integrity and availability of the networks are addressed by Task 3.

[2] A *Security domain* is defined as a collection of entities to which applies a single security policy enforced by a single authority.

Key management is crucial in mobile tactical networks. Coalition partners that move into an area need the proper cryptographic key in order to start communicating using a crypto device (e.g. an IPsec device). Task 3 addresses key management solutions that use digital signatures, which require a *Public Key Infrastructure* (PKI). PKI refers to the set of services, software and protocols necessary to manage a set of public key pairs (private/public key). Public key schemes may be desirable in a tactical network whose management must be as easy and flexible as possible. However, the PKI services are not designed for the communications constraints of tactical networks. Traditional PKI services require a high amount of system resources, in terms of networking capacity (data volume and number of protocol round trips) and networking connectivity (frequency of protocol invocations). Task 3 investigates scalability and optimization opportunities of the PKI protocols.

Task 3 addresses confidentiality, authenticity and integrity protection of the user traffic between coloured enclaves. We also address traffic flow confidentiality (TFC) solutions that may be used in scenarios where the data traffic traverses third party networks, i.e. public networks or military coalition networks not supporting traffic flow confidentiality.

Further, cross-domain solutions that support information exchange between different security domains are investigated and demonstrated. Interconnection between two different security domains has traditionally been achieved using solutions such as diodes, only providing a one-way information flow. However, a guard based approached may be used to provide a more flexible two-way solution. A guard has the potential to provide two-way information exchange between a classified military domain and a non-governmental organization (i.e., an unclassified domain) where messages are released from the military domain to the NGO based on confidentiality labels. Task 3 also investigates cross-domain solutions for exchanging management data between a classified domain (coloured enclave) and an unclassified domain in the transport network.

Lastly, protected and controlled communication between civilian and military networks is addressed. The presence of non-governmental organizations (NGOs) in a conflict zone is frequently seen, and their operations may be safer and more efficient through communication with military forces. The NGO requirements for a Civilian-Military communication arrangement are expected to be different from the military requirements for such an arrangement. Task 3 investigates and demonstrates security technologies that can meet both the NGO and military requirements.

1.3 Report structure

The remainder of this report is structured as follows

- Chapter 2 provides an introduction to the Transport Network concept

- Chapter 3 presents the work done on protecting the transport network; i.e. security mechanisms and methods needed in the transport network
- Chapter 4 presents the work on key management
- Chapter 5: presents the work done on protecting the user traffic; i.e. security mechanisms and methods in the coloured enclaves
- Chapter 5 presents the conclusions of our work

2 TRANSPORT NETWORK CONCEPT

The CoNSIS reference model consists of a core IP network to which user domains are connected via IPsec crypto devices.

The core network itself is composed of a number of interworking communication systems operated by different administrative authorities.

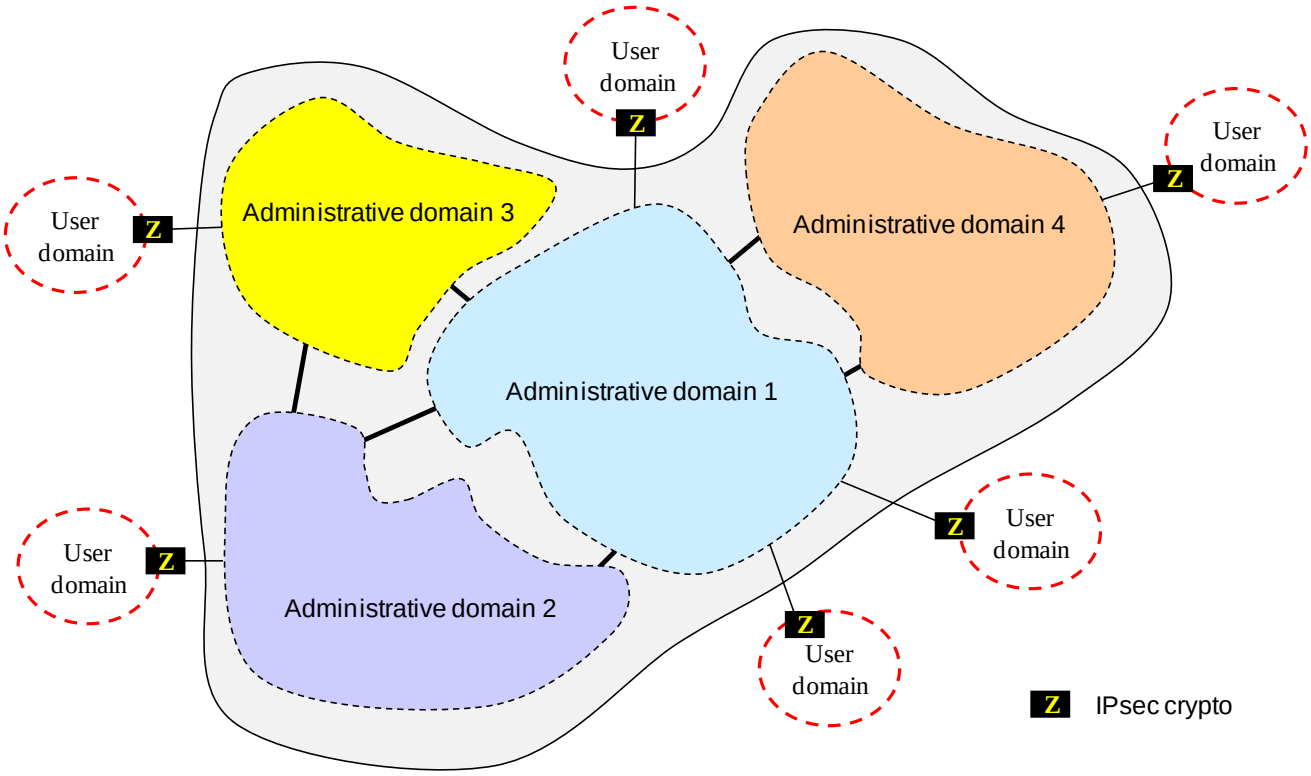

Figure 1: Main elements of the CoNSIS architecture

In the CoNSIS terminology, the core network as a whole is called the **Transport Network** (TN) and the individual administrative black domains are referred to as **Transport Network Segments** (TNSs).

A TNS is a set of nodes which run the same Interior Gateway Protocol. From the routing perspective, it is thus an IP Autonomous System. The most natural situation is one in which all network elements of such an AS belong to the same nation. A TNS of this kind is referred to as an National TNS, or **N-TNS**.

However, there can also be cases when for operational reasons the telecommunication assets of two or more nations have to be mixed within the same AS. This construct is then administered by a coalition operator and is called a Coalition TNS, or **C-TNS**.

The set of hosts and network elements located in one site and connected to the TN via an IPsec device is generically referred to as a **Coloured Enclave** (CE). CEs with different levels of classification and right to know can be connected to the same Transport Network since the traffic they transmit is made opaque by end-to-end IPsec encryption.

A Coloured Enclave can be embedded in another CE of a lower classification level. It is referred to as an **Inner Coloured Enclave** (ICE). An ICE is isolated from the CE in which it is embedded by its own IPsec device.

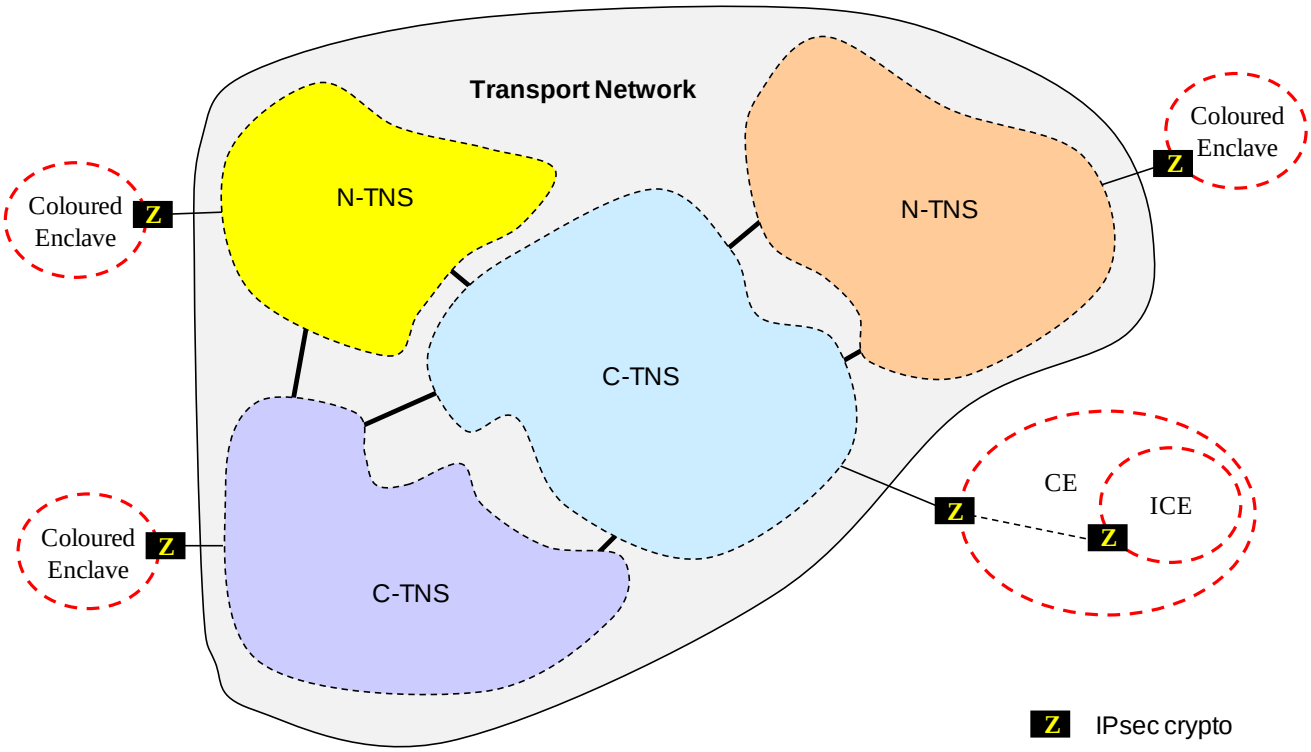

Figure 2: CoNSIS reference model and terminology

This architecture is close to that of PCN in which *Coloured Clouds* (CCs) are connected to a *Protected Core* composed of *Protected Core Segments* (PCSs). PCN CCs are the counterpart of CoNSIS CEs, the PCS is the counterpart of the TN and PCSs are that of TNSs.

However, the two reference models are not identical. In particular, CoNSIS administrative domains are not assumed to have exactly the functions as PCSs regarding e.g. security protection or the management of SLAs, and they interwork via interfaces which are not supposed to have the same features as the PCS-1 interface.

Likewise, the generic interface between CoNSIS user domains and the core network is not necessarily compliant with the PCS-2 interface.

3 TRANSPORT NETWORK PROTECTION

In this chapter we provide solutions, experimentation results, experience and future work for the transport network protection topics.

3.1 Core Network Protection

3.1.1 Introduction

Associating Transport Network Segments together to form a global coalition TN can pose certain threats to national networks. The goal of the core network protection study was to identify these threats and propose mechanisms and procedures to thwart them.

The output of the core network protection work is a theoretical study documented in [2]. In addition, several experiments on the core network theme were conducted during the CoNSIS project and a description of the experiments and of their major findings can be found in [3]. The core network protection work documented in these reports is briefly summarised in the following.

The analysis started from the fact that attacks from a TNS to another TNS are necessarily borne by flows exchanged between these autonomous systems. These inter-network flows fall into the following general categories:

- 1. User to user flows,
- 2. Flows whose both endpoints are border routers,
- 3. Flows exchanged between internal network elements on the two sides,
- 4. Flows between a network function on one side and a user on the other side.

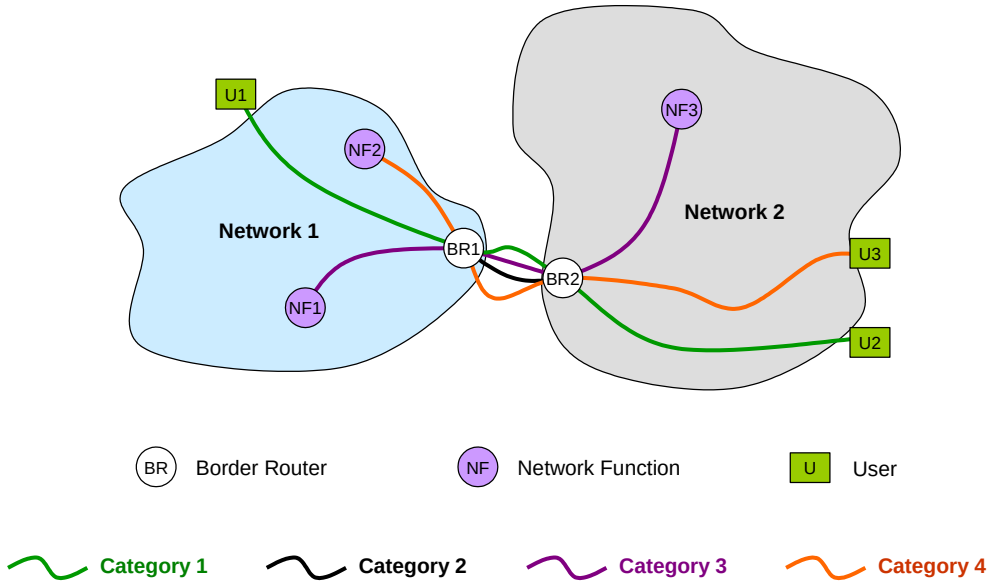

Figure 3: General model of flows exchanged between two TNSs

3.1.2 Category-1 flows

By definition, user-to-user flows are addressed to hosts in Coloured Enclaves and cannot directly interfere with elements in the black domain. However, by requiring an excess of network resources, they can be the vector of a Denial Of Service (DOS) attack.

A DOS attack can be easily detected and blocked at the border between two TNSs if it simply consists of trying to flood the target network with an excess of data. The inter-TNS SLA can be assumed to specify what amount of traffic can be accepted by the receiving network, overall and per class of service.

However, a DOS attack can be more subtle, and can be aimed at a particular segment of the receiving network without infringing the inter-TNS SLA. One or several illegitimate flows which all abide by the SLA can for example saturate a link in the target TNS.

Detecting such a subtle DOS requires measurements of the traffic mix over *each link* in the target TNS. Practical experiments have shown that these measurements could be performed with software probes which are compatible with the required compactness of tactical network nodes. However, they also showed that detecting the attack unambiguously was a daunting task, and above all that spotting which flows are illicit within a complex traffic mix was even more tricky. An alarm can be set when a DOS is suspected, but with the current state of the art, a man in the loop is still indispensable to make a decision as to which reaction should be carried out.

3.1.3 Category-2 flows

Transactions between border routers are essentially BGP advertisements. They can bear two major threats to this routing protocol:

- Exhaustion of the processing capability of a border router via the transmission of an abnormally huge number of messages,
- Deception of the routing process via the installation of false routes.

The first threat would typically result from a hostile border router masquerading as the equipment of a legitimate allied nation. It can relatively easily be thwarted by authenticating peer routers in the external networks and by ensuring the integrity of the messages they send.

The second threat is a lot more difficult to combat because the router which issues false advertisements can very well be a legitimate one, and have been deceived itself by another router further upstream.

Two IETF working groups dubbed Secure BGP (S-BGP) and Secure Origin BGP (soBGP) have tackled this issue. But both efforts resulted in complex schemes which left unsolved vulnerabilities, and were discontinued in 2006. In 2009, the Department of Homeland Security of the US government has resumed work on BGP security in the frame of the BGPsec project, but this new attempt had not come to fruition when the CoNSIS theoretical study was completed.

In the absence of a standardised solution, the following pragmatic approach was proposed:

- Prevent attacks on BGP via mutual trust based on the enforcement of engineering rules and good practices within a coalition,
- And as this form of prevention cannot be regarded as 100% safe, complement it with means to detect an attack and counteract it once it *has* happened.

The preventative engineering rules and good practices provide that each TNS participating in a coalition Transport Network must abide by the following conditions:

- Condition 1: as deceptive advertisements can originate from the interior routing process, each TNS must protect the operation of its own IGP. This implies mutual authentication between routers within a TNS, as well as protecting the integrity of routing messages.
- Condition 2: each TNS must protect its routers and all other functions which participate in the generation and transportation of IGP or EGP messages. This implies a physical protection of the relevant equipment, of management centres, as well as means to prevent tampering with management flows.
- Condition 3: each network must ensure its peer TNSs that it does not propagate wrong routes learnt from a third party. This means that a TNS must only exchange routing information with trusted TNSs (i.e. those which abide by conditions 1, 2 and 3) and that it must make its best effort to detect biased routing information using the complementary corrective measures detailed below.

It is actually a lot more practical to detect a false route *once* it is installed than when it is just advertised. If this false route just creates a “black hole” (i.e. does not lead to the advertised addresses), a simple connectivity test will spot it as soon as it is implemented.

However, a false route can indeed lead to the advertised addresses, but via an excessively long and meandering path. In principle, this can be detected by measuring QoS parameters (e.g. transit delay and packet loss rate) end to end to the advertised destination. Experimentations performed during the CoNSIS project have shown that a degradation of QoS indicators could be spotted and that the installation of a false route could be suspected this way. However, interpreting the results automatically proved difficult, and a man in the loop appears still necessary to distinguish between a natural cause and the effect of an attack on BGP.

3.1.4 Category-3 flows

Flows exchanged between internal network elements across the border between two TNSs are the most complex category. They include:

- Signalling flows,
- Management flows,
- And miscellaneous streams which will be collectively referred to as ancillary flows in the following.

3.1.4.1 Signalling flows

Signalling flows which can be exchanged between two ASes are for example RSVP, or PIM.

The attacks associated with PIM are very similar to those BGP can convey and require the same sort of protective measures. In contrast, the main threat RSVP can pose is an excessive consumption of network resources via illegitimate reservations. This can be thwarted by enforcing the inter-TNS SLA in the border router.

The difficulty with these signalling protocols is that they are potentially countless and that the type of attacks they can bear is dependent on each protocol. There is no doubt that specific protective steps can be devised for each of them, but in order to ensure the security of a black core, the first rule is that a protocol of this type must be prohibited at the border between two TNSs by default, and authorised only when appropriate countermeasures against the particular threats it can pose have been implemented.

3.1.4.2 Management flows

Likewise, there could in principle be a large variety of management flows, but a more in-depth analysis of this issue shows that the situation is actually simpler because there is no reason why a management centre in one TNS should monitor (and even less configure) network elements in another TNS. No management transactions between a

management centre and managed elements should therefore cross the border between two national systems, and the flows which can legitimately be exchanged between two administrative black domains can essentially be:

- Flows between two management centres (e.g. to give a foreign operator a summary status of a national TNS).
- Or flows intended for specific measurement purposes on an end-to-end path across two or more TNSs. These include the test streams on which measurements are actually performed and the flows which allow the coordination between probes.

Data exchanged between management centres could bias the decisions made by the receiving party if they are erroneous. They could also bear attacks at data processing level by including viruses or other forms of malicious software. To prevent this sort of attack, the establishment of mutual trust based on a set of engineering rules and good practices similar to those relative to BGP was proposed:

- Condition 1bis:
  - OSS functions must be protected against physical intrusions.
  - Management communications must be protected against intrusions and must guarantee the integrity of the transported data.
- Condition 2bis:
  - The OSS in each TNS must do its best effort to detect attacks against computers and to destroy their bearers.
  - It must only connect to trusted OSSs, i.e. network management systems which abide by conditions 1bis and 2bis.

Attacks associated with measurement flows include:

- Denial of service. Just like those issued by users, flows transmitted by measurement platforms could use up an excessive amount of network resources and prove detrimental to user traffic.
- Tampering with test packets in order to bias measurement results.
- And interfering with the coordination flows between probes, with the same goal of deceiving network operators by supplying them erroneous measurement results.

The first attack requires exactly the same protection as a DOS borne by user flows. The last two require a protection of the integrity of test flows and of flows which control the measurement process. They also require means to prevent the perturbation which could result from the action of a rogue measurement platform, i.e. precautions very similar to those recommended for routers:

- A protection against physical intrusions,
- And a protection against illicit management of the platform, either from a local position or from a remote management centre.

3.1.4.3 Ancillary flows

Ancillary flows are those exchanged between servers in the black domain, such as e.g. DNS, NTP or various directories (e.g. LDAP). Again, the possible types of these servers

are countless, but the following set of precautions can avert the attacks their flows may convey, whatever the nature of the protocol:

- Protection of the relevant functions against intrusions (either physical or via communication ports),
- Mutual authentication between servers or between servers and their clients,
- Guarantee of the integrity of data flows as they are transmitted over networks,
- Connection of servers to trusted servers only.

3.1.5 Category-4 flows

Finally, there is normally no reason why a network element in the black domain should exchange data with a host in a Coloured Enclave, and any attempt to do so will actually fail to get through the IPsec device which isolates the CE. However, there is no convenient means to prevent a router from sending ICMP messages back to an enclave which originated an erroneous packet. This sort of message cannot pose a threat since it will be discarded by the destination crypto device. There would thus be no point in filtering it out at the border between two TNSs, and as some network operators might want to accumulate statistics about the number and types of ICMP messages, it was eventually decided to let these flows through.

Practical experimentations showed that simple ACLs in the border routers could selectively block flows which were prohibited (e.g. packets sent from a CE to an internal network function) and at the same time let ICMP messages through. However, it must be pointed out that defining these ACLs might turn out to be particularly burdensome and error-prone if the addressing plan in use in the TN does not make a clear and straightforward distinction between the user domains and the various national black domains.

3.1.6 Future work

Two major areas for future work were identified during the study:

- Ways to trustworthily detect a DOS attack resulting in the congestion of an internal segment within a TNS, and to automatically identify which flows are illicit.
- Threats posed by the merger of assets from different nations into a single coalition TNS, and ways to thwart these threats.

3.2 Network Authentication Header (NetAH)

3.2.1 Introduction

Motivation: Standard IPsec leaves the mutable fields of the outer IP header unprotected. This makes the QoS signaling in the IP header prone to attacks. The Network Authentication Header (NetAH) is a modified version of the standard IPsec AH. It also protects the QoS signaling in the outer IP header.

The concept was first described in the article ”QoS signaling in IP-based Military Ad Hoc Networks” [4].

The purpose of the CoNSIS Task 3 NetAH-activity has been to implement the concept and gain experience with the scheme through experiments and demonstrations.

Summary: The output of the NetAH work is a software patch and a report describing the implementation [5]. In addition, two different experiments were planned for the international CoNSIS experiment in Greding. But due to time limitations, only one was conducted.

3.2.2 NetAH Implementation

The NetAH is implemented as a hop-by-hop extension header. It is compatible with standard IPsec in the sense that it can be used in addition to this. The SR600 and WM600 tactical communication nodes from Kongsberg were used as experimental platforms. The NetAH can be ported to other experimental platforms. The NetAH patch is available for Linux 2.6.39.4 and it comes with a patch for IPsec-Tools version 0.7.3.

3.2.3 Experiment 1 - NetAH protected QoS signaling

The experiment demonstrates how NetAH can be used to distinguish between two flows of the same DiffServ class and prioritize the one with correct NetAH. This is the main NetAH experiment.

Figure 4 illustrates the test setup. There is a soldier node carrying a camera, and a camera mounted in an NGO vehicle. Both parties forward video streams via two radio nets to vehicle 1. (Actually the NGO node was included in the Squad net during the experiment due to lack of radios). The NetAH is verified in Vehicle 1 before is forwarded to the monitor in Vehicle 2 via the Convoy C-TNS.

As long as there is capacity enough in the C-TNS, video streams from both the camera of the NGO vehicle and the Soldier are forwarded to the monitor in Vehicle 2. But when the bandwidth in the C-TNS is limited, video from the Soldier node that adds correct NetAH is prioritized over the NGO video stream without correct NetAH. At the monitor in vehicle 2 one could see that the stream from the NGO vehicle deteriorate while the other stream did not.

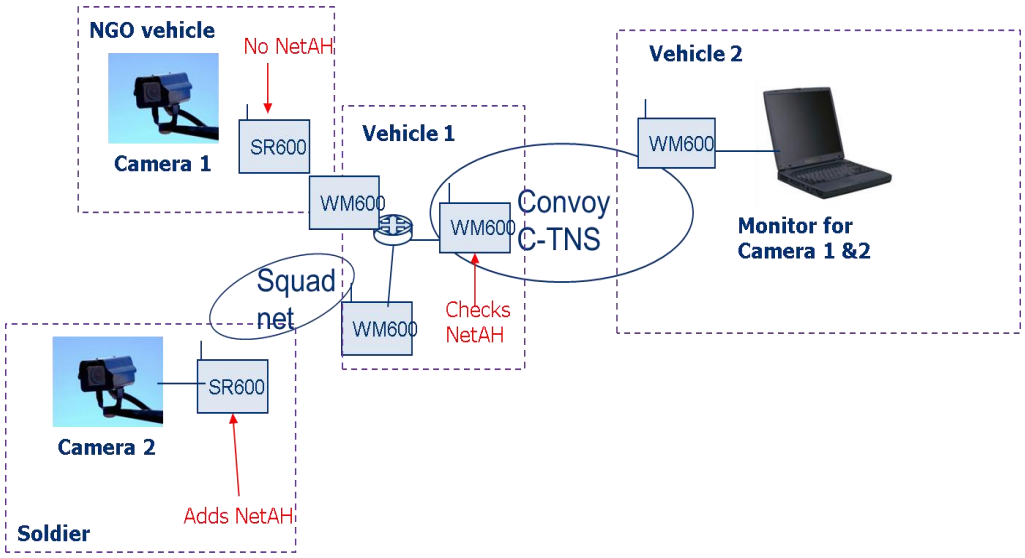

Figure 4 Test setup for demonstration of NetAH protected QoS signalling

3.2.4 Experiment 2 - NetAH to build an authenticated C-TNS

Whereas protection of QoS signaling is the main purpose of NetAH, it can also be used to build an authenticated network. Nodes with the correct NetAH are authorized to become part of the network. Those not presenting a proper NetAH are not included in the network. Only traffic with the proper NetAH is forwarded. Figure 5 shows the test setup. The tactical radios in vehicle 1 and 2 add and check NetAH, and build an authenticated network service. Vehicle 3 does not add a proper NetAH and is not included in the network. This experiment was not conducted in June 2012 due to time limitations.

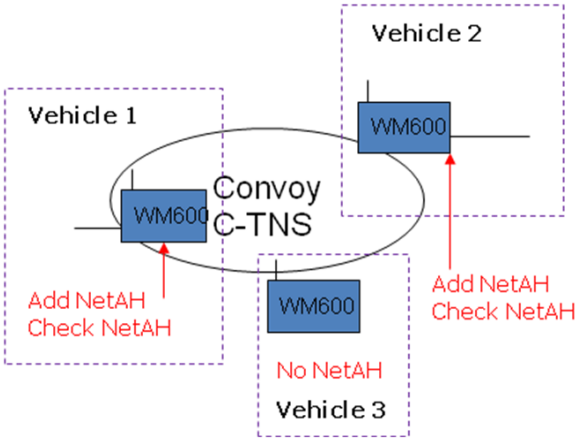

Figure 5 Test setup for using NetAH to build an authenticated C-TNS

3.2.5 Conclusions/future work

The experiment in Greding showed how NetAH can be used to prioritize a military data flow over an NGO one with identical QoS marking. The experiment also demonstrated

NetAH compatibility features. NetAH was added to a stream of data that was already encrypted by standard IPsec. The data were routed along a path that included both non-NetAH-aware as well as Net-AH aware nodes.

NetAH enables authorized nodes to detect if the QoS marking has been changed by unauthorized nodes, but does not prevent them for doing it. In some cases the packet should be dropped, in other cases it is better to accept the packet but send it Best Effort, as done in the CoNSIS experiment. Different policies for the treatment of datagrams with failing NetAH are a topic for further studies.

The use of NetAH to build an authenticated C-TNS represents another use of the NetAH concept. More ways to exploit the NetAH concept and enhancements of the scheme are other topics for further research.

3.3 MLS routing

3.3.1 Introduction

Secure information sharing is crucial in coalition operations. In such operations, actors may interconnect different networks and establish a coalition routing domain.

In military systems, multilevel security (MLS) has applied to user information. No specific MLS scheme has applied to network information[3]. In mobile wireless networks, such information is especially vulnerable to attacks. Network information should be subject to specific security policies, not necessarily similar to the policies regulating user information security. Further, coalition networks may include partners that do not fully trust each other, and interoperability requirements call for new solutions for protection of network information.

The concept of MLS (multilevel security) routing in coalition IP networks was first described in “Multilevel security for IP routing” [6]. It proposes a separation of routing information into different levels of security, and describes how routers may employ multilevel security.

The output of the work on MLS routing are a *proof-of-concept* router prototype developed by Thales NOR. The implementation extends the multi-topology router [7] with multilevel security in the three dimensions confidentiality, integrity and availability. In addition an experiment setup was described. Due to resource limitations the experiment was not conducted at the Joint Distributed Experiment in June 2012.

3.3.2 Experiment setup

[3] Network information includes routing information, QoS signaling and other management information

The objective of the experiment is to demonstrate a multilevel security (MLS) scheme for routing information, using a *proof-of-concept* MLS router prototype. The routing information is classified along three independent dimensions; confidentiality (C), integrity (I) and availability (A). There are two levels (*high*, *low*) for each dimension.

The actors initiating traffic on the network are the Multi National HQ (MNHQ), a national HQ, a military vehicle, an NGO office and an NGO vehicle. The experiment shows how the MLS routers can be used in the CoNSIS scenario with both civilian (NGO) and military actors.

Figure 6 shows the architecture of the experiment. An overlay network is established to interconnect the MLS routers, as the MLS routing protocol is not supported in the CoNSIS network. The overlay network reflects important properties of the CoNSIS network.

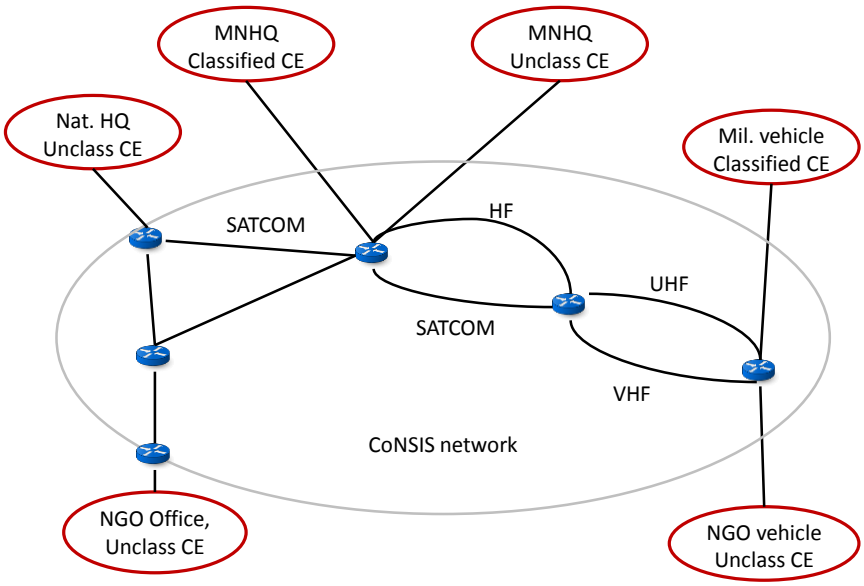

Figure 6: MLS overlay network

The overlay network consists of tunnels in the CoNSIS network. The tunnels can be GRE, SIT (IPv6 over IPv4) or IP6IP6 as supported by standard Linux. The MLS router supports configuration of these tunnels from the Router configuration interface.

3.3.3 Classification of router links

Each link is classified in three dimensions. The confidentiality classification C_{HIGH} (High Confidentiality) assigned to a link means that the link only is visible for the coalition routers, thus the link is advertised to coalition routers only. C_{HIGH} links are not visible for the civilian partners (i.e. the NGO in the CoNSIS scenario) and are not advertised to the

civilian routers. Only links that are assigned the C_{LOW} classification are visible to the civilian routers.

The availability classification at A_{HIGH} (High Availability) assigned to a link means that the link, or elements of the link, is a limited resource (e.g. limited capacity) or internal resource that is only made visible to authorized routers. The lowest level of availability (A_{LOW}) encompasses the largest number of links and is advertised to all routers.

The integrity classification I_{HIGH} (High Integrity) assigned to a link means that the routing messages exchanged on the link are trusted (e.g. only internal routers are involved) or the link is protected by source authentication and data integrity check. In the CoNSIS scenario only internal routers of the coalition is assumed to be trusted. Thus all links between the coalition routers are classified I_{HIGH}, and all links involving civilian routers are classified I_{LOW}.

With three dimensions (C, I and A) and two levels per dimension we have eight possible combinations. Note that some of these combinations are not relevant. It is for example meaningless to combine C_{HIGH} with I_{LOW}, as pinpointed in **Error! Reference source not found.**.

3.3.4 MLS privileges and routing of packets

The MLS privileges of a packet are encoded in the *Traffic Class* field of the IPv6 header. The MLS privileges are specified in all the three dimensions C, I and availability A.

The packets are routed according to the rules depicted in **Error! Reference source not found.**:

- Packets with C_{HIGH} privileges are allowed to utilize routes from both confidentiality levels.
- Packets with C_{LOW} privileges are allowed to utilize low level routes only.
- Packets with A_{HIGH} privileges are allowed to utilize routes from both availability levels.
- Packets with A_{LOW} privileges are allowed to utilize low level routes only.
- Packets with I_{LOW} privileges are allowed to utilize routes from both integrity levels.
- Packets with I_{HIGH} privileges are allowed to utilize high level routes only.

3.3.5 Conclusions

The output of the work on MLS routing are a *proof-of-concept* router prototype. Due to resource limitations the experiment was not conducted at the Joint Distributed Experiment in June 2012.

4 KEY MANAGEMENT ISSUES

In this chapter we provide solutions, experimentation results, experience and future work for the key management issues.

4.1 Scalability of a tactical PKI

4.1.1 Introduction

A Public Key Infrastructure (PKI) allows the use of certificates and the automatic distribution of crypto keys or other secret elements for such applications as mutual authentication, the signature or the encryption of data transmitted through a communication system.

Because it does away with the task to manually configure keys into e.g. IPsec devices, resorting to a PKI is very desirable in a tactical network whose management must be as easy and flexible as possible. However, the use of a PKI implies the transmission of overhead data which can prove detrimental when bandwidth is scarce.

The output of the work on *scalability of a tactical PKI* is a theoretical study documented in [8]. In addition, several experiments on the tactical PKI theme were conducted during the CoNSIS project and a description of the experiments and of their major findings can be found in [9]. The work documented in the theoretical study is briefly summarised in the following.

4.1.2 Overhead evaluations

As the overhead associated with the use of a PKI depends on a large number of operational and technical factors, it can only be estimated via scenarios. In the CoNSIS project, three scenarios were considered: two in which PKI users are IPsec devices for the exchange of crypto keys, one in which they are user hosts which authenticate with applicative servers. The first IPsec scenario assumes a rather conventional situation in which crypto devices in Coloured Enclaves exchange certificates to establish IKE or MIKE administrative security associations. The second IPsec scenario is a more futuristic one in which each individual soldier is an actual communication node outfitted with its own crypto device, which results in a much larger number of user entities.

The three CoNSIS PKI scenarios are summarised by the following table:

Characteristics	Scenario 1	Scenario 2	Scenario 3
Domain in which the PKI is used	Black	Black	Red
User entities	Crypto devices	Crypto devices (down to the fighter)	User terminals and application servers
Use of certificates	Authentication at IKE/MIKE SA establishment	Authentication at IKE/MIKE SA establishment	Authentication at the establishment of applicative sessions
Number of user entities	400	20 000	2 000
Number of peers per user entity	4 IKE devices	4 IKE devices	10 application servers
Number of multicast groups	20 MIKE groups	200 MIKE groups	Not applicable

It must also be noted that one bit of overhead will use more resources if it traverses 10 links than if it only travels over a single hop. To reflect this reality, a network of 200 nodes and 400 links was assumed, with an evenly distributed population of user entities within this topology.

Other operational assumptions can also have a strong bearing on the overhead. To quote but the most important ones, the following hypotheses were retained for the scenarios:

- Lifetime of an IKE or MIKE security association: 1 day,
- Rate at which certificates get revoked: 2% per day,
- Delay for the notification of a revocation: 1 hour (i.e. the NATO requirement),
- Certificate validity duration: 1 month,
- Average number of members in a multicast group: 10.

The analysis based on these scenarios showed that the overhead associated with the use of a PKI is mainly due to three processes:

- The exchange of certificate between user entities when they authenticate or transmit their public keys,
- The transmission of certificates from user entities to the PKI for online verification,
- The periodic dissemination by the PKI of the Certificate Revocation List (CRL).

Overhead evaluations when no specific precaution is taken to conserve bandwidth are as follows:

PKI overhead (average bps per link)	Scenario 1	Scenario 2	Scenario 3
Certificate exchanges between users (bps)	14	647	1 502
On-line certificate status validation between users and SCVP responder (bps)	38	1 692	3 931
CRL dissemination between publishing repository and users (bps)	105	190 942	2 044
Mean overhead per link (bps)	**157**	**193 281**	**7 477**

The level of overhead is tolerable in scenario 1 when the number of user entities remains moderate. But it is clearly not so for the other two scenarios, bearing in mind that the typical capacity of a link in a tactical network is in the order of magnitude of 100 kbps and that the figures given by the previous table are only averages. **The deployment of a PKI for purposes such as those of scenarios 2 and 3 cannot be considered without optimisations**.

4.1.3 Overhead reduction

Four optimisation avenues were considered:

- The transmission of a unique certificate identifier instead of the full certificate itself. Whenever two user entities of the PKI have to exchange their certificates, they can remember that these documents were already transmitted to their party, and just send it a pointer which will unambiguously identify the certificate.
- Having the PKI "push" the list of revoked certificates using multicast instead of waiting for user entities to individually "pull" it in unicast mode.
- The choice of an optimised location for the PKI: deploying it more or less "in the centre" of the network will reduce the average number of links traversed by the CRL and by messages associated with the online validation of certificates.
- Selective dissemination by the PKI of delta-CRLs (i.e. only changes which were brought to the list since the last distribution) instead of the complete CRL.

A fifth optimisation consisting of reducing the lifespan of certificates was also considered. This benefit brought by this method would be to shorten the CRL since expired certificates can be quicker deleted from the list. But a shorter validity period would inevitably pose an organisational problem on a theatre of operations where user entities cannot be assumed to be able to physically go to the PKI site for the renewal of their certificates. Besides, analyses showed that the benefit brought by this optimisation method was dramatically reduced when it was implemented in conjunction with the other two relative to the CRL (i.e. "push" mode and delta-CRLs). Shorter-lived certificates were therefore disregarded.

The gains in terms of overhead brought by the 4 above-mentioned methods are summarised by the following table:

Efficiency of optimisation measures	Scenario 1	Scenario 2	Scenario 3
transmission of a unique certificate identifier	26%	1%	57%
CRL dissemination in push mode	40%	59%	16%
Optimised PKI location(1)	31%	27%	35%
Selective CRL dissemination	50%	90%	24%

(1) As compared to a situation in which the PKI is located “in a corner” of the network.

The 4 optimisation methods can even be combined together, and although their respective gains do not exactly add up, the overall improvement in overhead is most significant. It is actually sufficient to make a PKI usable in scenario 3 which was otherwise unfeasible, as shown by the table below:

Cumulative effect of the 4 optimisation measures			
	Scenario 1	Scenario 2	Scenario 3
Mean overhead per link (bps)	16	5 090	877
Compression rate vs nominal case	90%	97%	88%

Tests conducted with two different PKI software packages (EJBCA and OpenSSL) demonstrated the feasibility of the proposed optimisation methods. The EJBCA product even includes a delta-CRL capability which, albeit not exactly designed in a way to bring the overhead to a strict minimum, could be used efficiently. However, the other two methods (use of a unique certificate identifier and dissemination of the CRL in “push” mode) are not natively supported by currently available software and required the addition of some scripts both in the PKI and in user entities.

4.1.4 Recommendations

At the end of the study, the following recommendations can be made:

- The feasibility of a tactical PKI use-case cannot be taken for granted. Before envisaging such a set of mechanisms, network designers should always conduct an overhead analysis.
- If a PKI use-case turns out to be unfeasible without specific precautions, the four optimisation measures proposed above should be considered. These methods do not imply revolutionary changes in existing software or in current operational procedures. None of them infringes the security policy.
- All four optimisation measures can be used simultaneously. However, not all of them will prove efficient in all forms of scenario. Prior to implementing a particular method, network designers should conduct an efficiency analysis comparable to the one briefly described in this document.

4.2 Optimization of PKI protocols

4.2.1 Introduction

The term *Public Key Infrastructure* (PKI) refers to the set of services, software and protocols necessary to manage a set of public key pairs (private/public key). The services offered by a PKI is normally associated with subject identification, key generation, key deployment, certificate generation, key revocation, certificate directory service and certificate status provision.

These services require a high amount of system resources, in terms of

- networking capacity (data volume and number of protocol round trips),
- networking connectivity (frequency of protocol invocations),
- software maintenance (cost of configuration and update)
- computing power (cost of necessary calculations)

A study of different optimization strategies for PKI operations has been offered to the CoNSIS project. The study has been published in [10]. This section contains a brief summary of the discussions and findings of that report.

4.2.2 PKIX standards and the use of COTS software

PKI services may be implemented using the PKIX standards. These standards describe protocols and data structures relevant to the PKI services. The PKIX standards are governed by IETF and describe important PKI properties like:

- Data structure of digital public key certificates (X.509)
- Data structure of certificate revocation lists
- Procedures for certificate validation
- Protocols for a revocation status service (OCSP)

In addition to the standards published by PKIX, the PKI operation may rely on a larger set of underlying IETF standards, or on standards published elsewhere:

- General IETF standards like HTTP, LDAP etc.
- Public Key Cryptography Standard (PKCS), published by RSA Data Security. A selection of these standards is used for transportation of key pairs and certificates.
- Standards for signature representation (XML-DSig, S/MIME etc.)

The reason why PKIX standards are important is because commercial off-the-shelf (COTS) software is using them for the purpose of message signing, message encryption and certificate validation. Alternative architectures may perform better, but that would require custom built end-user software if they were to replace the PKIX protocols. For this reason, the PKIX protocols are unlikely to be replaced.

4.2.3 Certificate revocation

Although certificates have an expiration date, circumstances may require that they are revoked before they expire (e.g. the person represented by the certificate is being reassigned). Below four distribution methods for revocation information are mentioned:

1. Certification Revocation List (CRL)
2. Delta CRLs
3. Partitioned CRLs
4. Online Status Checking

The revocation operation is the single most controversial mechanism in the PKI, since it has serious consequences for system scalability.

4.2.4 The functional components and operations of a PKI

The main functional components of a PKI are:

Certificate Authority (CA) The CA issues public key certificates. It may also generate key pairs on behalf of subjects. The signature of the CA is trusted by everyone in the domain, making it a trust anchor.

Registration Authority (RA) The RA plays a part during certificate generation by verifying the correct association between the identity of the subject and the identification of the certificate.

Validation Authority (VA) The VA is delegated (from the CA) the responsibility to decide whether a certificate is valid.

The operations of a PKI are:

- Certificate Issue.
- Certificate Revocation.
- Signature Generation (related to PKI, but not performed by a PKI component).
- Certificate Validation.

4.2.5 PKI operations - size of data units

A data unit will grow in size if it is signed. Several programs are able to sign documents and messages, and the amount with which it increases is widely different. Observations have reported increased sizes from 3,6 kBytes (MS Word) to 26 kBytes (Adobe Acrobat), using keys with size 2048 bits.

For validation purposes, status revocation may be checked through CRL distribution or with online validation servers. Observed sizes of CRLs have been 36 bytes per certificate reference, and an additional constant part of 700 bytes. The network traffic generated by a OCSP transaction (for online revocation check) is 2,8 kBytes.

Other PKI operations, like certificate issue, happens less often than signature validation and do not generate significant traffic.

4.2.6 PKI optimization opportunities

PKI operations can be made more resource efficient through the employment of relatively well-known optimization techniques like:

- Reduce the size of the signature by excluding the signer's certificate
- Employ a content-distribution network (CDN) for push-based distribution of CRLs
- Employ caching of recent OCSP responses for subsequent validation of the same certificate
- Use certificates of short lifetime so revocation checking is not necessary

The traffic and scalability analyses of the optimization alternatives made a number of assumptions regarding different operating parameters, like number of users and certificates, message frequencies, expiration times, revocation latency etc. An important assumption was that messages are distributed over a number of senders according to a *scale-free distribution*.

For reasons of brevity, the full analysis from the report cannot be shown here. The below table summarizes the findings with the following interpretation of the columns:

Client traffic - Traffic rates in each client related to validation, given as bytes per second

Server traffic - Traffic rates in central server (CA) given as bytes per second

Connectivity demand - An indication of the client's dependency on frequent connection to a central server in order to complete a validation or renew a certificate

Traffic variability - An indication of the server's tendency to experience high peaks in the request traffic.

Optimization alternative	Client traffic	Server traffic	Connectivity demand	Traffic variability
Pull based CRLs	7.15 Bps	2860 Bps	Medium	High
Push based CRLs	7.15 Bps	7.15 Bps	Medium	High
Delta CRL (push)	1.83 Bps	1.83 Bps	Medium	High
Basic OCSP	13 Bps	5190 Bps	High	Low
Cached OCSP	13 Bps	4670 Bps	High	Low
Short lived certificates	0.17 Bps	69 Bps	Medium	Low

4.2.7 Cross domain operation of a PKI

A relying party can validate certificates issued by a different PKI if a cross domain relationship between the two PKIs exists. Such relationships are expressed through *cross certificates*, through which one PKI certifies the public key of the other CA.

Through cross certificates, a relaying party can construct a certificate path from the foreign certificate to its own trust anchor, which would allow the validation of the foreign certificate.

A cross domain relationship also requires that revocation information is exchanged between the PKIs. The exchange of CRLs may have serious consequences for resource consumption in particular where a large and a small PKI enters a relationship. If revocation information is made available over online providers (OCSP) then the providers must be made available to the other PKI's relying parties.

The combination of revocation arrangement and cross domain operation seems unlikely to be successful.

4.2.8 Conclusion

The use of short lived certificate seems to be the optimization technique that consumes the least network capacity, has moderate connectivity demand, works in a cross domain environment and is a conceptually simpler mechanism for signature verification since the entire verification process relies only on one object (the certificate), whereas other alternatives rely on two objects (certificate and revocation status).

4.3 IPsec discovery protocol (IDP)

4.3.1 Introduction

This paragraph discusses the IPsec discovery protocol (IDP) and its Java implementation developed by Fraunhofer FKIE, Tiber. Tiber functionality has been extended beyond the original discovery task. Unless noted otherwise this paragraph describes the behavior of the Tiber implementation.

This paragraph also describes the experiment conducted at the Joint Distributed Experiment in June 2012.

4.3.2 Functionality

Tiber instances are installed on IPsec devices each connecting a Colored Enclave (CE) to the Transport Network (TN). The eponymous functionality allows the instances to detect each other. This is done by periodically sending a HELLO message to a previously agreed IP multicast address. As soon as another IPsec device is detected through receipt of their HELLO message, IPsec Security Associations (SAs) to that device are established. In its stand-alone configuration, a pre-shared encryption key is used. A set of connected IPsec devices will end up with pairwise SAs.

In order to facilitate routing between the CEs, another periodic message is sent by each IPsec device. This encrypted message announces the network prefixes of the CE it is connected to.

4.3.3 Usage with MIKE key establishment protocol

The Multicast Internet Key Exchange (MIKE) protocol developed by Fraunhofer FKIE allows establishment of a common group key which can be used to derive keys for use as symmetric encryption keys. Tiber can use MIKE to establish an encryption key without having to distribute keys out-of-band in advance. This requires a public key infrastructure, since MIKE instances use digital signatures to authenticate themselves. This means that each instance requires the certificates of all other participating instances which contain the public keys to verify their signatures.

The advantage of pre-shared certificates and MIKE-generated keys as opposed to “Tiber only”-mode pre-shared keys is generation of a fresh key upon connection. We only have to distribute signing keys with long-term lifetime out-of-band. We do not have to do this for encryption keys which have to be changed often based on session key lifetime requirements.

The cost of MIKE usage is the data rate required for MIKE traffic. We also have to cope with key change during operation when connectivity changes. The MIKE traffic occurs when the IPsec device group composition changes.

4.3.4 Experiment

During the experiment in Greding in June 2012 all IPsec nodes connecting CEs to the TN were running Tiber during the experiments. In order to perform the individual experiments without being influenced by MIKE operations, Tiber was used with pre-shared keys. It successfully set up the IPsec security associations.

In addition, during a dedicated Tiber experiment Tiber was run together with MIKE. Instances were run on four Norwegian and three German vehicle IPsec nodes, the vehicle gateway node and a node in the wired network. Then two German vehicles moved away from the others and subsequently returned to disrupt and then restore connectivity (see Figure 7). Data on MIKE group behavior was collected. It shows the formation and partitioning of MIKE groups as wireless connectivity changes.

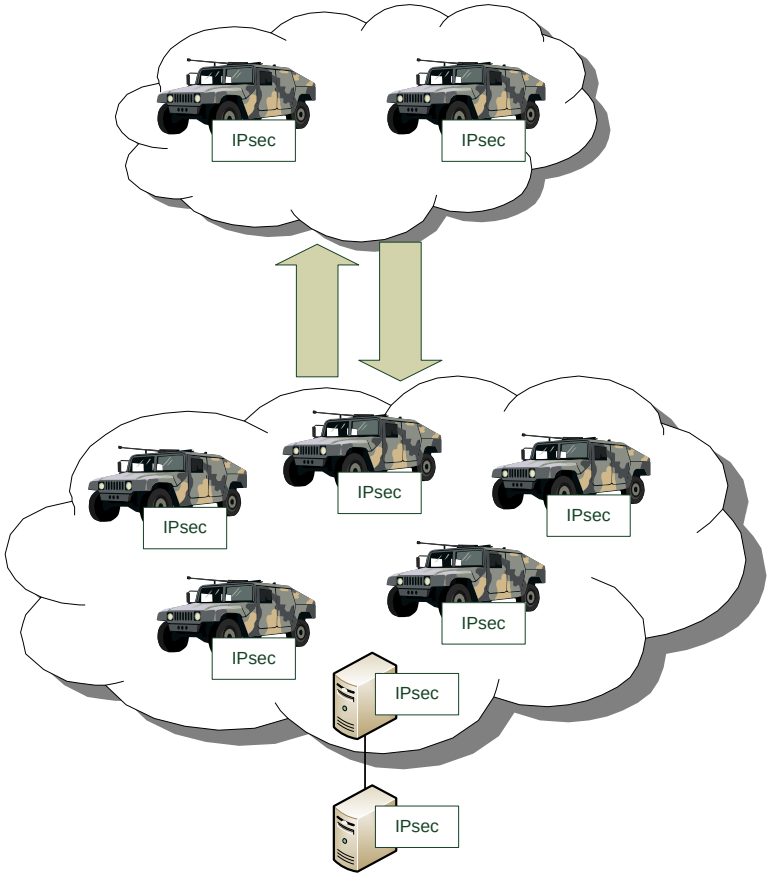

Figure 7: MIKE experiment

4.3.5 Lessons learned/future work

Future work on Tiber includes optimization of the Tiber data rate. Available data rate is a scarce resource in MANETs. Thus the data rate required by protocols has to be limited. Since the messages already use a compact representation, looking at the rate of HELLO messages is an important factor for optimizing the Tiber data rate.

A dynamic Tiber hello message rate can limit overhead by reducing the message rate in case of a stable group. Especially in case of using Tiber with MIKE choosing long timeouts for node removal is recommended. This limits group fluctuation caused by temporary connection problems.

In addition, dividing the IPsec devices into several groups using different IP Multicast groups improves scalability. For example, a set of devices connected by wireless links can be assigned its own multicast group. This prevents these wireless nodes from having to handle Tiber traffic of other nodes. Simply installing several Tiber groups provides IPsec tunnels between members of the same group only. This means that mechanisms for connecting the groups such as nodes participating in two groups are required. Routing and correct reencryption of packets have to be taken care of. Depending on the use case planning of Multicast group membership according to expected communication can limit such cross-group traffic.

4.4 MIKE study

4.4.1 Introduction

Motivation: Much of the communication at the lower tactical echelons is multicast radio traffic by nature. One example is position data for friendly force tracking. All data are confidentiality protected hop-by-hop: they are encrypted before are transmitted over the air. Coalition partners that move into an area need the proper cryptographic group key in order to start communicating. Current systems typically rely on pre-placed symmetric group keys (PPK). These systems have the drawback that they do not easily enable inclusion of new members ad hoc. The Multicast Internet Key Exchange (MIKE) scheme does. This triggered our interest in MIKE within Task 3 of the CoNSIS project, and gave rise to an assessment of MIKE for its use in tactical Ad Hoc Networks.

Summary: The outcome of the MIKE study is a report [11] and an article [12]. The main contribution of the work is an assessment of MIKE for its use in tactical Ad Hoc networks. An additional contribution is a number of suggested enhancements of MIKE.

4.4.2 Scenario

MIKE was assessed for its use in the scenario illustrated in Figure 8; a multi-hop mobile ad hoc network. The network consists of heterogeneous VHF or UHF wireless tactical communication nodes. Some are vehicle mounted. Others are battery powered and carried by dismounted soldiers. The nodes differ in level of mobility as well as in power resources and transmission range. There is a connection to deployable infrastructure, but connectivity cannot be guaranteed at all times.

Communication is protected by a group key. The group key can be pre-placed. But there is also a need for including new members ad hoc. The number of nodes in the wireless network is typically from 10 to 50.

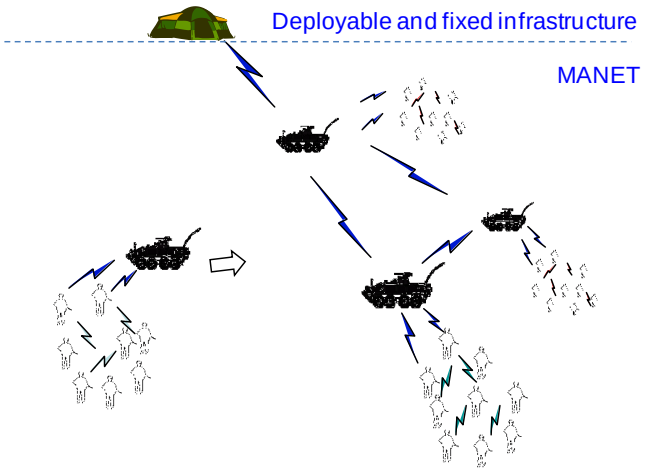

Figure 8 Scenario: Wireless communication at the lower tactical echelons

4.4.3 Outline of MIKE

Basically, MIKE does for multicast groups what the Internet Key Exchange (IKEv2) does for unicast peers: It provides automated negotiation of symmetric cryptographic keys for IPsec. IKEv2 is geared towards two communicating peers. MIKE provides a group key.

MIKE operates either in *Key Agreement Mode* or in *Key Distribution mode*. All group members must agree and contribute to the key in the Key Agreement mode. This is bandwidth consuming and limits the scalability. One member is appointed *Transaction Manager*. The Transaction Manager authenticates newcomers and assists in inclusions and exclusions from the group.

The Key Distribution mode is centrally controlled. The *Group Controller* has all power. It authenticates newcomers and generates and distributes the keys and decides who shall be excluded and included.

Forward and backward secrecy are obtained by changing the group key every time a member joins or leaves the group. Missing a key updated means the node may have to re-join in order to continue communication.

All newcomers are authenticated with the aid of their public key in a three-way handshake between the new node and the Transaction Manager or Group Controller. This means that MIKE demands a PKI or pre-shared public keys and certificates.

4.4.4 Assessment

Table 1summarizes the results. The table indicates to what extent the specific requirement is fulfilled. It also includes a PPK w/KDC column for comparison with the well known approach. The PPK w/KDC refers to a pre-placed group key and a Key

Distribution Centre (KDC). It assumes pre-shared unique symmetrical keys for protection of the communication between each member and the KDC and a pre-placed group key. See **Error! Reference source not found.** and **Error! Reference source not found.** for more details.

Criteria	Key Distribution	Key Agreement	PPK w/KDC
Secure Protocol	Partially	Partially	Yes
Forward Secrecy	Yes	Yes	Yes
Add members dynamically	Yes	Yes	No
Seamless key change	No	No	Partially
Seamless add new member	Partially	No	Yes
BW efficient	Partially	No	Partially
Robust to link loss	Partially	No	Yes
No single point of failure	No	Yes	No
Power efficient	Yes	No	Yes
Mature	No	No	Yes
Scope of Use	Small to large net Separation of COI Unsuitable as initial key	Small net Separation of COI Unsuitable as initial key	Small to medium net Separation of COI OK as initial key
Preconditions	Trust relation exists PKI Running network service	Trust relation exists PKI Running network service	Pre-distribution

Table 1：Result of the Assessment

4.4.5 Proposed enhancements

Table 2 lists the suggested enhancements and shows what aspect they improve. A closer explanation of each suggestion is found in **Error! Reference source not found.** and **Error! Reference source not found.**.

4.4.6 Conclusions

MIKE is not well suited for hop-by-hop protection of tactical ad hoc networks. It enables inclusion of new members ad hoc, but requires that the Transaction Manager or Group Controller is within direct transmission range of the joining node or a policy change that allows the other nodes to forward traffic of not yet authenticated nodes. MIKE furthermore assumes good connectivity and reliable multicast. Otherwise communication may be disrupted. It does not address seamless key changes. The assessment shows that the Key Distribution mode performs better than the Key Agreement mode in a tactical environment.

Whereas a number of enhancements have been proposed, they do not fully solve all problems. Further work is needed.

	Suggestion	Security	Seamless key changeg	BW efficiency	Robustness
General	**Retransmit last key**		X		X
	Allow key overlap		X		X
	Extend use of sequence numbers	X			X
	Skip LeaveConfirm message			X	
	Backup GC/TM				X
Key Distribution mode	**CRL only to GC**			X	
	Change key only on ejects				X
Key Agreement mode	**Collapse TM- and UpdateDistribute messages**			X	
	Add TM willingness			X	X
	Don't transmit entire Key-tree			X	

Table 2: Outline of possible optimizations of MIKE and their impact

5 PROTECTION OF USER TRAFFIC

In this chapter we provide solutions, experimentation results, experience and future work for the protection of user traffic.

5.1 Cross-domain information exchange

5.1.1 Introduction

Efficient information exchange is essential for the success of NNEC. While service-oriented architectures (SOA) have the potential to provide more efficient information exchange, by facilitating interoperability between nations and organizations, such interconnection cannot be allowed unless the applicable security requirements can be fulfilled. Interconnection between different security domains has traditionally been achieved using solutions such as diodes, only providing a one-way information flow, or through the use of air-gaps. However, a guard based approached may be used to provide a more flexible two-way solution.

To address this issue, the cross-domain information exchange CoNSIS experiment considered the use of a guard to provide interconnection between two different security domains in a service-oriented environment. In the experiment, the guard was used to provide two-way information exchange between a classified military domain and an NGO (i.e., an unclassified domain) where messages were released from the military domain to the NGO based on confidentiality labels.

Apart from the correctness of the guard itself, the security of such a solution is clearly dependent on the trustworthiness of the confidentiality labels. For this reason, the experiment also included the use of a MILS separation kernel based solution for attaching confidentiality labels, in order to assure that classified information was not labeled as unclassified.

The experiment was conducted at the Joint Distributed Experiment in June 2012 and is further described in the next section.

In addition, there has also been performed a study [6] of typical solutions for performing cross-domain access control. The implications of the architectural approaches and communication patterns used by these solutions are analyzed with regard to their applicability in a military SOA, considering availability and assurance requirements. This study also proposes alternative approaches to provide higher availability and assurance. An overview of this study is provided in Section 5.1.3.

5.1.2 Cross-domain information exchange experiment

As shown in Figure 9, the experiment consisted of two main components, the XML/SOAP guard and a workstation with a MILS separation kernel. In addition, the user within the NGO domain had a viewer application that was used for visualizing received messages.

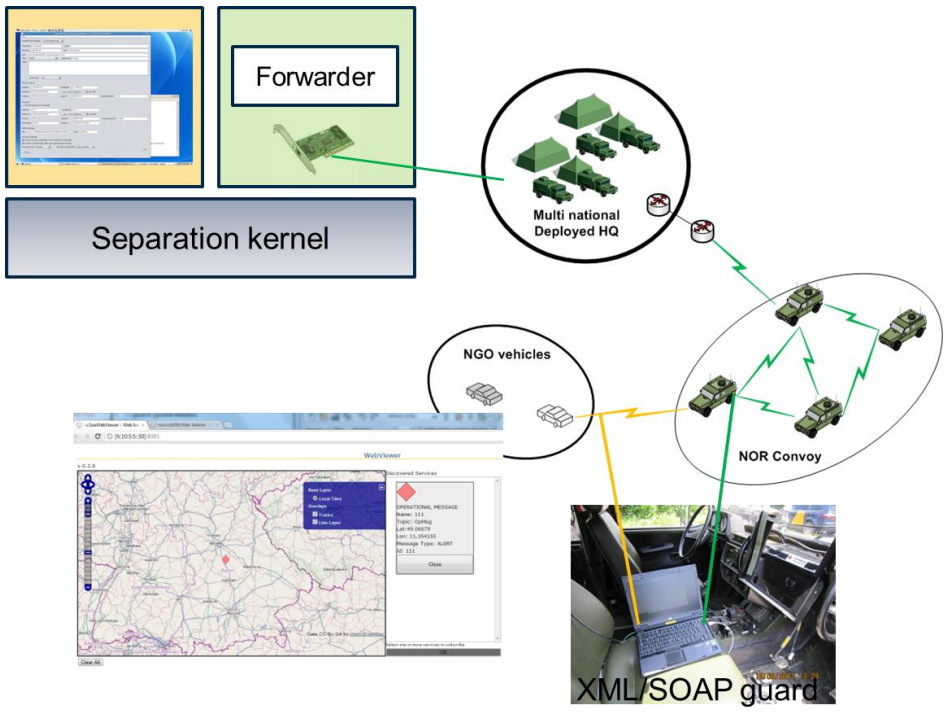

Figure 9 Cross-domain information exchange experiment overview

For the purpose of the experiment, information exchange was performed using the SOA infrastructure provided by Task 2. More specifically, the information exchanged during the experiment was using the Operational Message notification service, which was used in the experiment to send warning messages from the military domain to the NGO.

With the configuration used for the experiment, only messages labeled as Unclassified were allowed to be released to the NGO through the guard. This labeling were performed using the XML confidentiality label [13] and related binding [14] proposed by the IST-068 XML in cross-domain security solutions STO group. Each message and its associated label (including the binding between them) were also integrity protected by a digital signature.

The warning messages to the NGO were sent from a workstation within the Multinational deployed HQ. As shown in Figure 9, this workstation was running a separation kernel where warning messages originated within a partition without direct network access. The purpose of this was to demonstrate the use of such a solution to prevent classified information, originating from the classified domain, from being mislabeled. By originating unclassified messages within a separate partition without access to classified information, and performing labeling within this same partition, there is increased assurance that classified information is not mislabeled and additional protection of the cryptographic key used for creating the digital signature. It may be noticed that further protection of this key could be provided by using a hardware security module only accessible from within this partition, although such a module was not used for this experiment. The partition with a network connection did not play an active role in this experiment, apart from running an application forwarding messages from the “unclassified partition” onto the classified network. A prototype solution for providing the user access to multiple security domains on the same machine has previously been demonstrated at FFI, and the reader is referred to [15] for additional information.

The guard was positioned within a vehicle in the military convoy, and had one physical connection to a classified enclave within the military domain and one physical connection to an unclassified enclave connected with the NGO convoy. Messages to the NGO were forwarded through the guard, with the guard being accessed as an HTTP proxy. When processing a message, the guard checks both that the confidentiality label specifies a classification releasable according to policy (i.e., Unclassified in our case) and that the digital signature covering the message and label is valid. Otherwise, the message is blocked. The reader is referred to [16] for additional information on the guard.

5.1.3 Cross-domain access control study

The potential in information exchange and integration across administrative domains cannot be fully realized unless access control to the various resources within each domain can still be enforced. A cross-domain solution for access control is therefore required.

In order for access control to be performed in a cross-domain scenario, the attributes of the subject must be communicated from the administrative domain of the subject to the domain of the resource.[4] To this end, the Security Assertion Markup Language (SAML) defines an attribute statement by which the attributes of a subject can be expressed within a SAML assertion. Furthermore, mechanisms that can be used to exchange such attribute statements (i.e., assertions) are provided by standards such as WS-Trust, WS-Security, and the SAML protocols. An overview of these and related standards is provided in [17].

As such, SOA standards to realize cross-domain access control are readily available and their use in military systems would be advantageous considering interoperability and cost. The question considered in this study is to what extent these solutions, intended for use in civilian applications, are suitable for use in military systems. In particular, the architectural approaches and communication patterns used by these solutions are analyzed with regard availability and assurance.

The reader is referred to [18] for the full details of the study, but in general it was found that the considered existing solutions are not directly applicable to many military systems. In particular, the additional connectivity dependencies introduced by these solutions poses a critical issue in some military systems (e.g., tactical networks) where connectivity disruptions must be expected (or domains are only connected by one-way channels). Given that timely access to information and services may be of high importance, such degradation in availability may not be acceptable.

To ensure availability, it would be desirable that attribute statements have a long lifetime to prevent denial of service in situations where the subject is unable to renew its attribute statement(s), e.g., due to a loss of connectivity with a statement issuer (aka identity provider or security token service). On the other hand, to ensure validity, it is preferable that attribute statements have a short lifetime (which is typically the case for such solutions). Thus, the requirement for attribute statement freshness is in conflict with the requirement for availability.

Replicating the statement issuers close to the end-users (i.e., subjects) could mitigate the availability problem. However, the statement issuers are highly security critical and this would therefore pose a significant security risk, as the compromise of a statement issuer could enable unauthorized access to all resources subjects from that domain may be allowed to access (i.e., both the resources within the domain of the issuer and potential resources within other domains). The security critical nature of statement issuers is to a large extent due to the requirement for creating attribute statements upon

[4] As attributes can be used to specify both roles and identities, attribute based access control can be viewed to encompass both identity and role based access control, and may also be used to enforce mandatory access control according to the Bell-La Padula model. Hence, without loss of generality, it is assumed that some variation of attribute based access control is used.

request (i.e., a statement issuer is not simply a repository of pre-created attribute statements, thereby differing from a certificate repository).

The study [18] proposes an alternative approach to avoid this issue, where the statement issuer is split into two separate entities, providing separation of duty between one entity able to create attribute statements (e.g., specifying the attributes of a subject) while the other entity is only able to specify and renew the validity time of this exact attribute statement (within a relatively longer maximum validity time specified by the entity originally creating the attribute statement). This has the advantage that the second entity cannot be used to issue additional privileges (or renew attribute statements beyond the limit specified by the first entity) in the case it is compromised. Furthermore, the entity creating the attribute statements can be better protected (potentially behind a one-way channel) as it only needs to make the attribute statements available to the renewing/validating entity in some way. Because this latter entity is less security critical, it can also more easily be replicated closer to the end users, thereby improving availability. By not requiring subjects to obtain attribute statements from foreign domains, as detailed in [18], the connectivity dependencies can be further reduced (thereby improving availability).

5.1.4 Conclusion

The cross-domain information exchange experiment established that messages were correctly blocked or forwarded by the guard depending on the confidentiality label of the message (i.e., messages with a confidentiality label specifying Unclassified and having a valid signature were forwarded while messages for instance specifying a higher classification level or with an incorrect/missing signature were blocked). Furthermore, the applicability of attaching (and signing) confidentiality labels within a separate MILS separation kernel partition was validated in the experiment scenario. The experiment also served the purpose of providing experience on integrating such security solutions in a service-oriented environment.

The cross-domain access control study showed that established SOA solutions for access-control may not be directly applicable in military environments, due to the connectivity characteristics of some military networks and the differing requirements for availability and assurance. An alternative approach has therefore been suggested to improve availability and assurance.

5.2 Protected communication between military and civilian networks

5.2.1 Introduction

The presence of non-governmental organizations (NGOs) in a war zone is frequently seen, and their operations may be safer and more efficient through communication with military forces. Military information about safe routes, road conditions and observations regarding the situation for the population may be sent to the NGOs. Positions and

movements of NGO vehicles and personnel may be sent to the military forces in order to avoid inadvertent attacks.

The output of this work is a prototype and an article [19]. In addition an experiment was conducted at the Joint Distributed Experiment in June 2012. The summary given in this section is copied from the article [19].

5.2.2 Technical Requirements

The functional requirements for a Civilian-Military (CiMi) communication arrangement may be expressed in the following manner:

- *COTS equipment and protocols*: The NGO should avoid the use of military communication equipment from reasons of impartiality and cost.
- *Protection of communication channel*: The CiMi connection must be a black network, i.e. it can run through any unprotected link.
- *Robustness of separation (fail-close)*: The separation of the NGO and the military equipment should have the fail-close property (also called fail-safe).
- *Authentication of participants*: Participants in the communication should be fully identified before or during the service. A Cross Domain mechanism should be in place where a trust relation between the registration authorities allows mutual authentication across the interface without the need for multiple registration of identities.
- *Role-based access control*: Role Based Access Control should be the basis for the access control decisions, which enables the owner of a service to reserve its use for clients which possess certain roles.
- *Confidentiality labeling*: In the classification hierarchy found in military information management there is a need to decide if information kept in classified systems can be released for use on lower classification levels and even released to an NGO. Confidentiality labels are cryptographically bound to the information object and can be automatically inspected by a guard.

5.2.3 The Prototype Configuration

For an experimental evaluation of these principles a prototype was developed with the following services in mind:

Protected service invocation: A client in the NGO network should be able to invoke a positioning service in the classified network, and to receive the GPS coordinates of a mobile military unit.

Secure chat: The mobile client may write text messages to other users on a chat client program.

Configuration details
The figure below outlines the structure of the prototype. It consists of the following actors:

- An Android smartphone, acting as an NGO terminal for chat and protected service invocation.
- A chat server for the XMPP chat protocol. This server will forward both chat messages and service invocation messages.
- Two Identity Providers (IdP), one for the NGO domain and one for the military domain. They provide identity information for authentication operations.
- An application server, residing in the military domain, hosts application services or proxies for Web Services.
- A SOAP guard, which connects the military classified and unclassified networks. It ensures that only correctly labeled data is passed from the classified to the unclassified part.
- Other chat clients which use the XMPP protocol. They are connected to the XMPP server.

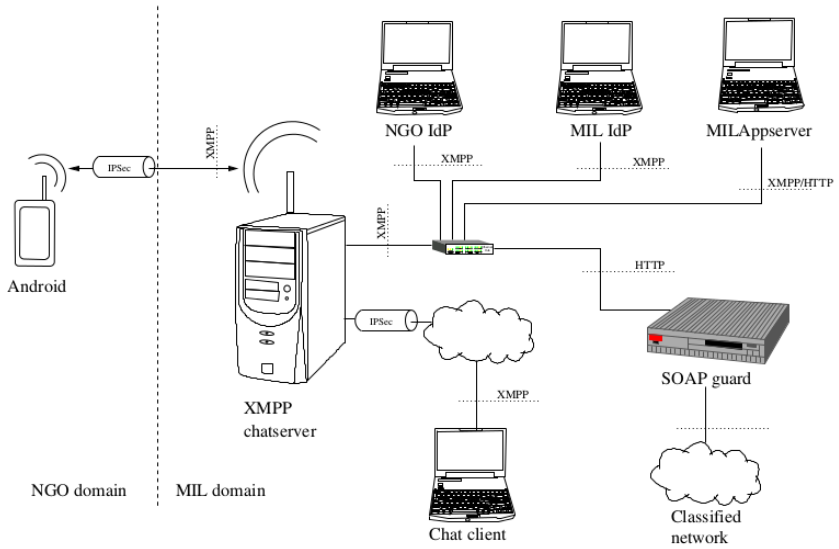

5.2.4 The GISMO IdM architecture

For the purpose of authenticated service provisioning in military tactical networks (meaning wireless, mobile, multi-hop, multicarrier networks), an Identity Management system has been developed under the project name “GISMO” (General Information Security for Mobile Operation).

- It uses short lived *Identity Statements* containing the subject’s public key and subject attributes. No revocation scheme is necessary. Identity Statements are issued by an Identity Provider (IdP).
- Cross COI (community of interest) relations are represented by ordinary identity statements issued from one IdP to another.
- IdPs can issue *Guest Identity Statements* when presented with an Identity Statement issued by an IdP with which it has a Cross COI relation. A guest identity statement contains the same information, but is signed by a different IdP.

- Authentication takes place either through a signature in the service request, or through the encryption of the service response.
- It supports Role/Attribute Based Access Control (RBAC/ABAC) through the subject attributes.
- Employs, but encapsulates existing PKIs. Clients never see X.509 certificates or revocation info.
- Identity Statements are cached and re-used during its lifetime. An IdP is invoked to issue Identity Statements, not to verify authenticity.
- There is loose coupling between IdP and services/clients, and between COIs. No redundant registration is necessary.

5.2.5 Service Invocation

IdP operations and service invocations are using serialized Java objects (called *POJO*) as PDUs which opens up interesting opportunities: The client may simply send a parameter object to the server containing the parameter values, and the class of the object identifies the service method. This arrangement eliminates the need for a separate scheme for service addressing and also eliminates the need for separate stub/skeleton compilation.

5.2.6 Messaging Protocols

In a wired private network where capacity and reliability suffice, and there exist IP routes between the nodes that wish to communicate, the HTTP protocol works just fine for IdP operations and service invocations. For mobile networks this is not necessarily the case: they are slow, unreliable and consist of several partitions connected with application level gateways (for reasons of security and traffic control).

In the context of this experimental study of the GISMO IdM, an XMPP (eXtensible Messaging and Presence Protocol) network was already in place for chat communication. Through the XMPP routers (working as application gateways) otherwise isolated networks (where no IP route exists between them) can exchange chat messages.

5.2.7 SOAP guard and confidentiality labeling

As can be seen in the figure, a SOAP guard connects military networks of different classification levels as an application gateway in the form of an HTTP proxy. It relies on *confidentiality labels* that are bound to information objects in a form that can be inspected and validated by the guard in order to make decisions whether to allow objects to be transferred from a high to a low classified network. The transport may be initiated by a client on the low side as an HTTP operation (e.g. a Web Services request), in which case the response will need a label in order to pass through. The request will need to be labeled if it is initiated on the high side.

5.2.8 Conclusion

This part of the CoNSIS experiment was conducted with the intention to study a range of security technologies for the separation of military and civilian networks, and to study how commercial mobile units (a waterproof Android smartphone) could be employed inside that security framework.

Most of the technologies (StrongSwan IPSec, serialized Java objects, homemade IdM, SOAP Guard) were working well. The use of Android was a bit over-ambitious, in the sense that IPv6, IPSec and network routing was implemented in a rather basic fashion.

The Android unit turned out to offer excellent portability of existing Java SE sources, and the XMPP stack was directly ported to Android without the need for any corrections. The low price, availability of development tools and the existence of waterproof Android units is promising for the future use of mobile COTS units in tactical networks.

5.3 Multilevel security and CoNSIS network management

5.3.1 Introduction

This paragraph discusses a concept for data transfer from a classified domain to an unclassified domain provided by Fraunhofer FKIE. The focus is on network management traffic in a CoNSIS network. More detailed information can be found in a report on Multilevel Security [20] in chapters 8 and 9 and in [21].

The transport network segments (TNS) of the transport network (TN) of the concept explained above have to be managed. When planning management access, security mechanisms restricting data flow have to be taken into account. The architecture prevents data in the colored enclaves (CE) from leaking into the TN by encrypting all outgoing traffic using IPsec devices. While this is intended for user data handled inside colored enclaves, it also means that management data cannot be sent from inside a CE to a TNS. Thus an administrator either has to have direct access to the TN or additional mechanisms have to be in place.

5.3.2 Current architecture

Due to the security mechanisms it is assumed that an administrator has direct access to the TN. This has the advantage of unmodified security infrastructure but it has practical drawbacks. If the administrator requires information from inside the CE, he either has to physically travel between workstations or a TN-connected workstation has to be placed next to his CE workstation. In the latter case TN connections have to be placed inside a secure facility next to CE devices. In both cases no digital data transfer between the TN-connected management workstation and the CE workstation is possible.

5.3.3 Concept

The alternative is providing a connection between a CE device and a TN device. This allows an administrator to use the CE infrastructure for network management. No additional devices or cabling are necessary. The administrator can, if necessary, have access to CE data when making management decisions.

Simply bypassing IPsec devices by connecting a CE to a TN network would break the security architecture. A more complex solution is required. We discuss three possible solutions and provide more details on the third one which we consider the most promising. We can use

- A data diode,
- An unclassified process or
- A Cross Domain Guard.

The first option is installing a data diode between CE and TN. A data diode is a hardware device which restricts data transmission to one direction only. Due to the low complexity of their task, highly secure ones are available. It allows data flow from TN to CE only. This means that its installation does not endanger the confidentiality of CE data. Even if malicious software got into the CE via the connection, it had no means to send classified information to someone outside a CE. If such a diode is used for network management purposes, network status information can be provided to workstations inside the CE. The obvious drawback is the lack of return channel to actually influence the TN. This still requires out-of-band mechanisms such as moving to a TN-connected terminal or phone calls.

The second option is having an unclassified process for managing the TN inside the CE and using the CE for connecting to the TN. The unclassified process can either be a standalone machine or an unclassified partition on an MLS system. An authenticated IPsec tunnel can be established between the unclassified process and the connection to the TN to securely mark the unclassified traffic exchanged between the management process and the TN.

The third option is installation of a Cross Domain Guard, a device which controls data flow between CE and TN (see Figure 10). Only messages conforming to specified rules may pass from CE to TN. This allows an administrator to do his job from a regular CE workstation. It is the most complex option, since traffic has to be treated differently based on its content. Legitimate management traffic may leave the CE without being encrypted, any other data must not. Other options treat all traffic from the same interface equally. Data flow from TN to CE can be handled by a data diode as in the first option. We now further discuss this option.

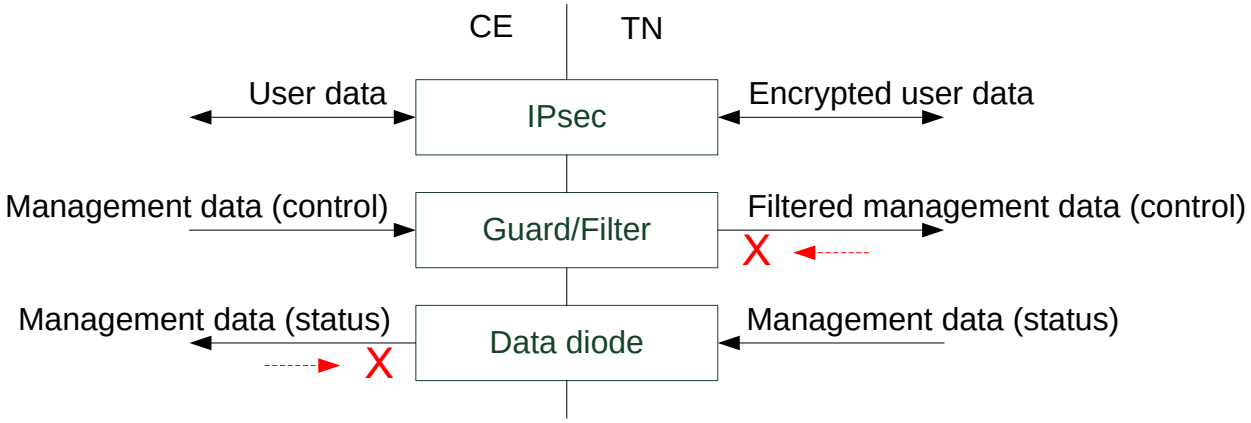

Figure 10: Guard data flow

5.3.4 Cross Domain Guard

When employing a Cross Domain Guard for enforcing rules on which traffic may leave a classified network such as a CE, several steps need to be taken. Legitimate traffic has to be defined precisely. The risk of abuse of legitimate traffic through steganography or covert channels has to be analyzed and limited. The acceptable covert channel data rate has to be determined based for example on the classification level of the data inside the CE and the utility of the connection.

We assume that legitimate traffic consists of XML messages. In this case definition of legitimate traffic can be done by specifying all required message types in a schema language such as XML schema.

Even if traffic is restricted to messages conforming to management message syntax, steganography and covert channels pose a threat to confidentiality. Steganography, the art of hiding data inside other data, allows an attacker to leak information by modifying legitimate messages. One of the strategies used to counter covert channels is using filter functions which remove unnecessary details from files such as images. We assume that management messages conforming to the schema do not have unnecessary content and do not apply additional filtering. We make the pessimistic assumption that all management messages sent may be part of a covert transmission. We then compute the amount of bits encoded in the messages assuming that they have to conform to the schema and throttle the message rate to limit possible information leakage to an acceptable level.

Covert channels are mechanisms used for covert data transmission which were not intended for communication. A covert channel relevant for Guards is timing of messages. Even if message content is strictly controlled, data can be encoded in the time delay between messages. In order to limit the possible covert channel, the suggested Guard enforces regular intervals between messages. It stores all messages which passed the schema filter in an internal buffer and forwards them at regular intervals. This can be integrated with the message rate throttling mechanism.

5.3.5 Management proxy

Limiting cross-domain management traffic from CE to TN to the minimum amount required is important. In addition to general bandwidth-efficiency considerations, the steganography risk is mitigated this way. The less traffic there is the less data can be embedded by an attacker without risking detection. Depending on the confidentiality of the data in the CE, a lot of effort may be justified.

Depending on the protocol a management proxy aggregating data inside the TN may be helpful (see Figure 11) to achieve this. If for example the administrator requires status information on every router, he can send this request to the proxy which in return contacts each router and returns the aggregated status to the administrator. This means that just one request instead of one per router has to pass the Guard. The same holds for a new configuration pushed to all routers and the like.

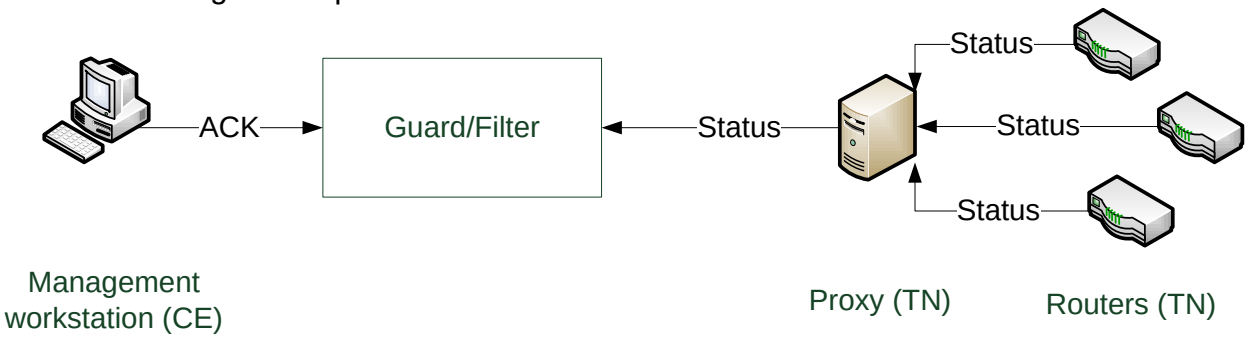

Figure 11: Proxy in TN

Our goal is to get by with such a low legitimate cross-domain traffic rate that we can severely limit the allowed message rate and thus the risk while still providing management functionality.

5.3.6 Conclusion

A concept for exchanging management data between a classified domain and an unclassified domain was presented. A Cross Domain Guard is suggested for filtering data leaving the classified domain. The main advantage of the concept is its simplicity, relying on secure devices between CE and TN only. The approach is limited to use cases which allow restricting the classified to unclassified traffic to a low data rate. If high amounts of data have to be transmitted from a classified to an unclassified domain, other approaches such as label based filtering are necessary.

5.4 Traffic flow confidentiality

5.4.1 Introduction

Traffic Flow Confidentiality (TFC) is a data confidentiality service used to protect against traffic analysis. This is done by masking the original traffic either by changing the data packets' frequency, lengths or origin-destination traffic patterns between network endpoints.

The currently available TFC solutions are based on link layer functionality. Hence they may not be used in scenarios where the data traffic traverses third party networks, i.e. public networks or military coalition networks not supporting traffic flow confidentiality.

The work on TFC focuses on military IP networks supporting differentiated quality of services provided using the DiffServ mechanisms. To offer a predictable end-to-end service quality, the DiffServ mechanisms must be supported through the use of Service Level Agreements (SLA). The SLA regulates the amount of traffic allowed for each DiffServ class and states penalties (drop, remarking) for traffic exceeding the maximum rates and bursts.

IPsec in tunnel mode provides integrity and confidentiality for data traffic run over unprotected infrastructure. However it is more difficult to offer traffic flow confidentiality. Different QoS classes will be run over dedicated tunnels and the QoS class is visible through the TOS marking of the tunnel IP packet. A basic method would be to add dummy traffic to each tunnel to ensure that the traffic level in each tunnel remains constant. However, this negates the statistical multiplexing of the underlying IP network, since each path would carry a constant level of traffic, and high priority dummy packets would delay real traffic running at a lower priority.

The TFC solution presented in this section provides end-to-end traffic flow confidentiality in heterogeneous QoS enabled secure IP networks. It is based on the TFC concept developed by Baseline Communications as [22]. The output of the work is an implementation of the TFC concept [22] and experiments conducted at the Joint Distributed Experiment in June 2012.

5.4.2 TFC architecture

Our TFC concept assumes that there will always be a mixture of traffic types being transmitted and the different traffic types are classified as different priority levels or DiffServ classes. The main idea is to promote traffic from lower priority levels if the total traffic of a higher priority class is less than the maximum allowable traffic according to the SLA. This ensures that the higher priority classes are filled with traffic – although with a combination of the original high priority traffic and traffic with lower priorities. This makes it difficult to identify variations in the high priority traffic. However the total amount of traffic is not changed, but it is impossible to deduct the amount of traffic that originated within each class. The underlying assumption is that normally the fluctuations

in the low priority classes like best effort (BE) traffic are large, and that there will always exist traffic that can be used to pad other classes.

The TFC mechanisms need to operate closely with the QoS to make sure that the traffic load is being adjusted according to the SLA limits. The promotion to a different QoS is always to one with a “better” QoS, and it is remarked only for the IPsec tunnel. Once decoded, the packet is marked with its original QoS. The packet size can be used to determine the application creating the traffic. Although the traffic in a TFC protected tunnel has the same QoS, the packet size may be sufficient to distinguish between “real” and promoted traffic. The promoted traffic must therefore be fragmented or padded to adhere to the packet distribution of the “real” traffic. In addition the solution supports traffic padding, i.e. the introduction of dummy packets. This is mainly used if there is no traffic to promote or if promotion into certain classes is not feasible. The potential benefit of this functionality was one to the targets for the experimentation.

We make no assumption about any structure in the QoS definitions. For each class, the network owner must define a sequence of classes the traffic can be promoted to if these classes have less traffic than specified in the SLA. The TFC mechanisms can be applied to a subset of the QoS classes, and the system is set up to accommodate up to 64 different promotion policies, which is the maximum possible QoS classes.

A policy based management solution will allow the operator to define which priority class traffic may be promoted, which priority classes should not be promoted, if dummy traffic should be allowed, and maximum traffic load of dummy traffic.

The final functionality is tied to the definition of SLA. In the initial release, we assumed that a fixed target existed for each QoS class. However, if the underlying network carrying capacity varies, the SLA might be dynamic, either defined as a percentage of total capacity or as varying target per class. In either case, the TFC incorporates functionality to estimate available capacity and lets the SLA target be adjusted by an associated module.

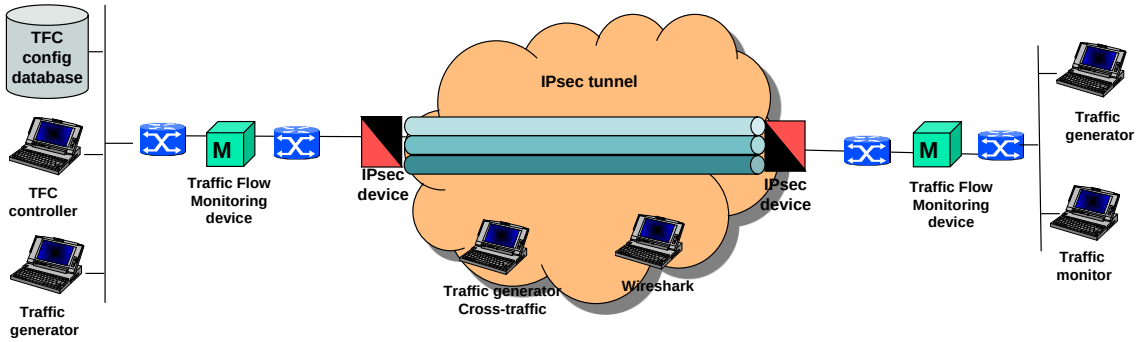

Figure 12 Demonstration architecture

5.4.3 Experiments

The experiments were set up to validate and demonstrate the benefit of TFC, the design trade-offs for allowing for additional dummy traffic and dynamic SLA support. The experiment setup is illustrated in Figure 12.

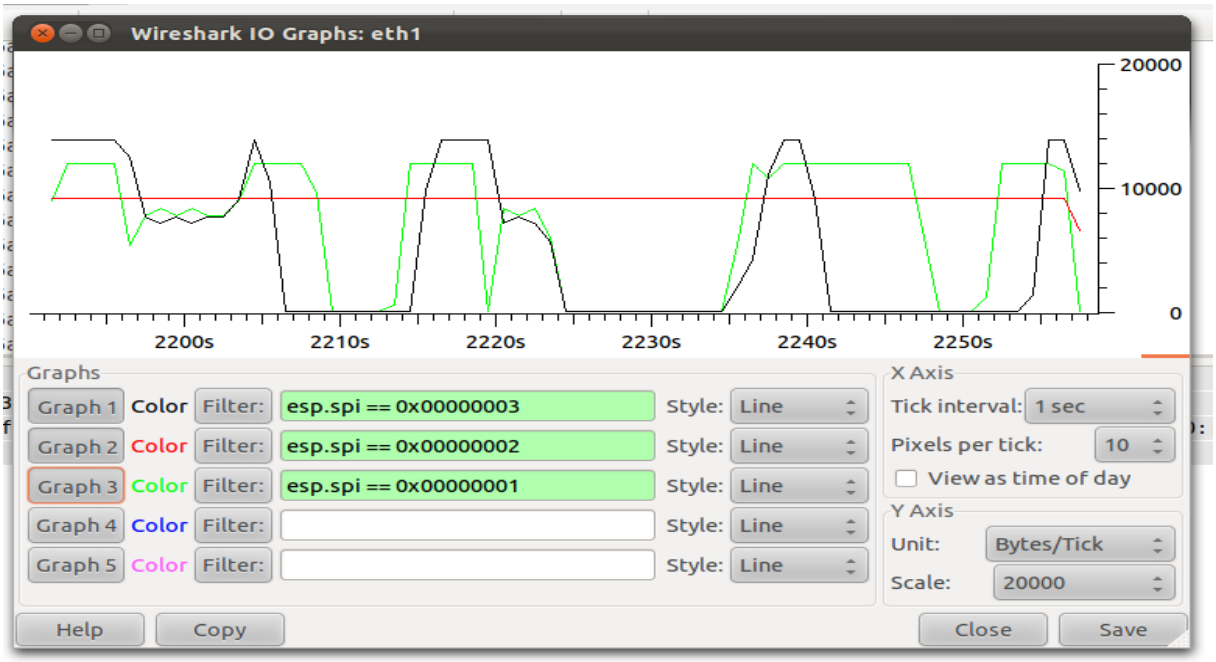

Figure 13 Traffic observed in black network without TFC

Three different QoS classes were used, VOIP, video and BE. The VOIP was not protected by TFC, while the video class used for videoconference was protected. For each experiment Wireshark was used in the black network to capture a copy of all traffic on an IPsec tunnel. In all figures, BE traffic is carried in tunnel 3 (esp=3), voice traffic in tunnel 2 (esp=2) and video in tunnel 1 (esp=1). To speed up the demonstration we used fast variable traffic sources. In a real operational setting the changes in traffic would occur on a longer time scale. As seen from Figure 13, it is easy to detect when a video conference is turned on and terminated.

Once TFC is turned on (see Figure 14), the traffic level in the video class remains constant, not reflecting the variations in the actual load at the red side. The underlying assumption is that there is always enough traffic available to promote.

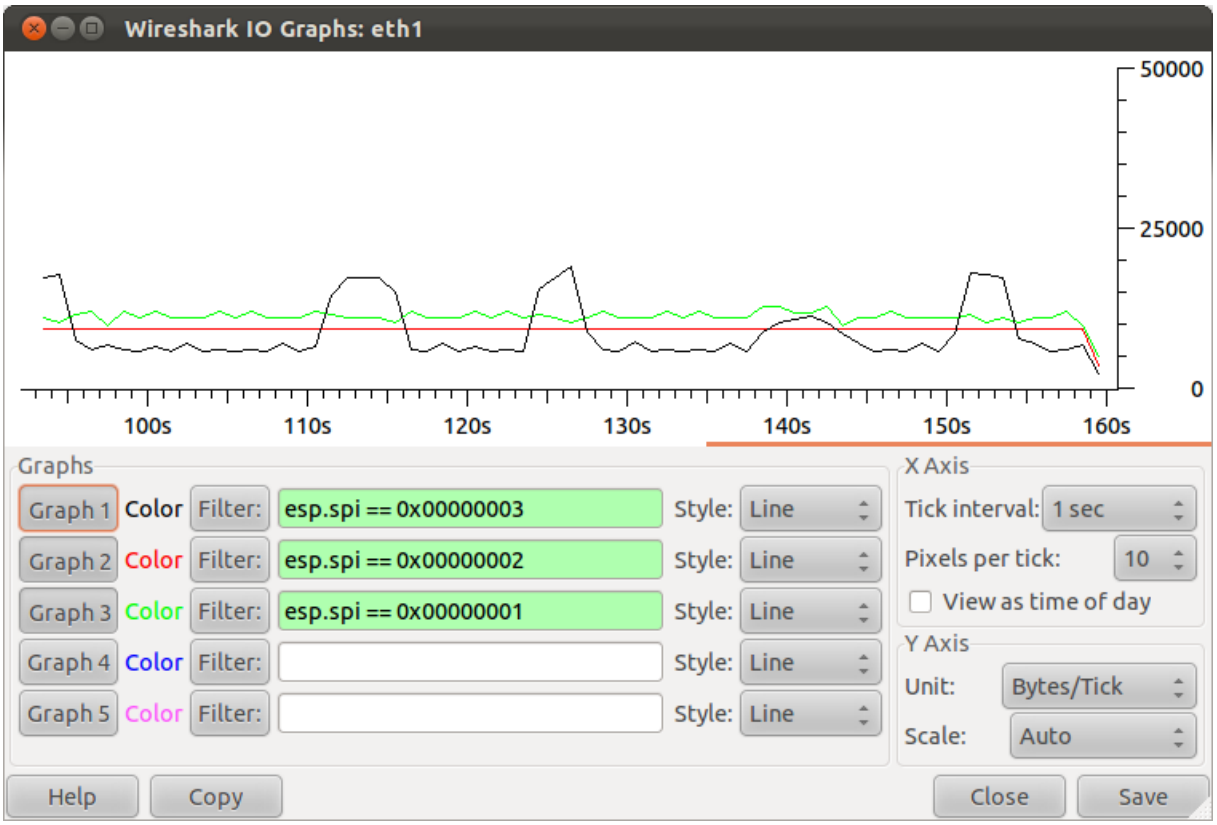

Figure 14 TFC is turned on

When there is minimal BE traffic available, the activity in the video class will be more transparent. Enabling dummy traffic when there is not sufficient BE traffic can be used to hide the activity in the video conference. The use of dummy traffic or not is a trade-off between the probability of minimal BE effort traffic against the transparency of the video traffic activity and the waste of network capacity used for dummy traffic. Figure 15 illustrates the functionality to adapt to changes in the black network. The capacity is reduced first to 8 Mbit/s due to link break and subsequent rerouting in the black network. After a while the link is up again and the capacity is back at 10 Mbit/s. The TFC system detects the change in capacity and according to the configured policy reduces the SLA and thereby the promoted traffic accordingly.

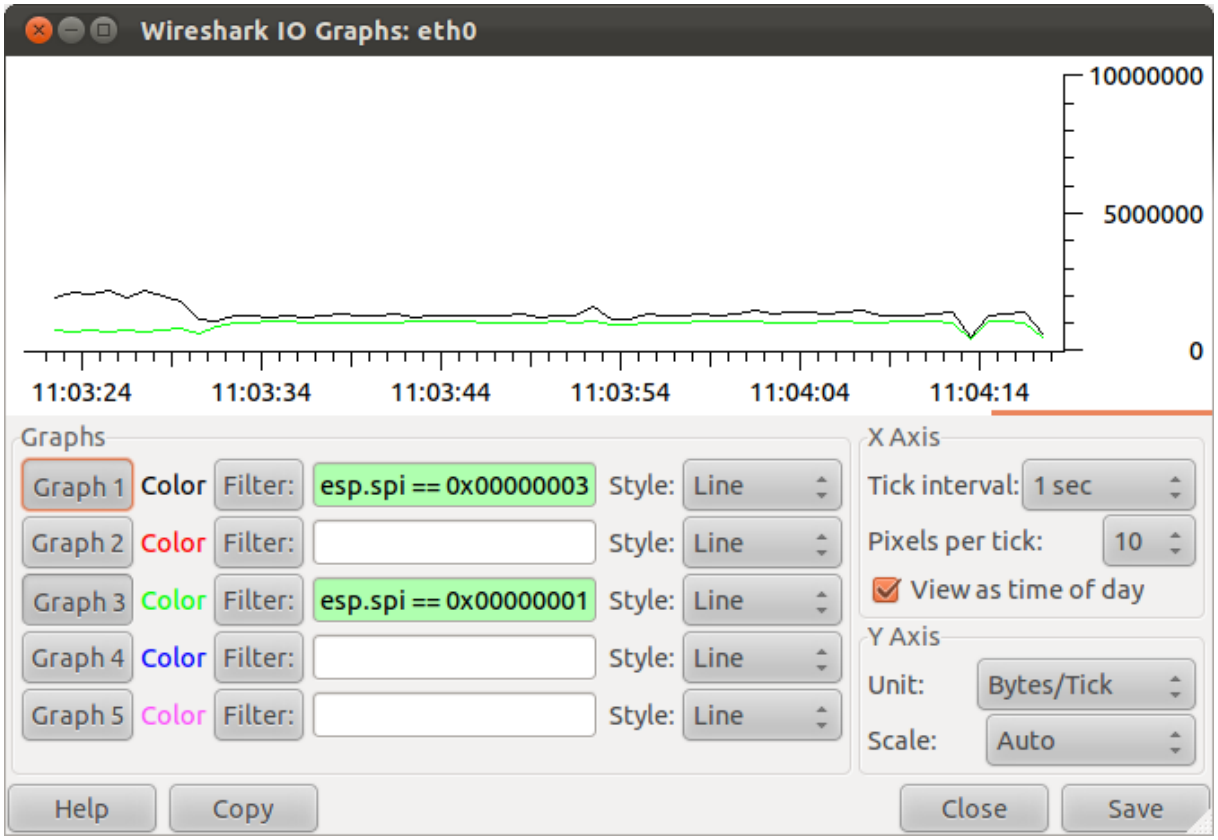

Figure 15 TFC detects changes in network conditions and adjust the SLA levels for video traffic

5.4.4 Conclusions

The TFC solution will provide an added layer of security compared to IPsec, but at some processing and minimal bandwidth costs (a few extension header's TLVs). Despite this added overhead we do believe that this is the most efficient way to provide traffic flow confidentiality in an IP network supplied by a combination of commercial and military providers. As seen in the experiments, the benefits are validated. The functionality is verified and the cost/benefits of the various functional options are illustrated.

6 CONCLUSIONS

Task 3 has addressed security topics in the three areas *transport network protection*, *key management* and *protection of user traffic*. Further, the work on security has covered confidentiality, integrity, authenticity and availability aspects.

Various theoretical studies have been delivered. The core network protection study proposes mechanisms and procedures to thwart threats to national networks. The study also identifies future work. Outputs from the key management area are two studies on

the scalability of a tactical PKI. Both studies identify optimization opportunities and make recommendations. The two studies have been developed independently and their conclusions have not been harmonized. A third *key management* study is on the Multicast Internet Key Exchange (MIKE) protocol. The main contribution of this study is an assessment of MIKE for its use in tactical Ad Hoc networks. Lastly, a study on multilevel security and use of cross domain guards for CoNSIS network management has been delivered.

Task 3 has conducted several security experiments. One group of experiments addressed the issue of automatic detection of threats to national networks, using software probes, and the findings were documented. A second experiment on the network protection theme was a demonstration on how Network Authentication Header (NetAH) can be used to prioritize a military data flow with NetAH over an NGO one with identical QoS marking, but without such protection. The NetAH also enables authorized nodes to detect if the QoS marking has been changed by unauthorized nodes.

In the *key management* area an experiment on the scalability of a tactical PKI was conducted. Two different PKI software packages were used to demonstrate the feasibility of proposed optimisation methods. The findings were documented. Another experiment in this area was on the MIKE protocol. MIKE combined with the IPsec discovery protocol was used to automatically configure IPsec devices. Further work within this area was identified.

In the last area (protection of user traffic) an experiment on cross-domain information exchange was conducted. The experiment showed that messages were correctly blocked or forwarded by the guard depending on the confidentiality label of the message. In addition the experiment provided experience on integrating new security solutions in a service-oriented environment. A second experiment in this area was on protected communication between military and civilian networks. It provided experience on a range of security technologies for the separation of military and civilian networks. Lastly, the Traffic Flow Confidentiality (TFC) experiment validated and demonstrated the benefit of TFC. In addition the functionality of the TFC solution was verified.

The approach of conducting experiments on individual security topics was advantageous. Such an approach allowed concentration on one topic at a time. It also gave a controlled environment for the experiments. However, a more comprehensive approach is needed in order to bring new security solutions into operational systems. This comprehensive approach should be considered for future work.

On two security topics both theoretical studies and experimentation was performed. This approach was useful as the experimentation served as a supplement to the studies.

References

[1] Coalition Networks for Secure Information Sharing (CoNSIS) (2011), "System and Experimentation Architectures", CoNSIS/Task 5/DL/002, Version 1.0.

[2] Simone P. (2010), "Core Network Protection", CG/UM-ESIO/IDRE/10.109/V1.1, Cogisys.

[3] Bret N., Ridard B., and Simone P. (2011), "Core Network Protection - Lessons Learnt from Demonstration Tests", CG/UM-ESIO/IDRE/11.238/V1.0, Cogisys.

[4] Hegland A. M. and Winjum E. (2008), "QoS signaling in IP-based Military Ad Hoc Networks", *IEEE Communications Magazine*, vol. 46, no. 11.

[5] Kongsberg Defence & Aerospace AS (2011), "Implementation of the NetAH (Military Authentication Header / AH*/QAH), Technical Report", 2/1559/2-FCPR10127 Rev A.

[6] Winjum E. and Berg T. J. (2008), "Multilevel Security for IP Routing", *Proceedings of MILCOM 2008*.

[7] Hauge M., Brose M. A., Sander J., and Andersson J. (2010), "Multi-Topology Routing for Improved Network Resource Utilization in Mobile Tactical Networks", *Proceedings of MILCOM 2010*.

[8] Engohan E. and Simone P. (2010), "Scalability of a Tactical PKI", CG/UM-ESIO/IDRE/10.110/V1.1, Cogisys.

[9] Bret N., Ridard B., and Simone P. (2011), "PKI Scalability – Lessons Learnt from Demonstration Tests", Cogisys.

[10] Fongen A. (2010), "Optimization of protocol operations in a Public Key Infrastructure", FFI-rapport 2010/02499.

[11] Hegland A. M. and Ellingsrud H.-A. (2011), "MIKE assessment, Technical Report", 1/1559/1-FCPR10127 Rev B, Kongsberg Defence & Aerospace AS.

[12] Hegland A. M. and Ellingsrud H.-A. (2012), "The Multicast Internet Key Exchange (MIKE) in tactical Ad Hoc Networks", *IST-111 Symposium on Information Assurance and Cyber Defense*.

[13] Eggen A., Haakseth R., Oudkerk S., and Thummel A. (2010), "XML Confidentiality Label Syntax - A Proposal for a NATO Specification", FFI-report 2010/00961 (NU).

[14] Eggen A., Haakseth R., Oudkerk S., and Thummel A. (2010), "Binding of Metadata to Data Objects - A Proposal for a NATO Specification", FFI-report 2010/00962 (NU).

[15] Nordbotten N. A. and Gjertsen T. (2012), "Towards a certifiable MILS based workstation", FFI-report 2012/00049.

[16] Haakseth R. (2012), "SOA Pilot 2011 – demonstrating secure exchange of information between security domains", FFI-report 2012/00117.

[17] Nordbotten N. A. (2009), "XML and Web services security standards", *IEEE Communications Surveys & Tutorials*, vol. 11, no. 3, pp. 4-21.

[18] Nordbotten N. A. (2010), "Cross-domain access control in a military SOA", *Proceedings of MILCOM 2010*, pp. 448-455.

[19] Fongen A. (2012), "Protected and Controlled Communication Between Military and Civilian Networks", *Military Communications and Information Systems Conference (MCC 2012)*.

[20] Steinmetz P. (2012), "Multilevel security and network management in CoNSIS", Technical Report CD-2012-02 Fraunhofer FKIE.

[21] Steinmetz P. (2012), "Use of Cross Domain Guards for CoNSIS network management", *Military Communications and Information Systems Conference (MCC 2012)*.

[22] Sorteberg I. and Kure Ø. (2010), "Traffic Flow Confidentiality - Architecture and mechanisms", Baseline AS.

Glossary

ACL: Access Control List
AS: Autonomous System
BGP: Border Gateway Protocol
CC: Coloured Cloud
CE: Coloured Enclave
CRL: Certificate Revocation List
C-TNS: Coalition TNS
DNS: Domain Name System
DOS: Denial Of Service
EGP: Exterior Gateway Protocol
ICE: Inner Coloured Enclave
ICMP: Internet Control Message Protocol
IGP: Interior Gateway Protocol
IKE: Internet Key Exchange
LDAP: Lightweight Directory Access Protocol
MIKE: Multicast Internet Key Exchange
NTP: Network Time Protocol
N-TNS: National TNS
OSS: Operations and Support System
PCN: Protected Core Networking
PCS: Protected Core Segment
PIM: Protocol Independent Multicast
PKI: Public Key Infrastructure
QoS: Quality of Service
RSVP: resource ReSerVation Protocol
SLA: Service Level Agreement
TN: Transport Network
TNS: Transport Network Segment