Security Shift in Future Network Architectures

Tim Hartog, M.Sc
Information Security Dept.
TNO Information and
Communication Technology
Brassersplein 2, 2600 GB,
Delft, The Netherlands

Tim.hartog@tno.nl

Harm Schotanus, M.Sc
Information Security Dept.
TNO Information and
Communication Technology
Brassersplein 2, 2600 GB,
Delft, The Netherlands

Harm.schotanus@tno.nl

Cor Verkoelen, B.Sc
Information Security Dept.
TNO Information and
Communication Technology
Brassersplein 2, 2600 GB,
Delft, The Netherlands

Cor.verkoelen@tno.nl

ABSTRACT

In current practice military communication infrastructures are deployed as stand-alone networked information systems. Network-Enabled Capabilities (NEC) and combined military operations lead to new requirements which current communication architectures cannot deliver. This paper informs IT architects, information architects and security specialists about the separation of network and information security, the consequences of this shift and our view on future communication infrastructures in deployed environments. The result of this paper is a proposal for a new architecture which addresses both security and flexibility requirements in deployed infrastructures as well as system management while retaining the "system high" mode of operation.

1.0 INTRODUCTION AND BACKGROUND

The availability of all relevant and supporting information to get a complete Common Operational Picture (COP) on which decisions are based is crucial for the successful execution of a military operation. However, these Common Operational Pictures and decisions made during military operations are not solely based on the information that is provided by each nation's own systems but increasingly depends on the information that is shared by others. The possibility to exchange information becomes increasingly important for the success of an operation. In this paper, we focus on the architecture for deployed military communication infrastructures. The aim of these infrastructures is to enable future interconnections of different communication infrastructures that belong to different nations and organisations. These interconnections support the exchange of information. However, the concept of controlled information exchange lies outside the scope of this paper. For more information about mechanisms to enable controlled information exchange see [1].

Nowadays military organisations, for example members in the NATO Response Force (NRF), have to react quickly to incidents worldwide. The ability to quickly deploy networked information systems suitable for the particular situation is therefore crucial to effectively fulfil their tasks. An important driver for future communication architectures is Network-Enabled Capabilities (NEC) which is based on an integrated and coordinated deployment of all capabilities, heavily leaning on controlled information sharing [2]. Controlled information sharing implies that the responsibility for sharing information lies by the owner of the information, who has to determine whether information is suitable to be shared with other members.

Current practice is that each nation or organisation deploys its own stand-alone networked information systems. To participate in different operations, an organisation has to deploy multiple stand-alone networked information systems. This has a severe negative impact on:

		Form Approved OMB No. 0704-0188

Report Documentation Page

1. REPORT DATE **NOV 2010**	2. REPORT TYPE **N/A**	3. DATES COVERED **-**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Security Shift in Future Network Architectures	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Information Security Dept. TNO Information and Communication Technology Brassersplein 2, 2600 GB, Delft, The Netherland**	8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADA564697. Information Assurance and Cyber Defence (Assurance de l'information et cyberdefense). RTO-MP-IST-091

14. ABSTRACT
In current practice military communication infrastructures are deployed as stand-alone networked information systems. Network-Enabled Capabilities (NEC) and combined military operations lead to new requirements which current communication architectures cannot deliver. This paper informs IT architects, information architects and security specialists about the separation of network and information security, the consequences of this shift and our view on future communication infrastructures in deployed environments. The result of this paper is a proposal for a new architecture which addresses both security and flexibility requirements in deployed infrastructures as well as system management while retaining the system high mode of operation.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT **SAR**	18. NUMBER OF PAGES **10**	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**			

1. the ability to deploy them quickly,
2. the ability to exchange information using the communication infrastructure,
3. system management,
4. flexibility and adaptability,
5. costs.

To overcome these challenges a shift to a different architecture of networked information systems is necessary. The capability to share infrastructural components with participating nations and organisations without a negative impact on the information security and network robustness may reduce the need to deploy stand-alone networked information systems. Sharing infrastructural network components does not imply that the information is also shared with the participating nations and organisations. This will lead to the principle of separation of a (shared) communication infrastructure and the information (systems) itself. The information system may still use the communication infrastructure but each having its own set of protection requirements.

Developments such as Protected Core Networking (PCN) and Trusted Operating Systems may enable this shift toward a new architecture. These developments will not be without consequences, especially in the areas of protection of the shared infrastructure, protection of the information and cryptographic technology a lot of challenges have to be solved. This article highlights the most important of these technical challenges.

2.0 CURRENT INFRASTRUCTURES

Current deployed infrastructures are mainly developed following national requirements, based on a generic deployment scenario. In case the deployed infrastructure participates in a coalition environment, coalition specific requirements are defined and must be implemented in the deployed infrastructure. Compared to the NATO Response Force, each member nation will have a six-month period for training and testing of their capabilities, including capabilities with respect to communication and information exchange.

The whole infrastructure a nation or organisation brings to a coalition may comprise more than one domain. We distinguish four different types of domains:
1. The information domain
   This entails the means by which the C2 functions are performed for a predefined information classification.
2. The technical infrastructure
   This entails all hardware needed to enable communication.
3. The security domain
   This entails the security requirements which apply to an information domain.
4. The responsibility for these domains.
   This entails the collection of information domains and technical infrastructures for which an organisation is responsible.

Traditionally a domain is bounded by its own type (classification) of information which is processed and the security requirements which apply, and consists of its own infrastructure. Hence the boundary of the infrastructure is effectively the same as the information and security boundary. This is a direct result of the system high mode of operation. These traditional domains are strictly separated. It is not uncommon to bring a Mission Secret, a National Secret, a National Restricted, and an Unclassified domain to the coalition. During a mission, each of the participating nations and organisations is the administrative owner (responsible) of one or more technical infrastructures.

In the current system-high approach, with different infrastructures for each security domain, we see that different domains are delineated by the technical network boundaries. This means that the security requirements for a certain information classification level apply to the whole technical network which processes that classified (unencrypted) information. And because these domains (the information and security domain and the technical infrastructure) are combined, the security measures which protect the infrastructure and the security measures which protect the information are integrated and mutually dependent, and therefore hard to distinguish and separate.

An improved situation would be where information can be exchanged over a shared, common infrastructure. Both the information domain and the communication infrastructure will have their own protection mechanisms to fulfil their own specific security requirements. For a communication infrastructure this may be the availability of communication between end-points, for information this may be to safeguard its confidentiality. This is a very important distinction.

2.1   Examples of current communication infrastructures

Two examples of currently used communication infrastructures, TITAAN and a Navy Ship, are described below to clarify these disadvantages.
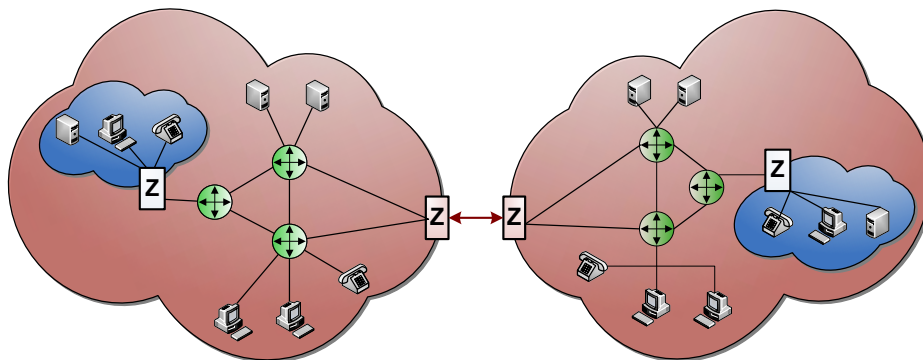

Figure 1 - Current practice in deployed environments

Figure 1 shows an example of a general overview of two interconnected Command Posts (CPs) in the Dutch TITAAN architecture. The TITAAN infrastructure is generally Mission Secret – the red clouds in Figure 1. Each CP is equipped with a basic set of infrastructure components such as routers, transmission devices, switches, and a number of workstations and VoIP telephones. The communication between CPs may use various transmission media, including radio, satellite and land lines. Usually the transmitted information and the CP infrastructures are protected by line encryption devices (Z) to ensure the confidentiality of the information. The interfaces are typically standardised and the red cloud is managed remotely from the home base in the Netherlands. Because network management is part of the mission domain each mission requires its own separate management infrastructure.

Besides Mission Secret, there can also be other system high environments in the TITAAN-based deployed infrastructure. For example a National Restricted cloud. Such a cloud is typically layered on top of the red cloud. This is shown as a blue cloud in Figure 1. This layer is constructed by using IP encryption devices. This IP encryption device creates a tunnel throughout the red cloud through which the traffic is forwarded. The leads not only to extra overhead, but also to the limitation in the interaction between the two clouds on the communication layer which in turn may restrict network capabilities (e.g. QoS).

Another scenario is one in which another nation or organisation uses the Dutch TITAAN infrastructure as a transmission medium. In this situation a similar approach to that of the "*blue*" cloud is taken. The difference is that a second layer of protection (tunnel) is added. The first layer is meant to ensure the

confidentiality of the other nation's data. The second layer of protection is meant to ensure the protection of the Dutch TITAAN red cloud. This is shown in Figure 2. It is clear that this incurs a significant overhead.
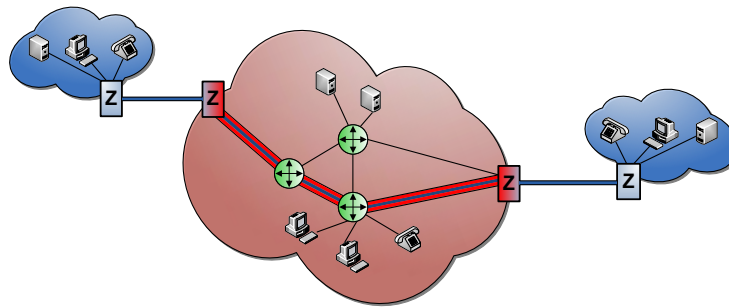
Figure 2 - Tunnelling over infrastructure from other nations or organisations.

The second example is the infrastructures onboard navy ships. These infrastructures typically consist of multiple separated networks in which the separation depends on the type of information that is processed. For example, on a ship there are separate networks for voice communication and for the management activities of combat and platform systems. There is no information exchange between these networks. The number of separated networks onboard of a ship depends heavily on the type of navy ship and its role.

The use of separated networks onboard a ship has several consequences. Not all communication networks can be used by other partners, e.g. partners that get onboard the ship for a (limited) period of time. This is caused by the physical limitations of these networks and network interfaces, and due to the classification that limits access to the infrastructure. Adding yet another network with basic communication services for other nations or organisation and including the possibility of external communication may solve this problem. However this solution adds yet another infrastructure into an already limited space.

2.2 Disadvantages of current architectures

The current approach in deployed communication infrastructures leads to highly varied infrastructures (1) between nations and organisations and (2) during different missions. This combined with the System High approach results in an abundance of (mostly unconnected) infrastructures. This has several consequences, including:

- Inflexible deployment;
  Deployment and use is inflexible since these infrastructures cannot be shared efficiently with others, both national as well as international.

- Inefficient use;
  There are several communication infrastructures that basically provide similar communication paths between end-points. All these infrastructures compete for the same, often scarce resources such as radio frequencies and satellite channels.

- Ineffective use of resources;
  When nations share communication infrastructures they implement an extra layer of protection to protect their own communication infrastructure and information. A security measure could be adding a self-defined and controlled tunnel throughout the used communication infrastructure through which all traffic will be forwarded. This leads to extra overhead.

- Inefficient network management;
  The system high approach and the abundance of communication networks require separate network management for each communication infrastructure. The security domain also entails the infrastructure. Thereby security requirements which are meant for the information domain also apply to the infrastructure domain, leading to an excessively high level of security for the infrastructure.

- High costs;
  Due to the amount of system-high environments and management networks which have to be deployed, transported, maintained, powered, and so on.

This chapter described how the current System High approach impacts communication infrastructures by merging different domains (information domain and the infrastructure) within the technical network boundaries and has shown the negative consequences thereof. The next chapter describes a new architecture for communication infrastructures which aims at the following goals:
- Improve flexibility.
- Reduce deployment time.
- Provide a basis infrastructure for future coupling to facilitate information exchange.
- Consolidating the security level.

3.0 DEVELOPMENTS AND INTEGRAL VISION

The main problems identified in the previous chapter stem from the fact that the information domain and infrastructure domain are combined within one security domain. This chapter describes an architecture in which these domains are clearly separated. This is realised by the creation of a shared 'protected black' core network for different organisations wherein protection of information (confidentiality) is moved from the network border towards the information domain (individual computer systems and even to logical partitions within a system).

Two specific developments give direction to this proposed new architecture for communication infrastructures, being Protected Core Networking (PCN) and Trusted Operating Systems.

PCN provides the design for the shared protected black core network, the Protected Core (PCore). This PCore forms the highly available transport network. This core network is by definition unclassified and enables the classified domains, also called Coloured Clouds (CC), to communicate over the core. Confidentiality of the exchanged information between classified domains is the responsibility of the Coloured Cloud owners and not the responsibility of the protected core.

The core network describes and publishes its capabilities to the different segments[1] and to the Coloured Clouds. These capabilities can for example be Quality of Service, or Traffic Flow Confidentiality. An important effect of a black core network is that many network components, such as routers and switches, are moved from the red domain to the black domain. This reduces the security requirements of these components significantly. The components however still need to be protected, but they do not have to comply with the more strict information confidentiality requirements. This less strict requirements also enables a more flexible manner of remote management for these components.

Protected Core Networking (PCN) pushes the system-high environments to the end-network and even further, to the information domain. This enables the shared usage of available communication infrastructures from different nations to provide a high availability transport network without the creation of individual communication links per nation. Figure 3 shows this new architecture which consists of a protected black core network which may be used by different nations and organisations. The term "protected black network" refers to the fact that the black network is protected by some basic protection

---

[1] Segments are parts of the protected core which are supplied and managed by different nations and organisations.

measures. This includes for instance some measures to implement authentication and authorisation before access to the network is granted. The term "black" means that the network does not contain any classified information. This is the opposite of a "red network" which contains classified information. Before classified information leaves a "red network" it is encrypted and only then enters the "protected black network". So the black core implements security measure to protect the infrastructure (availability) and the red network security measure focus on protection of the classified information.
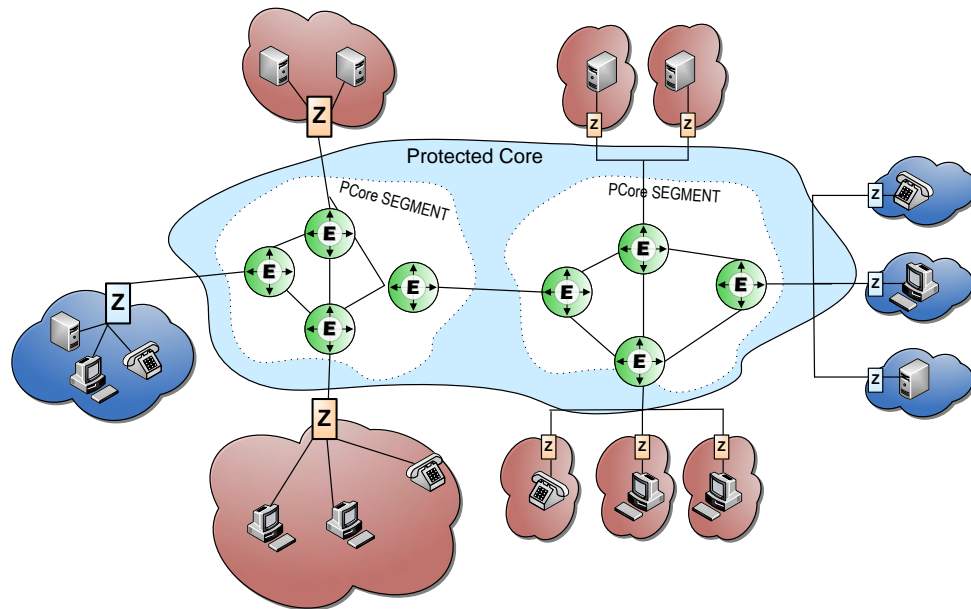

Figure 3 - New architecture for communication infrastructures

The components for the protected (black) core, known as the Protected Core Segments, are brought by the nations and organisations themselves and will be interconnected. These components may vary between a single network element to an entire sub-network. PCN has a more open character and can be used by all participants within a mission if they agree on the defined terms of usage and comply with the minimal requirements that apply for the PCN. Specifically PCN defines a set of well-defined interfaces that provide the means to interconnect these Protected Core Segments to each other and to connect to the "Red" clouds.

The classified domains will be connected through the protected black core network. Any classified domain wishing to use the PCN needs to have appropriate credentials to be able to connect to the PCore. Such a classified domain may vary between an entire network domain, consisting of multiple locations and computer systems, to a single computer system.

With the introduction of PCN the system high concept can be maintained. The classified domains are only interconnected through the PCN and still fall under the responsibility of the owner of the classified domain. However, this is not enforced by the PCN concept. How the classified domains are created, and how the confidentiality of the classified information is safeguarded is the responsibility of the owner of classified domain. This is most likely ensured by using different keys and cryptographic units for each domain, and thus creating different classified domains individually connected to and through the protected black network.

If we extend on this separation between the communication infrastructure and the information domain and the consequent shift of protection of information (confidentiality and integrity) from the network boundaries to the information domain we can utilise Trusted Operating Systems to provide a Multiple Independent Levels of Security (MILS) solution, we call this a fat client. Note that MILS is very different from Multi Level Security (MLS). See figure 4.
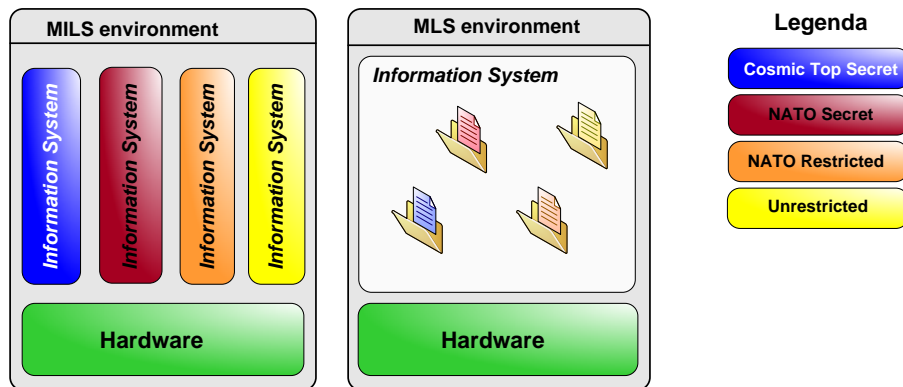
Figure 4 - MILS versus MLS environment.

MLS refers to a computer system which processes information with different classifications and permits simultaneous access by users with different security clearances. A MILS solution is a computer system where information of different classifications are strictly *separated in partitions of the system*. In both MLS as MILS access to the classified information also depends on the need-to-know of the user.

Such a Trusted Operating System thus provides the ability to run different system-high environments on one hardware platform. This is realised through virtualisation, which can create and run different (classified) partitions (virtual environments) securely while maintaining their strict separation. An additional advantage of such a solution is the improved flexibility in deploying hardware and reduced complexity of remote management. Virtualisation makes it much easier to deploy and restore an image of a system, without the need to access the actual system. Furthermore, hardware can be replaced without effecting the environment as long as the virtualisation layer supports it.
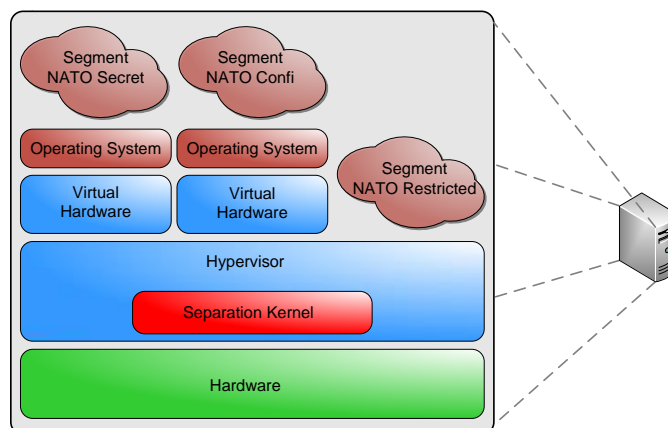
Figure 5 – Trusted Operating System used as MILS solution

A MILS solution is built around a separation kernel. The separation kernel forms the barrier between the virtual systems and the actual hardware. The sole purpose of the separation kernel is the separation of the different partitions to prevent uncontrolled information flow between these systems, see Figure 5.

The idea of a MILS solution is not new, see [4], but it has not been available yet in (approved) products. This may change with the appearance of several new commercial products. A MILS solution can be seen as an enabler for PCN to extend the protected black core and shift the red domain to the physical systems (information domain) itself and even to the 'inside' of these systems (such as servers or workstations). Note that this idea has large consequences for the infrastructure design and the way cryptography is used to separate the red domains.

In the first place the amount of Cryptographic Units (CU) will increase significantly. The CUs will move towards the servers and workstations and thus every server and workstation or cluster thereof will need a CU of its own. In the MILS-based Trusted Operating System, this is taken one step further because the CU has to be a component of the fat client itself and be controlled by the separation kernel to separate the communication of the different (virtual) systems. As a result many more users will have to interact with cryptography. The proposed architecture can, to a large extent, be accomplished with current IP CUs, but not by line encryption devices because they operate below the IP layer and hence cannot be used across IP networks.

The increased amount of CUs has a large impact on key management. Much more cryptographic key data will have to be generated, distributed, exchanged and managed. This will probably lead to more frequent key exchanges and more different keys to mitigate the risk of compromised keys. This leads to a strong requirement for electronic key management and distribution.

If we look back to the described developments in this chapter we can see a transition towards a 'protected' black core network which is shared by different organisations wherein protection of information (confidentiality) is moved from the network border towards the individual computer systems and even to logical partitions within a computer system. Summarising the changes with current practices we see:

- The creation of a protected black core network by and for different nations and organisations.
- The utilisation of a shared infrastructure which is flexible in use. This black core network ensures that communication requirements can be fulfilled.
- The separation of security measures for the infrastructure, focused on availability and integrity, and the security measures for the protection of information, focused on confidentiality.
- Application of the System-High approach on the information domain level instead of the infrastructure level. This entails the shift of information security towards the boundaries of the network (information domain).
- Multiple Independent levels of security (system high environments) within computer systems.
- Cryptography is shifted from the edge of the network towards the end-nodes within the network.
- Different application of cryptography and key management.

4.0   FUTURE WORK

Although the described new architecture offers many benefits there are several challenges left to be solved before this concept can be used in the field.

4.1   Architecture of the red domain

The network architecture in both the red and black domain changes significantly. One of the consequences is that the systems in the red domain have no knowledge of their point of (physical) attachment in the network, whether a system is in the same area or on a remote site. Note here that systems in the red domain may relocate to other sites (different points of attachment in the black domain), but this will not be visible in the red domain. This may result in large overhead, this is exemplified below.

Assume for example a red domain spread over two physical locations A and B, each contains a mail server, connected over a satcom link. If a client in location A needs to send an email, it is most logical to connect to the mail server also in location A. However if the client moves to location B it would mean that sending the mail to mail server in location A may incur excessive bandwidth waste on the satcom link and therefore it would be more efficient to connect to the mail server in location B. But the client must thus be capable of finding the right server

For example, a VoIP connection is setup between two clients. The clients are unaware of each others point of attachment. A VoIP connection requires a reasonable bandwidth and latency to work properly. So, when the VoIP communication uses a satellite link between two clients it must still result in a reliable connection. In the situation where both clients are connected to the local network this hardly forms a problem, but the application (VoIP client) cannot make that distinction.

As a result, service discovery and signalling between the red and black domain become of a very high importance. This requires interaction between the red and black domain to determine whether requirements from the red domain, can be granted in the black domain. Results should be communicated back to the red domain, without compromising the security of the red domain.

4.2   Protection of the black network

The black network also requires protection even though it is an unclassified network. The integrity of the components needs to be protected, to ensure the availability of the black network and hence the red domain. Nevertheless these protection mechanisms have other requirements than traditional cryptographic units provide, which focus on providing confidentiality, and where integrity of information and systems is usually included as a bonus.

4.3   Cryptographic units and key management

The requirements to the cryptographic units and its usage change. The increasing amount of cryptographic units and shift from the network boundary towards the end-nodes (information systems) themselves requires smaller (maybe even logical instead of physical) and more efficient cryptographic units. Additionally the cryptographic units are no longer applied at relatively fixed places but appear practically all over the network. This requires an efficient and secure key management solution.

5.0   CONCLUSIONS

To conclude; combining the development of Protected Core Networking with Trusted Operating System technology results in a very flexible, promising architecture which can decrease costs and effort associated with network management and can reduce the amount of components, especially the components in the red domain. However there are many open issues that will have to be addressed; especially the application and amount of cryptography as well as the network architecture of the red and black domain and the interaction between them.

REFERENCES

[1]   B.J. te Paske, D. Boonstra, D.H. Hut, H.A. Schotanus, *Information Labeling – Cross- Domain Solutions*, Intercom Vereniging Officieren Verbindingsdienst, 38e jaargang, nr. 2, June 2009.

[2]   Martis, E.R., et al. *Information Assurance : Trendanalyse*, TNO report TNO-D&V 2006 B312, Oktober 2006.

[3]   Hallingstad, G., Oudkerk, S., *Protected Core Networking – Initial concept description,* March 2007.

[4]   Rushby, J., *Design and Verification of Secure Systems*, ACM Operating Systems Review, Vol. 15, No. 5, pp 12-21, http://www.csl.sri.com/papers/sosp81/sosp81.pdf, December 1981