A SECURE NEC-ENABLING ARCHITECTURE DISENTANGLING INFRASTRUCTURE, INFORMATION AND SECURITY

Daniel Boonstra, Tim Hartog, Harm Schotanus, Cor Verkoelen
TNO
Brassersplein 2
2612 CT
Delft, ZH
THE NETHERLANDS
Tim.hartog@tno.nl Harm.schotanus@tno.nl Cor.verkoelen@tno.nl

ABSTRACT

The NATO Network-Enabled Capability (NNEC) study envisions effective and efficient cooperation among the coalition partners in missions. This requires information sharing and efficient deployment of IT assets. Current military communication infrastructures are mostly deployed as stand-alone networked information systems operating at the System High mode. This impedes the ability to support effective and efficient information sharing. This paper describes a security architecture for deployed military communication infrastructures based on new advanced security concepts. The objective of this security architecture is to facilitate an infrastructure that provides flexible and efficient use of technical resources and enables controlled exchange of information.

1.0 INTRODUCTION

Military operations are increasingly carried out by a coalition of different nations and organisations in order to reach the mission goals. Effective and efficient cooperation among the different coalition partners requires information sharing and efficient deployment of IT assets. This is also envisioned in the NATO Network-Enabled Capabilities feasibility study (NNEC FS)[1]. Current military communication infrastructures are deployed as stand-alone networked information systems operating at the system-high mode. This limits the current infrastructures in the ability to effectively share information. Hence, a change in these current infrastructures is required. The current military communication infrastructures are based on the concept of the 'duty-to-protect' the information. However current operations require a more 'duty-to-share' approach without compromising current security demands with regard to the protection of information. This change includes changes to network infrastructure set-up as well as changes in the security architecture. This shift in network architectures and the consequences are described in [2].

1.1 Problem description

The system-high approach shown in Figure 1 has consequences with respect to effective and efficient cooperation. First of all, stand-alone infrastructures limit the ability to interconnect these infrastructures for information sharing. This does not mean no interconnections are established at all or no information can be shared, however current means used to interconnect with, or share information with coalition partners are not always reliable or verifiable regarding security and are mostly specific for each interconnection.

Secondly, each environment has to provide the necessary communication means to transport the information within the system-high environment, resulting in possible inefficient use of available communication means among coalition partners.

Finally, each system-high environment has its own specific implementation of security requirements making this a cost-inefficient approach (stovepipes).

These consequences are the result of the tight coupling between the information security requirements and the security measures implemented as part of the infrastructure. They result in the inability to achieve effective and efficient support for information sharing in current architectures.

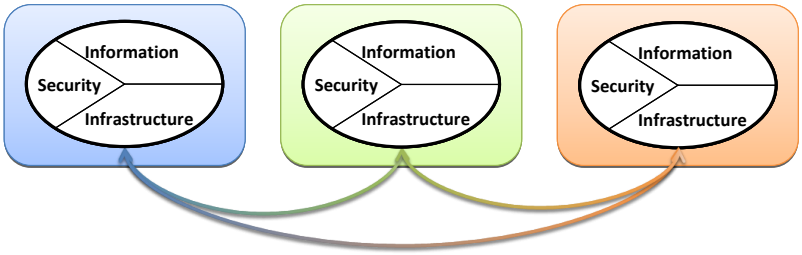

Figure 1: Interconnected System High communication infrastructures

1.1 Objective

The objective of the proposed architecture is to facilitate the creation of an infrastructure that enables:

1. flexible and efficient use of relatively scarce technical resources;
2. the controlled exchange of information between different classified information domains.

This has to be realised without compromising current security demands for the protection of information. An approach to reach these benefits is to disentangle the current cohesion between the information and infrastructure, as indicated by the NNECC FS [1]. The NNEC FS describes (a) Networking, (b) Information and (c) People as areas that should be further developed. This paper focuses on the first two aspects.

Networking is the need for a robustly networked force to enable improved information sharing. This means interconnecting the different coalition partners in order to be able to share information.

The area of information is about the ability to share this information, making use of the robust network. Despite the topics of networking and information are closely related, one cannot (effectively) share information without the network and the network is useless unless it is used to share information.

NNEC FS states further "The strategy for developing the networking and information sharing aspects of NNEC focuses on the 'joining together' of networking systems and core information systems from NATO and NATO nations, to form a Federation-of-Systems (FoS) capability that implements the Networking & Information Infrastructure (NII)." In order to support coalition-based military operation in the form of a federation of systems, the basic principle for our proposed architecture is the disentanglement of information and infrastructure as shown in figure 2.

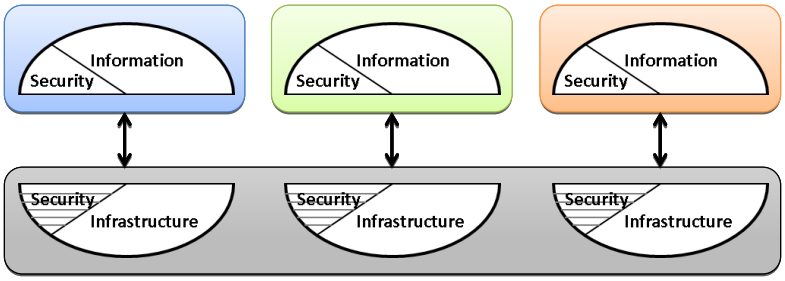

Figure 2: Disentanglement of information and infrastructure

2.0 PROPOSED ARCHITECTURE STEP BY STEP

This chapter describes the different building blocks of the proposed architecture individually. The description includes a brief overview of each building block, the impact on the information or networking objective of NNEC FS, and the relation with other building blocks.

2.1 A protected shared core network

NATO is developing a concept of a shared network architecture based on an unclassified network infrastructure [3][4]. The infrastructure's primary function is to provide connectivity among the different network nodes and a basic level of protection of this unclassified network. Unclassified means that the network does not provide any confidentiality measures of the information that is transported. The unclassified core network, known as the Protected Core (PCore) is a dynamic collection of parts brought and managed by the different participants in a coalition. These individual parts are known as Protected Core Segments (PCS).

The PCore provides connectivity and security measures to ensure the availability of the network. An important aspect of the PCore to be able to guarantee the availability and robustness of the communication network is access management – which and under what conditions can nodes connect to the PCore. This is meant by the term 'protected'. PCN therefore describes an interface between different PCS and between a PCS and users of the PCore.

A collection of nodes that operate at a certain classification and is connected to the PCore is named a coloured cloud (CC). As the PCore has no means to provide confidentiality, information confidentiality protection is left to the CC. Figure 3 shows the connection between the PCore and different CC's.

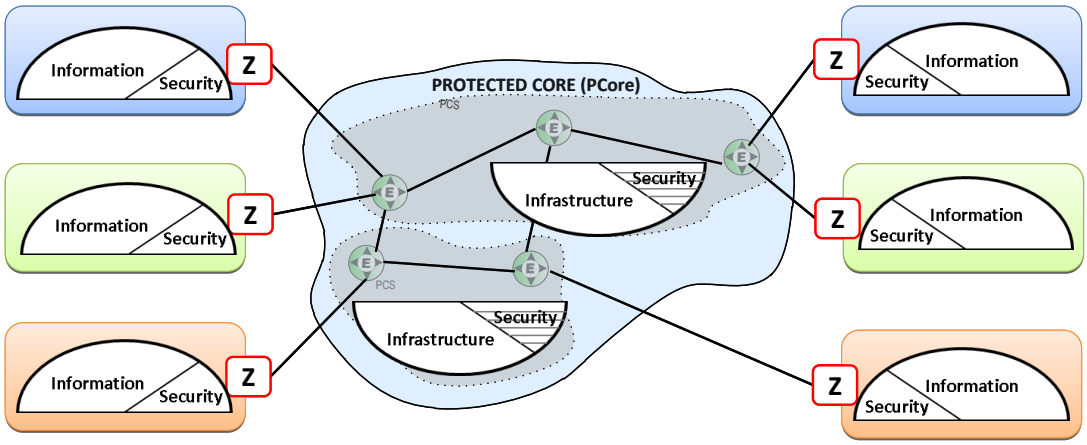

Figure 3: Protected Core connected to different coloured clouds

PCN will allow the use of scarce resources such as bandwidth and frequency ranges more efficiently, as they can be shared among participants. And in cases where one's own connectivity fails or is not present, connectivity of another participant can easily be used. The more elements the PCore is comprised of, the more efficient the core can be used by the CCs. This implies that the coloured clouds should be made as small as possible.

2.2 Compartmented workstations

Originally, workstations and servers are specific for a single security domain, e.g. one system-high environment. Recent improvements in the development of operating systems allow combining more than one of these security domains on a single physical system. These improvements are based on the concept of Multiple Independent Levels of Security (MILS) [5][6]. In this set-up a virtualisation layer provides a secure

basis for building differently classified compartments side by side on one physical system. These compartments can be seen as small system high environments, possibly containing differently classified information but processing all this information according to the security classification of the compartment. This is shown in Figure 4.

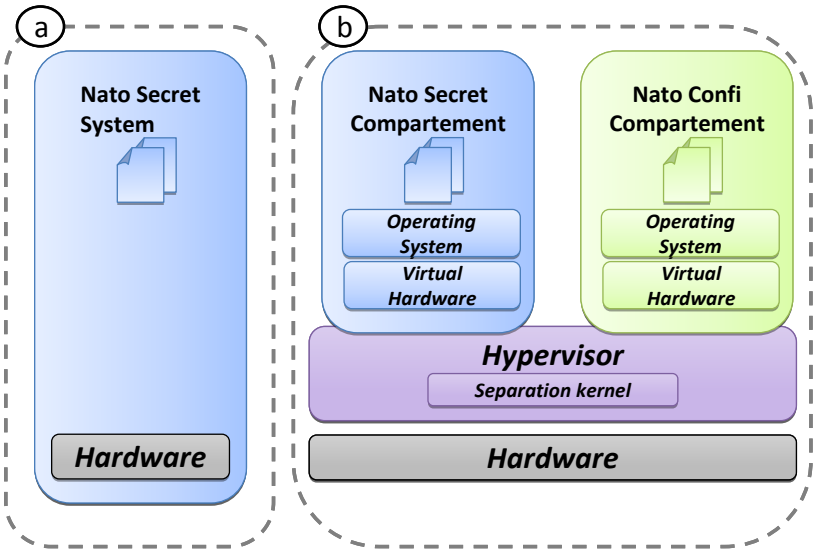

Figure 4: (a) non-MILS configuration, (b) MILS configuration

A general virtualisation layer however does not provide the required assurance that information in the compartments is strictly separated from each other. A more strict separation of the compartments can be realised by using a hypervisor in combination with a separation kernel [17]. This hypervisor enforces the separation and the virtualisation layer. The separation is needed for all physical resources that are shared by the workstation, such as memory, CPU, network access and storage. A set of requirements identified by the US Information Assurance Directorate for a separation kernel are defined in a protection profile [7]. Note that the concept op MILS is not the same the concept of MultiLevel Security (MLS). MILS is targeted at the separation of classified information, as where MLS is meant to control access to (differently classified) information in one system.

A system that contains different compartments can reduce the amount of physical systems needed to be brought to a mission, because there is no need for individual systems for each classification. Furthermore it provides more flexibility in deploying hardware due to the fact that hardware is not bound to a specific classification. This will also reduce the need to provide multiple infrastructures, one for each classification.

As a result, the compartments within a system can be seen as the Coloured Clouds of PCN. One CC can be diminished to merely one compartment within a system. That is, each compartment will become a CC. Internally, the compartments typically do not need any additional security measures in comparison to workstations as used in traditional deployed environments. All required additional security is provided by the separation kernel, essentially ensuring the separation of the information among compartments. The separation of compartments does not explicitly address the protection of confidentiality while it is transported. This is the subject of the next section.

2.3 Connecting the Networking & Information Infrastructure

As a compartment within a system has become a Coloured Cloud ensuring the confidentiality from other compartments, it also needs to ensure the confidentiality of the information when information is exchanged with another Coloured Cloud of the same classification using the unclassified network infrastructure PCN. To guarantee the confidentiality protection, the needed cryptographic systems need to be integrated as close to the compartments, especially in the situation where multiple compartments are created on one system.

This results in the integration of the cryptographic systems with the systems, as hardware or software, as shown in Figure 5. The cryptographic systems are illustrated as the [z], and indicate the classification of the information they protect, e.g. NATO Secret (ns). The classification level of the compartments determines the assurance levels of the cryptographic systems and the possible options of integration. As there can be compartments of different classifications on one system there may have to be different cryptographic systems and keys involved.

Typical workstations only contain one network interface card which must be shared by all compartments that need to access the unclassified network infrastructure – the PCore. Assuming that different classified compartments exist and thus multiple cryptographic systems and keys, the separation kernel must now also ensure that the correct cryptographic unit and key is used when information is shared over the PCore. Therefore the separation is not only needed for generic hardware but the separation kernel must also have direct control over the cryptographic unit. Current developments of new cryptographic units that support the flexibility required and meet the demands of future military operation are for example the NII IP Network Encryption (NINE) Interoperabilty Specifications [8] and the development of SCIP [9][10].

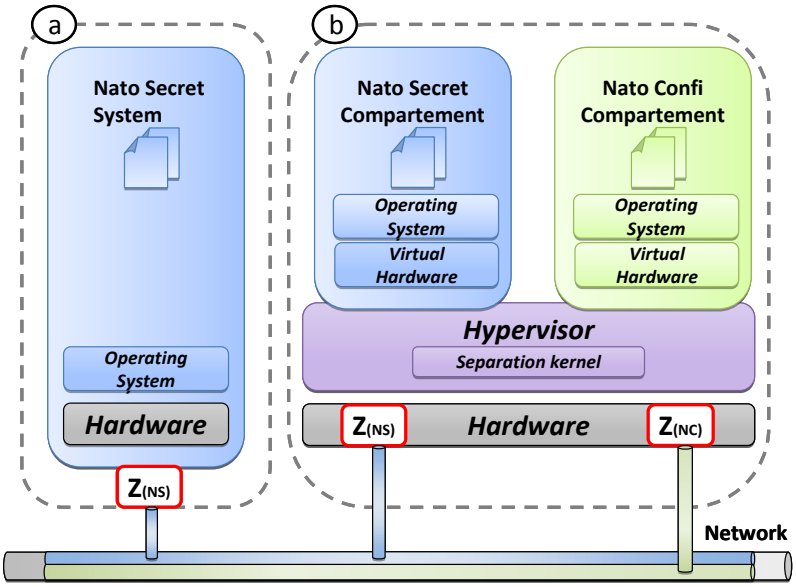

Figure 5: Cryptographic support to separate communication over a network

A consequence of having multiple classified compartments on one system, requiring cryptographic systems for each compartment, is the increase in the amount of cryptographic devices. Traditionally, systems of the same classification are grouped and share one cryptographic unit. The MILS concept connected to the PCore requires at least one cryptographic unit per system, grouping is not possible anymore. On the other hand, when combining compartments on one system and integrating the cryptographic unit with these compartments it is not longer required to have multiple infrastructures, one for each classification, is required. All compartments can use the unclassified PCore.

2.4 Extending NII with information release

Typically information exists within one classified compartment, and can be shared with other systems, compartments or servers that operate on the same classification. This can be realised using the MILS concept and the PCore. However if information must be shared with systems, compartments or servers operating at another (lower) classification level, we have to ensure that only the information that is permitted to be processed in the destination domain is exchanged. If for example information within a NATO Secret compartment that actually is classified as NATO Confidential needs to be shared with a compartment operating at the NATO Confidential level, one should determine that only the specified NATO Confidential

information is shared and no other information including NATO Secret.

A release mechanism can be used to determine whether the information is suitable for release to the destination domain. However, dealing with highly classified information requires a high level of assurance. Automatic interpretation of information is not yet deemed mature enough to provide these levels of assurance. Adding specific tags (meta-information) [11][12][13] to the information and determining the criteria that should be used by the release mechanism can be used to overcome this, especially for high classifications.

These specific tags and criteria must be used to determine whether the information is suitable for releasing to the target compartment [14][18]. The release mechanism needs to implement these criteria in a policy and use them as a filter on the specific tags that are added to the information. When information is presented to the release mechanism, it will determine whether it is in accordance with the policy. Only if that is the case the information is released to the other compartment. Hence the actual contents in the policy are the release criteria, the release mechanism facilitates the enforcement of these criteria.

To enable the exchange of information across two compartments on a system, there are two options regarding the location of a release mechanism. One, there is a centralised, dedicated system in the network at the same classification level that acts as a release mechanism. Or two, a dedicated compartment is created on the system itself where the functionality of the release mechanism is implemented, as shown in Figure 6. The applicable policy is illustrated as the red pentagon [P]. The first option has as a drawback that there is a strong dependence on a central system, which is not suitable for many situations such as the mobile domain. The benefit however is that the release mechanism can be used by multiple systems or compartments, reducing the number of release mechanisms that is required. The second option is more complex to implement due to the amount of policies that also have to be managed, but does not depend on a central system. However, the separation kernel for the compartmentalisation already provides a basis that can be re-used.

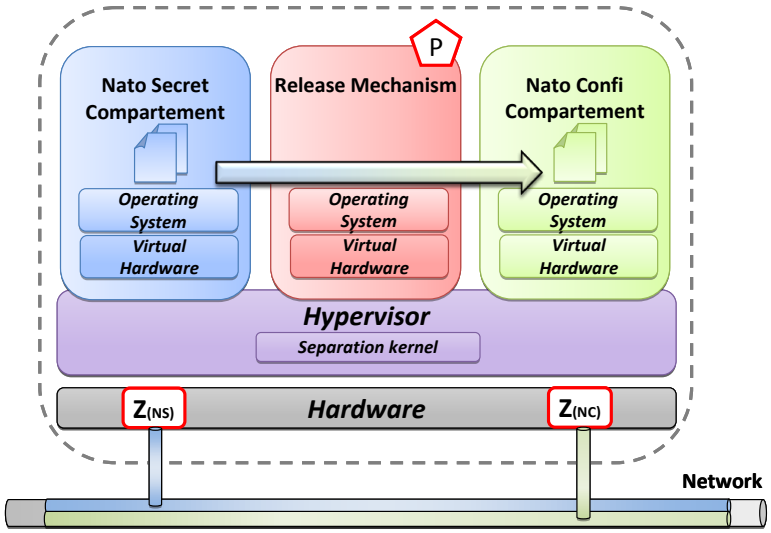

Figure 6: Conceptual position of a release mechanism in a MILS configuration

2.5 Labelling to support information release

One of the preconditions for information release is the availability of specific tags that are added to the information. We can use labelling to guarantee that these specific tags (labels) are securely tied to the specific information. That is, we enforce two requirements for the tags:

1. the tag nor the information cannot be altered without breaking the label,

2. the tags cannot be copied to another information object.

We propose to use a separate labelling object such as defined by NATO that is cryptographically tied to the information as a basis for the use in high classified information compartments.

A label can contain more tags than just the security classification of the information. The relevant set of tags depends on the release policy which has to be enforced. E.g. the resolution of a photograph or the GPS coordinates of a video can also be used as a criterion to determine whether the information may be released. These criteria could be defined if for example the capabilities of the sensors are classified.

The tags in a label can be used for more then determining if the information may be released. For instance for searching in data, archival or preselecting sets of information. This enables a growth path towards labelling as described in [15] but this also requires the alignment of labelling solutions.

The creation of the label is typically done by a user, determining the labelling information based on the information itself and the context of the information. However in some cases (some) labelling information can also be derived automatically. For example a camera attached to a UAV that take photos of a certain area. The capabilities of both the UAV and the camera can be used to automatically determine labelling information for the photos it takes. These capabilities are the resolution of the camera or the altitude of the UAV and can be included in the labelling information automatically.

Cryptography can be used to meet the two requirements regarding the binding between an information object and its label. E.g. Public Key Infrastructure (PKI) based signatures that are included in the label itself can be used to securely tie the tags to the information.

To be able to use the label in a release mechanism and to make a decision whether the information may be released, the release mechanism needs to be able to verify the label. This means that (1) it must verify if the label belongs to the information that is offered to the release mechanism; (2) it must verify whether the contents of the label conforms to the requirements stated in the policy; and finally (3) it must verify whether the label is still valid – i.e. the information and the label are not modified.

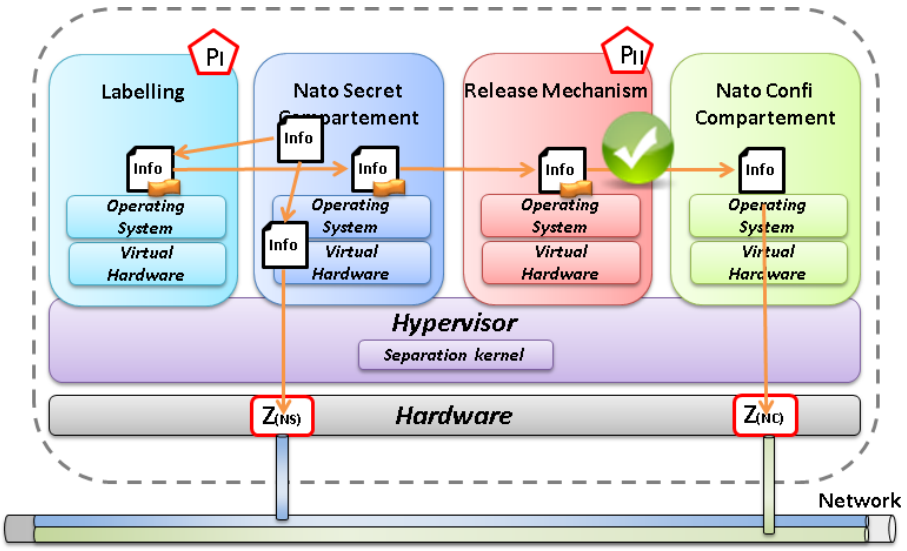

Figure 7: Labelling and release mechanism process

The label must be created in a secure environment where there are no adverse external influences to ascertain that the label is correctly bound to the right information. This means that we can use a separate compartment on the workstation with the sole purpose of the creation of labels and binding them to the information objects. This is further described in [16]. As shown in Figure 7 we have chosen for a decentralised implementation since the same benefits and drawbacks identified for the release mechanism apply.

3.0 MIGRATION OPTIONS

One advantage of the proposed architecture is that it supports a gradual realisation. E.g. creating a PCore first and connecting the existing system high networks on top of it as Coloured Clouds. Security measures with respect to protection of the confidentiality of the information will be based on the existing security measures, implemented on the borders of the system high network. In addition the compartmented workstations can be introduced gradually and can co-exist with current architectures and connect to other CCs or System High environments of the type. In this process the existing System High environments will in time be replaced by the systems with (multiple) compartments. The shift in confidentiality protection to individual systems with compartments does require a higher assurance level for these systems due to the added security functions within the systems.

Besides the increased assurance level, the security shift also introduces an increase of the amount of required cryptographic systems. Management of cryptographic systems and keys should therefore be part of the migration from traditionally environments towards the proposed architecture. Finally, labelling and release mechanisms can be introduced to enable information sharing. Figure 8 shows a possible migration scenario in which the modular and gradual characteristic of the proposed architecture is illustrated.

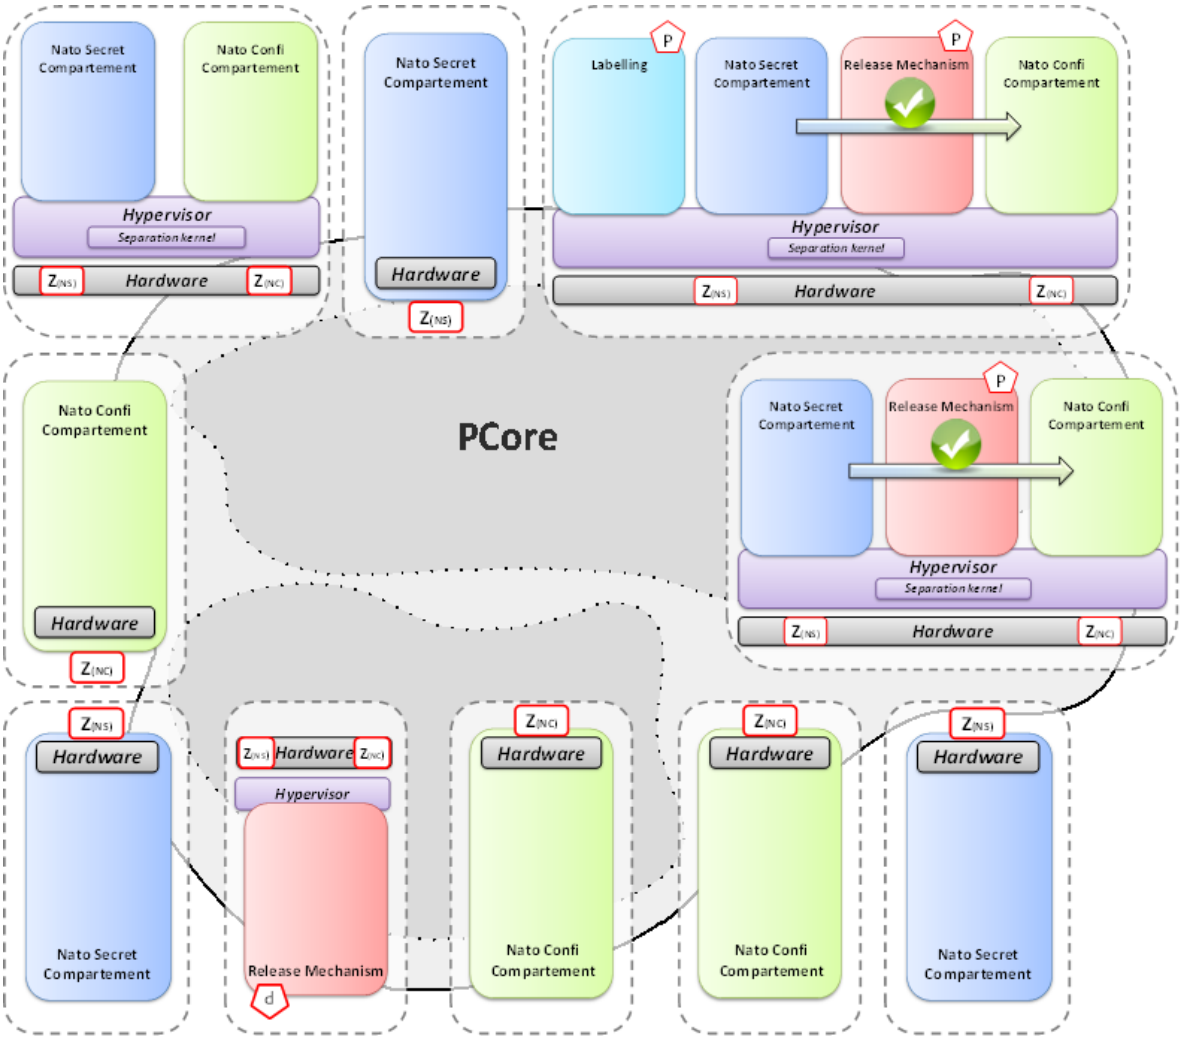

Figure 8: Flexible migration to the proposed architecture

Besides the gradual implementation of the proposed architecture, the sequence of implementing the

described modules is also variable. Thus one can choose to implement the compartmented workstations first using the current communication networks, or one could choose to implement the PCore first. Hence, the modular construction of the architecture leads to many possible scenarios for realisation in phases.

Additionally it is possible to realise this proposed architecture for national purpose only and at a later time offer the PCore to other nations and organisations within the coalition as well. This may simplify the technical and possible security-political problems associated with international sharing of a PCore. Nevertheless if the international PCore is still an objective, the interoperability is an important issue that will have to be addressed even if it is initially national oriented. This to prevent a (expensive) work-around or even redesigns of the concepts in a later stage when the national systems are used in an international setting.

4.0 FUTURE WORK

The described security concepts are useful to create a secure network and information infrastructure for future operations. To fully use the strength of the concepts additional research is required in the following areas:

1. Information management

   Since sharing of information is mandated, information management is increasingly important. Linking information and labels is a prerequisite and requires additional measures regarding the storage and retrieval. In deployed situations, a centralised approach will not always be applicable; hence a distributed approach is needed. The integration of Document Management Systems and Content Management systems extended with registry functionality can form a basis.

2. Policy management

   The described security architecture contains several components which act based on a security policy. These security policies are machine based rules derived from applicable organisational (security) policies and inter-organisational agreements (Memorandums of Understanding). The rules for the security components must be consistent and unambiguous for efficient and flexible operations. In general, policy management describes the theory to formulate policies, translate policies, the conversion of human readable rules into machine readable rules, management of policies, make policies (de)operational and phase out of policies. All this is required to deal with the complexity and to prevent security breaches.

3. Technology alignment

   The technological developments need to be further fine-tuned to user requirements and the interdependence of various developments. Currently not all of these developments can directly work together, for example Payload encryption and PCN currently cannot cooperate. For a compartmented workstation payload encryption may provide an important efficiency boost in comparison to traditional IP cryptographic devices.

4. Security requirements methodology

   Besides security concepts there is also the need for a structured methodology to define the security requirements and appropriate measures for each component or set of components in the communication infrastructure. In [19] a methodology is described to indentify and analyse inter-connection panes between information domains. This analysis is used to define security requirements for each indentified Security Policy Enforcement Point.

5.0 CONCLUSIONS

To reach the objectives of 'Networking and Information Infrastructure' (NII), the disentanglement of information and infrastructure is a necessity. This also implies that current security requirements have to be extracted and transformed as well. There are different security requirements for the infrastructure and for the information. We have proposed an architecture consisting of five basic components that can enable the efficient and flexible use of infrastructure independent of the classification of the information, over which the controlled sharing of information between different coalition partners can be facilitated.

An important aspect is how to arrive at said architecture – there are different migration strategies available. Each of the components can be introduced separately. Where international collaboration has to take place, it will require coordination and planning among the involved organisations.

Many of the described components are still in active development. Therefore to ensure that the results will be usable, nations need to address interoperability and standardisation of the components, but also focus on the interdependence of these technological developments. The creation of national deviations must be averted to ensure international deployment is possible.

6.0 REFERENCES

[1] Buckman, T.; Nato Network Enabled Capability Feasibility Study – Executive Summary; version 2.0, NC3A, http://www.dodccrp.org/files/nnec_fs_executive_summary_2.0_nu.pdf; July 2010.

[2] Verkoelen, C.A.A., et al.; Security shift in future network architectures; information assurance and cyber defence; NATO IST 091, 2010.

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

[4] Hallingstad, G; Oudkerk, S.; Selected aspects of Protected Core Networking; March 2008

[5] Rushby, J., Design and Verification of Secure Systems; Eighth ACM Symposium on Operating System Principles; pp. 12-21; Asiloma, CA; December 1981; (ACM Operating Systems Review, Vol. 15, No. 5).

[6] Timothy E. Levin, Cynthia E. Irvine, and Thuy D. Nguyen; Least Privilige in Separation Kernels

[7] Information Assurance Directorate, National Security Agency, Fort George G. Meade, MD. "U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness," Version 1.03, June 2007.

[8] NII IP Network Encryption (NINE) ISpec NATO C3 Agency; NINE Interoperability Level of Ambition, Edition 3.0, August 2008

[9] Daniel, E.J. Teague, K.A. Sleezer, R. Brewer, J. Raymond, J. Beck, W.J. Hershberger, J.; "The Future Narrowband Digital Terminal", The 2002 45th Midwest Symposium on Circuits and Systems, 2002. MWSCAS-2002, pp. II-589 - II-592 vol.2.

[10] M. Street, B. Bottesini, R. Russo, P. DeLaere, G. Elzinga, R. Murtland, "Providing Interoperable Secure Voice in Converging Heterogenous Networks", Military CIS Conference, Prague, Czeck Republic, September 2009.

[11] Oudkerk S., et al.; A proposal for an XML Confidentiality Label Syntax and Binding of Metadata to Data Objects, IST 091, 2010.

[12] A. Eggen, et al., Binding of Metadata to Data Objects – a proposal for a NATO specification,

Norwegian Defence Research Establishment (FFI) & NC3A, 22 April 2010

[13] A. Eggen, et al., XML Confidentiality Label Syntax – a proposal for a NATO specification, Norwegian Defence Research Establishment (FFI) & NC3A, 22 April 2010.

[14] Paske, B.J., et al.; Information Labelling – Cross-Domain-Solutions; Intercom 2009; volume 9, number 2, pp 47-50; June 2009.

[15] Hartog, T., et al.; Labelling: Security in Information Management and Sharing; 6th International Conference on Information Warfare and Security; March 2011.

[16] Hartog, T. et al.; High assurance platform for labelling solutions: TNO memo 2010 DEFENSIE 103, 14 December 2010.

[17] Sorrells, J.; WikiLeaks Shows the Need for Improved Separation and Isolation of Information; http://www.securityweek.com/wikileaks-shows-need-improved-separation-and-isolation-information; 03 November 2010

[18] Schotanus, H.A. et al.; Releasemechanismen voor gekoppelde netwerken; Report number TNO ICT 34369; 2007.

[19] Schotanus, H.A. et al; Decomposition of the Security Requirements for Connected Information Domains; Military Communications and Information Systems Conference; October 2011