Chapter 4 – COMMON CRITERIA AND RISK ANALYSIS

In this section, emphasis has been put on the Common Criteria (CC) as being a potential element of a risk management framework, where risk analysis are performed though Threat and Risk Assessments (TRAs).

4.1 COMMON METHOD USING COMMON CRITERIA

The Common Criteria (CC) is an internationally recognized approach to security evaluation. It provides a set of criteria, which can be used to set requirements and to guide the development of IT security features within a specific product. These requirements serve as a guide for the development, procurement and evaluation of IT security features and products. Using a set of defined "assurance levels", an accepted engineering standard can be applied to products under evaluation. The assurance levels are a graduated scale of documentation, and development and testing processes that appropriately grade the product's security functions.

In contrast, the TRA is a formalized process used to determine the risks to Information Technology (IT) assets and provide recommendations to mitigate the risks to acceptable levels. A TRA generally has a broader scope of investigation and will include elements that do not lend themselves easily to modelling or quantitative analysis. This characteristic of a TRA is generally seen as a significant rationale for the failure of automated commercial risk management tools. A successful TRA tool must be able to leverage current knowledge bases using a well defined inference engine to correctly reason with the fine points of contemporary IT network architectures and their security technology. Such a tool has not been refined to the point where it can replace the human involvement in the Risk Management process.

In a previous study[1], the preliminary findings show that the CC has the potential to relate to the following TRA aspects:

1) Structured terminology of controls;
2) Qualitative description of safeguards;
3) System architecture model;
4) Applicable threat model, including threat agent attributes (motivation, capability, opportunity, etc…) and threat scenarios;
5) Taxonomy of relevant vulnerabilities;
6) Classification scheme / sensitivity analysis of information assets;
7) Impact analysis of information assets, with respect to confidentiality, integrity and availability scenarios, and possibly mode of access;
8) Risk derivation model, the functional relation between risk and any of the above parameters;
9) Risk mitigation model linking safeguards and controls to threat scenarios; and
10) Risk acceptance of system operations is assessed based on CC evaluation results of security components of a system.

This concept needed some additional research to position the CC for use in a common framework for the NATO requirement. The CC is definitely a useful tool and can influence the risk management practice.

[1] Common Methods For Security Risk Analysis, prepared for DRDC by Cinnabar Networks Inc., 22 December 2004.

4.1.1 Similarities and Differences Between CC and TRA

The TRA methodologies and tools have the goal of estimating the degree to which the controls or safeguards of an organization / system are sufficient to prevent harm or damage from threats directed at the information assets. On the other hand, the CC was created for the purpose of establishing the functional and assurance characteristics of IT products that are systems or parts of systems. Links between CC assurance and TRA risks are indirect and sometimes fuzzy. While the purposes differ, these activities are relevant to one another insomuch as:

1) The scope of CC evaluation is generally a sub-component of a TRA-scope: Inputs and outputs of a CC evaluation of a component of a system can be relevant to a TRA of the greater system;
2) CC evaluations are generally deeper than TRAs; and
3) Common activities are involved in both CC evaluation and TRA.

The CC evaluation is generally a much more detailed analysis, focuses on a sub-set of threats, entails extensive documentation, is a lengthy process and leads to significant costs. On the other hand, the TRA has usually a broader scope, may require extensive analysis but can be completed in a relatively short period of time at a reasonably low price. In addition, the generally smaller scope of the CC evaluation makes it impractical to undertake full CC evaluations of virtually any of the components of a system in support of a TRA. It is therefore unlikely that the two activities would ever be combined in a joint objective but rather one be supported by the other.

4.1.2 Using CC with TRA

The focus of this section is on the application of the CC to TRA. The main areas in which this goal is achievable are as follows:

1) Threat mapping, where TRA threats are mapped to CC threats (note that there is no set of standard CC threats defined);
2) Safeguard mapping, where the TRA safeguard functionality are mapped to a set of CC SFRs;
3) EAL mapping of the TRA asset valuation and threat level of a specific threat scenario to an EAL (more specifically to an AVA component) required for safeguards mitigating the risk associated with the scenario;
4) Use of security assurance requirements (SARs) deliverables of a CC evaluation of a target subsystem of interest as inputs to the TRA; and/or
5) Use of Evaluation Technical Reports (ETRs). Section: Testing and Vulnerability Assessment.

The first three areas were introduced in a Canadian study published by the Communications Security Establishment (CSE) in 2002[2]. They were introduced in the context of the definition of a multiple-entry point modification of a standardized TRA process (ITSG-04). The entry points are threat mapping and EAL mapping processes that link to the TRA, as well as a post-TRA entry point that performs safeguard mapping. This leading edge approach provides the missing assurance component to the TRA methodology. While it is beyond the scope of this document to describe more fully the details of this schema, it is representative of what can be gained from mapping CC concepts to the TRA process.

The last two areas above are much more direct input of CC-evaluation documentation into a system TRA. Risk-relevant information exists in specific CC SAR document sources and ETRs of components that have been successfully evaluated. However, these documents may be available only through direct negotiation

[2] CSE – Threat and Risk Assessment Controls and Safeguards in Relation to The Common Criteria Report and Recommendations, Version 1.1, 27 March 2002, Conducted by: Cinnabar Networks Inc.

with the developer/vendor, presumably through nondisclosure agreement. Only in the case of the Security Target (ST) document, there is a general requirement for public disclosure, although this right can be withheld under appropriate conditions under the CC Mutual Recognition Agreement (MRA). A fundamental assumption is therefore that the documentation be available to the TRA team. The brief list of CC SAR documents that can serve as useful input to a TRA is shown in the table below.

Table 4-1: CC V 2.1 Documents Useful for TRA

CC Class	CC Family	Description	Relevance as input to TRA
Vulnerability Assessment	Vulnerability Analysis AVA_VLA	Analyses TOE vulnerability to threats based on level of assurance	Primary: Highly relevant to TRA Vulnerability Analysis
	Misuse Analysis AVA_MSU	Analyzes misuse by a user prevented by the TOE, based on level of assurance	Primary: Highly relevant to TRA Vulnerability Analysis
	Covert Channel Analysis AVA_CCA	Performs covert channel analysis based on assurance level	Primary: Highly relevant to TRA Vulnerability Analysis
Security Target	Target of Evaluation (TOE) Description ASE_DES	Basic description of TOE	Secondary: Potentially useful information for system description and architectural analysis in the TRA
	Security Environment ASE_ENV	Threat analysis, policy analysis, environmental assumptions	Primary: Highly relevant to TRA threat scenarios and threat analysis
	IT Security Requirements ASE_REQ	The evaluated TOE SFRs	Primary: Highly relevant to TRA risk analysis
Testing	Coverage Analysis ATE_COV	Demonstrates the coverage of SFRs by operational testing	Secondary: Potentially relevant as an evaluation of safeguard effectiveness
	Depth Analysis ATE_DPT	Demonstrates the depth of test cases by operational testing	Secondary: Potentially relevant as an evaluation of safeguard effectiveness
Other	Audit, Cryptographic Functionality, Design Documentation, etc. All SFRs	Describes TOE security functionality	Marginal: High level and detailed TOE information to clarify TRA

The AVA class deliverables provides indirect information on the vulnerabilities, misuses, and channels countered by the TOE and its environment. The TOE and its environment must mitigate all vulnerabilities, misuses and channels in the AVA analysis effectively. What does it mean? and What is the usefulness of the AVA class for a TRA vulnerability analysis?

The meaning depends of what is performed in the AVA analysis and what is effective. First, the AVA analysis is dependent of the SFRs and SARs. Some vulnerability types may be discarded by the analysis because they do not depend on the TOE SFRs and SARs. For instance, if no covert channels analysis is required, no vulnerabilities related to covert channels are analysed; if no protection of the TSF data is required, no vulnerabilities related to TSF data are analyzed; and if no low level design information is available, no analysis is performed on the vulnerabilities introduced in the refinement process from the TOE high level design to the TOE implementation. Research needs to be pursued to get a clear understanding of these dependencies.

Second, the AVA analysis is also dependent of the TOE environment. So, it may be difficult to assert if a countered vulnerability is countered by the TOE, the TOE environment, or both. However, the ASE_ENV may help for this by listing environmental assumptions.

Third, the notion of effectiveness can be better expressed in the context of the new release of the CC (Version 3,0), where the AVA class has been modified. Effective means that the TOE and its environment mitigate all discovered threat scenarios in which threat agents have attack potential lower than or equal to a specific value, which value is defined in AVA_VAN. The set of discovered threat scenarios is a representative set of threat scenarios that involves vulnerabilities discovered through the AVA analysis.

From this it seems the AVA analysis would be useful for a TRA vulnerability analysis of a CC product in its intended environment if the vulnerabilities of interest are those exploitable by threat agents of attack potential lower than or equal to the value specified in AVA_VAN.

4.1.3 Protection Profile and TRA

A protection profile (PP) is primarily an optional vehicle for user groups to publish evaluated security requirements for a specified class of products. PPs must be evaluated in the same manner that products are evaluated. Full compliance to the PP is mandatory if the PP compliance claim is made in the ST. There is no notion of partial compliance by the product.

While there is no formal connection between PP content and TRA processes, they do resemble the ST format closely, and many of the ST mappings to TRA described above would apply equally to PPs. Generally, however, the product ST duplicates this information and further instantiates those product-specific requirements in the PP that are unspecified (i.e., left as refinements, assignments, etc.). Therefore, the PP stands as a relatively generic requirement. As a catalogue of threats, the PP can play a significant role in threat analysis, specifying what type of protection is expected by informed users from a class of products. It is possible that no products on the market actually satisfy an evaluated PP. This may suggest caution on the use of PPs in a TRA safeguard selection.

4.2 CRITICAL ASPECTS OF COMMON CRITERIA EVALUATION

The CC in general is based on the notion of evaluating a given product, be it a network, operating system, database, etc., irrespective of the environment. This focus is further enforced in the notion of the boundary of evaluation of the Target of Evaluation (TOE). While much of this supports the practical process of evaluating a commercial product, it is not particularly well suited to the needs of risk analysis. The following critical observations further expand on this theme:

1) Although the Common Criteria does expect a vendor to indicate the type of environment the product should be used in, it discusses environments in a broad manner and leaves the level of detail to the developer discretion.

2) Many products are evaluated to co-exist within specific environment. Deploying them in other environments results in the invalidation of the evaluation. This characteristic places the onus on the part of the customer to check statements in the Security Target (ST) regarding Environmental Assumptions and Environmental Security Policy against actual deployment in order to obtain a sense of fitness for use. Such decisions are beyond the control or scope of the CC evaluation.

3) Although the CC is a formalized methodology, it may be inconsistently applied across the CC community. A product evaluated under two separate laboratories would be treated to differing levels of examination. The Criteria does not enforce a methodology upon the laboratories and therefore there is no current way to ensure evaluations are truly equitable. Nevertheless, the Common Evaluation Methodology (CEM) CC document provides a framework for the processes to be performed by these methodologies. Reciprocity agreements, such as the Mutual Recognition Agreement (MRA) accept evaluations from other nations but are not an indication of agreement to the methods used.

4) The CC focuses on how well (i.e., trustworthy) the functionality is implemented within a given product. It does not provide any direct answers as to the efficacy of said product within a given environment. As noted above, CC vulnerability analysis must demonstrate complete coverage by the product and its environment, without residual risk, but only for a set of vulnerabilities dependent of the SFRs and SARS and of the attack potential. Nothing is said with respect to vulnerabilities not related to SFRs and SARs. This approach is at variance with the more general notion of vulnerability used in TRA practice, and requires further analysis on the part of the TRA team to find potentially relevant vulnerabilities not covered in AVA_VLA (or AVA_VAN) documentation.

5) An evaluation is based on the product itself. Little consideration is given as to the type of information processed by the product. This lack of information identification is a major point separating CC evaluation methodology from that of TRA. There is no explicit classification of data process or statement of sensitivity in the CC process. In the end, if the product obtains a successful evaluation, it is solely on the basis of whether or not it met the minimum requirements for functionality and trust.

6) There is no assurance provided that two evaluated products can be comparable. Differing methods of evaluation can and do result in differences in whether or not a given product met a given requirement. Although reciprocity has been gained on a number of fronts, for some applications the US, for example, will not accept non-North American evaluations.

7) While the original goal of harmonization of security product evaluation has been actively pursued in the development of the CC, constant forces of a protectionist or exclusionary tendency exist among the signatory nations that favour evaluations of one or more nations over others. These tendencies are not always identified as such, nor is there an effective global mechanism to counteract them. Another continuing barrier to increased harmonization is the lack of progress in achieving an extension of the MRA to beyond-EAL4 evaluations. Recent developments in the US and elsewhere indicate that this is an unlikely expectation in the near future.

8) The final issue with CC evaluations is the time required for completion of the evaluation. Even for rudimentary evaluations the time taken to evaluate a product is excessive. An EAL-3 or 4 product may well have been under evaluation for two or more years. Hence, the product will be two or more revisions away from the currently commercial available version. In the commercial world the unforeseen length of the evaluation is a direct cost overrun. This timeliness issue continues to plague evaluations and has led to the cancellation of otherwise viable CC evaluation projects.

A major consideration in the decision to pursue CC evaluation is the total cost of the evaluation, both in monetary terms and in terms of developer time and effort. The direct cost of CC evaluation is a significant factor in the decision by organizations to pursue either a PP or product evaluation. The significant costs

are mostly due to the large volume of documentation that must be produced by the developer, evaluated by the CC laboratory and further corroborated by extensive testing. These costs are expected to be proportional to the assurance level being claimed by the developer, because more evaluation documentation is required as the EAL increases. Further costs are related to higher assurance, such as increased complexity of site visitation evaluation tasks, testing requirements, and the evaluation of formal modeling, specification and design that is mandated at EAL5 – EAL7. In addition, certain technologies may require more extensive evaluation and testing costs than others, e.g., biometrics, cryptography, complex protection systems.

Cost is not a controlled or fixed item, and is largely the result of open market competition on a country-by-country and laboratory-by-laboratory basis. Some countries provide financial support to laboratories under their CC Schema. In North America, an EAL2 evaluation may cost several hundred thousand dollars. Funding of such projects must be justified either by developer-perceived increased market share or level of support by funding agencies. The level of assurance required in many currently evaluated PPs is another cost factor. In many cases commercial developers face a significant technical challenge in justifying the development of products that satisfy particular PPs. In some cases, government-sanctioned PPs have attained the status of technical IT standards for entry into government and military markets. When based on a purely financial return on investment, CC evaluation in a competitive market remains a business risk issue.

While the above analysis provides some of the predominantly negative aspects of the CC in general, there are positive reasons to consider the use of CC concepts in a TRA context:

1) It is possible to consider a given existing product evaluation within a TRA. A product with an evaluation may be more trustworthy than one that has not been evaluated. Two similar products, for example two UNIX-based operating systems, can be fairly ranked on whether or not they have been evaluated. However, the usefulness of the lowest assurance levels has come into question. The fact that a product has undergone EAL 1 or 2 may not be of any greater intrinsic value than a product that has gone through a rigorous quality assurance process from a reputable firm. The jury remains out on that issue;
2) The most advantageous elements of the Common Criteria remain the taxonomy of functionality and assurance so that products can, more or less, be evaluated to similar security standards. The common grouping of functionality and the notion of interdependency between specific functionality components is something that would be beneficial to introduce into a formalized and automated TRA methodology;
3) A TRA can be augmented with a certain assurance level by using the CC threats and adding a non-CC level, by mapping the TRA safeguard functionality to CC SFRs (or SARs), or mapping the TRA asset valuation and threat level of a specific threat scenario to an evaluation assurance level (EAL); and
4) The use of security assurance requirements (SARs) deliverables of a CC evaluation of a target subsystem of interest as inputs to the TRA.

4.2.1 CC and TRA Summary

The use of the CC in TRA methodologies and tools is still at an early stage. This section has indicated some sources of information produced by a CC evaluation that would be relevant, to varying degrees, in a TRA context. The role of PPs has been discussed as a generic class specification in much the same way. The mapping of CC to IT security framework shows weakness in the common framework approach. Due to excessive expense and time required, it would not be recommended to pursue CC evaluation of a product simply as part of a TRA. The practical methods of applying CC technology to the TRA process are found to be:

1) The use of the common terminology, taxonomic and definitional structure of the CC to organize TRA methods and tools, and
2) The use of component CC evaluation deliverables as information sources in the analysis of threats, vulnerabilities, environmental assumptions and related security functional information.

The CC methodology could be leveraged for involvement in the TRA process by combining CC evaluated components into a higher order system, which encompasses (in part or in entirety) the scope of the TRA engagement. However, a method by which CC evaluated products can be combined in such a way as to retain the usefulness of their assurance level ratings is needed. Two such methods were proposed, but the need for extra research in this area is clearly indicated. Specifically, methods are needed which will allow such an analysis to be:

1) Effective in the face of emergent behaviour of such a composed system; and
2) Efficient in the reduction of manual effort so as to justify leveraging the work performed in the CC component evaluation.

In the absence of methods by which these goals can be achieved, the use of composed system to assess the risk level posed by a system under evaluation cannot be recommended at this time. Annex A provides an overview of the use of the composed system model as a potential mechanism by which the Common Criteria approach can be leveraged to be of use in performing TRA.

Recommendation # 1 The common methodology to risk management should leverage the CC evaluation of components and possibly use its structure terminology but should not rely on the CC framework as the model for the NATO common TRA framework.

4.3 SECURITY REQUIREMENTS AND CC

Risk analysis methods and tools often include the identification of security requirements that can be proposed to counter the identified risk. It seems then useful to rely on the part 2 of the CC for this work. This means:

- Using at a minimum the security functional class names (security audit, user data protection, security management, …) to classify the security requirements.
- If possible, use the CC components when writing the security requirements. This could save time by relying on the CC efforts to produce relevant requirements. This is also a guarantee that requirements issued from different tools or methods will be comparable.

At the moment, there is no satisfactory tool to select automatically security requirements consistent with the estimated level of risk. This could be a research area given that there will always be a need for a human arbitration as many other factors have to be taken into account such as the cost, the environment, the operational constraints, …