A Quantitative Analysis of Common Criteria Certification Practice

Samuel Paul Kaluvuri[1,2,3], Michele Bezzi[1], and Yves Roudier[2]

[1] SAP Labs France
[2] Eurecom Institute
[3] Eindhoven University of Technology

Abstract. The Common Criteria (CC) certification framework defines a widely recognized, multi-domain certification scheme that aims to provide security assurances about IT products to consumers. However, the CC scheme does not prescribe a monitoring scheme for the CC practice, raising concerns about the quality of the security assurance provided by the certification and questions on its usefulness. In this paper, we present a critical analysis of the CC practice that concretely exposes the limitations of current approaches and provide directions to improve the practice.

1 Introduction

With increasing number of cyber attacks and security issues, governmental organizations and private companies are striving to get security assurance for Information Technology (IT) products. In many cases, these organizations may not have the required knowledge or resources to assess whether a certain product has the appropriate security features nor can they rely only on the statements of vendors. This is due to the trust deficit that exists between consumers and product vendors. One way to bridge this trust deficit is through security certification of software. Security certification provide a practical solution to address the lack of security assurance when assessing and purchasing IT solutions. The Certification Authorities (*CA*) perform rigorous security assessments that a particular software system has certain security features, conforms to specified requirements, and behaves as expected [7]. A customer buying a certified product can rely on the "stamp of approval" by the *CA*. Clearly, the value of a certification depends on the reputation of the certification authority issuing it, as well as the quality of assessment performed. Ideally, software purchasers can then choose among different certified products which address common security requirements.

Common Criteria for Information Technology Security Evaluation (CC)(ISO / IEC 15408) [2] is the most popular security certification standard. It is a globally recognized set of guidelines that provides a common framework for specification and evaluation of security features and capabilities of IT products. At the heart of the CC scheme lies a "common" set of security functional and security assurance requirements. These common requirements enable potential consumers to compare and contrast the certified products based on their security functional and assurance requirements and to determine whether a product fits into their needs. The CC scheme allows the evaluation of the products at varying levels of evaluation rigor, called *Evaluation Assurance Levels* (EAL), in a range of 1 to 7 (7 being the highest assurance level).

Despite the wide use and economic success of Common Criteria scheme [18, 10](mostly driven by government regulation and government purchase) its current practice has been receiving significant criticisms.

1. *Comparability.* One of the main objectives of CC is to allow consumers to compare certified products on the market in an objective way from a security point of view. However, certification documents are filled with legalese and technical jargon. Hence, comparison is not straightforward nor easy.
2. *"Point in time" certification.* CC certifies a particular version of the product in certain configurations. Any changes to the configuration or any updates to the product that affect the *Target of Evaluation (TOE)*, which is the part of the product that is evaluated, invalidate the certification. This is not a desirable situation, given that products evolve and are updated at a frantic pace and the certification must not be "frozen" to a specific version of the product.
3. *Long and expensive.* CC evaluation life cycle is lengthy and expensive [15, 20, 19]. In fact, due to the complexity of the process and the high cost, vendors have to spend a large effort on preparation for the evaluation, which adds to the cost and time of the evaluation itself. High assurance level (as EAL4) certification can take $1 - 2$ years, and, often, by the time the process is completed a new version of product is already delivered.
4. *Concerns for Mutual Recognition.* Though the CC scheme is a widely recognized international standard, there are several concerns regarding the consistency of the assessments by the evaluating laboratories located in different countries, since the *Common Criteria Recognition Arrangement* (CCRA) does not prescribe any monitoring and auditing capability. In addition, the relevance of CC certification for governmental institutions, specific national interests can impact the impartiality of the assessment [11, 5].

Although, most of these shortcomings of the application of the CC scheme have, to the authors knowledge there is no quantitative study of the CC certificates, which provides the evidence that these criticisms are applicable to a broad category of CC certified products or are limited to just a few cases.

The major contribution of this paper is filling this gap, providing an exhaustive analysis of CC certificates. Systematically analyzing the certificates (in Section 5), we can quantitatively assess the relevance of the points 1 and 2 above. We show how these issues are well grounded and affect a large part of existing certificates. We will also present possible directions (in Section 6) to enhance the current situation, considering current evolution of CC scheme and practice under discussion, and recent research results addressing security certification for web services. The points 3 and 4 are out of scope for the paper, because: an analysis on cost and duration of CC certifications has been discussed in [15, 19] (addressing Point 3), and the mutual recognition issue (point 4) cannot be analyzed looking at certificates.

2 Common Criteria Certification Scheme

The *CC* scheme allows product vendors to describe the Security Functional Requirements (SFRs) for the product and to prove that the set of SFRs are able to counter the threats identified for a Target of Evaluation (TOE), which identifies the specific aspects of the product that will be evaluated. In addition, the *CC* scheme allows product vendors to choose particular configurations of the product that will be evaluated and these "golden" configurations are also part of the TOE. This information is captured in a document called "Security Target" (CC-ST) which can be seen as the *descriptive* part of the *CC* certification [3]. The product vendor then defines the set of Security Assurance Requirements (SARs) that specify actions to be performed by the evaluating laboratories. Based on the *SARs* selected for the product the certification authorities determine the *Evaluation Assurance Level*.

The drawback of this approach is that the EAL can only specify how thoroughly the evaluation has been performed, but it does not answer the question of "Is the software secure?". The answer to this question can be provided by the SFRs that are implemented in the product. The *CC* scheme classifies the SFRs into 11 high level classes as shown here:

SFR Classes	
Security Audit (FAU)	Communication (FCO)
Cryptographic Support (FCS)	User Data Protection (FDP)
Identification and Authentication (FIA)	Protection of TOE Security Functionality (FPT)
Privacy (FPR)	Security Management (FMT)
Resource Utilization (FRU)	TOE Access (FTA)
Trusted Path/Channels (FTP)	

An example of an SFR in the *Security Audit* class and an SAR in the *Security Vulnerability* class can be seen below:

"***SFR: FAU_GEN.2.1*** *For audit events resulting from actions of identified users, the TSF shall be able to associate each auditable event with the identity of the user that caused the event.*"

"***SAR: AVA_VAN.1.3E*** *The evaluator shall conduct penetration testing, based on the identified potential vulnerabilities, to determine that the TOE is resistant to attacks performed by an attacker possessing Basic attack potential.*"

The *CC* scheme is generic and does not impose specific requirements for different types of IT products. Hence product vendors can implement certain specific security functionalities (SFRs) and get specific parts of their system evaluated in a certain way (SARs) and consequently certified, which may not address the requirements of consumers. To address this issue, *CC* allows consumers to use Protection Profiles (CC-PP) that contain a combination of SFRs and SARs for a particular type of application, such as Operating System or Databases. When products conform to a specific protection profile, it is easier for the consumer to select and compare the best fit for their needs. But conformance to CC-PP by products is not mandatory, and there is a criticism that product vendors exploit this flexibility of the *CC* scheme, and choose not to conform to any protection profiles that could be applied for their products [17, 5].

3 Analysis Objectives

The fundamental aim of our analysis is to verify whether the *CC practice* fulfills the intentions of the *CC scheme*. The main goals of the *CC scheme* are: *a)* Enabling the comparison of the security characteristics among (certified) "similar" products; *b)* Providing meaningful and useful security assurance to the consumer.

We defined the possible checks to assess whether these objectives are reached by the current *CC practice* (Checks are indicated in **bold** in the following). For comparing products of the same category, for example databases, from the security assurance point of view, we need to evaluate them against a common set of security requirements (SFRs). To support that, CC proposed the Protection Profiles, that allow for describing a predefined set of requirement for a class of products. Accordingly, to assess if this objective is reached in the actual practice, we need to check:

- **C1**: Are Protection Profiles available for the different categories ? **(Protection Profile Availability in each category)**
- **C2**: Are Protection Profiles actually used? **(Protection Profile conformance by products per category)**
- **C3**: Do similar products (same category) address the same set of SFRs? **(Differences in the number of SFRs for a given category)**
- **C4**: Does the usage of a specific Protection Profile results in a actual common set of SFRs? **(Differences in the number of SFRs for a given class in PP conforming products)**

To provide meaningful assurance to the consumer, the certification issued should be valid along the product lifecyle. Considering the need to perform changes in the software (e.g., security patches to address new vulnerabilities) or in the landscape (e.g., configuration), CC scheme proposes a *product certification maintenance* under the Common Criteria Maintenance Agreement (CCMA). Under this scheme, a specific version of the product can be initially certified and any changes made to it in future will be localized to the aspects that have been changed instead of the whole product being reevaluated. So, our next objectives are to evaluate:

- **C5**: Is the CCMA actually used in practice? **(How many products are maintained under the CCMA?)**
- **C6**: Are CCMA certified products secure? **(How many CCMA certified products have disclosed vulnerabilities?)**

4 Research Methodology

We use data from two main sources: the Common Criteria website [6], that provides details about certified products; and the National Vulnerability Database (*NVD*) [1], that contains the list of disclosed vulnerabilities in products. In particular we considered the following data sources: *a)* Certified Products List [6] *b)* Protection Profile List [6]; *c)* Security Targets of certified products [6]; *d)* CC Part 2: Security Functional Requirements Document [2]; e) NVD database [1].

The data collected from these sources requires additional processing, in order to perform advanced reasoning. The steps we performed are presented here in a concise manner: *1)* The *certified products* and the *protection profile* CSV files were converted to SQL tables and stored in a database; *2)* The Security Target files (in PDF format) are downloaded for each certified product (URLs are contained in the CSV file of certified products) *3)* We stored the standardized SFRs contained in *CC: Part* 2 document in the database; *4)* Search the CC-STs for SFRs and stored the results; *5)* Cross-reference certified products against the *NVD* for disclosed vulnerabilities. Except the steps *3* and *5*, the rest of the analysis is automated.

The certified product list contains data of products that fall under three categories: a) Certified products; b) Certified products under *maintenance agreement*; c) Archived certified products. We consider only products with valid certificates (1971 certified products) and ignored the archived certificates for our analysis. Due to technical reasons, such as malformed URL or a digitally signed PDF document that could not be parsed into text, we could not process 95 certificates. Hence the data set that we considered in our analysis was 1532 security targets of certified products and 344 security targets of products under the maintenance agreement.

5 Analysis Results

Due to space constraints we present the most important results from our analysis. The results presented here are focused on products certified at EAL4+, since most products are certified at level (close to 40 % of the certified products).

5.1 Comparability of Certified Products

Products that conform to protection profiles are expected to have homogeneity both in terms of functionality and security features. Hence, we examined the availability of protection profiles compared with the number of certified products across various product categories and the results are shown in Figure 1. It can be noted that the availability of protection profiles is rather low across all categories of products except the *ICs and Smart Card* category.

Figure 2 presents the percentage of certified products that conform to at least one protection profile across various categories. The average PP conformance rate among certified products of our data set is 14 %, with standard deviation

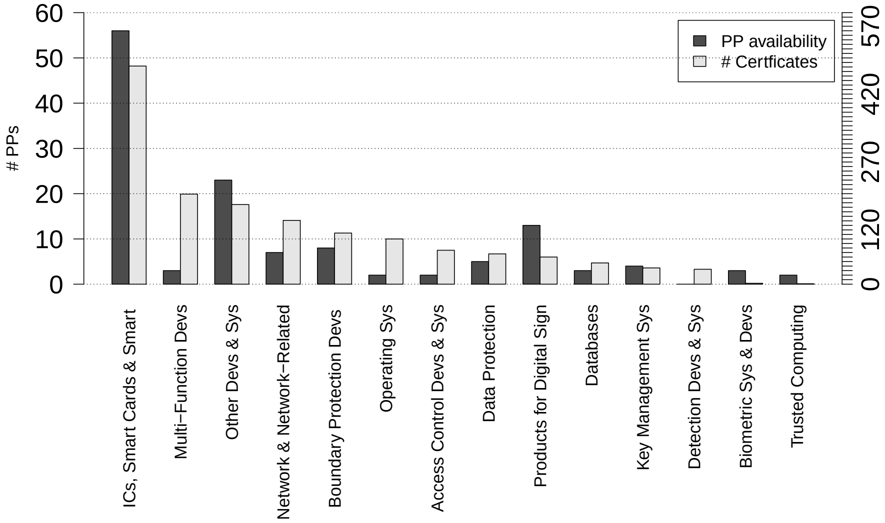

Fig. 1. Protection Profile Availability and Certified Products across categories

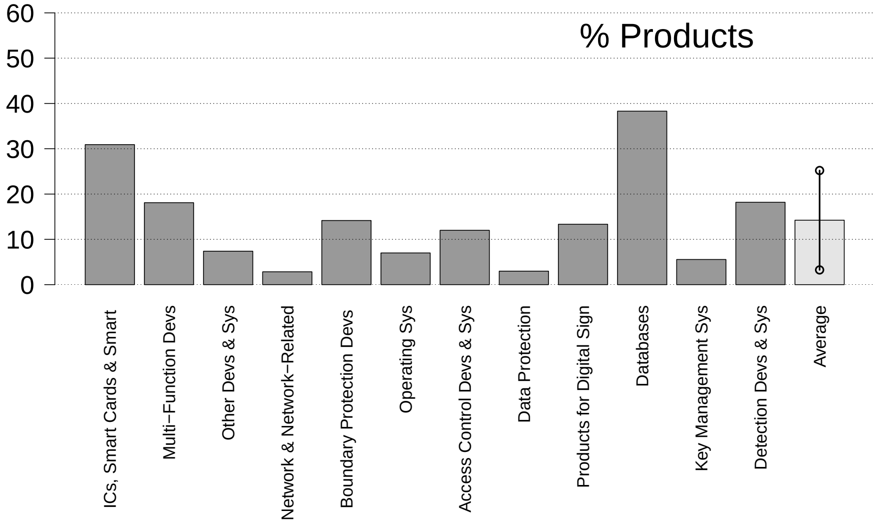

Fig. 2. Protection Profile Conformance among certified products (Right-most bar shows the average conformance to PP for our dataset, 14%, and corresponding standard deviation, 11%)Note that categories with less than 10 products are not shown

around 11 % (see Fig 2, rightmost column). This indicates that a relatively low number of certified product use CC-PPs with relevant differences among categories. Indeed, a closer inspection reveals that the products broadly related to *hardware or firmware* show higher conformance than products that fall under the software-only category. This low conformance could also be due to vendors finding it difficult to build products that conform to a particular CC-PP, while the products themselves are targetted for the general commercial market. Hence, to conform to a particular CC-PP, which is produced by specific consumer or a consumer group, does not provide any competitive advantage in the general market.

On the other hand, the low CC-PP conformance makes it difficult to compare and contrast the security requirements addressed by the certified products. In fact, the non-conformance to a CC-PP allows vendors to customize the scope of certification to features that are very different from other certified products. As an example, a product in a certain category could make claims that it addresses more SFRs related to data protection, while another certified product in the same category may have claims addressing SFRs related with access control. Furthermore, each certified product identifies different threats and makes various assumptions. Hence, comparison of certified products in such cases can become rather labour intensive and a very time consuming process.

Next, we compare products based on the number of *SFRs* that are addressed by each product in a certain category to understand the differences in certified products based on their security functionalities. Figure 3 and Figure 4 show the SFRs addressed in products for *Database* and *Operating System* categories certified at *EAL4* (and *EAL4+*) and conform to CC version *3.1*. Each shade of the bar in the figures 3 and 4 represents products that conform to a specific protection profile and the *white bars* represent products that *do not* conform to any protection profiles.

It can be observed from Figure 3 and 4 that even among products that claim conformance to a protection profile, there is a considerable difference between the SFRs addressed by the products. And products that tend to show little or no difference are either different versions of the same product or products from the same vendor. Among the products that do not conform to any protection profile there is a huge difference in the number of SFRs addressed.

5.2 Point in time Certification

The CC scheme certifies products at a point in time, that is, certification applies to a particular version of the product and in a certain set of configurations. But products do need to evolve - either to provide new functionalities or to fix problems or both. And in such cases, the CC certification does not apply to the new version and the whole product has to undergo the certification all over again which is once again a very time consuming and expensive process, especially when the changes made to the product are very minor. In order to avoid such situations, the CC scheme allows products to be under the CC Maintenance Agreement (CCMA) where only the changes made to the product are evaluated

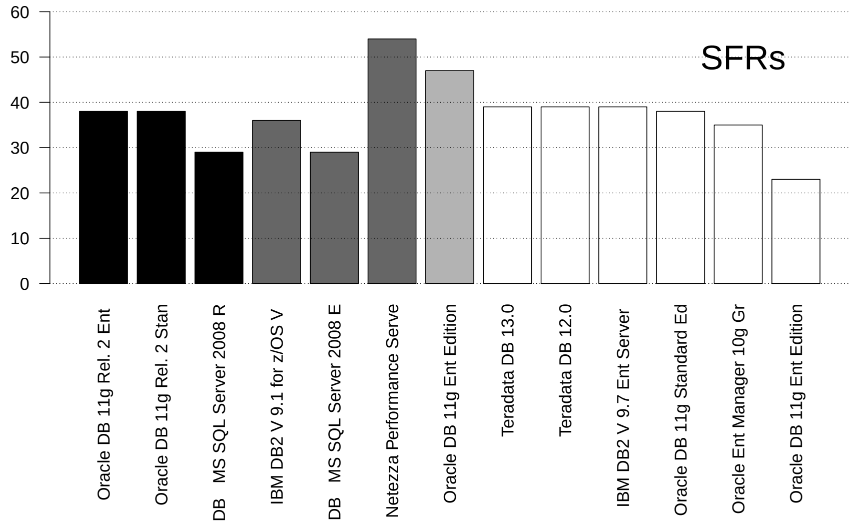

Fig. 3. SFR variation in Database Category for EAL4+ (Each shade of the bar represents products that claim conformance to a particular CC-PP, white bar implies non conformance to any CC-PP)

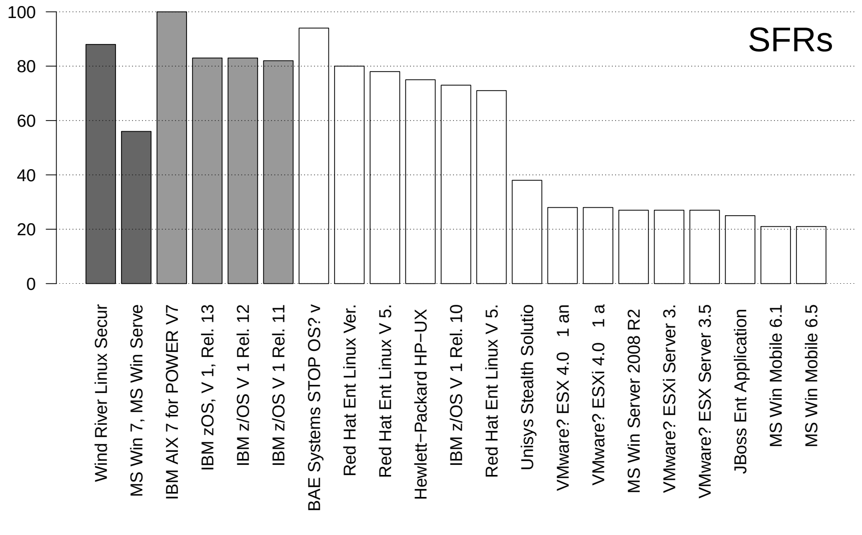

Fig. 4. SFR variation in OS category for EAL4+ (Each shade of the bar represents products that claim conformance to a particular CC-PP, white bar implies non conformance to any CC-PP)

and certified. This aspect of the CC scheme would allow the products to be certified over a *period of time* instead of a *point in time*.

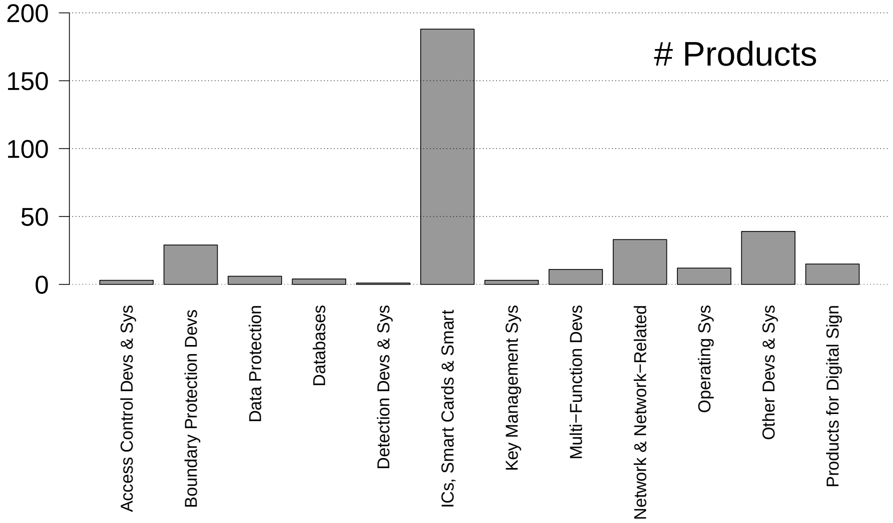

Fig. 5. Products under CCMA

We verified the Figure 5 shows the certified products that are under the maintenance agreement across the various product categories. It can be observed that the number of products under the CCMA scheme is high among *ICs and Smart Card* category when compared to the other categories. And indeed the total percentage of products that are under the maintenance agreement is just 22 % of all the certified products. And in fact, excluding the *ICs and Smart Cards* category, the percentage of products that are under the CCMA scheme comes down to approximately 15 %.

Such low numbers of products under maintenance raise an important question on the product's lifecycle, especially when vulnerabilities are found in the product which need to be fixed - can a product vendor issue a fix and technically *loose* the certification or keep selling the vulnerable version of the product to claim the certification?

In order to better understand this question, we have used the National Vulnerability Database (NVD), to cross-reference the certified products with products that have known vulnerabilities. Since we could not automate this step, we limited our analysis to the Database and Operating System categories certified at assurance level EAL4+. In the Operating System category, we found 22 % of the products under the maintenance agreement have disclosed vulnerabilities. And in the database category we found only 25 % products under the maintenance agreements are shown to have a known vulnerability. To contrast this, we cross reference products (in the database category at EAL4+) that are not under the maintenance agreement and 85 % of the products have shown to have a known vulnerability.

Though we do not claim that the vulnerability is in the certified "golden" configuration, these figures show that in practical usage of the products the issue of addressing new vulnerabilities must be discussed. And clearly, a point in time certification does not cope well with the dynamic landscape of a product's lifecycle.

6 Discussion and Conclusions

6.1 Findings about the CC certification practice

Comparability of certificates. Our results illustrate some reasons behind the lack of comparability of certificates. In particular, the security assurances sought in the certificates produced for the same class of products often exhibit large differences. When products conform to CC-PP, the variation between SFRs addressed by the products is not so large. However, the CC-PP conformance rate is rather low, particularly in software products. We believe that making products conform to *Standard Protection Profiles* in each product class could provide better comparability between certified products.

On a more fundamental level, we found out that without tool support it is not a trivial task to perform comparison between products based on their SFRs. In this regard, the CC-STs should be represented in a machine processable manner that facilitates automated reasoning to be performed on them.

One point in time certification. The low numbers of products under maintenance raise an important question on the product's lifecycle, especially when vulnerabilities are found in the product which need to be fixed - can a product vendor issue a fix and technically *loose* the certification or keep selling the vulnerable version of the product to claim the certification? It is rather obvious, that the product vendor will choose to fix issues and risk losing the certification.

Our results show that, despite the finding of new vulnerabilities, which are sometimes unknown at the time of initial certification, and the provisions made by the Common Criteria scheme to support incremental certification (CCMA), the certified products are overwhelmingly certified once and for all. While this is perfectly valid in itself, it shows that two certificates should be compared with respect to their time of issuance but also with the information from publicly available vulnerability databases (such as *NVD*).

6.2 Outlook

Contributions have been proposed in order to ease the comparability of the certificates produced by the Common Criteria evaluation process. Those approaches rely either on an extension of the CC certification scheme, or on tools to support a more homogeneous generation of certificates.

Common Criteria Framework Extensions. Countries that are members of the Common Criteria Recognition Agreement (CCRA) have recently agreed to develop internationally accepted Protection Profiles (known as Collaborative Protection Profiles - CPPs) for each class of products. Each product has to conform to the CPP that is applicable in its class, thus facilitating an easier comparison among certified products.

Computer Aided Certification. These approaches most notably aim at providing some guidance for evaluators in the production of certificates, and in making sure that their description is consistent. These approaches might therefore be extended in order to provide the necessary support to implement the recommendations we suggest above, in particular that of rendering certificates machine readable, with comparable SFRs and TOEs.

Certification Toolboxes have for instance long been designed. The CC Design Toolbox created by Tore Nygaard [14] aims at supporting the production of CC certificates. The toolbox aims at supporting the uniform definition of protection profiles, and at certifying those profiles themselves. Other proposals have extended such toolboxes with security ontologies. Ekelhart et al. [8] and Chang et al. [4] proposed to use an ontology as the core tool to manipulate Common Criteria certificates. The main improvement of this approach over plain toolboxes is that the definition of an ontology makes the relationships between the different concepts apparent. For instance, those relationships materialize consistency checks between the different sections of a certificate or of a protection profile.

Machine Processable Security Certificates. In [12] the authors propose a language that allows machine processable representation of the security certificates. Though the work focuses on the service environments, the certificate language proposed is capable of representing Common Criteria security targets in a machine processable manner. In [13] the authors present a machine processable language to represent protection profiles and a tool that automatically verifies the conformance of a certificate with its protection profiles.

6.3 Conclusions

We have presented the results from a thorough analysis of the certificates of CC certified products to concretely understand the drawbacks of the CC practice. We presented evidence on the variation of SFRs in products and that EAL should not be considered as the only metric to measure the security of a product. The low rate of conformance to CC-PP makes the comparison even more complex. In addition, we also discovered that very few products are under the maintenance agreement. We believe that the conformance to a standard (or basic) CC-PP for each product category could help in allowing easier comparison between products. Finally, we believe that machine processable CC-ST and CC-PP could help ease the burden to compare products on the consumer.

References

1. NIST National Vulnerability Database. National vulnerability database, 2012.

2. T. C. C. R. Agreement. Common criteria for information technology security evaluation part 1 : Introduction and general model july 2009 revision 3 final foreword. *NIST*, 49(July):93, 2009.
3. B. Beckert, D. Bruns, and S. Grebing. Mind the gap: Formal verification and the common criteria (discussion paper).
4. S.-C. Chang and C.-F. Fan. Construction of an ontology-based common criteria review tool. In *Computer Symposium (ICS), 2010 International*, pages 907–912, 2010.
5. Cisco and Intel. Common criteria embrace, reform, extend. discussion draft 1.0. 2011.
6. Common Criteria. Common criteria portal, 2012.
7. E. Damiani, C. A. Ardagna, and N. E. Ioini. *Open Source Systems Security Certification*. Springer, 1 edition, 2008.
8. A. Ekelhart, S. Fenz, G. Goluch, and E. R. Weippl. Ontological mapping of common criteria's security assurance requirements. In *SEC*, pages 85–95, 2007.
9. S. Fenz, G. Goluch, A. Ekelhart, B. Riedl, and E. R. Weippl. Information security fortification by ontological mapping of the iso/iec 27001 standard. In *PRDC*, pages 381–388, 2007.
10. D. S. Herrmann. *Using the Common Criteria for It Security Evaluation*. CRC Press, Inc., Boca Raton, FL, USA, 2002.
11. J. Kallberg. Common criteria meets realpolitik - trust, alliances, and potential betrayal. *Security Privacy, IEEE*, PP(99):1, 2012.
12. S. P. Kaluvuri, H. Koshutanski, F. Di Cerbo, and A. Mana. Security assurance of services through digital security certificates. In *Web Services (ICWS), 2013 IEEE 20th International Conference on*. IEEE, 2013.
13. S. P. Kaluvuri, H. Koshutanski, F. Di Cerbo, R. Menicocci, and A. Maña. A digital security certificate framework for services. *International Journal of Services Computing*, page 25.
14. T. B. Nygaard. Common criteria design toolbox. Master's thesis, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, 2007. Supervised by Professor Robin Sharp and Assoc. Professor Michael R. Hansen, IMM, DTU.
15. U. S. G. A. Office. Information assurance: National partnership offers benefits, but faces considerable challenges. Technical Report GAO 06-392, Report, March 2006.
16. V. Pretre, F. Bouquet, and C. Lang. Using common criteria to assess quality of web services. In *Software Testing, Verification and Validation Workshops, 2009. ICSTW '09. International Conference on*, pages 295–302, 2009.
17. J. Shapiro. Understanding the windows eal4 evaluation. *Computer*, 36(2):103–105, 2003.
18. R. E. Smith. Trends in security product evaluations. *Inf. Sys. Sec.*, 16(4):203–216, July 2007.
19. H. Yajima, M. Murata, N. Kai, and T. Yamasato. Consideration of present status and approach for the widespread of cc certification to a private field˜ cases in japan˜.
20. C. Zhou and S. Ramacciotti. Common criteria: Its limitations and advice on improvement. *Information Systems Security Association ISSA Journal*, pages 24–28, 2011.