

ANATOMY OF THE COMMON CRITERIA

Testimony before the House Government Reform Committee

**Subcommittee on Technology, Information Policy, Intergovernmental Relations and
the Census**

September 17, 2003

Submitted by J. David Thompson

**Director of the Security Evaluation Laboratory of CygnaCom Solutions, an Entrust
Company**

I. Brief description of Common Criteria

The motivation for a product testing capability is derived from the US military Certification and Accreditation (C&A) process for systems. Most systems include at least one computer, each employing an operating system that had to have its security functionality identified and assessed. Operating systems are complex and implement many key security functions, so considerable effort is required to do an appropriate security assessment of one. As computers became commodities, the notion of performing these difficult evaluations once and using the results in many C&As took hold.

The Orange book's set of five operating system criteria, the Rainbow Series of supporting methodology and interpretation documents written to supplement it, and the TPEP infrastructure, for the most part, accomplished this task. The Orange Book example was implemented in other countries, although in different and evolving ways. One of the evolutions was the European notion of a catalog of security assurance and functional requirements that could be used to specify security criteria for various types of IT components, not just operating systems.

In the early 1990's the Orange Book's narrow focus on operating systems became a problem, as was the expense to vendors of having to have their products evaluated to different security criteria in order to sell to the governments of different countries. A common criteria was sought that incorporated the best of the various existing programs

and that all of the major Western governments could endorse. We are still in the early stages of implementing the resulting Common Criteria, but all the originators are still enthusiastic participants and additional governments are signing on to recognize CC conformance in their procurements and to help produce product evaluations.

Industry also sees the value of a common security conformance process and is using the CC's processes and flexible criteria to fit its purposes. The Trusted Computing Group (TCG), for example, has adopted the Common Criteria for specifying conformance to the security components of its Trusted Processing Module (TPM). The TPM standard specifies a chip that can be installed on a PC motherboard to support secure e-commerce, in effect turning PCs into smart cards. The TCG is a consortium of most major PC hardware and software manufacturers and TPM conformant chips are destined to be built into nearly all PCs. Some are already available today. The TCG specifies the CC evaluation process as the way for vendors to demonstrate the conformance of their chips to the security parts of the TPM standard.

The Banking Industry Technology Symposium (BITS) is in the process of rewriting its security criteria to be Common Criteria conformant and replacing its ad hoc product testing facilities with testing in accredited CC testing labs. Their major criteria have been converted, and CC evaluations that include BITS conformance are underway.

II. Anatomy of the Common Criteria

The Common Criteria consists of two menus of security requirements -- one for security functionality (SFRs) and one for security assurance (SARs) -- and a process for using these components to evaluate products. The security assurance and security functional requirements are used to create documents that specify security criteria. The CC specifies two such documents: Security Targets and Protection Profiles.

A Protection Profile (PP) specifies security criteria for a class of products. A PP is intended to be written by a consumer or a group of consumers to specify the security requirements they want to see in a class of product they want to buy. The NSA has developed a large set of PPs for firewalls, operating systems, smart cards, and many other product types to specify security criteria for products for DoD use. BITS and the TCG are other examples of consumer groups producing PPs.

A Security Target (ST) specifies security criteria for a specific product. It also describes the security functions in the product and provides a rationale that the product's security functions meet the specified security functional requirements, among other things. It may also include an argument that its security criteria conform to one or more Protection Profiles. An ST evaluation confirms any PP conformance claims and the validity of correspondence arguments. A product evaluation confirms the underlying conformance of the product to the criteria in the ST.

The security assurance requirements identify a set of methodology that the CCTLs must execute satisfactorily in order to show that the product conforms to the security functionality in the security target. The security assurances used in an evaluation provide the consumer with a level of assurance that the product performs the security functions that the target says it will perform.

The CC defines sets of security assurances, called Evaluation Assurance Levels (EALs) that specify coherent groups of assurances roughly corresponding to those in the earlier Orange Book and ITSEC criteria. There are seven assurance levels defined by the CC, identified as EAL1 through EAL7. EAL1 specifies the least assurance and EAL7 the most. Other sets of assurances are viable, as well. The NSA has identified three sets of assurances associated with its Basic, Medium, and High Robustness levels. The authors of PPs and STs may, and often do, augment one of the EALs with additional SARs, as meets their needs. These are often specified as EAL3+ or EAL4+.

The most commonly used assurance levels are EAL2 and EAL4. EAL4 is often specified for products that are employed in the first line of defense, such as firewalls and operating systems, where untrusted users have full access to a feature-rich interface and the rewards of break-in are high. Lower assurances are appropriate for products with a lesser security role. Higher assurances are often required when information at different sensitivities must be separated or critical missions or assets are being protected.

The security assurances are grouped into the following classes:

- Configuration Management
- Delivery and Operation
- Development
- Guidance Documents
- Life Cycle Support
- Maintenance of Assurance
- Testing
- Vulnerability Assessment

The lower EALs do not include components from all of these classes, but it should be clear from this list that testing (trying to break into the product) is only one aspect of an evaluation, and only then in the context of an analysis of the design and implementation of the product. It is what we call white-box testing, requiring that the evaluator see inside the box and understand its functionality before running tests. Vendor cooperation is required. The alternative is black box testing, where nothing is known about the inside and the interfaces are tested against their specification. Black box testing is generally cheaper but provides significantly less assurance.

The result of a successful CC evaluation is a published Security Target that precisely documents the security functions that the product claims to meet and provides a precise expression of the assurance that has been applied to confirming the claims are true. The

ST can be used to determine the product's suitability to a particular security use. It can also be used to compare the security functionality of competing products in a way that vendor marketing information makes difficult, if not impossible.

It is theoretically and practically impossible to determine that a product of any complexity will be secure regardless of its configuration, or that security will mean the same thing in all the situations in which the product will be used. What CC testing does show is that the security functions the vendor claims the product to have work as described and that a coherent and mutually supportive set of security functions is available.

III. Misconceptions

Two common misconceptions are worth addressing. The first arises from the fact that many products that have been successfully evaluated are later found to have vulnerabilities. This leads to the conclusion that the evaluation process must be somehow flawed. In truth, the evaluation accurately demonstrated that the product can be used to implement secure systems in a specific configuration. The problem arises when vendors or users of the product do not use it in a secure way, choosing instead to misuse the security features to maximize utility or ease of use at the expense of security. The NSA, among others, publishes guidance for securely configuring critical widely used products on its website, at www.nsa.gov/snac.

Another misconception is that a product that has been evaluated should be free of bugs. The CC evaluation process is primarily focused on design and development process issues, not finding bugs. Higher CC assurances, such as those that comprise EAL7, do reduce the possibility of bugs significantly since they require that product design and development proceed in parallel with the evaluation and that important design documents be expressed mathematically and proven to correspond to each other. Further precise documentation and thorough correspondence between documentation layers is provided down to the code level, and anomalies of compiler and processor design are taken into consideration. But an EAL7 evaluation is only practical on small programs (less than 10,000 lines of code) when several million dollars are available. Developing perfect code, to perform security or other even moderately complex functions, is exceedingly difficult and out of the price range of even an EAL7 evaluation by several orders of magnitude.

NSA is developing protection profiles that specify small operating system kernels with simple security functions for use in special purpose high-risk applications that can be evaluated to a high assurance level. The most complete example is the Partitioning Kernel PP for Real Time operating systems. These systems must isolate data at multiple security levels (e.g., Top Secret and Unclassified). That PP specifies EAL7+ and several vendors are pursuing the development of such systems and their evaluation against it.

IV. Strengths

The Common Criteria evaluation process has several strengths, which are listed below:

- It provides consumers with an independent and well-monitored assessment of vendor security claims. These claims are often difficult to determine from marketing literature that touts security features, or even from independent reviews that compare products.
- It provides a precise expression of a product's security features that is readily comparable to those of other evaluated products. This description is similar to that used in legal documents, with carefully defined terms. Consequently, it allows comparisons between different expressions to be made more accurately.
- It assesses the ability of a product to be used to build secure systems. It clearly identifies the security functions and the limits of their implementation.
- It demonstrates that at least one configuration of a product meets the claimed security requirements. As part of the evaluation, the vendor must specify a secure configuration of the product. That configuration then becomes the basis for the vulnerability analysis and the vendor and evaluator testing.
- It allows precise tailoring of the criteria to the security capabilities of products. The flexible nature of the CC's menu of security functional requirements allows the specification of nearly every security function and its customization to the precise method implemented. The ability to augment the CC requirements with modified or completely new requirements allows the complete specification of any security function any product might have.

- It uncovers design flaws and, sometimes, software bugs. The CC process is best at uncovering design flaws. In some cases, the perspective of a CC evaluation often leads vendors to see security design flaws that they didn't recognize as flaws before. Sometimes, it also uncovers bugs.
- It focuses vendors on security issues. Some vendors do not spend much time worrying about security. A CC evaluation directs their energy into security and makes them defend their security designs to an independent third party.
- It constitutes the most rigorous and thorough independent product testing process commercially available. There are other independent testing processes that are cheaper and less intensive, but the CC is the most fundamental. Without it we would be much less able to select the right products to build secure systems or to understand the risk remaining in those systems.
- It provides International Mutual Recognition, so that vendors only have to pursue one evaluation against a single criterion. This is an important advantage of the CC over its predecessors. It provides a larger market over which to amortize evaluation costs.

V. Weaknesses

The Common Criteria evaluation process also has some weaknesses, which are described below.

- It creates an additional expense for product vendors. CCTL fees range from \$30K or less for an EAL1 evaluation to over a \$1M for an EAL7 evaluation. The cost

to the vendor to support the evaluation may be in the same range as the CCTL fees, effectively doubling these estimates. The necessary EAL is determined by customer requirements and by competition.

- The evaluation specifies a precise version of the product and a precise hardware environment. Other versions and hardware platforms are not strictly evaluated. This is due to the fact that letting those parameters vary makes it nearly impossible to reach a meaningful conclusion about the security of the product. Some consumers require strict conformance. The risk assessment part of the C&A process must deal with residual risk imposed by deviations from the evaluated configuration.
- As products protection profiles evolve, products must be re-evaluated over their life cycle.
- Because the CC evaluation process is complex and time-consuming, it requires a lot of vendor understanding and participation.

VI. Conclusion

The CC product evaluation process is a very effective tool for a very important purpose. It is critical for countering the growing threat arising from the convergence of global software development and international terrorism. Its wide international support and precise specification of security attributes minimizes the problems inherent in integrating systems and components built in different countries and services into effective and secure systems---systems whose security attributes are well understood. It does not, however,

serve every purpose. The fact that attempts are made to apply it to situations for which it was not designed shows how great the need is for other kinds of security testing and how many challenges face the available security evaluation services.