Cray UNICOS/mp Common Criteria Evaluation

Bryan Hardy, *Cray Inc*, *and* **Krystyne Supplee**, *Cray Inc*
May 18, 2004

ABSTRACT: *The Cray UNICOS/mp operating system is undergoing a Common Criteria evaluation. The U.S. NIAP (National Information Assurance Partnership) will oversee and approve the certification. This paper will give some background of the Common Criteria and NIAP, and give a status of the UNICOS/mp evaluation.*

1. What are NIAP and the CCEVS?

The National Information Assurance Partnership (NIAP) is a collaboration between the National Institute of Standards and Technology (NIST) and the National Security Agency (NSA) to promote and evaluate the security of IT products. Both NIST and NSA have a long history of developing technology, metrics, and standards to protect information critical to U.S. economic and national security interests. NIAP replaces the Trusted Computer System Evaluation Criteria (TCSEC) – widely known as the 'Orange Book' – with an evaluation program intended to be more adaptable to the evolving nature of technology and more cost-effective to operate. A long-term goal of NIAP is to "...help increase the level of trust consumers have in their information systems and networks through the use of cost-effective security testing, evaluation, and validation programs".[1]

The NIAP program to accomplish this is officially known as the NIAP Common Criteria Evaluation and Validation Scheme (CCEVS). In the U.S., independent testing labs certified as a Common Criteria Testing Laboratory (CCTL) conduct the majority of an actual evaluation – with the role of the NIAP CCEVS being to certify the CCTL, provide them with technical guidance, then validate the results of their evaluation. The first four CCTLs were certified in 2000 and the number has now grown to eight certified CCTLs in the U.S. Approximately 50 products have since been validated under NIAP and many others are formally in evaluation. Evaluated products come from a variety of areas including: biometrics, firewalls, network management, operating systems, and many more.[2]

2. What is the Common Criteria?

The Common Criteria (CC) document is a key resource in NIAP evaluations. Its official name is "Common Criteria for Information Technology Security Evaluation".[3] It is also known as ISO international standard 15408; however it is generally referred to as the 'CC'. The Common Criteria is just that – a common, multi-national, and mutually recognized set of criteria for evaluating the security of computer products and processes. It began in 1993 when three sponsoring organizations announced plans to align their separate criteria into a single set. These organizations were:

- CTCPEC – Canada
- ITSEC – Europe
- TCSEC – United States ('Orange Book')

Version 1.0 of the CC was completed in 1996. Based on public review and trial evaluations, it evolved into version 2.0 in 1998. With only minor changes, that became version 2.1 and ISO standard 15408 the following year. CC V2.2 was recently adopted in January 2004. (The evaluation of the Cray UNICOS/mp operating system began under CC V2.1 and will continue under that version.) Government agencies from a number of countries had similar need to evaluate IT security – and a similar need to reduce duplication of effort. Recognizing the value of a common set of security criteria when procuring products in a global economy, the list has now grown to eight full members. Through the Common Criteria Recognition Arrangement (CCRA)[4], signers have agreed to accept the results of CC evaluations performed by other CCRA members. These countries may have their own CCTLs or

1 http://niap.nist.gov, "Introducing NIAP"

2 http://niap.nist.gov/cc-scheme/vpl/vpl_type.html, "Validated Products List (by Technology Type)"

3 http://niap.nist.gov/cc-scheme/cc_docs/index.html, "Common Criteria for Information Technology Security Evaluation", Parts 1-3, Version 2.2, January 2004

4 http://niap.nist.gov/cc-scheme/mutual-rec.html, "Common Criteria Recognition Arrangement"

they may conduct evaluations by a governmental body. An additional eleven members currently lack evaluation capabilities themselves, but recognize certificates produced by the full members. See Table 1 for a complete list of CCRA members.

3. Comparison of NIAP to 'Orange Book'

Even prior to the formation of NIAP, an effort was underway in the U.S. to shift trusted product evaluation from the government to the private sector. This becomes apparent when comparing NIAP CCEVS to TCSEC:

3.1 CCEVS Evaluation

- In almost all cases, a commercial testing lab (CCTL) performs the evaluation.
- The vendor funds the CC evaluation.
- The CC conformance certificate is only valid for a specific HW/SW release.
- All CCRA members accept the evaluation.
- Assurance level is independent of functionality.

3.2 TCSEC Evaluation

- The NSA conducted the TCSEC evaluation.
- The NSA funded the TCSEC evaluation.
- Through on-going inspection, the concept of 'ramping' could keep the rating current across HW/SW revisions.
- A TCSEC certification was valid only in the U.S.
- TCSEC specified security functionality.

4. Cray's Pre-Evaluation Experience

4.1 Motivation

For Cray, the primary motive to evaluate the UNICOS/mp operating system was customer request. However a secondary incentive was the National Security Telecommunications and Information Systems Security Policy (NSTISSP) No. 11. In most cases, it requires evaluated, validated products for all systems used for national security information. Lastly, evaluation provided an opportunity for independent review of software development and support processes.

4.2 Choice of a CCTL

The role of a CCTL is two-fold. Some personnel act in a consulting mode: providing guidance throughout the preparation and evaluation process. Others act in a strictly evaluation role: analysing evidence provided by the vendor and ruling on its ability to satisfy criteria. Separate personnel are used to avoid any conflict of interest. Cray selected Science Applications International Corporation (SAIC) because they were the only CCTL to have completed an operating system evaluation and they had conducted the SGI IRIX evaluation (thus they already had a good understanding of UNICOS/mp roots). A contract and mutual non-disclosure agreements were signed in July 2003. (Non-disclosure agreements are essential because vendor-supplied evidence is often proprietary in nature while consulting documents created by the CCTL often contain information proprietary to them.)

4.3 Choice of a TOE

The Target of Evaluation (TOE) may be any of hardware, firmware, software, or other components of the product to be evaluated. For an intangible such as software, an evaluation is only meaningful if it includes the hardware on which that software runs. So while Cray's interest was in an evaluation of the UNICOS/mp operating system, the TOE also includes much of the Cray X1 hardware.

The choice of Target of Evaluation is a critical one. A TOE that is chosen too broadly could result in such a lengthy evaluation period that validation would be completed too far into the product's life cycle to be useful. Also, such an evaluation could become inordinately expensive and/or become a distraction from the software development and support processes. However a TOE that is chosen too narrowly may not result in a useful evaluation. Attempting to strike a balance here, Cray's choice of a TOE is the following:

- Cray X1 Mainframe – this includes items such as enforcing separation between user mode and privileged mode CPU instructions, enforcing memory access restrictions, etc
- RAID Disk Arrays – this includes the entire hardware path from the mainframe to the media
- Cray UNICOS/mp Operating System

On the hardware side, the TOE includes the Cray X1 mainframe but not support processors such as the Cray Network Subsystem (CNS), Cray Programming Environment Server (CPES), and Cray Workstation (CWS). An attempt to evaluate multiple platforms and operating systems was considered impractical given the time and resources available. Also, similar versions of much of this hardware and software have been evaluated or are being evaluated by their respective vendors.

On the software side, the TOE includes the UNICOS/mp operating system only. Asynchronous products such as the Cray Programming Environment, Cray Open Software (COS), and Portable Batch System (PBS) are not included.

The TOE is very specific to a particular software and hardware revision. If a fix package release is applied or a node module replaced with a follow-on version, that would no longer be considered a validated configuration. Since the precise UNICOS/mp fix package release chosen for the TOE will have been superseded several times prior to validation, it is questionable whether a customer would choose to install the validated configuration over a recent release.

It should be noted that the evaluation is still in progress and changes to the above TOE might occur before its completion.

4.4 Security Targets and Protection Profiles

The Security Target (ST) is a document defining the intended environment for the TOE and any security requirements that must be satisfied. It also states the set of security functionality and assurance level of the TOE. A CCEVS certificate is essentially a statement of conformance and assurance of the TOE to the ST. The ST becomes publicly available on the NIAP web site upon successful validation of the product – thus becoming a resource to users in selecting products for their environment.

A Security Target should specify a Protection Profile if one exists for the product type at the EAL sought. A Protection Profile is a common, approved set of security requirements for a key technology (eg, an operating system). Protection Profiles specify functional requirements and assurance requirements. They also specify one of three categories of robustness: basic, medium, or high. No Protection Profile existed applicable to the Cray UNICOS/mp evaluation (EAL2+ of an operating system) and so none is specified in the ST.

4.4.1 Security Target Details

The ST for the Cray UNICOS/mp evaluation includes:

- TOE Description (see section 4.3)
- Security Functionality
  - User Data Protection – Discretionary Access Control (DAC) policy, Access Control Lists (ACLs)
  - Identification and Authentication of Users
  - Security Management – Administrative tools and the underlying system calls to manage user accounts
  - Protection of the TSF – Ability of the TOE to protect itself from accidental or malicious compromise
- Security Assurance Requirements (see section 4.5)

It should be noted that the evaluation is still in progress and changes to the above ST might occur before its completion.

4.5 Choice of an EAL

Another difficult choice is the Evaluation Assurance Level (EAL) to have the TOE evaluated against. A larger number indicates a higher level of assurance. Note that a particular EAL is not a direct measure of the security of an end product. An EAL can be considered more a measure of the assurance in the evaluation itself – that is, the depth or rigor of the evaluation of the TOE against its Security Target. Commercial products can be evaluated at EAL1 - EAL4; internal government products can be evaluated up to EAL7.

Cray has chosen EAL2+ for the Cray X1 UNICOS/mp evaluation. The "+" refers to a Flaw Remediation augmentation. (Flaw Remediation is the handling of software security problems discovered after product release and would include items such as the Software Problem Report [SPR] processes and database, "trusted vendor" status and early receipt of CERT and other advisories, etc.). This augmentation is not required for EAL2 and thus raises the level of the evaluation.

It is difficult to compare EALs to the older TCSEC 'Orange Book' security levels and there is no generally accepted mapping between the two.

4.6 Initial Assessment

Research and discussion on the choice of a TOE and EAL had begun in the last half of 2002. Collection of evaluation evidence began in earnest in January 2003 with one person working full-time on the project and a second working one-quarter time. But the first official event was the Initial Assessment occurring in August. The purpose of an Initial Assessment is to gauge the vendor's readiness for an evaluation. Consulting staff spent a week at the Mendota Heights facility, reviewed evidence collected so far, and met with representatives from the X1 project team, Software Division, and Publications. Based on this, more detailed discussions were held on the practicality, time, resources required, and cost of various TOE and ST options. Also during the assessment, an outline for a Security Target was developed. (As is generally done, Cray chose to contract for the CCTL to write the actual ST.) Lastly, the SAIC consultants advised on the strong and weak points of the evidence gathered so far and what else would be required before entering the formal evaluation phase. Although the goal is an evaluation of UNICOS/mp 2.4.x, a "dry-run" using the 2.3 evidence available at the time proved helpful until 2.4 evidence became available.

5. Cray's In-Evaluation Experience

5.1 Evaluation Kick-off

The Kick-off is a meeting or teleconference involving the vendor, their CCTL, and representatives from NIAP. Following this, the vendor can formally say their TOE is "in evaluation", and it is listed on the NIAP "Products and Protection Profiles in Evaluation" web page[5]. The Kick-off for the UNICOS/mp evaluation occurred in September 2003.

5.2 Evidence

It is important to note that a CC evaluation of a TOE is not only of the end product but also of the processes and procedures used to develop and support that product. Evidence must be provided for the following assurance

5 http://niap.nist.gov/cc-scheme/in_evaluation.html, "Products and Protection Profiles in Evaluation"

components – with the scope and depth of that evidence dependent upon the EAL level sought:

5.2.1 ACM – Configuration Management

Configuration Management is defined much more broadly than just source code control; it describes the full path a release takes from initial inception, through design, coding, testing, packaging, and finally to media or pre-installation on a newly manufactured system. The evidence submitted by Cray is a combination of: existing procedures already available in written form, newly created documents based upon interviews with the personnel involved to document the processes they used, physical evidence such as manuals and release media, output from custom queries to various databases used within the company, etc.

From the perspective of Release Planning, ACM evidence describes what tools and procedures are used to track release plans, customer feature commitments, new hardware feature commitments, and the timing and contents of upcoming release packages.

From the perspective of UNICOS/mp Source Code Control, ACM evidence describes the Ptools source code management system and includes an overview description, man pages, and examples of how Ptools operates. This includes its ability to track mods, inform "owners" of each area of the OS when any mod is checked in affecting their area, enforce a "two person rule" in modifying source code where any mod requires at least one reviewer, and internal policies which further reduce anyone's ability to bypass Ptools to modify source code. ACM evidence also describes mechanisms controlling who has read access to UNICOS/mp source, who has write access (via Ptools), and the very small number of administrators who have root access to the Ptools server. ACM evidence also describes how changes to documentation - manuals and man pages - are controlled through similar source code management procedures.

From a Packaging perspective, ACM evidence describes the process used after code freeze to "split" the upcoming release from the development source tree so that it is isolated from further changes during the build, package, and installability test process. It also describes the procedure used should critical fixes need to be pushed back into the upcoming release - subject to management approval.

From a Distribution perspective, ACM evidence describes the Distribution Center "picking ticket" giving the recipient the ability to match part numbers and labels and help ensure that the correct components are received. Other procedures document this from the aspect of a system with a pre-installed OS done in Manufacturing.

Lastly, the ACM component requires that all evaluation evidence must itself be under configuration management control. So evidence includes a Ptools audit listing of the 'niap' evidence tree.

5.2.2 ADO – Delivery Operations

Delivery Operations are intended to ensure that the recipient receives the TOE that the vendor intended to send. ADO evidence includes steps taken by Cray to minimize any opportunity for a release to be tampered with in-transit (for example, use of well-established carriers and graphic designs on media which could not be easily counterfeited). Evidence also describes some flexibility the Distribution Center makes available to customers with unique delivery needs.

Another important element of the ADO component is installation. ADO evidence verifies that any unique security requirements are documented in the product's installation guide. For example, if system security depends upon performing a certain action during installation, then the install guide must document this.

5.2.3 ADV – Design and Development

Design and Development deals with information much more internal to the product being evaluated. ADV evidence for the UNICOS/mp evaluation includes a high-level overview of the OS, functional specifications for each functional interface, and an analysis of each functional interface's responsibility and relevance in each of four areas of security.

Only a subset will be listed here, but the breakdown of functional interfaces for the UNICOS/mp evaluation includes:

- Architectural components such as a node, an IO drawer - Each is considered a functional interface because of the operating system's dependence on their design for its own secure operation
- CPU instructions - This expands into a large number of interfaces (any user could potentially attempt execution of any instruction); but the privileged instructions are of the most interest
- System calls - This also expands into a large number of interfaces (as above, any user could potentially make any system call)
- Security-relevant commands - Examples include start-up commands, daemons called by inetd, setuid commands, sestgid commands, and commands referenced in installation and administration manuals.

For a non-setuid, non-setgid command executed without root privileges, the responsibility for security rests at the system call layer rather than with the command. (The issue of whether that command is genuine or a Trojan horse is a separate administrative concern.) Libraries were also excluded for this same reason that enforcement is (and should be) at the system call layer. (Most libraries are part of the Cray Programming Environment release and thus outside the scope of the TOE as well.)

The ADV component requires functional specifications for each of the functional interfaces listed above. (The remaining commands were considered to be functional interfaces but did not require functional specifications.) All

of the hardware interfaces could be satisfied from excerpts of existing engineering design documentation. Man pages were used as functional specifications for the remaining interfaces. This required efforts from the Publications department because some internal system calls (not intended to ever be called by a customer) did not have man pages. Also some existing man pages required revision because they did not address some of the information required by a functional specification.

Using both the functional specification and source code analysis, each functional interface then had to be evaluated for applicability to any of the following security functionality categories (as described in the Security Target):

- User Data Protection (UDP) - For example, the open(2) system call falls into this category because the kernel must check the file or device owner, permissions, and possibly an Access Control List (ACL) before completing the system call. However the read(2) system call would not fall into this category because UDP checks would already have been performed on the file descriptor supplied as the argument. Commands rarely fall into this UDP category because the burden is on the system call layer. An exception would be a setuid or setgid command, which then must take on this UDP responsibility while executing with a privileged effective id.
- Identification and Authentication (I&A) - Very few system calls fall into this category. A number of commands have this responsibility to verify user identity (typically through a login/password prompt) before proceeding. Examples are ftpd(8), login(1), passwd(1).
- Security Management (SM) - Restricts the ability to modify security attributes to the object owner or a user executing with root privileges. For example, the chmod(2) system call restricts changing of permission bits; the passwd(1) command restricts who can change what information in the password and shadow files. Both fall into this SM category.
- Protection of the TSF (TSF) - This category includes the ability of the TOE to protect itself from accidental or malicious acts. Most of the privileged instructions fall into this category because the TOE (in this case hardware rather than software) cannot allow non-privileged execution of these instructions without jeopardizing its own security as well as that of users. Many system calls fall into this category.

5.2.4 ALC – Life Cycle

No EAL level requires Life Cycle evidence. However it is available as an augmentation. Because of the well-defined and well-tested Software Problem Report (SPR) database and procedures going back to Cray Research Inc days, a Flaw Remediation augmentation was considered relatively easy to achieve and a good addition to the evaluation.

5.2.5 AGD – Guidance Documents

The Guidance Document component requires evidence for both administrative and user level guidance. For this evaluation, evidence consists of the entire UNICOS/mp manual set, man pages, and several custom documents addressing points specific to the evaluation.

5.2.6 ATE – Testing

The ATE component requires the vendor to provide evidence of test coverage for security related aspects of the TOE. This includes: a test plan, test suite(s), build and run instructions, expected results from each test, and actual results from runs performed on the TOE configuration. The Independent Testing phase has not yet occurred; but when the evaluation team is on-site, it is expected that they will rerun some or all of these tests to confirm the actual results. The team may also bring tests of their own. It's also expected that they will do some installability and administration testing using dedicated system time and the documentation supplied with the TOE.

5.2.7 AVA – Vulnerability Analysis

The Vulnerability Analysis and the Security Target were the only components of the evaluation contracted out (both to SAIC). Contrary to what its name might imply, the vulnerability analysis is not a "recipe to break into a system". It is a very high-level overview of generic vulnerabilities present in most System V and BSD derived operating systems. An example would be: given the minimum password length imposed by the TOE as released, the number of legal characters available for use in a password, the restrictions imposed by the TOE such as not allowing passwords that are permutations of login names, a guess rate of X guesses per minute, and a lock-out period of X seconds following an incorrect password -- then using no specialized equipment the average time to guess a user's password would be X years.

6. Current Status

6.1 Timeline

- Last half of 2002 – Initial planning for evaluation, resource and staffing requirements, etc
- January 2003 – Collection of evaluation evidence began
- February 2003 – Began CCTL selection process via phone interviews and networking at conferences
- July 2003 – Contract and NDA signed with SAIC
- August 2003 – Initial Assessment meeting on-site with SAIC
- September 2003 – Evaluation Kick-off teleconference, now "in evaluation"
- September 2003 – Version 0.1 of Security Target delivered

- October 2003 – First collection of evaluation evidence delivered to CCTL
- March 2004 – Release of UNICOS/mp 2.4
- On-going – The process of developing evidence is iterative: submit evidence, wait for feedback in form of an Evaluation Technical Report (ETR), refine, and resubmit
- (Planned) Summer 2004 – Independent Testing phase on-site in Chippewa Falls and Mendota

6.2 Estimated Completion

Cray's goal is to complete the evaluation and be validated by the end of 2004 – and preferably by the end of 3Q04. At that time, NIAP CCEVS will present a validation certificate and the Cray X1 UNICOS/mp entry on the NIAP web site will move from the "Products in Evaluation" list to the "Validated Products" list. The Security Target will then become publicly available on that web site.

6.3 Costs

By the time of completion, the total cost for the evaluation is expected to approach $1 million.

6.4 What Has Been Learned

Few conclusions can be drawn until the evaluation is completed. However the experience makes it clear that the process is a significant cost, time, and resource burden to a vendor. The process has confirmed that the development and support processes – some new, some inherited from Cray Research days – serve well.

Table 1. CCRA Participants

CCRA Scheme Title – Country

- Australian Information Security Evaluation Program (AISEP) Defence Signals Directorate – Australia
- Communications Security Establishment – Canada
- Bundesamt fur Sicherheit in der Informationstechnik – Germany
- Service Central de la Securite des Systemes d'Information – France
- Japan Information Technology Security Evaluation and Certification Scheme (JISEC) – Japan
- Government Communications Security Bureau – New Zealand
- Communications-Electronics Security Group and Department of Trade and Industry – United Kingdom
- National Information Assurance Partnership (NIAP) Common Criteria Evaluation and Validation Scheme (CCEVS) – United States of America

The following nations do not have a national scheme for conducting evaluations but have agreed to accept the certificates produced by the nations listed above.

- Federal Ministry of Public Servic and Sports – Austria
- Ministry of Finance – Finland
- Ministry of Public Order/National Information Service – Greece
- Presidenza del Consiglio dei Ministri Autorita Nazionale per la Sicurezza CESIS III Reparto – UCSi from Italy
- Ministry of the Interior and Kingdom Relations – The Netherlands
- HQ Defence Command Norway/Security Division – Norway
- Ministerio de Administraciones Publicas – Spain
- SWEDAC (Swedish Board for Accreditation and Conformity Assessment) – Sweden
- Ministry of IT and Telecommunication – Hungary
- Turkish Standards Institution (TSE) – Turkey

Table 2. Acronyms

- **CC** – Common Criteria – generally refers to the "Common Criteria for Information Technology Security Evaluation" documents (parts 1-3); however it is sometimes used to refer to this overall strategy of mutually recognized evaluations using a single set of evaluation criteria
- **CCEVS** – Common Criteria Evaluation and Validation Scheme – a NIAP program
- **CCPSO** – Common Criteria Project Sponsoring Organizations – this multi-national group of government organizations provided the framework for the CC document and also hold its copyright
- **CCRA** – Common Criteria Recognition Arrangement –signers agree to accept CC evaluations done by other CCRA members
- **CTCPEC** – Canadian Trusted Computer Product Evaluation Criteria – predecessor to the CC, it provided significant input to the first CC document
- **CCTL** – Common Criteria Testing Lab – third party commercial testing lab; U.S. based CCTLs are certified by the NIAP CCEVS
- **CEM** - "Common Methodology for Information Technology Security Evaluation"; similar to the CC document but written from the perspective of the evaluator detailing what they should look for in evaluating specific CC criteria
- **EAL** – Evaluation Assurance Level – the CC specifies levels of 1-4 for commercial products and 1-7 for internal government products
- **ETR** – Evaluation Technical Report
- **IT** – Information technology

- **ITSEC** – Information Technology Security Evaluation Criteria, European predecessor to the CC, it provided significant input to the first CC document
- **NDA** – Non-Disclosure Agreement
- **NIAP** – National Information Assurance Partnership
- **PP** – Protection Profile
- **ST** – Security Target
- **TCSEC** – Trusted Computer System Evaluation Criteria, predecessor to the CC, also known as the 'Orange Book', it provided significant input to the first CC document
- **TOE** – Target of Evaluation

References

http://niap.nist.gov

http://csrc.nist.gov/cc

http://www.niap.nist.gov/cc-scheme/ccra-participants

"Common Criteria for Information Technology Security Evaluation", Parts 1-3, Version 2.2, January 2004", *htpp://niap.nist.gov/cc-scheme/cc-docs/index.html*

"Common Methodology for Information Technology Security Evaluation", Version 2.2, January 2004, *htpp://niap.nist.gov/cc-scheme/cc-docs/index.html*

About the Authors

Krystyne Supplee is manager of OS Testing, Integration, and Installation, Cray Inc. Her e-mail address is: krys@cray.com. Bryan Hardy is a systems analyst in the OS User Interfaces group, Cray Inc. His e-mail address is: beh@cray.com. Both can be reached at: Cray Inc, 1340 Mendota Heights Road, Mendota Heights, MN, 55120.