

Understanding the Windows EAL4 Evaluation

Jonathan S. Shapiro, Ph.D.
Johns Hopkins University Information Security Institute

By now, you may have heard that Microsoft has received a Common Criteria certification for Windows 2000 (with service pack 3) at Evaluation Assurance Level (EAL) 4. Since a bunch of people know that I work on operating system security and on security assurance, I've received lots of notes asking "What does this mean?" On this page I will try to answer the question. For the impatient the answer is:

Security experts have been saying for years that the security of the Windows family of products is hopelessly inadequate. Now there is a rigorous government certification confirming this.

Since that's a pretty strong statement, bear with me while I try to explain it in plain English.

How a Security Purchase Should Work (In Abstract)

At the risk of telling you something you already know, here is how a purchaser ought to proceed when buying a security product:

1. Assess your needs. Determine what your requirements are.
2. Decide which product you are most confident will meet those needs.
3. Buy and deploy it.

Each of these is potentially an involved process, and most customers don't have the expertise to do them effectively. Even if you did, Microsoft (or any other vendor) isn't likely to let you examine their code and design documents in order to evaluate their product.

The purpose of the Common Criteria process is to develop standard packages of commonly found requirements (called Protection Profiles) and have a standard process of independent evaluation by which an expert evaluation team arrives at a level of confidence for some particular software product.

As a customer, this makes your life simpler, because you can compare your needs against existing requirements constructed by experts and then see how well the software you are buying meets those requirements. Security requirements are fairly hard to write down correctly, but if the resulting document is annotated properly they aren't all that hard to understand.

Obviously, if you don't know your needs (requirements) you don't stand much of a chance of getting them met. Likewise, if you don't know what requirements a software product was evaluated against, the evaluation result isn't terribly useful to you in practical terms.

How Common Criteria Works

From the customer perspective, a Common Criteria evaluation has two parts:

1. A standardized requirements specification called a *Protection Profile* that says what the system is supposed to do. Sometimes there will be more than one of these -- usually a general baseline protection profile and then some others describing additional, specialized requirements.
2. An evaluation rating. This is basically an investigation by well-trained experts to determine whether the system actually meets the requirements specified in the protection profile(s). The result of the evaluation is an "Evaluation Assurance Level" which can be between 1 and 7. This number expresses the degree of confidence that you can place in the system.

In order to understand the result of an evaluation, you need to know both the evaluation result, which will be a level between EAL1 and EAL7, and the protection profile (the requirements that were tested). Given two systems evaluated against the same protection profile, a higher EAL rating is a better rating *provided* the requirements meet your needs.

Knowing that a product has met an EAL4 evaluation -- or even an EAL7 evaluation -- tells you absolutely nothing useful. It means that you can have some amount of confidence that the product meets an unknown set of requirements. To give a contrived example, you might need a piece of software that always paints the screen black. I might build a piece of software that paints the screen red with very high reliability, and get it evaluated at EAL4. Obviously my software isn't going to solve your problem.

The Windows 2000 Evaluation

Microsoft sponsored an evaluation of Windows 2000 (with Service Pack 3 and one patch) against the *Controlled Access Protection Profile* (plus some enhancements) and obtained an EAL4 evaluation rating. This is most accurately written as "CAPP/EAL4".

Problem 1: The Protection Profile

The **Controlled Access Protection Profile (CAPP)** standard document can be found at the Common Criteria website. Here is a description of the CAPP requirements taken from the document itself (from page 9):

The CAPP provides for a level of protection which is appropriate for an assumed non-hostile and well-managed user community requiring protection against threats of inadvertent or casual attempts to breach the system security. The profile is not intended to be applicable to circumstances in which protection is required against determined attempts by hostile and well funded attackers to breach system security. The CAPP does not fully address the threats posed by malicious system development or administrative personnel.

Translating that into colloquial English:

Don't hook this to the internet, don't run email, don't install software unless you can 100% trust the developer, and if anybody who works for you turns out to be out to get you you are toast.

In fairness to Microsoft, CAPP is the most complete operating system protection profile that is presently standardized. This may be the best that Microsoft can do, but it is very important for you as a user to understand that These requirements are not good enough to make the system secure. It also needs to be acknowledged that commercial UNIX-based systems like Linux aren't any better (though they *are* more resistant to penetration).

Note that the "Don't install software" part means that you probably shouldn't install a word processor. On several occasions Microsoft has unintentionally shipped CD's with viruses on them. A CD with a virus qualified as "malicious system development."

Problem 2: The Evaluation Assurance Level

Having described the requirements problem, I now need to describe the problem of the EAL4 evaluation assurance level that Windows 2000 received.

As I mentioned before, EAL levels run from 1 to 7. EAL1 basically means that the vendor showed up for the meeting. EAL7 means that key parts of the system have been rigorously verified in a mathematical way. EAL4 means that the design documents were reviewed using non-challenging criteria. This is sort of like having an accounting audit where the auditor checks that all of your paperwork is there and your business practice standards are appropriate, but never actually checks that any of your numbers are correct. An EAL4 evaluation is not required to examine the

software at all.

An EAL4 rating means that you did a lot of paperwork related to the software *process*, but says absolutely nothing about the quality of the software itself. There are no quantifiable measurements made of the software, and essentially none of the code is inspected. Buying software with an EAL4 rating is kind of like buying a home without a home inspection, only more risky.

The Bottom Line for Windows 2000

In the case of the CAPP protection profile, there actually isn't much point to doing anything better than a low-confidence evaluation, because the requirements set itself is very weak. In effect, you would be saying "My results are inadequate, but the good news is that I've done a lot of work so that I can be really *sure* that the results are inadequate.

In the case of CAPP, an EAL4 evaluation tells you everything you need to know. It tells you that Microsoft spent millions of dollars producing documentation that shows that Windows 2000 meets an inadequate set of requirements, and that you can have reasonably strong confidence that this is the case.

Conclusion

Security isn't something that a large group can do well. It is something achieved by small groups of experts. Adding more programmers and more features makes things worse rather than better. Microsoft has been adding features demanded by their customers for a very long time.

It is possible to do much better. [EROS](#), a research operating system that we are working on here in the [Systems Research Laboratory](#) at Johns Hopkins University, should eventually achieve an EAL7 evaluation rating, and is expected to provide total defense against viruses and malicious code. It won't be compatible, because the most important security problems in Windows and UNIX are *design* problems rather than implementation problems. In fact, *none* of the viable research efforts toward secure operating systems are compatible with existing systems.

It remains to be seen whether EROS or one of the other attempts to build secure operating systems will prevail, but better solutions are coming.

Jonathan Shapiro is an Assistant Professor in the [Department of Computer Science](#) of [Johns Hopkins University](#). He has been working on operating system security and assurance since 1991. His past research has yielded both formally verified security properties and dramatically improved performance results in secure operating systems. His current research focuses on tying

these results together into a complete, usable system, and on evaluating and testing the correctness and reliability of the resulting system.

Dr. Shapiro is also member of [JHUISI](#), the Hopkins Information Security Institute.