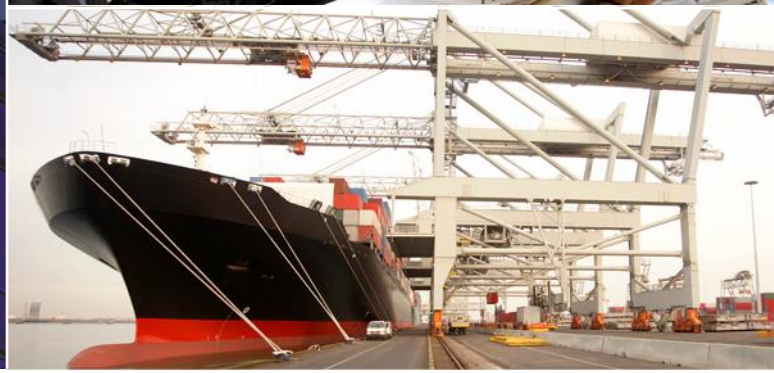

NDPP – Is It Better?

Tony Apted, James Arnold, Tammy Compton

19 September 2012

SAIC
From Science to Solutions

Synopsis

- Background
- Question
- Assessment
  - Consistency
  - Transparency
  - Scalability
  - Improved "time to market"
  - Leveraging industry expertise
  - International perspective
- Conclusions and recommendations

Background

The U.S. Common Criteria Evaluation and Validation Scheme (CCEVS) has been publishing a new series of so-called ‘Standard’ Protection Profiles (PPs), the first of which was the Security Requirements for Network Devices protection profile (NDPP). The NDPP and other standard PPs have been subject to review and debate, but only recently have the first of the NDPP evaluations approached completion.

This presentation will review the work (consulting and evaluation) related to an actual NDPP evaluation and relate the findings and work performed generally to historical evaluations (with and without historical PPs) with emphasis on end-user assurances resulting from evaluations.

Question

Is the Network Devices Protection Profile (NDPP) Better?

To answer this, it is best to first step back a bit.

- The U.S. Scheme is driving the change from the previous paradigm, based on the CC and CEM, to the currently evolving era based on Technology Community-driven Protection Profiles and Tailored Assurances. The guiding principles/objectives are: (ref. http://www.niap-ccevs.org/evolution/communities/)
  - Consistency
  - Transparency
  - Scalability
  - Improved "time to market"
  - Leverages industry expertise
  - International perspective

Leaving aside discussion about the stated objectives themselves, the interesting question is whether these objectives are being achieved and, further, whether they are being better achieved with the new approach.

Assessment - Consistency

There are a number of factors that could affect consistency. including developer understandings, consultant understandings, evaluator understandings, and validator understandings.

- One idea for the evolving era is that functional requirements will be fewer and more focused.
  - Developers are asking just as many questions about the functional requirements as they have for the retired Protection Profiles, leaving consultants and evaluators to make educated guesses or seek clarification from the certifiers/validators.
  - While simply having fewer requirements might serve to improve consistency, it turns out that in some cases specific requirements are being replaced with fewer more general requirements leading to more questions and uncertainty.
- Another idea is that explicit assurance activities will serve to better guide evaluators.
  - The assurance activities are more specific and serve to provide better direction in general. However, the same assurance activities performed by different experienced evaluators on essentially the same material are providing results that largely intersect, but also have interesting differences. It turns out that certain verbs, such as ‘describe’, are treated subjectively beginning the question of how much description is needed. There are also some requirements buried in unexpected places, like application notes, that might not be enforced consistently.
  - At first glance, it seems like there is less consistency in evaluation work based on assurance activities, but upon reflection assurance activities certainly have more potential to improve consistency. There is simply not much case law or community experience to smooth out the rough edges at this point.

Assessment - Transparency

The concept of transparency seems important because of perceived doubt or lack of understanding about evaluation results. This stems from the fact that vulnerabilities continue to be found even after products are evaluated. There is also apparent lack of trust between or confidence in various stake holders.

- The evolving era requires more definitive evaluator guidance, in the form of assurance activities, offering insights into what will actually occur during an evaluation.
  - While the evaluator instructions are public it remains to be seen whether any benefit is realized by users. The assurance activities remain subjective and in some cases extensive leaving comprehension a challenge.
  - It is not clear this is better than the CC and CEM standards in terms of conveying any benefits that may result from evaluations.
- The intent is also to publish an actual evaluation report.
  - Since no evaluation report has been published at this point, it can only be surmised that the requirement to publish will naturally result in better reports and an accurate reflection of the actual evaluation activities.
  - The production of such reports will necessarily introduce added cost and time to complete the evaluation and, again, it remains to be seen whether it has any actual value to any users.

Assessment - Scalability

Changes in policies over the past several years have served to reduce the U.S. evaluation volume on the order of 75%. Since its inception the CC has been predicated on a commercial laboratory concept to address scalability. In theory, the growing set of certificate producing Schemes should also serve to help with scalability. As such, perhaps this objective is really just to ensure that scalability is not lost.

- The reduction of required custom evaluation artifacts serves to reduce the need for developer labor and CC consultants.
  - This serves to generally improve scalability. Of course this really just reduces cost since there is an endless supply of consultants.
- The reduction of the overall assurance level for evaluations serves to reduce evaluation effort while the addition of assurance activities serves to increase it.
  - The level of effort to perform evaluations using new PPs was expected to generally decrease.
  - It remains to be seen whether that is truly the case given that PPs are starting to include onerous assurance activities (e.g., dozens of RFCs or extensive test coverage) and also given the intent to product a publicly consumable, detailed evaluation report.
- In the U.S. evaluation oversight was limited to three oversight review activities and that is now changing to a similar number of informal synchronization meetings.
  - While the oversight might be decreased, there is no evidence that will actually be the case so far.

Assessment - Improved "time to market"

It is assumed that this means that products should complete evaluation close to being available to users. This is a tricky concept since it is mostly dependent on good planning, timing, and execution,

- In the U.S. Scheme, the oversight body has traditionally not been involved in the overall project planning and is not committed to any schedules or milestones. As a result, the Scheme assumes no responsibility for evaluation schedules or completion dates.
  - This appears to remain the case, though the Scheme has alleviated the formal oversight milestone intended to occur before hands-on evaluation testing.
  - The U.S. Scheme did impose a formal final oversight review milestone that served to add 1-2 months, at a minimum, after testing was completed in order to wrap up the evaluation. While the formal milestone has been replaced with a wrap-up, it remains to be seen how long it will take after the evaluators are essentially done.
- Time to market has been and remains limited by hands-on testing of the product to be available. As such, the delay from product availability and evaluation completion is minimally the summation of the time it takes to test the product, the time it takes to finalize the test report, and the time it takes the oversight body to wrap-up and issue a certificate.
  - The evolving process has not improved, and may have extended, the time for testing given that is the primary evaluation focus and has extended the time to finalize the test report, given the requirement to publish it. It is unclear whether the oversight wrap-up could counteract those increases, but it seems unlikely.

Assessment - Leverages industry expertise

In the U.S., this seems indirectly related to scalability given that Protection Profiles were almost exclusively developed by the oversight body. Regardless, it is reasonable to expect that better and more palatable PPs could be developed with the involvement of industry experts.

- To date, some of the new PPs have been developed in vacuums or in limited communities.
  - While this does not serve the objective, it has apparently been done in some cases for expediency.
  - Regardless, there is undoubtedly improvement in potential to participate as an industry expert.
- There is a clear shift to more industry involvement in arising Technical Communities.
  - A problem is that with broader involvement brings more diversity and potential for inconsistency.
  - Care needs to be taken to ensure that objectives don't serve to work against each other.

Assessment - International Perspective

The international perspective presumably includes the notions of acceptance of the evolution to this new era of evaluations, mutual recognition and acceptance of results, and ability to participate as an international certificate producing Scheme.

- It is not clear all Schemes are on board.
  - The U.S. Scheme seems to be changing policy ahead of all other Schemes and it remains to be seen whether all Schemes are willing to follow that course.
- Mutual recognition remains limited to CC assurance requirements.
  - It remains to be seen whether EALs and potentially the CC assurance requirements in general can be replaced by explicit assurance activities as a basis for mutual recognition.
- Presumably other Schemes will be able to oversee evaluations against the new PPs.
  - However, the NDPP includes requirements for NSA evaluation outside the scope of a CC testing laboratory. This is akin to the explicit FIPS requirements found in a number of retired PPs.
  - Scheme-specific requirements need to be excluded from PPs in order to make real progress on this objective.

Assessment - Scorecard

To recap the assessment so far:

Objective	Much worse	Worse	No change	Better	Much improved
Consistency			→		
Transparency				?	
Scalability			?		
Improved "time to market"		?			
Leverages industry expertise					?
International perspective		?			

Assessment – Other Factors

Looking back over the stated guiding principles, it seems a couple of obvious ones are missing.

What about cost?

- The National Information Assurance Partnership (NIAP) Technical Community webpage doesn't even mention the word cost.
  - While NIAP might not have diminishing cost as a driving factor, the concept of cost-effectiveness has certainly been used. Just like time-to-market, Technical Communities should be directed to work toward cost-effective solutions.
- The paradigm shift has shifted and changed the costs of evaluation work, as depicted on the next slide.
  - It is envisioned that the total cost of evaluation has tended to decrease relative to EAL2 or higher evaluations, though it would be hard to predict the precise difference given that the new process is still evolving.

Assessment – Other Factors

The costs associated with evaluations moving from EAL2 or higher to an NDPP-like case:

Evaluation Aspect	Developer/Consulting	Evaluation
Security Target (ST)	More effort due to more required design information	More effort due to TSS assurance activities
Design	Much less effort – shifted to published documents and the ST	Less effort – shifted focus to published documents and the ST
Guidance	Slightly more effort due to specific required content	Slightly more effort due to Guidance assurance activities
Life Cycle	Less effort due to eliminating a need for a specific document	Less effort due essentially from the shift to EAL 1
Testing	Much less effort do to a shift of responsibility to the evaluators	Much more effort due to a shift entirely to the evaluators and test related assurance activities
Vulnerability Analysis	No change	Less effort due to limiting the work to published vulnerability data
Total	Substantially less effort	More effort, but probably less than the effort saved in developer/consulting

Assessment – Other Factors

So, given that evaluations will tend to be somewhat less costly,

What about security?

- While not called out as a guiding principle, the Technical Community concept addresses security in the following ways:
  - Technology evolution should be tracked to ensure they satisfy U.S. government protection needs.
  - Identify specific threats.
  - Identify minimal security functionality to address the identified threats.
- It seems as though this should be explicitly identified as a primary goal.
  - Reducing the assurance level to EAL 1 certainly means that evaluations would be less thorough. Adding assurance activities could possibly add assurance, perhaps even to surpass some higher assurance levels, that is uncertain as there has been no published analysis or guidance that might tend to ensure that is the case.
  - Given that new era evaluation are not currently concerned with architectural analysis or a complete understanding of product interfaces, it seems that evaluation findings will tend to be less conclusive about the actual security of evaluated products.
- At the end of the day, the question is whether the change in resulting assurance is cost effective.

Assessment – Other Factors

The change in relative assurance or security resulting from evaluations:

Evaluation Aspect	Evaluation
Security Target (ST)	In theory there is more assurance due to transparency as well as the added focus on the core security functions of a given Technology (note that while the NDPP doesn't focus on any actual business threats, it is assumed that subsequent PPs will do so.)
Design	There is less overall assurance the products meet their requirement since architecture isn't examined and there is no intent to ensure that all interfaces have been addressed – in effect the principles of non-bypassability and tamper are missing.
Guidance	The assurance here is pretty much the same – the assurance activities just serve to identify specific topics to be covered that should otherwise have been determined by the evaluators.
Life Cycle	The assurance here might seem like it is less, but in reality there is little difference.
Testing	This is a hard one, but replacing developer testing with evaluator-developed tests would not likely produce any more assurance in the end result.
Vulnerability Analysis	The assurance here is likely similar or only slightly less, while evaluators were required to consider evaluation evidence in addition to public information that likely will continue to happen informally.
Total	In total, it seems like there is less assurance given the rather limited analysis and understanding of the product architecture and more detailed design.

Assessment – Other Factors

Considering that it seems both costs and resulting assurances have decreased and an assessment against the overall objectives indicates only marginal improvement, it is not clear that the new paradigm is a success at this point.

This is not meant to imply that the new paradigm won't become better as it matures.

In order to improve the overall effectiveness:

- Do not impose unnecessary work – meaning that all of the work should yield assurance for the end user. For example, requiring substantive content to be added in a Security Target with no associated tests does not yield meaningful assurance.
- Do not include claims that will not be meaningfully evaluated – again, this will serve to create ineffective work. For example, security claims that do not result in visible behavior likely cannot be meaningfully evaluated.
- Always define claims, assurance activities and other material to be as clear, concise, and objective as possible to mitigate mistakes and disagreements.

Conclusions and Recommendations

In summary, we seem to have some improvements, some minor setbacks, and the likelihood of added improvements with time and experience.

It seems the Technical Communities need to have more well defined and carefully thought out objectives with applicable guidance, rules, etc. to ensure they are as effective as possible.

It also seems as though the National Information Assurance Partnership (NIAP) needs to work on better adherence to its guiding principles in order to serve as a better role model for Technical Communities.

- For example, to suggest that technical Communities should specifically seek to satisfy US government protection needs is not very ‘international’. Rather, Technical Communities should be looking out for the protection needs of any and all stakeholders.

Contacts

Tony Apted
James Arnold
Tammy Compton

SAIC Accredited Testing & Evaluation Labs

http://www.saic.com/infosec/testing-accreditation/common-criteria.html