
UK IT SECURITY EVALUATION & CERTIFICATION SCHEME

UK CC Interpretation - UK/2.3/014

13 February 2009

Status: Endorsed for use in UK Scheme

Subject: ALC_FLR.3 Requirement for a Timely Response to Security Flaws

REFERENCES CC, Version 2.3, August 2005
 CEM, Version 2.3, August 2005

Issue

1. The CC requirement ALC_FLR.3.9C [CC, Part 3] states the following:

“The flaw remediation procedures shall include a procedure requiring timely responses for the automatic distribution of security flaw reports and the associated corrections to registered users who might be affected by the security flaw.”

2. The question at issue is what the CC means by a “timely response”. Neither [CC Part 3] nor [CEM] offer any elaboration or explanation of the term that would enable an evaluator or developer to determine, in a truly objective manner, whether the responses are indeed **timely**. This requirement is one of the characteristics that distinguishes it from ALC_FLR.2, and hence a common understanding of the meaning of the term is essential if evaluators are to reliably and consistently assess claims of conformance with ALC_FLR.3.

Interpretation

3. In order to satisfy the ALC_FLR.3.9C requirement for a *timely response* to security flaws, the developer must provide an explanation of their flaw remediation process that details the procedures for:
 - a) Ensuring that sufficient resources are allocated to the investigation of security flaws and to the implementation and testing of temporary workarounds and actual fixes.
 - b) Assessing the perceived risk posed by security flaws¹, prioritising on the basis of that risk, and determining whether (in the case of high impact flaws) more immediate action is required to develop and promulgate interim workarounds.
 - c) Determining and setting appropriate target timescales for the resolution of each security flaw².
 - d) Tracking the progress of a security flaw against the defined timescales, and invoking appropriate escalation procedures to prioritise ahead of other development work in the event that progress is not satisfactory (dependent on the perceived risk).
 - e) Ensuring that the fix is sufficiently thorough to resolve the issue (including any related flaws that have the same underlying cause) rather than (for example) simply masking the behaviour that leads to the flaw.
 - f) Ensuring that fixes are thoroughly tested, such that the delivered fixes will completely resolve the reported flaw and any interim workarounds will be effective in minimising the risk.
 - g) Distributing fixes and interim workarounds to users in a manner and frequency that balances the need for ensuring a high uptake of the fix by users against the need to avoid unnecessary delays.
4. The above interpretation is based on a consideration what constitutes reasonable (and practical) best endeavours by a developer to respond to security flaws. A response to a reported security flaw can only be said to be **timely** if the process avoids *undue delay*. Such *undue delay* might arise because of:
 - a) A failure by the developer to assign adequate resources (in terms of manpower and expertise) to the investigation, resolution and testing of a fix.
 - b) A failure to prioritise those resources in an appropriate manner. Given that resources are finite, the risk posed to the users of the TOE should be minimised by effective

¹ The term ‘perceived risk’ means the risk to the customers’ assets as perceived by the developer. This may include factors such as ease of exploitation, impact and number of customers affected.

² The purpose of setting target timescales is to provide a basis for tracking progress (so as to avoid any ‘undue delay’). It is accepted that some fixes may turn out to be more complex than originally envisaged, such that it takes more time than expected to get the fix right. The important requirement is that if the initial targets are not met, the reasons for the delay will be investigated, the targets revised, and other appropriate action (if any) taken. Note also that since ALC_FLR does not require *evidence* of application of the procedures, evaluators cannot demand *evidence* that target timescales are being set and met.

prioritisation by the developer. This would ensure that the resolution of those vulnerabilities that present the highest risk to users is not unduly delayed because the necessary resources have been diverted to fixing low risk vulnerabilities³.

- c) A failure to adequately track the progress of a reported flaw through its lifecycle towards resolution.
 - d) A schedule of delivery of fixes to the users that is too infrequent, such that it introduces needless delay into the process.
5. The term *undue delay* recognises the fact that there may be some cases where a delay in providing a response may be *reasonably* expected. Legitimate reasons for ‘delay’ (whereby some security flaws may take longer than ‘normal’ to fix) are as follows:
- a) Engineering complexity, for example, resolution requires an architectural change to the product.
 - b) A fix may require testing on many different hardware and/or software platforms or with many other dependent products. The degree of testing required will obviously impact on timescales (the more testing that is needed, the more time that is required). Curtailing testing simply to ensure that arbitrary timescales are satisfied would clearly be unacceptable, because of the adverse impact it would have on the quality of the fix that is delivered to the end-user. Indeed, as described above, an inadequate fix could pose an even greater risk than no fix at all.
 - c) Investigation of the security flaw could be delayed because the person or organisation that discovered the flaw has provided incomplete, misleading or unclear information to the developer.
 - d) Combining of fixes to vulnerabilities in one product or component usually has a practical ‘economy of scale’ for both the developer and the TOE user. Customers often wish to avoid repeated testing of fixes that would occur if the fixes are too frequent.
6. A *timely* response will therefore ensure an appropriate trade-off between these factors to provide the most effective risk reduction possible. The appropriate timeframe will vary from case to case, but it is important for the developer to set a target for resolution so that progress can be tracked. Hence a good guideline of “timeliness” is to provide the quickest response that is practical, commensurate with the need for a thorough investigation, the quality of the fix, and the measures taken to ensure high uptake of the fix. The frequency of delivery of updates may take into account the customer’s ability to ‘consume’ the patches (thus ensuring as high an uptake as possible); equally, this frequency should also be reviewed regularly in the light of the threat environment.

³ This might, for example, happen if the developer has set (or has imposed on them) target timescales for the resolution of security flaws that do not take into account the risk posed by the vulnerability.

Rationale

What CC and CEM have to say about Timeliness

7. As has already been stated above, neither CC nor CEM offer any explanation to the evaluator or developer as to how to determine when the flaw remediation procedures provide for a timely response, and when a procedure does not meet this requirement for timeliness. [CC Part 3] offers no hints, but a useful starting point for discussion is the CEM guidance on the associated work unit (ALC_FLR.3-11):

“The issue of timeliness applies to the issuance of both security flaw reports and the associated corrections. However, these need not be issued at the same time. It is recognised that flaw reports should be generated and issued as soon as an interim solution is found, even if that solution is as drastic as Turn off the TOE. Likewise, when a more permanent (and less drastic) solution is found, it should be issued without undue delay.

It is unnecessary to restrict the recipients of the reports and associated corrections to only those TOE users who might be affected by the security flaw; it is permissible that all TOE users be given such reports and corrections for all security flaws, provided such is done in a timely manner.”

8. It can be seen from this guidance that there is no hint as to what constitutes a “timely response” other than that it should not involve any “undue delay”.
9. In order to develop this interpretation, it was therefore necessary to go back to first principles, and consider the purpose and contribution to assurance of the flaw remediation procedures within the context of a CC evaluation. We then consider what is commonly understood by the word “timely”, and how this applies to flaw remediation procedures.

Flaw Remediation Procedures within the context of a CC Evaluation

10. In order to fully understand what is meant by timeliness we must first attempt to gain an understanding of how the evaluation of flaw remediation procedures fits into a CC evaluation. What does it actually *mean* when we say that the developer has evaluated flaw remediation procedures?
11. [CC Part 3, 17.2] has this to say about flaw remediation:

“Flaw remediation requires that discovered security flaws be tracked and corrected by the developer. Although future compliance with flaw remediation procedures cannot be determined at the time of the TOE evaluation, it is possible to evaluate the policies and procedures that a developer has in place to track and correct flaws, and to distribute the flaw information and corrections.”

12. [CC Part 3, 17.2] goes on (in its Application Notes) to provide a more detailed introduction to the Flaw Remediation (ALC_FLR) family of assurance components, which contains the following extracts of particular note:

“This family provides assurance that the TOE will be maintained and supported in the future, requiring the TOE developer to track and correct flaws in the TOE. Additionally, requirements are included for the distribution of flaw corrections. However, this family does not impose evaluation requirements beyond the current

evaluation.”

“Some flaws may not be reparable immediately. There may be some occasions where a flaw cannot be fixed and other (e.g. procedural) measures must be taken.”

“Once the evaluation of a TOE is complete, it is no longer the target for evaluation. Furthermore, any changes to this evaluated TOE result in the original evaluation results being no longer applicable to the changed version. The phrase “release of the TOE” used in this family therefore refers to a version of a product or system that is a release of a certified TOE, to which changes have been applied.”

13. And when we come to examine the ALC_FLR requirements themselves, the striking thing to note is that, whilst it requires an evaluation of the documented procedures themselves, it does not require the evaluators to check that the developer actually *follows* those procedures. This is in sharp contrast to the requirements within the ACM_CAP, ADO_DEL and ALC_DVS families, where an evaluation of documented procedures is always followed by a check that there is evidence of the application of those procedures. This places a limitation on the value of claims of conformance with ALC_FLR requirements: the consumer gains confidence that there are documented procedures, but no assurance is given that the procedures actually reflect what really happens (or will happen) in practice. **This applies equally to the requirement for timely responses as it does to any other ALC_FLR requirement.**
14. So why is it that there is no check that flaw remediation procedures are or will be actually applied? We might note that (unlike those other ALC requirements) it is quite feasible for such procedures to have never have been applied in respect of the TOE. However, the last quote from [CC Part 3, 17.2] gives us a hint as to why there is no check on the application of the procedures, even if evidence did actually exist. Flaw remediation procedures are unique amongst the CC assurance requirements. Application of those procedures to the evaluated TOE will result in changes that invalidate the evaluation results. In the context of the CC as it is currently defined, a check on the application of the procedures has little meaning. In some future version of CC, in which evaluation results were no longer tied to a specific TOE version, such checks would make sense; for CCv2.3, this situation obviously does not apply.
15. It is beyond the scope of this interpretation to call into question the value of the flaw remediation family within CC (as it is currently defined), but it is legitimate to ask what value does a claim of conformance with an ALC_FLR requirement (and ALC_FLR.3 in particular) actually give to a consumer of the TOE.
16. Customer expectations are very important when it comes to the extent of what is being claimed, and that applies equally to flaw remediation procedures as with any other functional or assurance claim. However it has to be qualified in this case. First it is very clear that the CC does not imply that the assurance in the evaluated TOE can be transferred to a patched release. All it does is offer the purchaser confidence that the TOE will be maintained in future (though there is no commitment to any particular timescales) - which means in practice that there are procedures in place to resolve flaws.
17. The fact that there are procedures in place has little meaning in itself, and is of no value to the consumer if they do not at the same time provide some assurance that security flaws will in practice actually be fixed. From the consumer’s perspective, it does (in general)

makes sense to apply any fixes to security flaws (notwithstanding any potential adverse impact on the validity of the evaluation results) as these will reduce the risk posed by the vulnerabilities to the assets that need to be protected⁴.

18. This risk reduction does (of course) need to be balanced against the risk that application of a patch to a security flaw invalidates the evaluation results. Whilst on a technical level any change does invalidate those results⁵ (in the sense that the assurance obtained in the unpatched version does not automatically transfer to the new version), from the consumer's perspective the only meaningful risk impact is that the change introduces a new security flaw⁶. Even then, it is interesting to note that it is only at ALC_FLR.3 that there is an explicit requirement for the flaw remediation procedures to ensure that no new flaws are introduced by any changes.
19. These considerations help provide a measure of consumer expectations from an evaluated claim against ALC_FLR. It is unlikely that consumers will appreciate distinctions in the "small print" of the individual CC requirements, especially given that most STs will include the CC assurance requirements by reference only (e.g. by stating the assurance requirements as "EAL4 augmented with ALC_FLR.3").
20. With that in mind, it is worth considering how CC chooses to characterise the ALC_FLR.3 at a high level. This is how [CC Part 3, paragraph 399] describes the objectives of ALC_FLR.3:

*"In order for the developer to be able to act appropriately upon security flaw reports from TOE users, and to know to whom to send corrective fixes, TOE users need to understand how to submit security flaw reports to the developer, **and how to register themselves with the developer so that they may receive these corrective fixes.** Flaw remediation guidance from the developer to the TOE user ensures that TOE users are aware of this important information."*
21. The emboldened part represents the text that differs from that presented for the ALC_FLR.2 objectives. It is interesting to note from this that this summary does not mention timeliness or the need for a timely response as being one of the key characteristics of ALC_FLR.3, nor does it suggest that this is an important difference between ALC_FLR.2 and ALC_FLR.3. Indeed, it is arguable that a consumer would expect that a timely response to security flaws is a pre-requisite for any ALC_FLR claim, and not one that is restricted to ALC_FLR.3. Certainly, CC does little to disabuse the reader of this notion – unless they do indeed have the time and energy to "examine the small print".
22. That does not of course mean that it is not important to interpret the timeliness requirement correctly and consistently; but it does mean that when formulating an interpretation it is

⁴ An exception to this general rule is in those cases where the application of a patch interferes with important operational requirements. Also a consumer might choose not to apply a patch if it is clearly not relevant in their environment.

⁵ In general, in this interpretation we think of a "change" to the TOE as involving a patch that results in a new version of the product. However even an interim procedural workaround constitutes a change that could invalidate the results, as it involves adding (or changing) guidance that is unevaluated and could (for example) be unclear, incomplete, or conflict with other guidance.

⁶ Even this may not be so clear cut. A new security flaw may actually be acceptable if it is minor in impact, and represents the "cost" of fixing another security flaw that has a major impact.

important not to adopt a “heavy-handed” approach that over-emphasises its importance. We should also not assume that even a significant minority of consumers will necessarily appreciate that ALC_FLR.3 means that a “timely” response is guaranteed when ALC_FLR.2 does not.

Conclusions from the CC and CEM

23. The conclusions we can draw from this analysis of CC and CEM are as follows:
- a) Viewing the flaw remediation procedures in the light of what it has to say about the risk management decisions that are involved when deciding whether to apply a patch to a security flaw would appear to form the only sensible basis on which to construct an interpretation of “timeliness”. A “timely” response will serve to reduce the risk to the consumer’s assets by reducing the window of opportunity to an attacker who might attempt to exploit that security flaw.
 - b) Whilst the concept of a “timely” response is one of the differentiators between ALC_FLR.2 and ALC_FLR.3 it does not appear to be a key differentiator, based on the high-level description of ALC_FLR.3 presented in CC Part 3.
 - c) The lack of any check on the application of flaw remediation procedures means that any interpretation of “timeliness” cannot require actual evidence that timeliness criteria are being satisfied in practice. This is because to do so would go well beyond the scope of the ALC_FLR.3 requirement, and as such cannot be used to justify a Fail verdict against this requirement (which would undermine the fundamental principle of *reproducibility* on which CEM and the CCRA is based).
 - d) Neither CC or CEM successfully reconciles the contradiction between the notion that evaluation results are invalidated once changes are made to a TOE, and the need for users of that TOE to minimise their risk exposure by applying fixes to subsequently discovered security flaws. Any interpretation of “timeliness” must therefore respect this fact, and be based on where CC and CEM currently are rather than (necessarily) where we in the UK might like it to be (or at least to where we might like it to go in the future).

Timeliness: purpose and understanding

24. As noted above, the stated purpose of flaw remediation is to provide TOE users of some level of confidence in the developer’s commitment to fix security flaws, with the implication that with evaluated flaw remediation procedures the risk that evaluation results are invalidated by applying fixes to discovered flaws is minimal when compared with the risk of not applying the fix. In that context, a *timely* response serves to minimise the risk posed by vulnerabilities to the assets that the TOE is protecting, by helping reduce the window of opportunity that is available to an attacker.
25. So what does timeliness mean? The general (dictionary) meaning of the term is:
“*coming or occurring at a suitable or at the correct time; well-timed; in good time.*”
26. We might also add to this definition “*without undue delay*”, this being the nearest CEM gets to trying to define what constitutes a *timely* response.

27. In the world of statistics (for example) *timeliness* refers to the elapsed time between the gathering of the data and the delivery to the consumers of that data. The analogy with vulnerabilities is obvious: *timeliness* relates to the elapsed time between discovery or notification of the vulnerability and dissemination of interim workarounds and patches to users.
28. It is instructive to focus on the “suitability” of the time of delivery of a fix. It is clearly the case that a fix cannot be “well-timed” if its delivery is *too late* (however we choose to interpret what we mean by “late” in this context) – because the window of opportunity presented to the attacker would then be unacceptably large, and the risk of actual exploitation consequently too great.
29. However, it is equally possible that delivery of a fix is made is *too early*. In particular, a fix will be too early if it does not **correctly** or **completely** address the vulnerability. This may be because of inadequate investigation, or inadequate testing for example. Distribution of such a fix may actually in some cases be worse than no fix at all, because of the false sense of security that it generates. In the worst case the window of opportunity to the attacker may be indefinitely extended, without either developer or user realising this.
30. It is also possible that some customers may actually prefer not to receive updates that are *too* frequent. Customers will often wish to carry out their own testing of fixes before applying them, in order to ensure that they do not interfere with critical operational requirements. The ability of customers to ‘consume’ patches is also an issue where the vulnerability has not been made public. In such cases the issuance of a patch may actually increase the risk to customers, since attackers may then be in a position to reverse engineer the patch to determine the vulnerability that has been fixed.
31. This also emphasises the point that the time taken to *deliver* a fix is not the whole story: we need also to consider the elapsed time from delivery of the fix to a customer *applying* it. This is more difficult to measure, but it will be influenced by the above factors.
32. We highlighted above the importance that a fix is not delivered *too late*. This is of course being over-simplistic: when we say *too late* we really mean *unduly delayed*. This is because we are dealing here with an assessment of risk, and we cannot (in generating an interpretation of the requirement) justify the imposition of arbitrary timescales for resolution of fixes, because this would be tantamount to saying that that we can identify the precise point (in any case) where an acceptable risk becomes unacceptable to all users.

Timeliness: external and internal perspectives

33. When judging whether the flaw remediation procedures ensure a timely response to security flaws, it is possible to view the question from one of two perspectives:
 - a) Internal perspective: this focuses entirely on the developer’s procedures. The procedures will be considered to provide a timely response if it can be argued that it would not be *practical* for the developer to deliver a fully quality assured fix in any significantly shorter timeframe. Product complexity (taking into account matters such as the number of platforms and versions that are supported) can be considered to be a legitimate mitigating factor in making this determination.

- b) External perspective: this focuses entirely on the needs of the user for minimisation of the risks to their assets. Internal factors such as product complexity do not come into consideration: they are not a sufficient “excuse”.
34. Focusing wholly on the **internal** perspective may be considered unacceptable, as this would take insufficient account of the needs of the users. However, a focus entirely on the **external** perspective could be seen as unfair to some product vendors; because “internal” mitigating factors such as product complexity may not be used to justify a delay in delivery of a fix to a security flaw in this instance, this would (in effect) preclude the vendors of such products from ever claiming ALC_FLR.3.
35. For this interpretation, we therefore take the view that it is important to keep a balance between the **internal** and **external** perspectives, and not to arrive at an approach that focuses exclusively on one to the detriment of the other. Nonetheless, we do not believe that these two perspectives are necessarily incompatible or in conflict. It should be clear from the discussion above about the contribution that ALC_FLR makes to assurance, that a consumer would not *reasonably* expect that conformance with ALC_FLR.3 means that the TOE (and future releases) will be protected in future sufficient to keep the risk to their assets to an absolute minimum. Such an expectation would be placing far too great an emphasis on the assurance contribution that ALC_FLR provides, and cannot be justified in the context of a CCv2.3 evaluation.
36. This interpretation therefore takes the view that *reasonable* consumer expectations would be that the developer is committed to making best efforts to resolve security flaws as quickly as practical, taking into account practical internal constraints. Hence the interpretation at paragraph 3 attempts to strike an appropriate balance, recognising the need to avoid ‘undue delay’ (as described at paragraph 4) whilst at the same time also recognising there may be factors at play that might legitimately cause delay in providing a response (as described at paragraph 5).