Compartmentalization and Privilege Management (CPM)

HR001123S0028

Frequently Asked Questions (FAQ)

As of May 16, 2023

Q38: Are conference travel and associated conference registration fees allowed?

A38: Yes, as long as the conference is within scope of the program.

Q37: Is the BAA number shown on p.24 of Amendment 1 of the BAA a typographical error?

A37: Yes. It should be HR001123S0028.

Q36: Would a TA1 proposal based on a Unix variant other than Linux be in-scope?

A36: This is addressed with Amendment 1 of the BAA. A TA1 proposal based on a widely used Unix variant other than Linux (e.g., FreeBSD) would be in-scope, though some additional considerations apply. The reason for focusing on a single legacy OS for test and evaluation is pragmatic. Requiring the TA3 performer to develop attack campaigns for more than one OS is likely to involve extra effort that does not directly advance the goals of the CPM program. However, a proposer who believes that a different Unix variant would provide substantial benefits may propose it, but must allocate the extra resources to work with the TA3 team in making sure that their attack campaigns work against the proposed variant.

As of April 10, 2023

Q35: Are bibliographic references included in the page limit?

A35: A bibliography can be included as an Appendix, which would not count against the page limit.

Q34: Would an approach that uses a privilege enforcement engine that is connected to one of the traditional processor buses be in scope?

A34: Yes.

Q33: Is it acceptable to include a placeholder in the budget for the engineer(s) to be hired, or do we need to identify the engineers by name in our budget proposal?

A33: Key Personnel should be identified by name. Other positions can be identified in the cost proposal by labor categories.

Q32: In the metrics table, for Phase 1-- we see compartment size = individual function but no metrics about Principals are mentioned. The Proposers Day presentation said something like

"'each instruction as its own Principal' is too fine grain." Is the intent then to treat each of the functions as its own Principal (in addition to actual users counting toward the total count of Principals)?

A32: To clarify the terminology used in the BAA, the term "compartment" is used to refer to an aggregation of data and, since instructions are themselves data, we also refer to a function as a compartment. However, a Principal represents on whose behalf the processor is running (see Q31, below); it is not a compartment. The assertion is, "If the Principal changed on every instruction, the overhead would be prohibitive". Part of the research task is to determine the appropriate boundaries for changing the Principal to obtain a good trade-off between performance and protection. If, as in the abstract model in the BAA, Principals can only change through gate calls, then part of the task for TA1 could be to refactor the program into appropriate functions and insert gate calls where appropriate.

Q31: In the BAA's Background section, the specification says that what is desirable is the ability to arbitrarily associate "Principals" with "Compartments" (where a Principal is code and a Compartment is data) in an access control matrix. There are also access controls on a Principal calling into another Principal. However, this approach only uses *static* time association with access controls. This seems to admit of the Confused Deputy Problem.

A31: The BAA contains an abstract model (pg. 5) that might help visualize how a compartmentalization and privilege management scheme would behave. It is not a "specification" and proposers are free to explore any approach to the goals of the program; some may resemble this abstract model, others may not. Explaining how a proposed approach avoids the Confused Deputy problem would be a plus.

A Principal in this abstract model is not synonymous with a chunk of code. A Principal is a representation of who the processor is running on behalf of at a particular time; it is part of the internal state of the processor that cannot be modified except through a "gate call". Many different principals may execute the same body of code at different times. On completion of the called routine, the Principal is restored to that in effect before the call.

Q30: Can Foreign Nationals or Foreign Organizations participate in the program.

A30: Per section III.A.2 of the BAA, "Non-U.S. organizations and/or individuals may participate to the extent that such participants comply with any necessary nondisclosure agreements, security regulations, export control laws, and other governing statutes applicable under the circumstances."

Q29: Can an organization submit multiple abstracts for multiple approaches?

A29: Yes.

Q28: Is there an expected size/grouping of a strong team for a proposal that attempts all three tasks?

A28: Each proposer must develop their own statement of work and should therefore size their team as appropriate. TA3 performers cannot work on any other TA. Per section III.D of the BAA, if a performer is selected for the TA3 award, that performer cannot be selected for the other TA(s) either as a prime or subcontractor.

Q27: How important is it that the work produced is accepted for upstreaming into the existing open-source projects? Is a route to technology transition considered favorable to this program?

A27: Desirable but not important.

Q26: What is the total program budget? What budget would be appropriate for a successful proposal?

A26: Proposal values should reflect a reasonable and realistic amount needed to perform the research solution. Please review Evaluation Criteria (Section V.A) in the BAA. DARPA recognizes that undue emphasis on cost may motivate proposers to offer low-risk ideas with minimum uncertainty and to staff the effort with junior personnel in order to be in a more competitive posture. DARPA discourages such cost strategies.

Q25: Are there any application areas of particular interest? Is the focus on general-purpose system and/or is there interest in cyber-physical or mission-critical systems?

A25: There are no specific application areas of interest, though applications should be relevant to DoD.

Q24: Experimental TA2 platforms may naturally also integrate (and depend on) memory safety. To what extent will the TA3 performer depend on, and engage with, memory-safety vulnerabilities?

A24: Memory safety is within scope for TA3 but TA3 is not limited to memory safety attacks.

Q23: Will software supply-chain vulnerabilities be in scope for TA3?

A23: No.

Q22: Is there any guidance on how many different attack campaigns should be developed?

A22: No. More is better.

Q21: Other than the traditional exploit method of memory errors (like ROP, ret2libc), are new exploit methods in scope, like advanced ROP or data-only attacks?

A21: Yes; however, all exploits must be appropriate to a Basic Research program (i.e., not classified and not controlled unclassified information).

Q20: It seems memory errors are the focus of CPM, what about other vulnerabilities, like side channel attack?

A20: Memory errors are not the only focus. Protection against side channel attacks is not the primary focus but certainly of use.

Q19: Do you envision that TA3 would create a virtual network of devices and move laterally/vertically through multiple machines or will it primarily be focused on individual devices?

A19: Program focus is on a single machine.

Q18: Are you also interested in *nix on IOT devices?

A18: No.

Q17: Would a microkernel-style approaches (e.g., a special version of Linux that runs atop of a microkernel) be considered in-scope?

A17: CPM's goal is to develop capabilities to compartmentalize a legacy system not replace it.

Q16: Is TA2 looking mostly for the combination of HW & SW approaches? Are mostly-software approaches that use some recent hardware extensions discouraged?

A16: We are looking for any approach that is performant and effective.

Q15: Is it in-scope under TA2 to address physical separation for particularly sensitive compartments -- that is, offloading the execution of operations with access to that compartment onto a separate CPU or device altogether?

A15: Yes. In general, we want to provide the most flexibility.

Q14: Will TA2 performers be able to budget to support TA1 performers in order to provide access to hardware and associated tools? How many TA1 performers should a TA2 performer expect to support?

A14: A combined TA1/TA2 proposal would be expected to support at least one additional TA1. A TA2-only proposal would be expected to support at least one TA1.

Q13: Where experimental hardware is provided by a TA2, who will cover the cost of that hardware (e.g., FPGA platforms for use by TA1 and TA3), and where might it be hosted?

A13: The TA2 team should propose a budget that covers the cost of the hardware.

Q12: High-performance simulators, such as extended versions of QEMU developed by TA2 teams, may provide better performance than FPGA for software and adversarial work. To what extent will TA1 and TA3 need to use actual hardware in much of their work?

A12: We are looking to get the best measurements we can. If simulators like QEMU provide better estimates against some objectives than pure hardware approaches, proposers are free to use it where appropriate.

Q11: For a compartmentalized program, should the switch between compartments be abstractable to a function call for purposes of using an existing ABI without compiler changes?

A11: Compiler changes are in-scope for TA1 and there is no requirement that compartmentalized code be compatible with the existing ABI, only that any new ABI is backward compatible with uncompartmentalized code.

Q10: Should TA2 act as a strict superset of an existing ISA? Should it be compatible with the existing application binary interface (no recompilation required to run uncompartmentalized software)?

A10: Since our goal is to compare features of the protected system to those of an unprotected system it would be required to run unprotected software on the same hardware as the protected system. This would require backward compatibility to the existing ISA and existing ABI.

Q9: Is there a specific architecture that is preferable, such as RISC-V or ARM?

A9: Only that it is a widely used instruction set.

Q8: What sorts of Linux subsystems are being considered for compartmentalization? Are all performers expected to work on the same chosen subsystem(s)?

A8: This has not been determined yet. The decision will be made by the government and TA3 performer, considering input from the other performers. Selected subsystems must be exploitable by attack campaigns yet also have code base sizes suitable to the evolving capabilities of the TA1 approaches.

Q7: LLVM is a frequent compiler foundation for research, but is not able to build some key elements of Linux (e.g., glibc). Will depending on LLVM for compartmentalization tooling be acceptable?

A7: Use whatever tools you find useful or needed.

Q6: Are there any limitations on how the software may be transformed by automated compartmentalization as long as performance against metrics is sufficient?

A6: No.

Q5: "Legacy systems" often implies that source code is not available. Is CPM addressing legacy systems for which source code is not available?

A5: CPM focuses on systems and applications for which source code is available. Analysis on binaries is within scope but not required. The goal is to provide researchers flexibility in their approaches. TA1 proposals that employ binary analysis should be prepared to build their own binaries from source.

Q4: What is "over privilege ratio"? How is it measured?

A4: Over privilege ratio is defined in the SCALPEL document (reference [6] in the BAA). There it is defined as "the ratio of the privileges allowed by a particular compartmentalization compared to the least-privilege minimum."

Q3: The BAA anticipates that there will be overhead associated with compartments and in switching between compartments; however, memory and CPU metric targets are given without a compartment count. Roughly how many compartments should be supported within the given metric targets?

A3: Part of the research is to explore the tradeoff between compartment size and frequency of context switching vs. overhead. CPM does not have fixed targets for either of those.

Q2: Are other types of software targets beyond Linux and User-Space (e.g., hypervisors) of interest to the program?

A2: Phase 1 will deal with Unix-like systems such as Linux. The BAA says, "In Phase 2, which will focus on application code running in user space, TA1 performers should be prepared to work on user-space systems […]. Proposals should identify at least one such system the proposer is prepared to work with." A hypervisor is not an example of an application running in user space.

Q1: for Figure 1 of the Proposers Day slides, please clarify if the labels "C-1, C-2, C-3,.." refer to individual compartments, or if the collection of C-1,..C-n comprises an individual compartment?

A1: Each was an individual compartment.