Automatic Implementation of Secure Silicon (AISS)
Frequently Asked Questions (FAQ) Document
April 24 – May 10, 2019

1. Question: Will DARPA provide a detailed list of threats and adversary capabilities in RE, supply chain, and cyber?

   Response: No, the performers should include their methodology for establishing such lists based on public sources.

2. Question: Will the IV&V development start before contract awards start? The AISS development team needs a metric to develop to and then later it can be checked against.

   Response: It is anticipated that an IV&V team will start earlier than the performers and will be announced.

3. Question: Will the IV&V team provide verification services to performers or will it be completely separate? Need to understand how to estimate verification effort.

   Response: It will be separate. Proposers should plan for verification within their proposed effort.

4. Question: Who will provide the firmware for TA1/SE? Is this the role of TA1 performer? What license is this expected to be under? Same question for TA2.

   Response: Any firmware required for the TA1/SE must be provided by the TA1 performer, to support the TA1/SE implementations being developed. As stated in the BAA, what is developed under TA1 is expected to be available immediately to TA2 performers under ACA at time of contract.

   For TA2, the firmware would be based on the ARM or RISC-V based path pursued. Licensing would be based on the core selected, with open source or commercial licensing based on that core.

5. Question: Is there any GFE assumed or available, e.g. EDA tools or IP for security system design and testing?

   Response: No.

6. Question: Can you provide more specifics on the cloud computing resources that will be available to performers?

   Response: Contact Government Cloud Supplier (Nimbus).

7. Question: Is the intent of this program to make the end product as open source and commercial?

   Response: Commercial for the ARM based path; open source for the RISC-V based path.

8. Question: Is the final integrated product a cloud-based design environment?

Response: Yes.

9. Question: Does the security engine need to perform vulnerability analysis on the attack vectors at multiple levels such as at gate, transistor, and active silicon?

   Response: It has to perform analysis at the RTL level. It is acceptable to add capabilities if better results could be achieved at lower levels of abstraction,

10. Question: Do you anticipate the embedded tracker communicating to the cloud through the chip normal interfaces or a special interface? Wired or Wireless, either or both?

    Response: It is anticipated that the interface would be through normal chip interfaces. Even though the use of standards such as IJTAG is desirable, no restrictions or limitations are imposed.

11. Question: Can the obfuscation use advanced fab techniques such as 3dic?

    Response: No, implementation-neutral RTL is the end product.

12. Question: For deliverables of tools, do you mean new tools created in this program, or modification/improvement of EDA tools which unfortunately are not open-source?

    Response: Commercial tools enhanced to incorporate AISS capabilities do not have to be delivered in the source form.

13. Question: For the RTL deliverable, what type of RTL is expected? For example, Configurable RTL is used to generate final RTL.

    Response: Final RTL ready for Logic Synthesis.

14. Question: Is a demonstration of Cloud-based functionality required for Phase 2 and Phase 3?

    Response: Cloud-based functionality will be required for TA2 Phase 3. However, cloud based demonstrations as early in the program as possible are encouraged.

15. Question: For remediation of attacks, would detecting an attack and providing an error code and system interrupt be sufficient?

    Response: Even though this is acceptable as a bare minimum, more effective strategies are preferable.

16. Question: How does AISS define "usability" of deliverables?

    Response: The basic question to answer is: "Would a major design house use the tool in a production flow".

17. Question: How and when in the program will the interface(s) between the security engine and the processor platform be defined and specified?

    Response: This should be defined by the performer teams early in Phase 1. Interaction and coordination between TA1 and TA2 should be a continuous and ongoing process, enabled by ACAs at contract award time.

18. Question: Do you have indications from the VC community that the AISS effort will indeed be attractive to numbers of venture investors?

    Response: AISS tools and methodology are not relevant to the VC community, however, ability to get to first working silicon with a typical Round A investment is of value.

19. Question: Given interest in generating VC demand/investment, how are discoveries arising from AISS efforts and the associated IP rights going to be treated?

    Response: From a patent rights perspective, any subject inventions arising from the AISS efforts will be owned by the performer, and the Government would obtain certain rights to practice those inventions. From a data rights perspective, any data/software delivered under the AISS efforts will be owned by the performer, and the Government will obtain certain rights to use (preferably Unlimited Rights, but other rights could be negotiated depending on the nature of the project and the applicable data deliverable/s, such as Government Purpose Rights or Limited Rights). As noted in the BAA, the proposer is required to appropriately identify any potential restrictions on the Government's use of any Intellectual Property (IP) contemplated under the award instrument in question, and that any such IP restrictions will be taken into consideration during proposal review to assess the extent to which they may potentially impact the Government's ability to transition the technology (prevent the Government from meeting the overall program goals and objectives). The standard Patent and Data Rights clauses (to include associated definitions) associated with procurement contracts are: 1) DFARS Clause 252.227-7013, DFARS Clause 252.227-7014, FAR Clause 52.227-11 or DFARS Clause 252-227-7038 (as applicable). Similar types of IP rights language would be applicable to Other Transactions; however, there is far greater flexibility to craft rights that are specific to the project being proposed.

20. Question: If we are not already on the IV&V team, are you looking for novel metrics frameworks?

    Response: Yes.

21. Question: Can you identify the IV&V team?

    Response: Not yet.

22. Question: Do TA1 proposers have to address all three tracks in the proposal?

    Response: Yes, per the BAA all three tracks within a technical area must be fully addressed within a proposal to yield a conforming response to that technical area.

23. Question: Outside of design companies, do you imagine a role for other private companies in the semiconductor supply chain to participate in AISS?

    Response: All compliant proposals will be considered.

24. Question: Is the $75M funding for phase 1 only, or for all phases?

    Response: $75M is for all phases and all performers selected from this BAA.

25. Question: Do you expect companies to be involved with universities for TA1? Can a team be composed by universities only? Same question for TA2.

    Response: Yes and yes. Please do note the type of award instruments that are available per the BAA.

26. Question: What is the purpose of the multiple tracks within a TA? Is a proposal that addresses only 1 track of 1 TA compliant/responsive?

    Response: Per the BAA all three tracks within a technical area must be fully addressed within a proposal to yield a conforming response to that technical area. At a minimum, each TA must meet listed deliverables at specific metrics as defined in the BAA.

27. Question: Should proposals address just phase 1, or the anticipated full phase completion?

    Response: The proposals should cover all phases.

28. Question: Can performers choose which cloud-based FPGA approach to use? Can it be a private vs. public cloud provider? Clarify if this is Phase 2 or 3 deliverable.

    Response: Yes, though the Government Cloud platform is recommended, cloud based operation should be demonstrated as early as practical.

29. Question: Can a foreign national working in US be a PI/Co-PI for AISS?

    Response: Yes.

30. Question: Will DARPA support new security clearances for non-traditional contractors?

    Response: No.

31. Question: The BAA specifies TA1 metrics for all three phases for supply chain, reverse engineering, and malicious hardware - but not for side-channel?

Response: As noted in the BAA, a side-channel attack simulator will be available in Phase 2 of the program. This development is a separate, but related effort which will also yield metrics. In the interim, independent metrics development is encouraged.

32. Question: What does an ideal team look like? E.g. is there a combination of non-traditional, FFRDC, large prime, product cos etc. that you're looking for?

    Response: Teams should have the capability to deliver usable solutions into production design flows. The mix should be determined by the proposer to produce the ideal response to the BAA.

33. Question: If third-party services are needed as part of a proposal, should they be "baked in" to the proposal?

    Response: Yes.

34. Question: Would a proposal be considered "non-responsive" if the order of deliverables does not exactly match what is called for in the BAA?

    Response: The deliverables should be met within each corresponding phase, as detailed in the BAA. If key capabilities can be delivered in an earlier phase, such plans should be clearly articulated in the proposals.

35. Question: Is it acceptable to participate on multiple AISS proposals?

    Response: Yes.

36. Question: Is the use of CEP a requirement?

    Response: No, however, it is highly recommended to those proposers who do not have a test platform already implemented. The AISS program will not fund creation of new evaluation platforms.

37. Question: Are there restrictions as to where work can be performed (i.e., OCONUS)?

    Response: There is no restriction on where work can be performed. However, capabilities developed must be made available to support DoD interests.

38. Question: Are all the slides going to be available to attendees?

    Response: Yes. Please see at darpa.mil “Opportunities” page.

39. Question: Will you publish the list of attendees and their affiliations/competencies so as to facilitate collaboration opportunities?

    Response: Yes.

40. Question: Is the video live stream going to be available later as a recording?

    Response: Yes. Please see at darpa.mil “Opportunities” page.

41. Question: Are timing side channels in addition to power and EM channels within the scope?

    Response: Yes.

42. Question: Are FPGA targeted proposals in the scope (possibly extending to ASICS)?

    Response: Yes, though it should be noted that the SoCs are the primary implementation target.

43. Question: Is investigation of more scalable and effective block level protections as part of the automated platform allowed and encouraged?

    Response: Yes.

44. Question: Are special computing modalities in the scope? e.g. high temp ICs, mixed signal

    Response: No.

45. Question: Could the proposal extend the software stack implementation flow? (Software stack: root of trust -> secure boot -> OS -> application layer)

    Response: Yes, but only if it provides truly unique and differentiated capabilities relative to the current state of the art.

46. Question: Does the program aim for a cloud-based IP repository that contains different versions of an IP with various PASS metrics so the platform can choose from them?

    Response: Yes.

47. Question: Are software implemented Trojan detection mechanisms (e.g. bus sniffing / register scanning firmware) out of scope?

    Response: In scope.

48. Question: Can you elaborate on passive and active Trojan detections that are in the scope of the program?

    Response:
    Passive Example: monitoring busses for deviations from memory mapped I/O rules;
    Active Example: generating bus transactions with intent of triggering hidden functionality.

49. Question: If the attacker doesn't have test vectors, then SAT attacks would not be in scope, correct?

    Response: SAT attacks are in scope.

50. Question: Are synthesizable analog solutions for protections allowed?

    Response: No.

51. Question: Does the on-chip security require the capability to detect "zero day" attacks?

    Response: The Security Engine should be upgradable to enable rapid responses to zero day attacks.

52. Question: Are analytics in scope or out of scope?

    Response: Tracking parts through the supply chain and using that data to drive meaningful analytics would be a compelling differentiator.

53. Question: Is it OK to target both ARM and RISC-V architectures?

    Response: Yes.

54. Question: Does manufacturing have to be included as part of the proposal?

    Response: Presence on the manufacturing floor in some form is needed to support enrollment and possibly other supply chain protection steps.

55. Question: Are Machine Learning solutions desired/responsive?

    Response: Only if they are truly useful and relevant to the solution.

56. Question: Are the performers responsible for defining the vulnerability database?

    Response: Yes.

57. Question: Are defense techniques against physical attacks (e.g., FIB tool) desired?

    Response: No.

58. Question: Can proposals include the use of non-volatile memory, OTP memory, etc.?

    Response: Yes.

59. Question: Are "watermarking" techniques responsive?

    Response: Yes.

60. Question: Can a proposal include gate-level analysis to support the RTL deliverable?

    Response: Yes.

61. Question: Would "shared accelerators" be considered as part of an AISS solution?

    Response: If they improve design productivity or security they may be viewed as compelling differentiators.

62. Question: The assumption that a Trojan has already been inserted seems to exclude Trojan prevention and design for trust techniques. Is this correct?

    Response: Trojan prevention techniques are NOT discouraged. If effective, they can be viewed as differentiators

63. Question: For reverse engineering, do we assume the attacker has access to input/output behavior of the design (in addition to design files)?

    Response: No. The assumption is that they have the netlist only.

64. Question: Can we assume that designs will be manufactured in a trusted foundry?

    Response: No.

65. Question: What TRL level do you expect the security techniques to be leveraged are at?

    Response: Sufficient for commercial chip companies to deploy in production design flows.

66. Question: For crypto-blocks do you intend to follow the classical ones such as AES, RSA, or interested in next gen crypto such as quantum resistant or lightweight crypto?

    Response: It is not a goal of this program to produce advances in encryption, but novel ideas in that area will be considered.

67. Question: How do you distinguish between “estimating” security and “measuring” security?

    Response: "Estimation" is a heuristic based approach that can quickly lead to one architecture being assessed as better than another in a particular security dimension;
    “Measurements" are derived through painstaking process that may require actual hardware and an exhaustive staging of attack scenarios.

68. Question: Who will have access to the government cloud? Does this include defense contractors, commercial chip makers, and universities?

Response: All AISS performers and their subcontractors will have appropriate access to the government cloud. For the duration of the program, AISS data stored on the cloud will only be accessible to AISS performers and their subcontractors, DARPA, the AISS IV&V team and government contractors supporting AISS. After the program ends, it is anticipated that the performers will maintain their solutions on the Government Cloud for some period of time, enabling evaluation and initial design activities by the AISS performers and members of the Government community (agencies, Government labs and FFRDCs), although such post-program activity is not currently part of the BAA (program) and should not be included in AISS proposals.

69. Question: What are the characteristics of an ideal transition partner?

    Response: A party that can deliver capabilities in form sufficient for commercial chip companies or defense contractors to deploy in production design flows.

70. Question: What is the blue sky vs. industry baseline for security architecture?

    Response: We anticipate that the performers will articulate clear and compelling vision.

71. Question: Can you clarify the "time per estimation" metric?

    Response: If the estimation is to be used in the context of combinatorial optimization, it needs to be computed very quickly as the optimizer is expected to traverse billions of possible solutions.

72. Question: Can the AISS BAA proposal due date be extended?

    Response: No, the due date will not be extended.

73. Question: For a team that submits a joint TA1/TA2 proposal, do you need a single Summary Chart, or one for each technical area?

    Response: Please provide a Summary Chart for each technical area. That will prevent excessive material in a single chart and allow a better representation of each technical area.

74. Question: The BAA mentions logic locking techniques that are "SAT-resilient." Does it also mean locking techniques that assume that the attacker does not have access to an oracle (and are naturally SAT-resilient) under consideration?

    Response: Yes, please also address techniques that assume the attacker does not have access to an oracle (and are naturally SAT-resilient) under consideration.

75. Question: Malicious HW Detection (Table 3 of the BAA) references MITRE's CVEs. That db (along with the errata documents from a lot of companies) contains hw bugs (exploitable by sw) - those bugs are design mistakes/errors. Are we to treat design/logic errors as a malicious hw?

    Response: Design/logic errors that provide exploitable security threats should be treated as malicious hardware, in the sense that it needs to be addressed for the overall security of the system.

76. Question: We are a prime applicant and would like to partner with a for-profit firm with security expertise who is a non-traditional defense contractor and has little or no experience with FAR and DFAR requirements and is not setup to comply with those. a) Is it possible that if sub-contracted the subcontractor can be exempt from these requirements or be awarded an “Other Transaction for prototype” or a “Technology Investment Agreement.” b) Alternatively, can DARPA/gov provide assistance in managing those requirements?

    Response:
    a) A subcontractor cannot be exempt from applicable FAR/DFARS flowdown requirements. However, what flowdowns are required is dependent on the type of subcontract being negotiated - for example, firm-fixed-price or cost-plus-fixed- fee. Firm-fixed-price procurement subcontracts do not have any FAR/DFARS clauses dealing with cost accounting/systems, as an example, so they would not come into play.
    b) The prime cannot submit a proposal that asks the Government to award separate direct instruments to one or more team members in place of the prime awarding subcontracts to those team members due to instrument type issues (or for any other reason - other than FFRDCs or Government entity team members). It is the prime's responsible to fully understand each team member’s contractual requirements/limitations and to propose a prime instrument that fits accordingly - failure to do this could result in a proposal being deemed non-compliant.

77. Question: Are proposals scored lower if one or more of the sub-contractors would prefer to be administered by an “Other Transaction” in pursuit of supporting the prime contractors mission of research and development for the DARPA BAA ?

    Response: Subcontractors can only receive a sub-award that is of the exact same type as the prime (meaning only one set of regulations can be used throughout the performer team). So - in the context of this question, the sub can only receive an OT type subaward if the prime is receiving an OT prime award. To stretch this discussion point a bit further - it is important to understand that a prime cannot submit a proposal that asks the Government to award separate direct instruments to one or more team members in place of the prime awarding subcontracts to those team members due to instrument type issues (as previously mentioned, the exception being FFRDC’s and Government entity team members). It is the prime's responsibility to fully understand the wants and needs of its team member/s and to propose a prime award instrument that fits accordingly - failure to do this could result in a proposal being deemed non-compliant.

78. Question: If we need to provide our proprietary IP to other awardees in support of the overall program, can we provide encrypted RTL to awardees we deem competitive to our commercial business, or does encrypted RTL detract from the overall goals of the program?

    Example: We bid as a subcontractor to TA1 and license a side-channel resistant crypto core to our prime. During integration for TA2, a separate prime awarded DARPA money needs to integrate this side-channel resistant crypto core to meet AISS program goals. This separate prime is a competitor to our company and, as such, we do not want to reveal our side-channel resistant crypto core source code to them. Can we deliver in encrypted netlist form, or does this detract from the overall purpose of the AISS program?

    Response: No encrypted IP will be accepted as an AISS deliverable.

79. Question: If we are a sub and we submit our commercial proprietary IP to our prime for integration into the program, we would also submit standard commercial costs to that prime for the use of our IP in this program. However, it's unclear how we should cost this if other awardees (with whom we are not in a business relationship under this program) need to use this IP in support of the overall program. In a standard commercial relationship, we would charge for every use of our IP. In the case of bidding on a government program, we are not sure if this would be acceptable, do not understand who would pay for it (DARPA? Our prime? The other awardees?), and don't understand how to cost this appropriately.

    Example: We bid as a subcontractor to a prime for TA1 and license a side-channel resistant crypto core to our prime. We then find that our side-channel resistant crypto core needs to be integrated as part of TA2 with another prime and separate awardee on the AISS program. Given that we are not in a relationship with this second prime, nor are we bidding on TA2, how should we cost the use of our proprietary IP on TA2 with a separate prime that we do not have a business relationship with?

    Response: While the end result of AISS are generators (i.e. programs that automatically produce IP unencumbered by licensing), pre-existing (background) IP may be exchanged among performers to enable generator development. Rights to background IP are to be free to all AISS performers in the course of program execution, but could be subject to inter-party confidentiality agreements. If background IP, ends up realized in silicon by a performer, it is expected that any applicable fees would be a matter or arrangement between the rights holder and a licensee and would not involve the Government.

80. Question: In the subject BAA, I do not see any allocated section for References. Typically, they would go under Section III (page 41), but I do not see them being mentioned there. Can you please clarify where References should go? Also, is Section III exempt from the page count?

    Response: References can be added to Section III of a proposal. There is no page limit for Section III.

81. Question: Can you clarify the discrepancy between the TA2 metric stating ">300 cycles/second simulation time" in phase 1 on page 21 and "1,000 cycles per second" for the integration event requirements in phase 1 on page 25?

    Response: The higher simulation performance applies specifically to the subset of Track A in TA1 while the lower simulator performance metric applies to the aggregated three tracks.

82. Question: We are a small business with traditional DoD performers as subcontractors. Section 8 of the BAA (below) indicates some limits that are confusing. Does it mean that as a small business we can only bid contracts up to $700K? Does this limit also apply to OT TIA and prototypes?

    Response: The "Small Business Subcontracting Plan" is applicable to procurement (FAR/DFARS) contracts only – and applies, for such award instruments, to prime contractors who are large businesses (per FAR Part 19) if the award (base and options) is at or above $700,000. If your organization is a small business (per FAR Part 19) then, if serving as a prime or subcontractor, this requirement would not apply.

83. Question: In general, our initial experience is that prospective primes are not willing to take as subcontractors any small businesses with pre-existing IP ahead of the industry, but would take them as consultants and work for hire. In this way the prime is sure that all the newly created IP is theirs. How can innovative small businesses still contribute to DARPA projects?

    Response: Small businesses are encouraged to team with other entities with whom they share common project goals/objectives and with whom there is a fit technically and contractually. As stated in the AISS BAA, an acceptable proposal must address all goals, objectives, metrics, phases and tracks of the respective TA. It is up to the proposing team to determine how a small business can contribute and participate.