

Software Defined Hardware program FAQs - drawn from September 19, 2017 proposers day Questions and Answers, questions from the SDH mailbox, and CMO clarifications – November 27, 2017:

-----October 03, 2017 responses-----

Question 1: Even though “HW and SW” is mentioned throughout the discussion of SDH, it sounds as though the interest in HW is limited to HW design? Is this correct?

Answer 1: Primarily yes, and also interested in the technology that enables the design. Unique architectural techniques and technologies that enable SDH goals are of interest. The proposer might want to consider other ERI areas depending on the specific architectural research being pursued.

Question 2: Is DARPA looking for a universal, homogeneous HW fabric? Is any heterogeneity in this fabric acceptable? There is a spectrum between a homogeneous fabric and DSSoC in terms of heterogeneity. What parts of this spectrum are acceptable for SDH?

Answer 2: This is a proposer’s judgement call. DARPA wants to realize as much of the specialization as possible later, closer to runtime, and to support run time optimization. But not all optimizations can be realized that way. The proposer needs to make an argument for the approach proposed and its use of homogeneous or heterogeneous components and heterogeneous specialization to support late binding and optimization.

Question 3: Is the SDH intent to dynamically recompile executable in <1 microsecond or will executable contain enough flexibility to support SDH optimization?

Answer 3: Reconfiguration of the HW is to occur in <1 microsecond. The microsecond does not include reconfiguration and optimization re-calculation. SDH will dynamically recompile executable and support optimization in parallel to system operation to support reconfiguration, and is not part of HW reconfiguration time.

Question 4: For a reconfiguration time of hardware is <1 microsecond, does the actual compilation and reconfiguration time need to be negligible with regard to runtime?

Answer 4: Most programs in D3M run for minutes. The relative overhead of compilation and reconfiguration depends on the application.

Question 5: What is the difference between reconfiguration versus missing icache?

Answer 5: This would appear to be a question regarding the choice of reconfiguring different access patterns on or off chip. The access of memory off chip is a big differentiator to be considered. This is up to the proposer where and how to address data access.

Question 6: Is the point of TA2 is to decide specialization offline, but still automatically?

Answer 6: If a proposer's wants to do this manually, they would have to decide on which "components" will be manually optimized and propose a way to discover those components. A goal of SDH is to specifically address real time optimization.

Question 7: Is there interest in TA2 tasks that can leverage SW optimization of standard HW (e.g. FPGAs) that can enable enhanced capabilities without TA1 HW development? Results may not be near-ASIC but could improve programming as well as performance of FPGAs?

Answer 7: Yes, that could be a reasonable starting point.

Question 8: What languages will the benchmark program be written in?

Answer 8: Python and NumPy.

Question 9: Can performers recode the benchmarks using SDH languages?

Answer 9: Yes. Re-implementation is acceptable.

Question 10: Data dependency - to what extent is this known a priori, versus to what extent is it blind to the data dependency?

Answer 10: Can a performer do work offline to assess/optimize code and data? Yes, a performer can. A performer can develop up front optimizations based on specific data patterns and flow. A performer will also see examples with and will need to be capable of optimizing for changing data at run time. It still will behoove the performer to also do runtime optimization.

Question 11: For programmability, NumPy hides target-specific efficiency underneath the hood and is the result of hundreds of programmer years of effort. How can performers equalize for these factors? Is hiding the complexity from application programmers sufficient?

Answer 11: For base linear algebra, this is true, aspects of current algorithms are highly optimized, but not for algorithms realized at higher levels (e.g. ML algorithms). Hiding optimization in libraries is reasonable, the programmer for instance should not be implementing SpMM.

Question 12: How much beyond stock linear algebra is expected: complex tensors, iterative solvers that operate on moving frontiers, streaming graphs?

Answer 12: Performers are encouraged to consider all of these, as well as examples such as stochastic search, machine learning, graph problems including sparse data problems. The performer is encouraged to provide as much breadth in solution space capability as possible.

Question 13: Will the problem space be primarily a very large problem (Facebook class) versus a lot of small problems (Imagenet)?

Answer 13: A proposal is desired to span as large a solution space as possible, including large and small problem spaces.

Question 14: In the BAA (page 11) the corpus of programs is stated to be from two sources (D3M and HIVE), however, in the PM's proposers day presentation only D3M was mentioned. Please clarify dense + sparse linear algebra only or also graph problems?

Answer 14: Yes, it is HIVE and D3M sources. The BAA is always the authority. Regarding the problems – yes – dense, sparse, and graph problems.

Question 15: Can the top 500 programs be prioritized to: top 50 that most care about; the 100 optimized for ASIC that will be used for benchmarking SDH? Having 500 programs is too many for design optimization. (clarification from audience – how will specialization be determined?)

Answer 15: The point of the TA2 research is to determine the specialization to be performed. Manual optimization can be considered, but must be coupled with automated optimization, to support optimization at run time. Any plan for manual optimization needs to be clearly described and the total optimization approach clearly described. The proposer needs to determine their best approach to optimization, and specifically address run time optimization based on changing data patterns.

Question 16: Can a proposer get access to codes during the proposal preparation phase?

Answer 16: No, the BAA states that the codes will be released at selection, but there are lots of publicly available code for proposers to consider. Open source data is available (look at the D3M program BAA).

Question 17: What scenarios are anticipated that require fast reconfiguration in a microsecond or less?

Answer 17: Examples include signal processing applications, stochastic search, language applications, radar and sensor processing. Proposers are encouraged to think not just about the HW, but to also think about the tool chain and what will be required to support run time reconfiguration.

Question 18: Will there be multiple TA1 and TA2 performers?

Answer 18: The intent is that there will be multiple TA1 and TA2 performers, but that will be based on the strength of actual proposals.

Question 19: How does the notion of a “prime” proposer mesh with the TA1/TA2 distinction? Does it matter who is the “prime”?

Answer 19: TA1 and TA2 can be bid independently. For a combined TA1/TA2 proposal it is up to the team who within the team should be the “prime.” The “prime” contactor should be a clear team leader and organizer, and that specifically would include the ability to organize and manage the overall program fiscally, contractually, and technically.

Question 20: What is envisioned as the team size of SDH performers?

Answer 20: There is no required or discouraged team size or specific team composition. It is totally up to the proposers to form the team size and select the team members they feel is appropriate to support the proposed SDH solution. A proposal should incorporate the resources necessary to accomplish the deliverables and program tasks.

Question 21: It is assumed that the prime contractor is fully responsible for delivering the proposed work at a negotiated price. The work can be done by full time employees, consultants and subcontractors. A performer will pay their subcontractors. Why does DARPA need to know who our subcontractors are? What if a subcontractor is added after the contract is negotiated, but they do work for hire? Are contractors/consultants considered subcontractors?

Answer 21: Your team composition is important information for DARPA to infer achievability. If a subcontractor or consultant’s contribution is important in the proposed solution and they are not in the proposal there will be a missing element that DARPA cannot consider in the proposal evaluation. The elements that compose a proposed solution and how they will be provided need to be proposed to be considered by the government evaluators. The government needs to be able to judge if an effective team has been formed to meet the SDH goals. Any change in team composition after contract negotiation should be reported to the government, since it could impact the ability of the team to meet program goals.

Question 22: Will there be down-selects?

Answer 22: The program is not designed to specifically include down-selects. However, based on funding limitations and contractor performance, the government may choose to exercise an option or may chose not to exercise an option at key points in the program.

Question 23: How is SDH and DSSoC relate to HIVE?

Answer 23: They are not directly related. SDH will use common problems from HIVE and there will be an overlap of benchmarks with HIVE. DSSoC does not have a relation to HIVE.

Question 24: If there are common tasks between SDH and DSSoC proposals, how should that be addressed?

Answer 24: SDH and DSSoC proposals should be totally independent. Joint SDH and DSSoC proposals will not be accepted or reviewed. However, it is recognized that there are common elements between the two programs. A proposer should not make any effort to differentiate common tasks between SDH and DSSoC proposals. SDH and DSSoC proposals should be written independently based on what is needed for the most complete response. However, it is asked that a proposer note tasks that may be common between SDH and DSSoC proposals so that this can be examined and coordinated at contract negotiation time.

Question 25: Should proposers plan on including studies to include human subjects to evaluate programmability?

Answer 25: NO. SDH performers will not be doing the evaluations. A USG evaluator will perform any programmability evaluations. As a result, SDH performers will not need an IRB.

Question 26: Can a prime propose a traditional DFAR contract with OTA subcontractors?

Answer 26: No, subcontractors under a procurement contract must receive a subcontract that is fully compliant with the terms and conditions of the prime contract (FAR and DFARS clauses), to include all of the mandatory flow-down clauses. Traditional defense contractors intending to propose as a prime contractor, with a non-traditional defense contractor as a subcontractor, could mitigate the procurement contract flow-down concerns by proposing use of an OT as the prime award instrument. The prime contractor would have to be mindful of the statutory requirements associated with use of an OT (such as non-federal cost share) but, as discussed during the proposers day event, the amount of such cost share - even in the case of a TIA - is negotiable (after selection) given the unique circumstances associated with each proposal/project.

-----October 11, 2017 responses-----

Question 27: Is the TSMC/CRAFT MPW (shuttle) only for academic proposals?

Answer 27: No, the CRAFT shuttle fabrication run is available to any DoD or DARPA contracted performer. If a proposer intends to use a CRAFT MPW shuttle run, they need to specifically include that in their proposal.

Question 28: What CRAFT tool access will be provided, and what is the process to obtain access to the CRAFT tools? How long will the SDH program have access to CRAFT (after the current program manager leaves or the program ends)?

Answer 28: Proposers will not have direct access to the CRAFT performers or research design tools. The Availability of the CRAFT developed tools will be through commercial vendor tools, so no special access will be required. CRAFT shuttle runs are already set, so that should not change. Access to the CRAFT MPW shuttle runs must be requested via the CRAFT program manager. Once access approval is in place the process and contacts to use the CRAFT MPW shuttle run will be provided.

Question 29: Will performers have access to CRAFT EDA tools in addition to the shuttle runs?

Answer 29: Performers will not have direct access to CRAFT performers or specific CRAFT program tool research. Availability of the CRAFT developed tools will be through commercial vendor tools, so no special access will be required.

Question 30: What are the CRAFT shuttle run dates?

Answer 30: The CRAFT shuttle dates are currently April 2018 and January 2019. Permission to use the CRAFT shuttle run will require the permission of the CRAFT PM (Litton Salmon). Once permission to be on a CRAFT MPW shuttle run is obtained, the process and contacts to use the shuttle run will be provided.

Question 31: What is the maximum size for a device using the CRAFT MPW runs? What process will the shuttles be in?

Answer 31: The maximum size available on a CRAFT MPW shuttle run is 5 mm x 5 mm. The CRAFT shuttle runs will be TSMC 16nm FFC.

Question 32: Can you provide a list of IP blocks that will be provided by the shuttle runs in order to build the full system? In particular, does the fabrication provider allow access to standard cell libraries, DRAM IP and PCI-e (or any other system interconnect) IP blocks?

Answer 32: No unique CRAFT or third party IP will be provided. Once permission to be on a CRAFT MPW shuttle run is obtained the designer will be provided access to a standard design development kit.

Question 33: For SDH is there a GPU baseline that will be used for evaluating our designs, particularly if our architecture is based on GPU like designs?

Answer 33: No, a specific GPU baseline is not currently planned. The specific program comparisons are with ASIC and CPU performance. In order to evaluate SDH against potential ASICs, the USG team will

develop baseline solutions on FPGA proxies and extrapolate these solutions to ASIC power efficiencies. CPU baselines will be based on Intel E7-8894 v4.

Question 34: Does the leading PI (or organization) have to be a person from a company, not a university faculty? I am not sure if this is the case, and I cannot see such a requirement in the solicitation. Could you please clarify this issue?

Answer 34: The lead PI or organization is not required to be from or to be a company. The lead organization can be a commercial entity, academic institution, or DoD contractor. The PI should be from the lead proposing organization. If a team is proposed, the lead organization would be selected as the organization that provides the strongest ability to organize and manage the overall program fiscally, contractually, and technically.

Question 35: I understand that DARPA expects TA-1 and TA-2 teams will need to cooperatively design both TA-1 hardware and TA-2 languages, and the program requires that TA-2 performers co-locate language design experts with TA-1 teams. Should the TA-2 team build the collaboration with a TA-1 team after the teams have been selected by DARPA for funding or before the proposal submission (i.e. the collaboration plan between TA-1 and TA-2 teams should be included in the proposal)?

Answer 35: It is the option of the TA1 or TA2 team if they desire to establish a collaborative relationship. If separate proposals are prepared, but a collaboration is desired, the proposed collaboration can be discussed and included in the proposals, and used to strengthen both the TA1 and the TA2 technical merit. Provided the collaborative pairs are both selected for funding, the efforts can be pursued as a collaboratively. If the collaborative pair is not selected, or T1 and TA2 proposals are independently proposed, the government plans to create TA1 – TA2 team pairs deemed to have the most value. Joint TA1 and TA2 proposals will also be accepted.

Question 36: Is there any guideline regarding the range or caps for the budget or the suggested team size?

Answer 36: No specific guidance for the budget will be provided. The proposer's budget should be appropriate for the proposed level of effort and activities being provided to meet SDH BAA requirements and goals.

Question 37: The instructions in the BAA for the SDH program, and under the abstract section on page 40 for "deliverables", refers the reader to see Section IV.B.10 for forms to be completed regarding intellectual property. The current topic at section IV.B.10 is Small Business Subcontracting Plan (page 53). I believe it was intended to refer to IV.B.11 Intellectual Property. I thought you might want to include a note on that in your FAQ.

Answer 37: That is correct, there is an error in the BAA. Under the abstract section on page 40 for "deliverables" it should refer to section IV.B.11 for forms to be completed regarding intellectual property, not to section IV.B.10 addressing Small Business Subcontracting Plan (page 53). This has been corrected via an amendment to the BAA.

-----October 19, 2017 responses-----

Question 38: Would an architecture that uses alternative on-chip memory approaches be considered that does not use memory blocks dispersed throughout the computational array as described within the BAA?

Answer 38: No specific architecture, or memory implementation approach, is mandated. Proposed solutions will be evaluated based on performance and capability to meet the described SDH goals and the ability to efficiently utilize data and data flow. Any architectural descriptions in the BAA are provided as representative examples only and are not intended to dictate a specific architectural approach.

Question 39: Would the SDH part of ERI be interested in a 10s of GHz FPGA that could be designed to permit rapid (one nanosecond) parallel configuration bit-programming?

Answer 39: Any proposed FPGA solution would have to clearly demonstrate how to overcome FPGA associated limitations and how the FPGA solution would meet SDH goals.

-----October 30, 2017 responses-----

Question 40: We are a small company with cash-basis accounting, and it is our understanding that our best option is to apply as TIA. Assuming a teaming approach where a university is the prime and we (a for-profit nontraditional) are the subcontractor, are there 2 proposals to be submitted (one from the prime and one from the subcontractor) or is the subcontractor proposal to be submitted within the prime's proposal (in Volume II, Section 2, G. Subcontractors)?

Answer 40: It is important to understand that a proposer team must submit a single proposal in response to the BAA so as to avoid having the individual proposals (in your example) deemed non-compliant with the BAA. The single proposal, in this example, is submitted by the organization who will be serving as the prime contractor. The award instrument type is dictated by the contractual relationship between the prime contractor and the government - all subcontractors would receive a sub-award that conforms to the prime contractor's award instrument (for example, a FAR-based procurement contract to the prime contractor would involve FAR-based procurement subcontracts, an OT agreement to the prime contractor would involve only those flow-downs necessary to ensure the prime meet their contractual obligations to the Government. It is worth noting that a subcontractor under a procurement contract could receive a FFP procurement subcontract - which, if the subcontractor is a non-traditional, mitigates any cost accounting system concerns and allows for milestone-based payments.

Question 41: According to Volume 2, Section 2 of the BAA, the subcontractors have to prepare their proposals at the same level of detail as the required by the prime. For the title sheet of the subcontractor's proposal is the subcontractor itself the leading organization? Does the subcontractor

proposal identify its own PI? Do we include in the SOW (or TDD) the tasks that have to be performed by the prime as part of our team effort?

Answer 41: Just as a point of clarification before directly responding to the above questions, it is important to note that the subcontractor is submitting a proposal directly to the prime contractor in their team – not to the Government per se, and the prime contractor is responsible for incorporating the information received by the subcontractor(s) into their prime contractor submission to the Government (only one formal submission is being submitted by each team – that submission including all information required by the solicitation).

The subcontractor/s would be a member of a team and, contractually speaking, subordinate to the prime contractor. Each subcontractor would submit to the prime contractor (typically in response to a formal request for proposal from the prime contractor if operating more formally and if the prime contractor is considering bids from multiple potential subcontractor organizations) a fully compliant subcontract proposal tied to their specific scope of work (as defined by the prime contractor - meaning the subcontractor SOW will include only those tasks that the prime has asked them to perform as a member of a single team). Generally speaking, the PI is a role filled by the prime contractor - although it's not improper for a subcontractor to also refer to their own technical lead as a PI. It is important to clarify that although this information (SOW, cost proposal, etc.) is submitted as part of the prime contractor's proposal to the Government per the solicitation, none of these items will be included in the award instrument between the Government and prime contractor (the contractual relationship will be solely between the Government and prime contractor).

Question 42: In case of proposing as loosely connected performers (one of the teaming models discussed at the ERI/Page 3 Workshop this past summer), is it possible for each team member to propose separately (for example, the prime submits a procurement contract proposal and the subcontractor submits a TIA proposal) but bound together by an ACA?

Answer 42: No - as stated in the above two responses, a single team must submit a single proposal in order to be compliant with the solicitation. Just for further, related, clarification, it is important to understand that the ACA model discussed in the ERI/Page 3 Workshop is something that is structured after a conforming proposal has been selected for negotiation - it's a model that is put together by the Government after proposals have been selected and when broader, loosely connected, teams need to be formed to meet the overall program goals and objectives (especially relevant when the solicitation allows proposers to propose to discrete elements/task areas versus to having to propose to all of the program elements/tasks). A single team needs to be captured in a single proposal and only a single award can result from that proposal - splitting a team into multiple proposals could very well result all of the proposals being deemed noncompliant since they must be reviewed independently per the solicitation.

Question 43: Do TIA proposers need to submit information required by Volume 2, Section 2, Detailed Cost Information? It is our understanding that according to Note 1(d) of this section "Cost or Pricing data" are not required if the proposer proposes an award instrument other than a procurement contract (i.e., cooperative agreement, grant, or other transaction). In this case, it is our understanding that as OT proposer we will need to submit information described in Volume 2, Section 3, instead of information required in Volume 2, Section 2. Please confirm if we understand this right.

Answer 43: All proposers need to submit the information required by Volume 2, Section 2 "Detailed Cost Information" to the maximum extent possible. This information is needed by the Government reviewers in order to fully understand what is being proposed - specifically, in order to assess cost realism, which is one of the evaluation criteria. The more cost details that are not provided the greater chance that the Government reviewers will not fully understand the proposal or will find weaknesses and or risks pertaining to cost realism. However, for those proposing a TIA, there is no requirement for the submission of "certified cost and pricing data" because this term is a procurement contract (FAR/DFARS) term. So - the lack of cost details in an OT proposal will not rise to the level of proposal non-conformance, like it might for a procurement contract whereby the requirement for certified cost and pricing information does exist (by regulation), but - as noted - the lack of any cost details could become an issue in terms of the Government reviewers not being able to fully understand the proposal, as whole, and/or in assessing cost realism (remember this is a competitive solicitation process).

Question 44: In regard to cost share, when providing the cost breakdown by phase and task do we list the full project cost or just the Government's share?

Answer 44: The cost proposal build-ups (Summary Build-up and Detailed Build-up) must capture the full cost (budget) of the project being proposed (all costs required to accomplish the technical approach being proposed). The cost share is simply applied at the bottom-line. The Government needs to be able to fully understand what is being proposed, and a full project budget is part of the information data set that is assessed. It's important to note that before agreeing to a proposed cost share amount/percentage, the Agreements Officer needs to first understand and agree with the total proposal amount (for example, 50% performer share in a total budget that is deemed to be twice as expensive as it should be for the project being proposed would not be reasonable).

Question 45: Regarding the SOW, do TIA proposers submit a TDD instead of SOW? According to DARPA ERI, Page 3 Investments, TIA Model Companion Guide, Attachment 1 "TDD should include both program and project background sections that can be understood by a non-technical reader" - does this mean that in the TDD we have to briefly outline the project objectives, as well as our proposed solution?

Answer 45: The SOW and TDD are actually the same documents - they serve the same purpose. It's just a terminology distinction. Statement of work (SOW) is the scope document for procurement contracts, Research Description Document (RDD) is the scope document for assistance instruments (grants and cooperative agreements), and Task Description Document (TDD) is the scope document for Other Transactions. They have different names but all are the same document, generally speaking, and include basically the same content.

Question 46: According to BAA Volume-II Section 2, G. Subcontracts: "The prime contractor is responsible for compiling and providing, as part of its proposal submission to the Government, subcontractor proposals prepared at the same level of detail as that required by prime." Does this mean that the prime will include their subcontractor's SOWs and Technical Approach, and Cost Estimates, etc. as part of their proposal to the Government, listing specifically which tasks will be done by the prime and which will be done by the subcontractor(s)? In order for the prime to prepare the proposal the subcontractor has to provide the prime with its own proposal, correct?

Answer 46: The prime will only be required to include as part of its cost proposal Volume to the Government each subcontractor cost proposal and the associated SOW. There is no requirement for the prime contractor to include, as part of its proposal to the Government, technical volumes for the subcontractors - the prime contractor will speak for the entire team when it comes to the technical volume (to include the prime contractor's SOW), when it comes to the prime contractor's own cost proposal build-ups (the subcontractor cost proposals, and SOWs, are included in the prime contractors cost volume as backup documents per the solicitation. But - yes - the prime contractor must have full cost proposals from each sub in order to build their own separate submission to the Gov't. If a subcontractor were to be unwilling to provide pricing to the prime contractor that includes proprietary rates and factors then that subcontractor (as discussed in the solicitation) would then, instead, provide their cost build-ups (summary and detailed) with certain elements reflecting full wrap rates (but still having all of the other required cost elements/information - labor categories (with wrap rate), labor hours per labor category, travel, ODCs, materials, equipment, etc.

Question 47: There seem to be an overlap between SOW, Technical Approach and Schedules, and Milestones.

Question 47 (a): Is SOW essentially like Executive Summary?

Answer 47 (a): The SOW is reflective of the technical approach. The technical approach is solely for proposal purposes - it does not generally get incorporated in to the resulting contract. The SOW is specifically defining the scope of work and does, in fact, become part of the resulting contract. If selected, the Contracting Officer will pull the SOW out of the tech volume and begin fine-tuning it, with assistance from the DARPA Program Manager, Contracting Officer's Representative (COR), and Principle Investigator (PI), for incorporation into the contract. Therefore, the technical approach and the SOW do generally capture the same thing - that being the project that you are proposing but the SOW is focused on tasks and deliverables by Phase in order to later serve a part of the resulting award instrument, while the technical approach is mechanism by which the proposer fully describes the project to the Government.

Question 47 (b): According to DARPA ERI, Page 3 Investments, TIA Model Companion Guide, Attachment 1 " TDD (SOW for OTs) should include both program and project background sections, that can be understood by a non-technical reader". Does it mean that in the TDD we have to briefly outline the project objectives as well as our proposed solution?

Answer 47 (b): The TIA Model and Companion guide are solely for information purposes only – their purpose to is to bring new performers up to speed on what an OT is and how they are negotiated - they are not part of the BAA/solicitation and proposers have no obligation/requirement to comply with them as part of the proposal preparation process. Please, just follow the BAA instructions when preparing your proposal. Details discussed in the TAA Companion Guide will be added to the SOW (TDD) during negotiations if your proposal is selected (information such as program background, project background, Government Furnished Information or Property expectations, interaction with other project performers, etc.).

Question 47 (c): According to BAA the SOW has to describe tasks, durations, dependencies and deliverables, detailed description of the approach, responsible parties, completion criteria, etc. and it

seems to overlap with the information to be included in Technical Approach, and Schedules and Milestones (tasks, durations, dependencies among the tasks). How is it better to approach these sections, in order not to repeat the same information in 2 places?

Answer 47 (c): It is fully expected that there will be some overlap between sections of your technical proposal as necessary to specifically address the BAA proposal preparation requirement; however, the page count allotted to the technical volume (and to specific sections of the technical volume, if applicable) takes this into account.

Question 48: There are several considerations regarding the Fundamental Research. How do I know if the program I'm intending to proposal to is considered fundamental research?

Answer 48: Please consult page/slide 9 of M. Blackstone's Proposers Day slides which are posted on the ERI Webpage for clarification. The clarification will help but, as stated in the BAA, this is a decision ultimately left to the Contracting Officer depending on the circumstances at hand taking into consideration the nature of the research project, nature of the proposer team, type of funding being used (6.1 vs 6.2), etc.. Please also consult the BAA covering the program you are intending to propose to.

Question 49: What is your foreign partner policy? I was trying to find this information in the BAA but could not. I am intending to include a foreign organization on my team – is this permitted by the BAA?

Answer 49: This subject is covered in each of the ERI/Page 3 BAAs. Please consult "III. Eligibility of Information." There it states that "All responsible sources capable of satisfying the Government's needs may submit a proposal that shall be considered by DARPA." Also, please consult slide 6 of the Contracting Officer's Proposers Day slide deck which is available on the ERI webpage, and which states "Foreign participants/resources may participate to the extent allowed by applicable Security Regulations, Export Control Laws, Non-Disclosure Agreements, etc." Regarding the statement made in the Proposers Day slides, there is a related question captured in one of the ERI/Page 3 BAA FAQs that you might find helpful as well. Also, please be sure you are coordinating often with your organization's contracting office as you prepare your proposal – they can help you understand BAA compliance issues.

-----November 27, 2017 responses-----

Question 50: In previous FAQ responses, it is noted that CRAFT MPW TSMC 16FFC runs occur pre-Phase 1 and during Phase 1, with the last noted CRAFT shuttle run scheduled for January 2019. Given the anticipated alignment of SDH Phase 3 it is not clear how SDH Phase 3 fabrication could use the CRAFT MPW run scheduled for January 2019. Can you clarify if there will be another MPW shuttle run at a later date that may align with SDH performance deadlines and objectives?

Answer 50: No specific SDH or additional MPW shuttle run is currently planned as part of the CRAFT program. Currently only existing, planned MPW shuttle runs are available via CRAFT.

Question 51: FAQ entry 31 states that the maximum size available on the CRAFT MPW is 5mm x 5mm.

- a) If a TA1 proposer is planning a larger chip, can they purchase additional area on the MPW run?
- b) How will that additional area be priced? Can we assume the CRAFT price structure (\$50,000 per 2.5mm x 2.5mm block)? What documentation will be required in the Cost Proposal for additional area purchase on a MPW run?

Answer 51a: We have been informed that the maximum size available is 5mm x 5mm. No access to larger fabrication size has specifically been made available.

Answer 51b: No access or pricing for larger sizes have been made available. The proposer should provide the best possible cost estimates they can and detail all assumptions made for proposed device fabrication.

From question 30: Permission to use the CRAFT shuttle run will require the permission of the CRAFT PM. Once permission to be on a CRAFT MPW shuttle run is obtained, the process and contacts to use the shuttle run will be provided.

From question 31: The maximum size available on a CRAFT MPW shuttle run is 5 mm x 5 mm. The CRAFT shuttle runs will be TSMC 16nm FFC.

Question 52: The BAA indicates “proposals can make use of DARPA/MTO fab shuttle runs (USG-furnished) as part of their Phase 3 efforts.” Phase 3 for SDH starts in Dec 2020 based on the June 2018 program start date. The posted Q&A from October 11th talks about CRAFT shuttle runs in April 2018 and Jan 2019, before the Phase 3 start.

Answer 52: The CRAFT MPW runs are independent of the SDH program. The CRAFT program will not be adjusted to support the SDH schedule.

Question 53: For Phase 3 fabrication, the YouTube video mentioned TSMC fabrication runs; the BAA mentions DARPA/MTO fabrication runs. Is the government providing a different shuttle run in Phase 3 for SDH? If so, what process will the shuttles be in and what is the maximum size of the device allowed? If so, how many wafers are provided via each government shuttle run? For Phase 3, how much die area will be allowed, and is there a nominal target process technology (e.g 16nm, 10nm, or 7nm; TSMC or GF)?

Answer 53: Specific details on a government provided MPW or fabrication shuttle run to support SDH are not available. Proposers should provide their best estimates possible, and provide all assumptions made, for proposed fabrication to support Phase 3 technical activities. The proposer should provide their best estimates for the cost of a shuttle run to be borne by the performer.

Question 54: Are the POSH/IDEA shuttle runs available for SDH?

Answer 54: No specific access to POSH/IDEA shuttle runs has been made for the SDH program.

Question 55: If letters of support are to be included, what section should they be included in?

Answer 55: Letters of support should be included in Section 3, Additional Information, in the in Volume 1, Technical and Management proposal.

Question 56: How should performers interested in teaming pursue potential teaming opportunities within the ERI and SDH?

Answer 56: Proposers can pursue teaming via the ERI site: <http://www.cvent.com/events/eri-teaming-partner-list/event-summary-b69b9f99098f48669d6e9eea8a382497.aspx>

Question 57: For travel cost estimating, what is the planned SDH program review frequency and location?

Answer 57: No specific SDH program review frequency or location/s were provided in the SDH BAA. As such, each proposer should price program review meetings based on their own assumptions. The proposer should clearly define the review meetings being assumed for pricing purposes. An example could be something like quarterly reviews with bi-annual reviews held at DARPA and with performer site visits held for the other two quarterly reviews.