

DARPA-BAA-16-52

Hierarchical Identify Verify Exploit (HIVE) Frequently Asked Questions (FAQ)

August 18, 2016

DARPA-BAA-16-52

Hierarchical Identify Verify Exploit (HIVE) Frequently Asked Questions (FAQ)

August 18, 2016

TA1

1. Q: The BAA notes scaling to 20W, is that just for the graph accelerator chip?

A: Power scaling is important for applications that TA3 identifies to be on the network edge, and in those cases the chip itself is desired to consume no more than 20W.

2. Q: Is the bandwidth goal (1TB/s) in terms of compressed data or the equivalent uncompressed data?

A: An IO bandwidth of 1TB/s is an estimate of IO performance necessary to achieve orders of magnitude acceleration of graph applications. Whether to compress (or not) is up to the performer.

3. Q: Is the HIVE chip envisioned as a pure digital design? Is non-digital processing anticipated to be in conformance with the BAA, at least for TA1? If so, will some of the primitives be evaluated for precision requirements?

A: Once the primitives are decided in phase 1, TA1 is welcome to solve/address them in the best manner possible.

4. Q: Is this seen as a standalone system or could it be attached to a conventional system?

A: The proposed hardware solution could be a co-processor attached to a conventional system.

5. Q: Is fabrication in overseas foundries a problem in HIVE?

A: No, fabrication in overseas foundries is acceptable in HIVE.

6. Q: Will a complete GDS layout specification of the chip be needed by the end of phase 2 or can that spill over to phase 3?

A: To meet the program timeline, performers should expect to deliver the GDS2 specification at the end of phase 2.

7. Q: How important is it to consider very low power solutions, for portable & in-field deployment?

A: Low power solutions (up to 20W chip power consumption) are application-dependent, though they are expected to be important for edge applications.

8. Q: Will the delivered phase 3 HIVE system be an accelerator connected/hosted by a CPU system? If so, does interconnect to the host system come into scope for TA1?

A: The I/O connection system is in scope for TA1, and co-processor designs are in scope as well.

9. Q: By which metric must the graph processor achieve a 1000x improvement?

A: The 1000x improvement is meant to be relative to power efficiency. However, the scale at which the 1000x improvement must be demonstrated is not defined. Solutions that can demonstrate 1000x performance improvement across the problem sizes of interest are preferred. However, solutions that only achieve 1000x improvement for a portion of the problem sizes of interest will still be considered.

10. Q: Are the HW performers for TA1 creating HW architectures exclusively designed to accelerate sparse computation and data movement?

A: No, the goal of TA1 is to create HW architectures and a graph processor that accelerate DoD/US government mission needs in graph computation. Graph computation (both streaming and static) is characterized by sparse computation and heavy data movement of a random nature, but can include some stages in the algorithm that involve significant dense computation and data movement. Thus, HW solutions will likely need to efficiently handle dense workloads as well as sparse workloads. The government furnished challenge problems are designed to characterize the computation and data movement characteristics of DoD and US government graph challenges.

11. Q: Are we allowed to make changes to DRAM design? If so, do they need to be compatible with DDR standards?

A: Yes, it is acceptable to make changes to DRAM design as part of a proposed solution. No, compatibility with DDR standards is not required.

12. Q: Can you clarify the 1TB/s BW per processor to memory requirement, as this seems to be on commercial roadmaps today, what size memory would be connected here?

A: This will be determined based on evaluation of the applications in phase 1.

13. Q: The BAA specifically calls out micro-code as a development expectation for TA1. The term "micro-code" has a very specific meaning in the industry. Is there an expectation that micro-code models will be exposed to other parties, or that micro-code needs to be used at all? What is the specific set of expectations around micro-code, particularly for those platforms that do not use micro-code?

A: "Microcode" was used as a generic term in this BAA. Anticipated information needed to be released to other performers is ISA, performance data, and details of the ISA sufficient to model behavior and performance.

14. Q: While the BAA calls out desired peak bandwidths to memory and between nodes, it does not discuss per-node peak computational capability in FLOPS or IOPS, nor memory capacity. What are the target goals for these?

A: Success is evaluated based on overall application acceleration, not component performance

TA2

1. Q: Will the HIVE graph analytics processor require a front-end with a particular operating system and programming environment w/ debugging and performance analysis tools for graph applications development?

A: The HIVE anticipates using standard environments such as Python and C along with associated graph libraries. No restrictions on operating system are anticipated in this program.

TA3

1. Q: Does TA3 have to address all 5 problem classes?

A: It is anticipated that a TA3 proposal will identify DoD usage scenarios from across the five areas with both static and streaming cases for each resulting in 10 application areas.

2. Q: Will TA3 teams be able to do research into making synthetic datasets more realistic?

A: The HIVE program is focused on realistic workloads from TA3 applications and not synthetic benchmarks. Any synthetic data must accurately reflect real workloads. It is the role of TA3 performers to determine how to develop unclassified surrogate data sets that preserve useful properties of the applications.

3. Q: What is the expected role/level of resiliency expected, especially for streaming cases where data is persistent?

A: The level of resilience will be determined by the application areas identified by TA3 performers.

4. Q: Are cyber scenarios important to the HIVE program or should TA3 look at different scenarios?

A: Cyber scenarios are viable application areas in HIVE, for example the Anomaly Detection use case given in the Proposer's Day briefing is an example from the cyber domain.

INTER-TA

1. Q: GPUs today can already provide a compute limited 16 node system with BW's that will soon approach 1TB/s, can you clarify what 1000x performance means for this system?

A: Performance improvements are with respect to the applications identified by TA3 over current CPU/GPU solutions. HIVE seeks to demonstrate approximately a 1000x improvement in power efficiency over current GPU/CPU solutions.

2. Q: How does a 100x improvement in evaluation framework relate to the 1000x goals in the BAA?

A: The 100x improvement is based on the direct (non-optimized) use of the TA2 tools on TA1 hardware, while the 1000x improvement goal should be attainable by optimized use of the tools by expert software developers.

3. Q: Please clarify importance of HIVE systems also having to perform well on traditional (dense) applications compared to conventional architectures (GPUs, etc.).

A: HIVE systems will be evaluated based on their performance across the five applications areas listed in the BAA, which will have a mixture of sparse and dense applications.

4. Q: What will be included with the government furnished challenge questions?

A: Two challenge problems will be provided to guide TA1 and TA2 development efforts and judge performance. One problem will represent a static graph problem, the other a streaming or dynamically updating graph problem. Each problem will consist of:

- a. A problem statement
- b. A data set (or tool that generates the data set) scalable from toy problems (millions of edges) to data center scale problem (trillions of edges)
- c. The computation that must be performed to solve the problem. The complexity of the computation will be scalable in some manner (ex: resolution or accuracy)
- d. The desired solution

5. Q: The BAA mentions baselining performance on a GPU. Is there a particular GPU that should be used?

A: No, however it is envisioned that proposers will compare to recent hardware so as to provide a meaningful baseline of comparison.

6. Q: What is the range of Graph sizes that are of interest? A 16 node system is still relatively small, to what size do you anticipate a fully scaled system of interest to be?

A: The goal of the 16 node system is to provide evidence that the system could be scaled in size to handle problems of at least 1 trillion edges (to as high as 100 trillion edges). The overall range of problems scales of interest are 1 million edges up to a trillion edges (with a stretch goal of a 100 trillion edges).

7. Q: BAA states that size, weight and power constraints are critical for TA1 performers to understand but it appears that this information is not provided until the end of phase 2 by TA3 performers (which is too late). Will initial guidance be given?

A: Size, weight and power constraints (SWAP) are important for a subset of applications and not all applications will be driven by SWAP constraints. Though the constraints are not finalized until the end of phase 2, we expect ongoing communication to occur during phase 2.

8. Q: The BAA briefly mentions Causal Modeling, but previously there seemed to be much more interest in dense probabilistic graph models (Bayesian Nets, Belief Networks, Ontological reasoning etc.), is that still the case? (And if so, is this only for TA3 performers?)

A: This is only one of the five general graph analytic use cases. HW and SW solutions should be broadly applicable to all five graph analytic use cases.

9. Q: In the TA3 phase 3, it calls for 16 nodes all interconnected. I am confused by the nomenclature here.
- Is the intention to be 16 systems which can interoperate with each other?
 - Or is it to have 16 of the final chips interconnected into a single graph processor to allow a larger data set to be evaluated?
 - Or is it simply to have 16 of these devices which can be used on a single training session with 16 operators?

A: TA3 is responsible for testing and evaluating application performance of the 16 node system in phase 3. TA1 is responsible for fabricating it. The intention is to have 16 nodes (each containing a graph processor) interconnected working on a single graph problem (generally of larger size than a single node could handle). This is similar to a cluster used in HPC computing.

PROCEDURAL

1. Q: Can you please clarify what is meant by "Proposer's reference number"? Page 28 of the BAA.

A: The "Proposer's reference number" is normally used as an internal control number by larger companies/organizations that send/receive large amounts of correspondence.

2. Q: What is meant by the term "lead organization"?

A: The "lead organization" refers to the "Prime Organization" that will be responsible for submitting the entire proposal. This is different from the sub-contractor, who may do work or provide support but are not the lead organization on the effort. The prime contractor will typically manage the subcontractor(s).

3. Q: What is meant by "Other team members"?

A: This refers to other organizations (subcontractors, consultants, etc.) that will contribute to the proposed effort/work.

4. Q: Will the Proposer Charts be provided?

A: Yes, they are posted on the DARPA Opportunities website under the DARPA-BAA-16-52 section: <http://www.darpa.mil/work-with-us/opportunities?PP=1>.

5. Q: Is there an Excel template for the cost buildup?

A: No, there is not.

6. Q: If TA2 performers must open source the valuable and strategic SW technology, will TA1 performers be required to open source the HW design completely? Do TA1 proposals gain advantage by open sourcing the HW in their plan? If HW is not to be open sourced, why?

A: An open source HW design is desirable but not required in this program; data rights will be negotiated on a per-performer basis. Neither TA1 nor TA2 performers are required to open source their design.

7. Q: Will you be posting responses to questions before the final deadline in case we realize we need to rephrase them if they are not understood as we intended?

A: Questions will be answered as received until the October 5th 2016 FAQ submission deadline listed in the BAA. Responses can be expected within 24-48 hours after question submission.

8. Q: Can we submit to TA3 if not all PIs have clearance? As an example, can university participants with no clearance be included as subcontractors for a TA3 prime?

A: It is permissible for a TA3 prime to work with subcontractors who do not hold security clearances; however, TA3 primes will need to be able to handle classified information and are still required to follow all applicable rules and regulations pertaining to the access and control of classified information.

9. Q: Could you clarify/explain the expectations for use of classified data in TA3?

A: TA3 requests DoD-relevant applications be identified from across the five areas listed in the BAA, to ensure realistic applications are developed, real DoD system data will be used.

10. Q: Is an abstract required to participate in HIVE? What is the expected government response time to abstract submissions?

A: As stated in the BAA, abstract submission is not required and government response on the abstract will be provided 15 days after submittal date.

11. Q: For TA1: can the abstract/proposal be for phase 1 only, or does it need to cover all 3 phases?

A: The government is seeking complete solutions across all 3 phases of the program in any of the TAs.

12. Q: When are deliverables/milestones expected?

A: Major deliverables listed in the BAA are expected to be delivered by the end of the phase in which they are described. Please see the BAA for timing information on reporting etc.

13. Q: Who pays foundry costs?

A: Foundry costs need to be included in the proposal. Ultimately, DARPA pays the costs by funding the performer during phase 3, but the costs must have been priced out and included in the proposal.