

Dispersed Computing

Stuart Wagner
Program Manager
Information Innovation Office (I2O)
DARPA

June 30, 2016





Agenda

TIME	EVENT	SPEAKER
8:30 AM - 9:30 AM	CHECK-IN / REGISTRATION	
9:30 AM - 9:35 AM	DARPA SECURITY BRIEFING	ALICIA HIRALDO
9:35 AM - 9:40 AM	WELCOMING COMMENTS	STU WAGNER
9:40 AM - 10:00 AM	CONTRACTING & BAA BRIEFING	MARK JONES
10:00 AM - 10:20 AM	DISPERSED COMPUTING PROGRAM BRIEFING	STU WAGNER
10:20 AM - 11:00 AM	BREAK	
11:00 PM - 11:30 PM	DISPERSED COMPUTING PROGRAM Q&A	STU WAGNER

Dispersed Computing

Mark Jones
Contracting Officer
Contracts Management Office (CMO)
DARPA

30 June 2016





Disclaimer

**If the BAA contradicts any information in these slides,
the BAA takes precedence.**



BAA Overview

BAA follows procedures in accordance with FAR 35.016.

BAA Amendment 1 (as well as any future amendments) is posted on FEDBIZOPPS at www.fbo.gov and Grants.gov at www.grants.gov

Proposals due by 12:00 noon ET on September 7, 2016

BAA covers all info needed to submit proposals. Follow instructions for proposal preparation and submittal.



Potential Award Information

DARPA anticipates multiple awards for Technical Areas (TA) 1 and 2. TA 3 will be an option occurring in Program Phase 2.

Program structured into a four-year program comprising two 24-month phases.

Single proposal may address either TA1 or TA2, but not both. Can submit multiple proposals.

Proposal for either TA must contain 2 options – TA3 Phase 2 effort and field exercises effort.

Awards may be Procurement Contracts, Cooperative Agreements or Other Transaction Agreements – Grants do not appear to be suitable award mechanism.



BAA Eligibility

All interested/qualified sources may respond subject to the parameters outlined in the BAA.

Foreign organization/individuals – check all applicable Security Regulations, Export Control Laws, Non-Disclosure Agreements, and any applicable governing statutes.

FFRDCs/UARCs and Government entities

- Subject to applicable direct competition limitations
- Must clearly demonstrate eligibility per BAA

Real and/or Perceived Conflicts of Interest

- Identify any conflict
- Include mitigation plan



Proposal Preparation Information

Proposals consist of two volumes – Technical and Cost

Volume 1 - Technical and Management:

- Maximum 32 page limit
- Includes mandatory Appendix A – does not count towards page limit.
- Includes optional Appendix B – does not count towards page limit

Volume 2 – Cost - No page limit

The BAA will describe the necessary information to address in each volume:

- Make sure to include every section identified.
- If a section does not apply – put “None” (e.g., Animal Use – None, OCI - None)
- Include a working/unprotected spreadsheet as part of your Cost Volume submission.
- Review individual TA descriptions, IP rights, and the deliverables section for submittal information



Statement of Work (SOW) Preparation Tips

Write a SOW as if it were an attachment to an award

- Don't use proposal language (e.g. we propose to do . . .)
- Break out work between any phases/time periods identified in the BAA
- Succinctly and clearly define tasks & subtasks
- Mark TA3/Field Exercise efforts as Optional
- Identify measurable milestones and define deliverables
- Do not include any proprietary information!

NOTE: For cooperative agreements: SOW = RDD or Research Description Document



Proposal Preparation Tips

Substantial Time Commitment

- Propose substantial time commitment for key personnel
- If PI is committed to multiple projects, consider co-PI(s) or document mitigation efforts to make up for PI's lack of commitment to effort

Risk – Do not be afraid to address Risk in Technical Volume

- Identify risk(s) to show an understanding of technical challenge(s)
- Discuss potential mitigation plans / alternative directions

Awareness of New Terms & Conditions

- DFARS Clause 252.203-7997 Prohibition on Contracting with Entities that Require Certain Internal Confidentiality Agreements
- DFARS Clause 252.204-7012 Safeguarding of Covered Defense Information and Cyber Incident Reporting



Proposal Preparation Tips Cont'd – Intellectual Property Rights

Government desires, at a minimum, **Government Purpose Rights** for any proposed noncommercial software and technical data. (See DFARS 227 for Patent, Data, and Copyrights)

Since DC will emphasize creating and leveraging open architecture technology, IP rights and software licenses asserted by proposers are strongly encouraged to be aligned with this goal.

Data Rights Assertions – IF asserting **less than Unlimited Rights**:

- Provide and justify basis of assertions
- Explain how the Government will be able to reach its program goals (including transition) within the proprietary model offered; and
- Provide possible nonproprietary alternatives

IF proposed solution utilizes commercial IP – submit copies of license with proposal



Items To Note

Work expected to be fundamental research

Understand and comply with SAM, E-verify, FAPIIS, i-Edison and WAWF.
Links are found in the BAA.

For planning purposes - anticipated Program Start Date is February 1, 2017

Subcontracting Issues

- Non-Small Businesses: Subcontracting Plans required for FAR-based contracts expected to exceed the applicable threshold.
- Subcontractor cost - Proposals must include, at a minimum, a non-proprietary, subcontractor proposal for EACH subcontractor. Include any internal price/cost analysis of subcontract value in proposal.
- If utilizing FFRDC/UARC, Government entity, or a foreign-owned firm as a subcontractor, submit their required eligibility information, as applicable.



Items To Note Continued

Proposals must be valid for a minimum of 120 days

If a prospective proposer believes a conflict of interest exists or has a question on what constitutes a conflict - promptly raise the issue with DARPA

Document files must be in .pdf, .odx, .doc, .docx, .xls, and/or .xlsx formats.

Submissions must be written in English.

Proposers are not required to hold or obtain security clearances, but may be required to have personnel with a SECRET clearance to participate in field exercises. Proposers must provide a CAGE code and security POCs, as well as the name of personnel that would attend any field exercise in the proposal.



Proposal Submission

All submissions will be completely UNCLASSIFIED.

Submit FAR based contract and OT proposals via DARPA's web-based upload system for unclassified portion of proposal. Submission must be in a single zip file not exceeding 50 MB.

Submit cooperative agreement proposals via Grants.gov.

Follow submission procedures outlined in the BAA. DO NOT submit proposals except as outlined in the BAA (e.g., email/fax submissions will NOT be accepted).

DO NOT wait until the last minute to submit proposals – the submission deadlines as outlined in the BAA will be strictly enforced!

DO NOT forget to FINALIZE your proposal submission in the submission tool!



Evaluation / Award

No common Statement of Work - Proposal evaluated on individual merit and relevance as it relates to the stated research goals/objectives

Evaluation Criteria (listed in descending order of importance) are: (a) Overall Scientific and Technical Merit; (b) Potential Contribution and Relevance to the DARPA Mission; and (c) Cost Realism.

Evaluation done by scientific/technical review process. DARPA SETAs with NDAs may assist in process.

Government reserves the right to select for award all, some, or none of the proposals received, to award portions of a proposal, and to award with or without discussions.



Communication

Prior to Receipt of Proposals – No restrictions, however Gov't (PM/PCO) shall not dictate solutions or transfer technology. Unclassified FAQs will be periodically posted to this BAA's DARPA web page.

After Receipt of Proposals – Prior to Selection: Limited to PCO – typical communication to address proposal clarifications.

After Selection/Prior to Award: Communications range from technical clarifications/revisions to formal cost negotiations. May involve technical as well as contracting staff.

Informal feedback for proposals not selected for funding may be provided once the selection(s), if any, are made.

Only a duly authorized Contracting Officer may obligate the Government



Take Away

Submit proposals before the due date/time - Do NOT wait until the last minute to submit.

Read and understand the BAA - Follow the BAA when preparing proposals.

Be familiar with Government IP terms from the DFARS Part 227.

Submit working/unprotected spreadsheet(s).

The Contracting Officer is the only Government official authorized to obligate the Government.

Dispersed Computing

Stuart Wagner
Program Manager
Information Innovation Office (I2O)
DARPA

June 30, 2016





What Are We Trying To Do?

Develop algorithms and protocols that harness physically dispersed computing capabilities to boost application and network transport performance by orders of magnitude.

- Enable the network to become the cloud
- Automate strategic movement of data to code
- Catalyze a migration away from problematic aspects of current Internet architecture

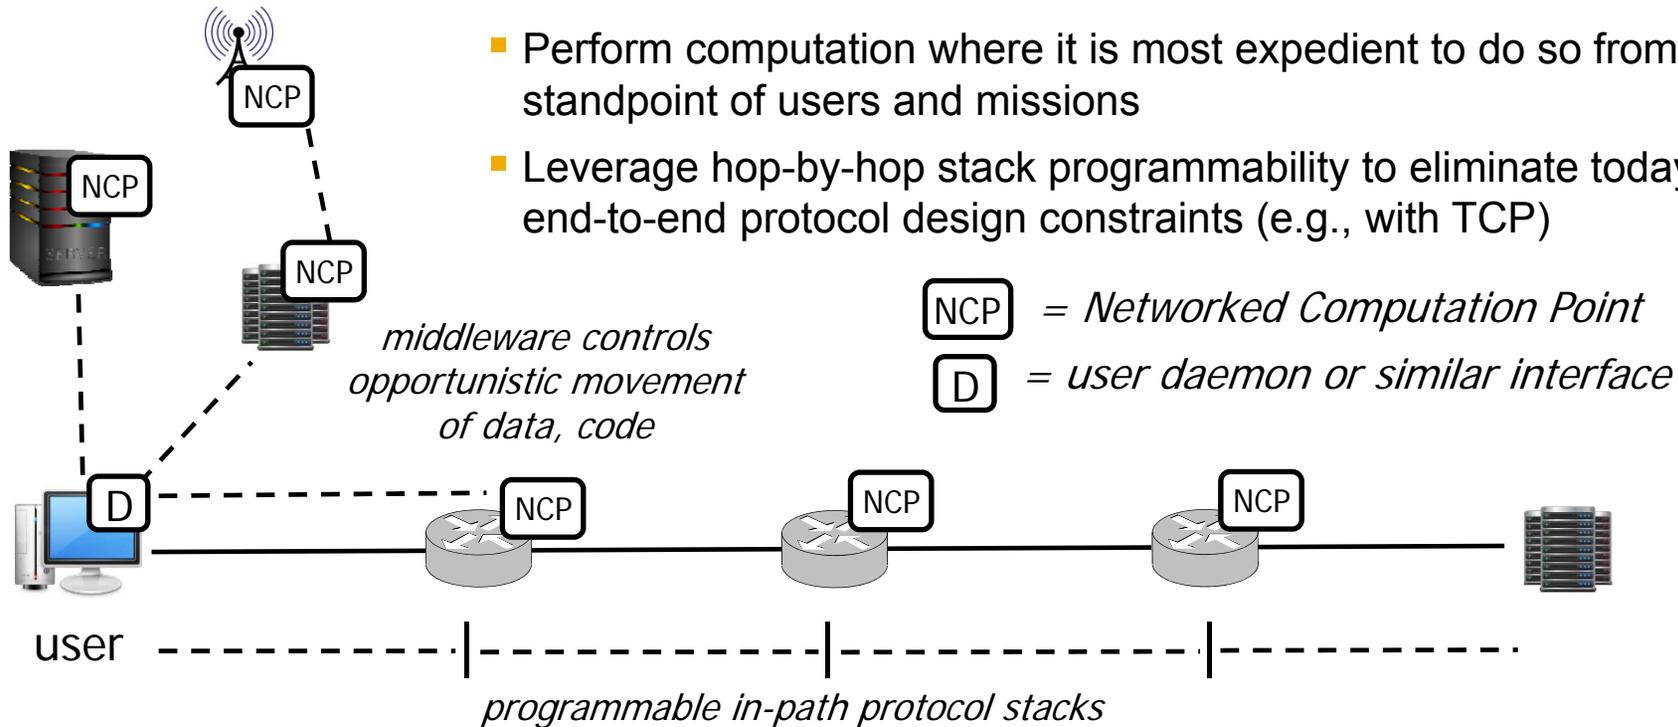


Limits of Current Art

Criterion	Current Art	Dispersed Computing
Code and data movement consistent with DoD mission success	<ul style="list-style-type: none">• Latency can far exceed mission or app tolerance• High network loads (backhaul)• Manual control	<ul style="list-style-type: none">• Latency reduction 10x – 100x• 90% reduction in network load• Automated control
Computational Mission Awareness (priority, deadlines)	<ul style="list-style-type: none">• Minimal to no prioritization• No deadline awareness• Centralized control	<ul style="list-style-type: none">• Prioritized tasking• Deadline-aware resource usage• Distributed control
Network Awareness of current research proposals	<ul style="list-style-type: none">• Algorithms assume static, favorable network conditions	<ul style="list-style-type: none">• Algorithms respond rapidly and efficiently to network dynamics
Efficient, mission-responsive network transport	<ul style="list-style-type: none">• Transport and app-layer logic confined to end points• Difficult diagnosis and mitigation of path degradations• Protocol logic statically defined and positioned	<ul style="list-style-type: none">• Transport and app-layer logic within network nodes, supporting improved app performance and diagnostics• Logic dynamically formulated and instantiated



The Dispersed Computing Vision



- Perform computation where it is most expedient to do so from the standpoint of users and missions
- Leverage hop-by-hop stack programmability to eliminate today's end-to-end protocol design constraints (e.g., with TCP)

Networked Computation Point Features:

- Programmable execution environment incorporating Dispersed Computing middleware
- Enables secure, mission-responsive resource sharing and code ↔ data movement
- Users (or their proxies) interface with NCPs via daemons or similar transparent software on end hosts
- Examples: programmable network elements, sensors with embedded programmable processors, micro/nanoclusters, smart phones (if their resources are accessible by other end points)



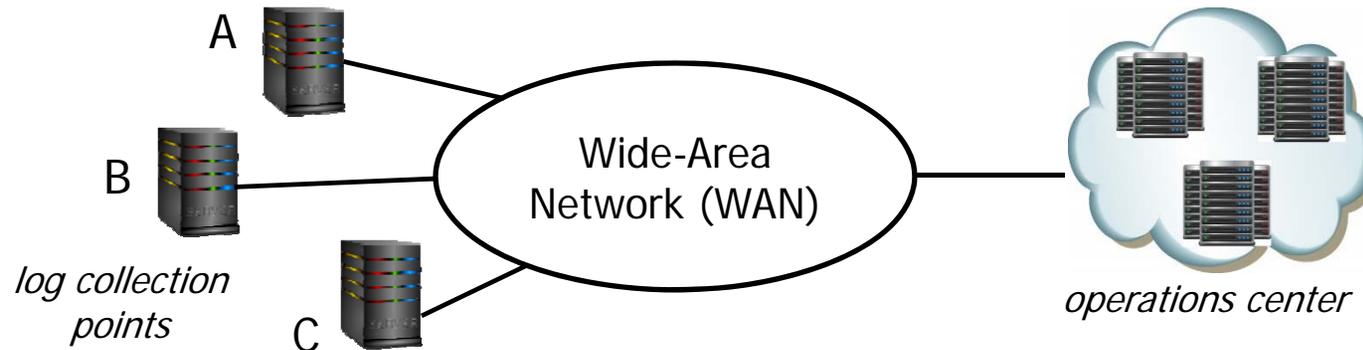
Dispersed Computing Program Structure

TA	Focus	Objectives and Outputs
1	Algorithms for Dispersed Mission-Aware Computation	<ul style="list-style-type: none">• Develop distributed algorithms, in software, that jointly optimize movement of data and placement of tasks across available NCPs• Instantiate algorithms in NCPs and users' end hosts• Test the algorithms at scale; demonstrate responsiveness to varying network conditions and mission constraints (prioritization, deadlines, security restrictions)
2	Programmable Nodes and Protocol Stacks	<ul style="list-style-type: none">• Develop end-host and in-path stacks tailored to applications and dynamic path properties• Enable remote, on-the-fly instantiation of stacks on NCPs• Verify fairness and stability across flows and users without centralized control
3	Integration and Demonstration (Phase 2 Option)	<ul style="list-style-type: none">• Integrate synergistic outputs across TA1 and TA2 performers during second half of program• Demonstrate integrated systems to stakeholders

Note: Development and demonstration of novel applications is in scope for all TAs but is not the main focus of any TA



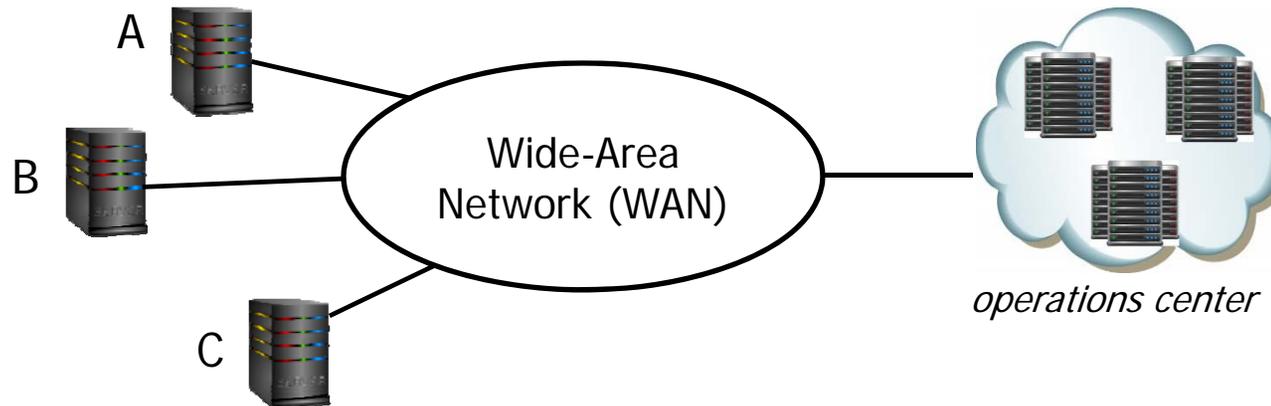
An Example Scenario for Dispersed Computing



- Points A, B, C and the operations center are all interconnected over the WAN but with different available bandwidths
- Points A, B, and C have computing capabilities of different capacities and loads; machines at A, B and C and at selected points in the WAN are NCP-enabled
- Administrator at operations center has a daemon on her host that can interact with NCPs
- A network attack is suspected; goal is to identify and characterize the attack ASAP
- Network log files $\{L_A, L_B, L_C\}$ contain critical traffic data but are each over 1 TB
- The necessary analytics require sequential execution of three computing operations $\{O_1, O_2, O_3\}$ to each of the log files
- Each operation produces a data reduction of 10x – 100x
- Overall time to complete the analytics = time to move data + time to execute the operations + time to transmit the resulting outputs to the operations center



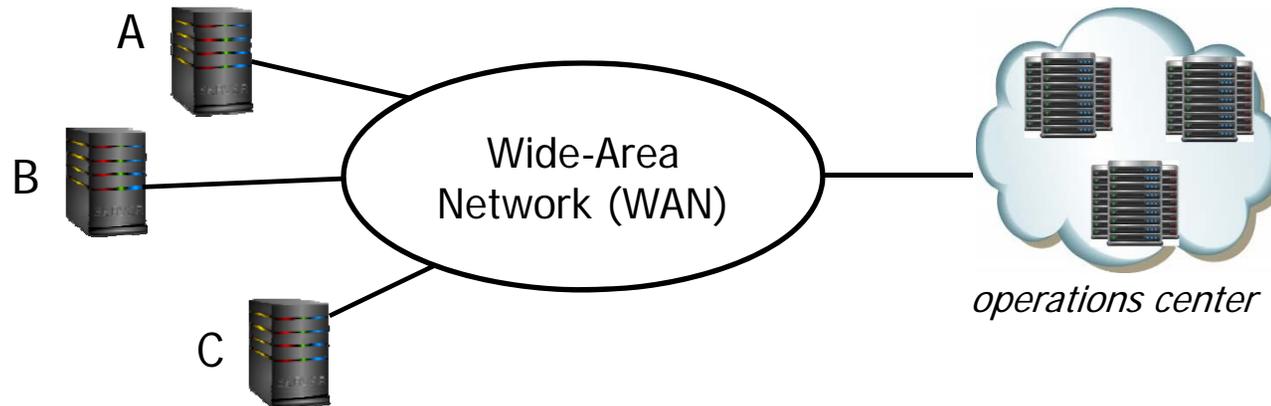
Example Part 2



1. Daemon at operations center polls NCPs at A, B and C to determine available computing power and bandwidths of WAN connections; task priority specified as high
 - Point C determined to have substantial computing power but WAN connectivity to operations center, A, and B is severely limited due to attack
 - Points A and B determined to have lower available computing power than C, but less-degraded WAN connectivity
2. User daemon optimizes over alternative strategies for code and data movement
 - Code for O_1 and O_2 pushed to Point C; code for O_1 pushed to A and B
 - NCPs at A, B, C execute these operations on the log files at their locations



Example Part 3



3. In parallel with step (2), user daemon at operations center polls NCPs within WAN to obtain more-specific measurements of network state
4. With data from (3), user daemon tasks NCP-enabled computing resources to synthesize transport protocol(s) optimized for maximum goodput, for the observed network conditions
5. Synthesized protocols pushed to WAN NCPs in anticipation of transferring data from points A, B, and C
6. Outputs of O1 (and O2 at point C) are transported to operations center, which executes O2 and O3 to complete the analytics; daemon tasks NCP-enabled computing resources, if needed, to complete this final step



TA1 Challenges and Issues

- Identify suitable objectives and objective functions (latency, network loads, power dissipation, accuracy, robustness to node loss)
- Perform timely optimization over a vast space of alternatives, given limited, highly distributed, heterogeneous computing resources
- Scale: How many nodes, users and tasks can be handled?
- Effectively handle competing tasks and users of differing priorities and deadlines, in a distributed fashion
- Distribute the decision-making
- Incorporate network state uncertainties and variabilities into decision-making
- Note: focusing on specific applications and scenarios may be useful and valuable for demo purposes, but solutions that have potentially broad applicability are preferred



TA2 Challenges and Issues

- Tailor protocol logic to the needs of the application and mission, and to the state of the network
- Perform effective protocol and node programming in presence of network dynamics
- Make holistically favorable programming decisions across all NCPs on the path (as opposed to point optimizations)
 - Preferably without centralized orchestration of the programming
- Avoid onerous processing and measurement overheads (probing) by NCPs that might degrade node throughput
- Ensure node-safe and network-safe programming of protocol stacks
 - Network-safe: preservation of fairness and stability across flows
- Note: be specific about the context of your proposed solution and planned demonstrations – applications, wired vs. wireless, scale, etc.



Example Dispersed Computing Program Metrics

	Metric	Phase 1 Goal	Phase 2 Goal
Computing	Architectural Scale	100 NCPs, 500 simultaneous users	1500 NCPs, >1000 simultaneous users
	Bandwidth Consumption	50% reduction w/r to backhaul approach	90% reduction w/r to backhaul approach
	Mission-Aware Computation	20x avg., 50x worst-case improvement in aggregate utility	50x avg., 100x worst-case improvement in aggregate utility, inc. deadlines
Protocols	App Performance Improvement	20x aggregate gain in application utility	50x aggregate gain in application utility, or within 10% of path capacity
	Response Time to Network Events	1 sec	100 msec

Impact: greatly enhanced user experience (latency, reliability) geared to mission needs (priorities, deadlines), demonstrated at scale

Project-specific metrics encouraged (e.g., power consumption)



Program Approach, Schedule and Milestones

	Phase 1										Phase 2													
Fiscal Year	17				18						19				20						21			
Program Month	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48
Kickoff and PI Meetings	◆		◆			◆			◆			◆			◆			◆			◆			◆
Proposed Test Plans		◆					◆						◆						◆					
Project Demonstrations				◆				◆				◆				◆				◆				◆
Integration Plan										◆						◆					◆			
Multi-Project Demos																◆					◆			◆
Field Exercise Proposal						◆													◆					
Range Testing									◆													◆		
Field Exercises													◆											◆

- Iterative development, prototyping and experimentation within and across TAs addressing program-level and project-specific metrics
- Base program (TA1 and TA2) is a four-year effort. TA3 is an optional effort of up to two years, during Phase 2.
- Contexts (hardware platforms, network scenarios, applications) will vary across projects
- Range testing and field exercises provide controlled environment for stakeholder demos and collection of feedback



Important Points From BAA

- Areas out of scope for Dispersed Computing:
 - Development of hardware to support NCPs
 - Evolutionary extensions to current art
 - Node programming via executable code encapsulated within user data packets
- Proposals must clearly state the context for the proposed technical approach - what kind of networks, applications, and scale are to be addressed?
- Proposals must define metrics relevant to their approaches and contexts, and describe the envisioned experimental plan and environment
 - Should clearly support the main elements of the technical approach
 - Strong proposals will specify target values and how they will be met
- Individual TAs have additional specific requirements – see BAA
- Heed the BAA evaluation criteria including (but not limited to):
 - Do you provide evidence that the approach is feasible and achievable (technical details, substantiated claims, clear description of how and why it works)?
 - Have you identified technical risks (some aspect of your approach that may not work) and mitigation plans?
 - Does the approach produce a revolutionary, high-payoff result? Why is it innovative?
 - Are the costs realistic with respect to the proposed effort?



Break

- The Dispersed Computing Program Q&A session will begin at 1100.



www.darpa.mil