

# FastNICs

---

Jonathan M. Smith

Proposers' Day

7/10/2019





## Goal

---

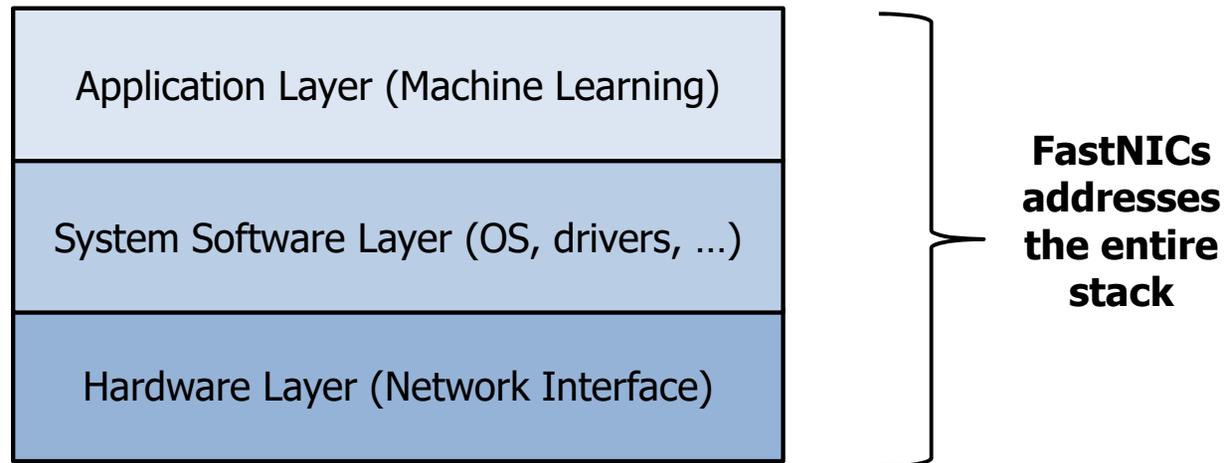
Speed up distributed applications by  
100x by creating new networking approaches



## The Stack: Layered Architectures

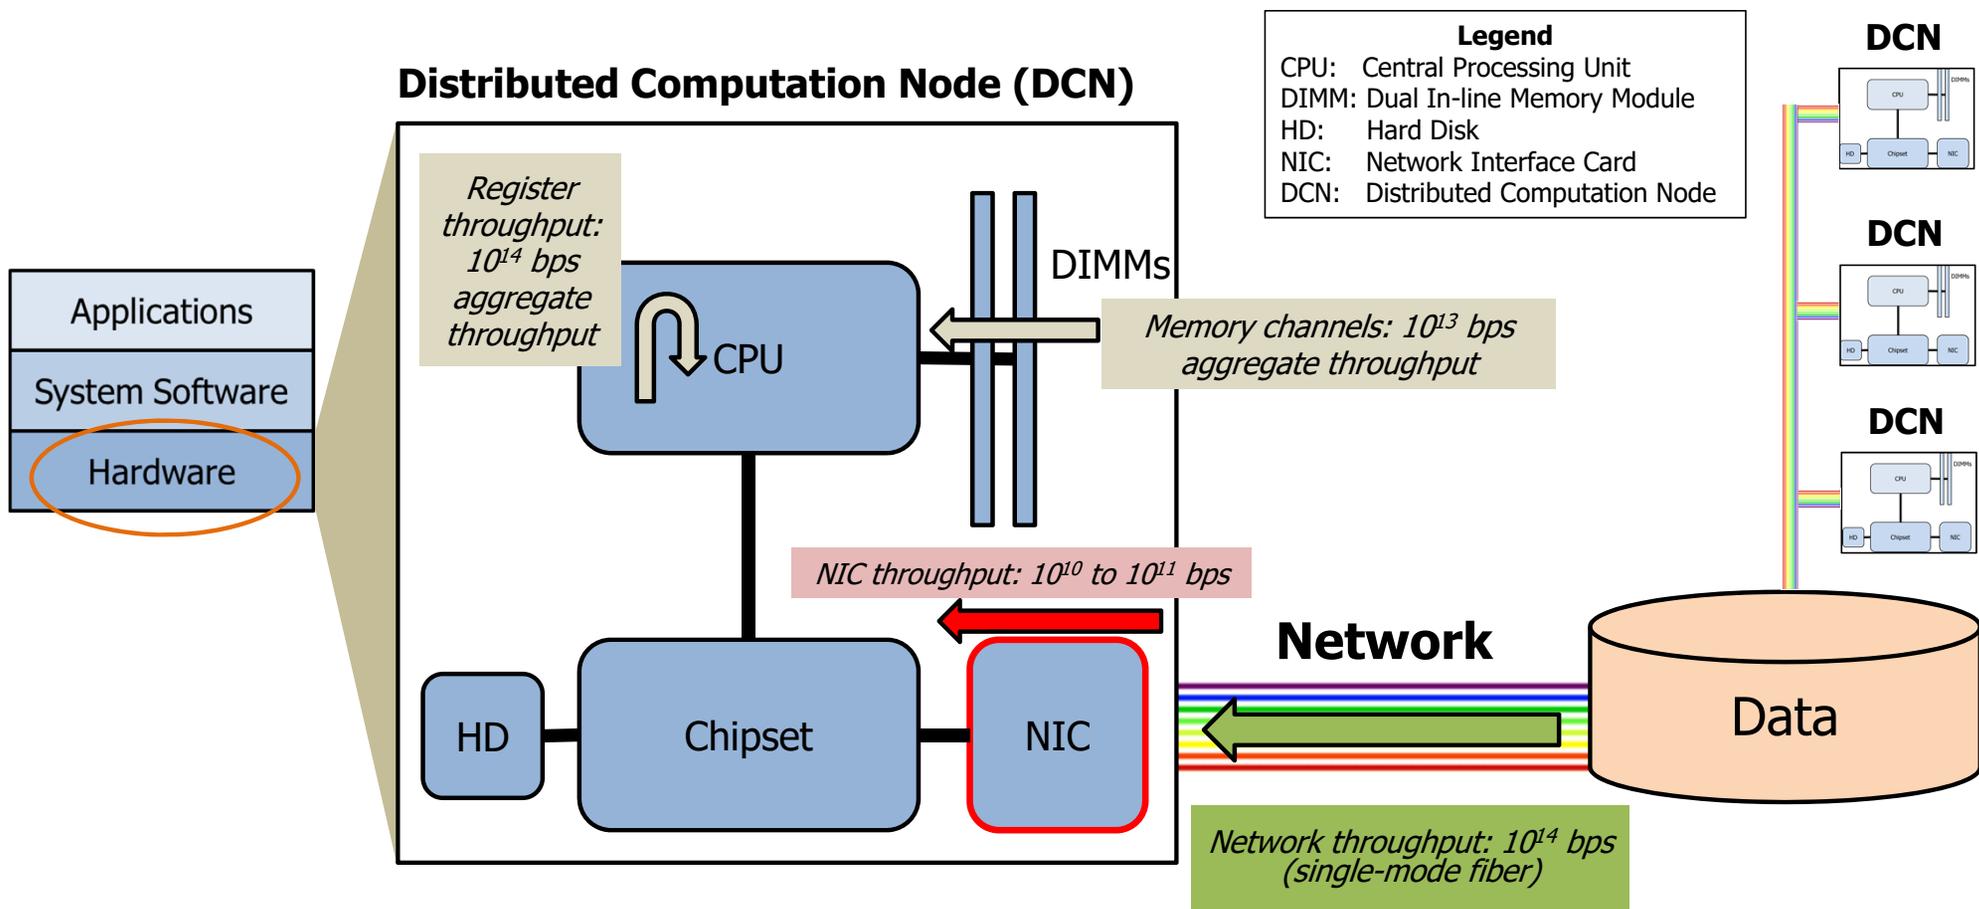
---

Higher layers are built on top of lower layers, creating strong inter-layer dependencies and assumptions



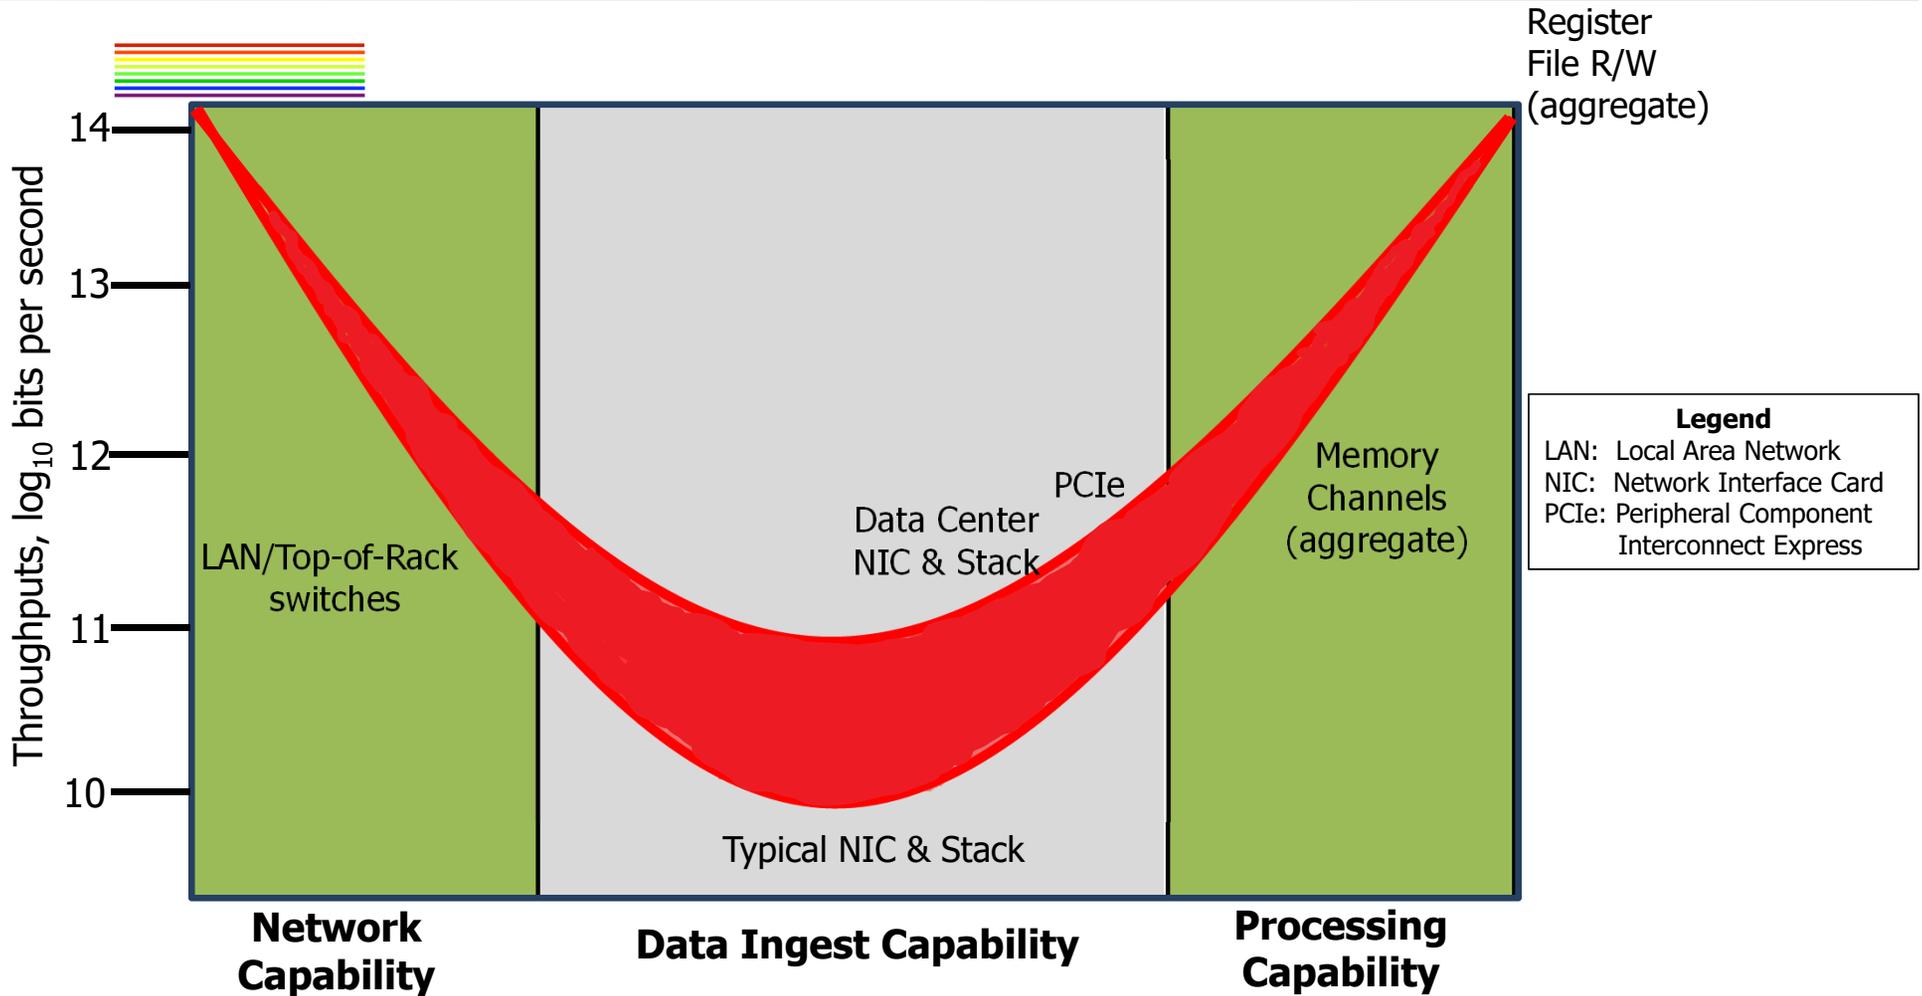


# Hardware Layer Problem: Data Ingest Bottleneck





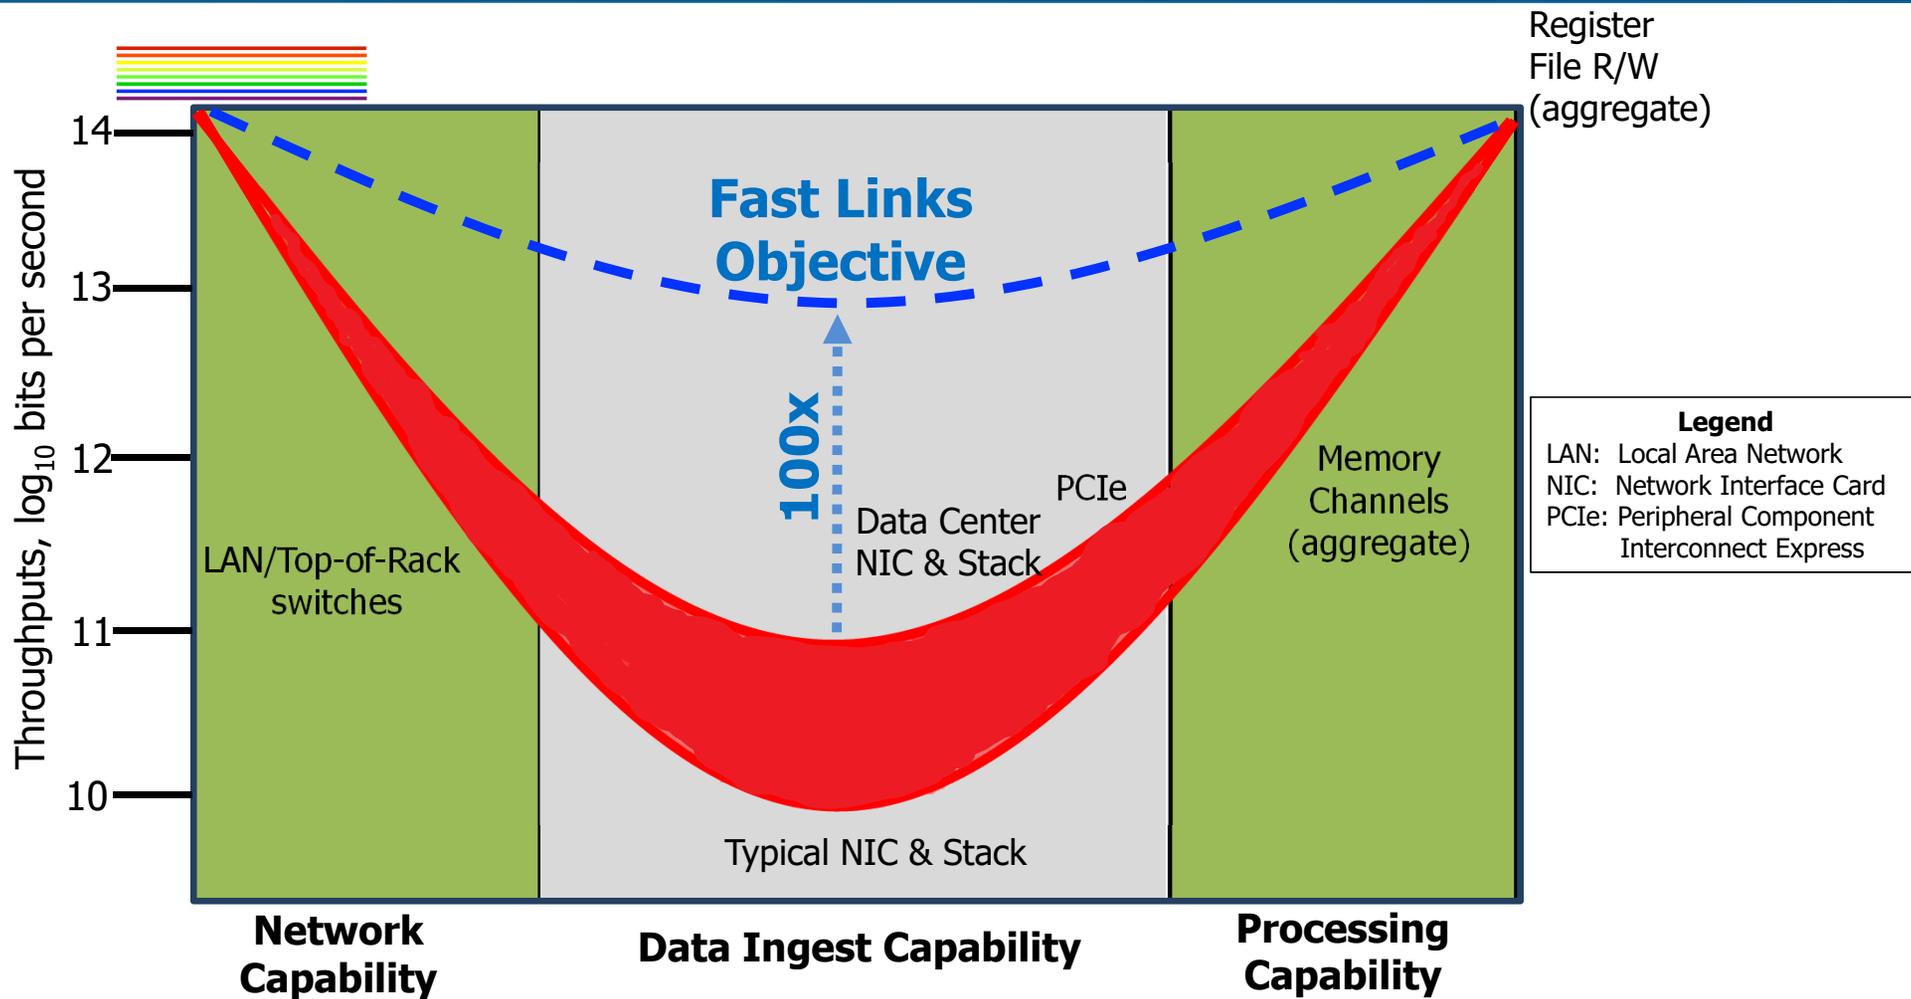
# Another Perspective: Current State-of-the-Art



Distribution Statement A: Approved for public release, distribution is unlimited.



# Another Perspective: Current State-of-the-Art



Distribution Statement A: Approved for public release, distribution is unlimited.

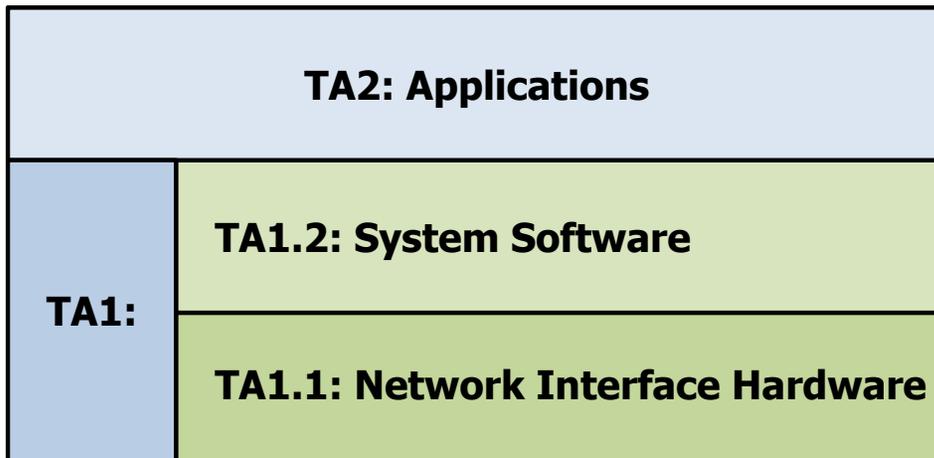


## Program Structure

---

### 48-month program

- Phase 1: 24 months
- Phase 2: 12 months
- Phase 3: 12 months





## TA1.1: Network Subsystem Hardware

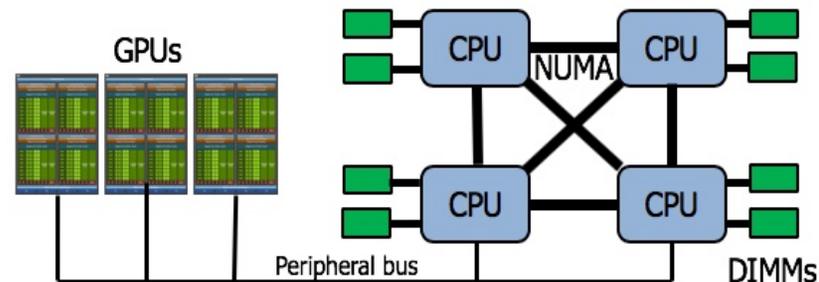
Goal: Design, implement, and demonstrate 10Tbps network interface hardware

Challenges:

- Increase throughput of network attachment by 100x
- Use existing / road-mapped hardware interfaces
- Increase parallelism with multiple interfaces [Traw96]

Potential Approaches: attach to host using

- Intel Compute Express Link (CXL)
- Peripheral Component Interconnect Express (PCIe)
- Intel Ultra Path Interconnect (UPI) Non-Uniform Memory Access (NUMA)
- Dual Inline Memory Module (DIMM) slots





# TA1.2: Network Subsystem Software

Goal: Design, implement and demonstrate system software managing the efficient use of TA1.1 hardware for TA2 applications, collaborate with TA2 to provide interfaces, APIs, and programming tools

## Challenges:

- Develop novel general programming interfaces that support TA1.1 hardware
- Modify OS to support TA1.1 hardware: latency-aware scheduling and task placement, virtual address management, coherence, protection
- Provide programming language libraries / tools to support TA2 (Applications)

## Potential Approaches

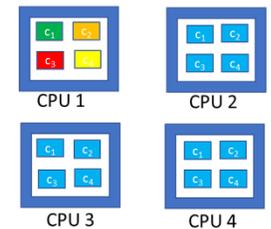
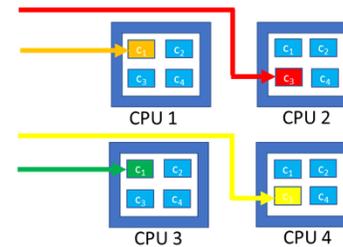
- Develop and implement computing surfaces
- Implement light-weight kernels
- Implement hardware resource disaggregation

## Computing Surfaces

```
Surface Tile[2][2];
```

```
Surface_Read(Tile,  
Remote_Tile,  
THROUGHPUT);
```

```
Surface_Process(Tile,  
LATENCY|DYNAMIC);
```





## TA2: Applications

---

Goal: Design, develop, and implement at least one application that demonstrates 100x speedup when executed on the TA1 stack

### Challenges

- Collaborate with TA1.2 to drive the design and development of novel, general interfaces and APIs
- Develop parallel algorithms using TA1.2 programming models
- Implement algorithms using TA1.2 languages and tools
- Demonstrate use on relevant DoD/IC problems (e.g., image processing)

### Applications of interest:

- Deep learning training
- Imagery/video processing



## TA3: Independent Test and Evaluation

---

Goal: Test and evaluate the performance gains accruing from TA1.1, TA1.2, and TA3, as well as overall integrated performance

### Challenges

- Define performance baseline
- Develop test and evaluation environments
- Refine evaluation metrics and evaluation methodology
- Find or generate data required by TA2 (i.e., labelled training data)
- Conduct 1-socket system-level test and evaluation at Phase 1
- Conduct 2-socket system-level test and evaluation at Phase 2
- Conduct 4-socket system-level test and evaluation at Phase 3



## Metrics

TA	Metric	Phase 1 (24 months)	Phase 2 (12 months)	Phase 3 (12 months)
1.1	Datapath, link to cores	>2 Tbps	>5 Tbps	10+ Tbps
	Latency, link to cores	100 nsec max, 4 cores in 1 CPU	40 nsec max, 8 cores on 2 CPUs	20 nsec max, 16 cores on 4+ CPUs
1.2	Sustained SW throughput	12TB in 100 sec	12TB in 30 sec	12TB in 16 sec
	Software RTT (3m fiber)	1000 nanoseconds	400 nanoseconds	200 nanoseconds
2	Training time, ImageNet (50 minute baseline)	600 seconds	120 seconds	30 seconds
	Training time, labelled 5GB EOS images (90 minute base)	1200 seconds	180 seconds	40 seconds
	Image processing throughput, 95% accuracy classifier, single FastNICs-equipped server, EOS-sized images	25 images/sec	75 images/sec	150 images/sec



## Proposal Information

---

DARPA anticipates a single TA1 award, multiple TA2 awards, and a single TA3 award

Proposers should submit one proposal per TA

The TA3 performer cannot perform on a TA1 or TA2 team, either as a prime contractor or a subcontractor



[www.darpa.mil](http://www.darpa.mil)

# Fast Network Interface Cards (FastNICs) Program

---

Mark Jones  
DARPA Contracts Management Office

Proposers Day  
Arlington, VA

June 10, 2019





## **DISCLAIMER**

**If DARPA publishes an FastNICs Broad Agency Announcement (BAA) and it contradicts any information in these slides,**

**the BAA takes precedence!**



### **BAA OVERVIEW**

- BAA follows procedures in accordance with FAR 35.016.
- Any BAA (as well as any future amendments) will be posted on FEDBIZOPPS at [www.fbo.gov](http://www.fbo.gov) and maybe Grants.gov at [www.grants.gov](http://www.grants.gov)
- Proposal due dates will be identified in the BAA
- BAA will cover all info needed to submit proposals. Follow instructions for proposal preparation and submittal.



### **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
  - BAA will identify a maximum page limit
  - Includes mandatory Appendix A – will not count towards page limit.
- Volume 2 – Cost - No page limit.
- Level of Effort Summary by Task Spreadsheet – Example an Attachment to BAA
- Proposal PowerPoint Summary Quad Chart – Example a second Attachment to BAA
- 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”
  - Include a working/unprotected spreadsheet as part of your Cost Volume submission.
  - Review individual TA descriptions, IP rights, and any deliverables for submission information



## **STATEMENT OF WORK (SOW) PREPARATION TIPS**

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

- Don't use proposal (e.g. we propose to do . . .) or marketing language
- Break out work between any phases/time periods identified in the BAA
- Succinctly and clearly define tasks & subtasks
- Identify the primary organization responsible for task execution
- Identify measurable milestones and define deliverables
- Do not include any proprietary information!

**NOTE:** For grants/cooperative agreements: SOW = RDD or Research Description Document. For Other Transactions: SOW = TDD or Task Description Document



## PROPOSAL PREPARATION TIPS

- **Heilmeier's Catechism – Learn it. Know it. Live it.**
- **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 metrics / potential mitigation plans / alternative directions
  - If conducted prior research, use data to support why approach will work
- **Page Limits** – Depth better than breadth
  - Focus on most critical/beneficial aspects
  - Don't restate problems at expense of explaining solution
  - Don't sacrifice SOW



### **PROPOSAL PREP CONT'D – INTELLECTUAL PROPERTY RIGHTS**

- Government typically desires, at a minimum, **Government Purpose Rights** for any proposed noncommercial software and technical data. (SEE DFARS 227 for Patent, Data, and Copyrights)
- Data Rights Assertions – IF asserting **less than Unlimited Rights**:
  - Provide and justify basis of assertions (e.g. privately funded under IRAD project XYZ)
  - 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

- Fundamental vs. non-fundamental research
- Understand and comply with SAM, E-verify, FAPIIS, i-Edison and WAWF. Links can be found in the BAA.
- Subcontracting Issues
  - Non-Small Businesses: Subcontracting Plans required for FAR-based contracts expected to exceed the applicable threshold.
  - Subcontracting plans with <5% SDB goal – provide an explanation why
  - 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 typically must be valid for a minimum of 120 days –recommend putting in a longer time period
- TFIMS-2 Usage
- Document files must be in .pdf, .odx, .doc, .docx, .xls, and/or .xlsx formats
- Submissions must be written in English



### **ABSTRACT/PROPOSAL SUBMISSION**

- Abstracts are submitted via DARPA's web-based upload system.
- FAR based contract and OT proposals: Required to be submitted by via DARPA's web-based upload system for unclassified portion of proposal. Submission must be in a single zip file not exceeding 50 MB.
- Assistance Instrument proposals: Required to be submitted 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 make submissions – the submission deadline(s) as outlined in the BAA will be strictly enforced!
- DO NOT forget to FINALIZE your proposal submission in the DARPA 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) at a minimum will be: (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.



### **TAKE AWAY**

- Submit abstracts/proposals before the due date/time - Do NOT wait until the last minute (i.e. hour) 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.



[www.darpa.mil](http://www.darpa.mil)