

Harnessing Autonomy for Countering Cyberadversary Systems (HACCS)

Angelos D. Keromytis
Program Manager
Information Innovation Office (I2O)
DARPA

July 31, 2017





Agenda

TIME	EVENT
1:00 PM - 2:00 PM	Check-in
2:00 PM - 2:05 PM	Welcome – Angelos D. Keromytis, Program Manager (PM), DARPA/I2O
2:05 PM - 2:10 PM	HACCS Security – DARPA Security
2:10 PM - 2:30 PM	HACCS BAA – Mark Jones, DARPA Contracting Officer
2:30 PM - 3:15 PM	HACCS Program – Angelos D. Keromytis, PM, DARPA/I2O
3:15 PM - 3:30 PM	Informal Teaming Discussions/Turn-in questions
3:55 PM - 4:05 PM	Question & Answer – Angelos D. Keromytis, PM, DARPA/I2O



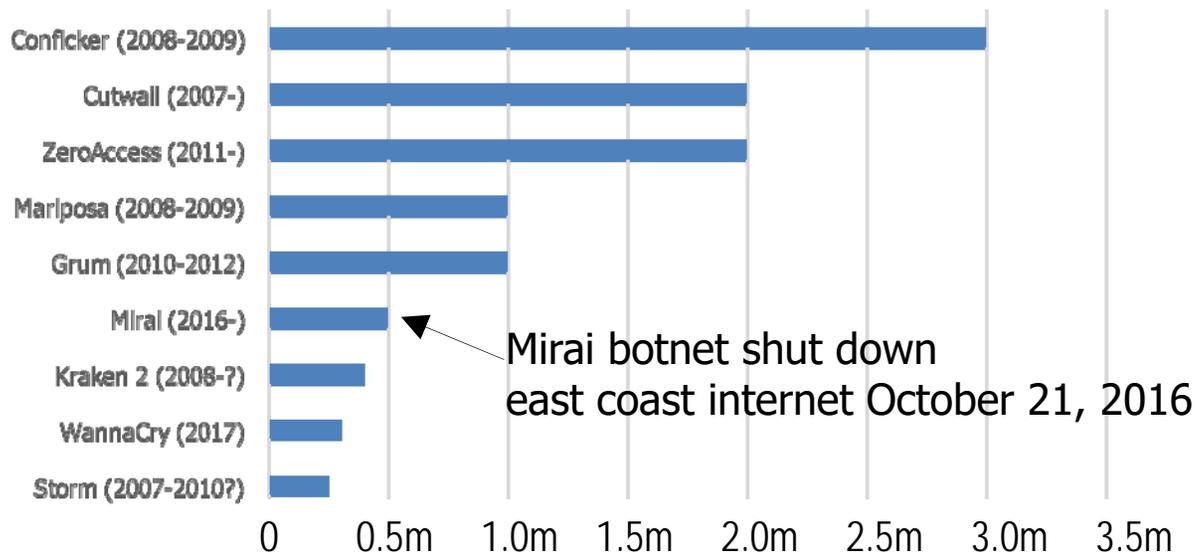
Program Goal

Develop safe, reliable, and effective capabilities for conducting Internet-scale counter-cyber operations to deny adversaries' use of neutral (gray) systems and networks (e.g., botnets)



Cyber Attackers Can Muster Massive Botnets

Botnet Sizes Observed on the Internet,
in millions of compromised devices



State and non-state adversaries can compromise and conscript large numbers of gray (neutral) networks and systems

- Gradual or rapid buildup through compromise and purchase of resources
- “Botnet for hire” services
- Botnets can DDoS networks, provide pivot points for operations, impede the flow of information, circumvent defenses, and amplify influence operations via social media



Current Countermeasures Are Slow and Ineffective

Computers are not patched reliably, configured properly, or used safely, allowing widespread exploitation

- 99.9% of exploited vulnerabilities has been publicly disclosed over a year earlier (Verizon Data Breach Report, 2015)

Incident response is slow and costly when possible

- Most botnet nodes are outside US jurisdiction

Adversaries have adapted to countermeasures

- e.g., from centralized to peer-to-peer or social network-based C2

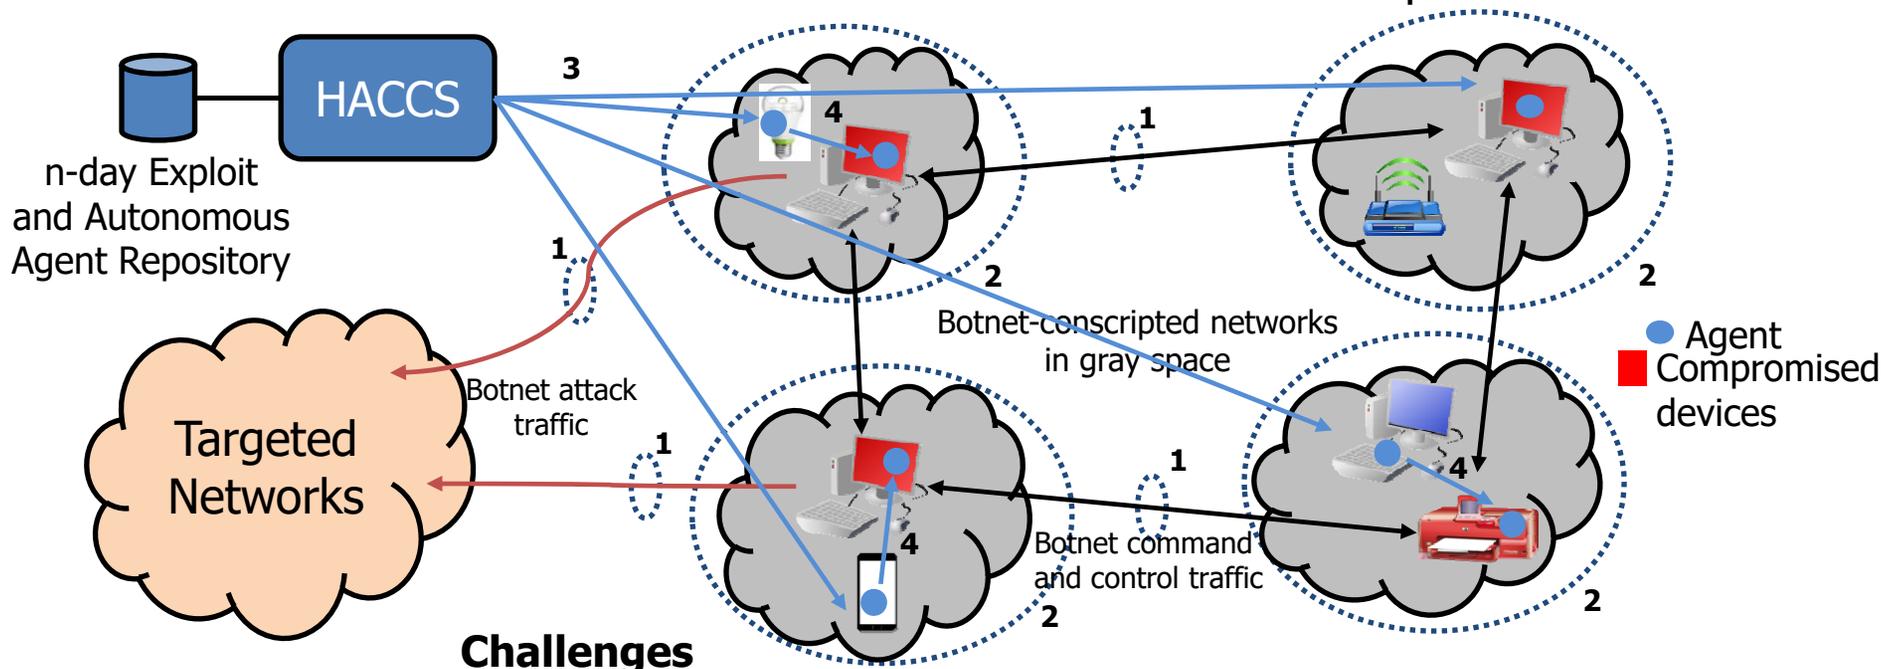
Active defense cyber operations against individual botnet nodes are difficult

- Feasible in principle but unreliable and unsafe
 - Welchia, Santy, Hajime
- Risky and illegal for the private sector, with no reward structure



Harnessing Autonomy for Counter Cyber Systems

Develop safe and reliable autonomous agents that can be introduced into gray networks at scale to counter botnets and similar adversarial implants



Challenges

1. Find botnet-constripped networks **TA1**
2. Fingerprint botnet-constripped networks
3. Exploit n-day vulnerabilities to insert agents **TA2**
4. Identify and safely neutralize botnet implants **TA3** at scale, according to verified rules of operation

Why Now?

Recent Technical Advances in:

1. Multi-dimensional network analytics
2. Cyber Reasoning Systems
3. Autonomous software agents leveraging AI



TA1: Find and Fingerprint Botnet Infrastructure

Key Research Challenges

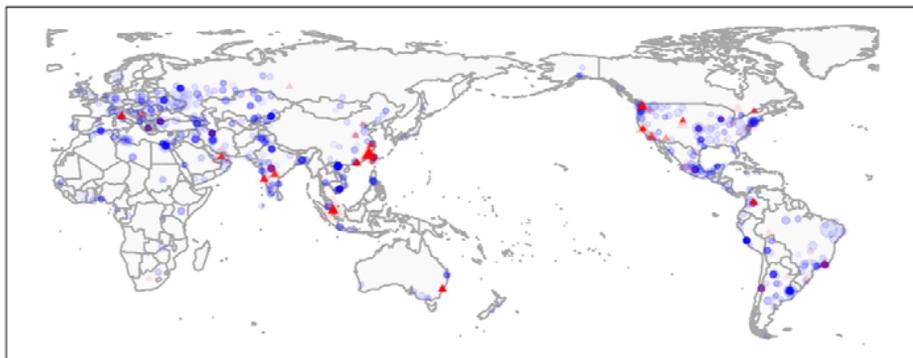
1. Internet-scale real-time botnet detection in the presence of evasive/covert C2
2. Accurate fingerprinting of devices and software in compromised networks

Possible Approaches

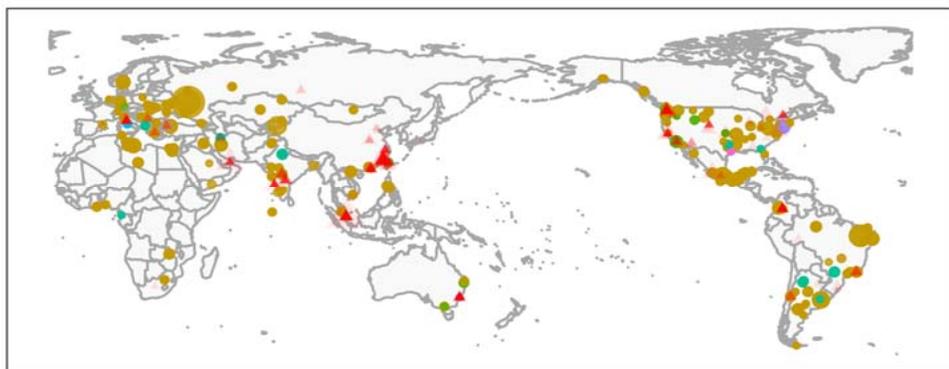
1. Automated traffic analysis using disparate and noisy data sources
2. Efficient and scalable black-box characterization of device network behavior
3. Precise white-box analysis of network-observable software behavior using information flow

Metrics

- Accuracy
- Percentage of devices characterized across the Internet
- Speed/work factor of fingerprinting new device/software



Hidden Cobra (DPRK)



Hidden Cobra co-resident IoT devices

Type of IoT device

- Backup
- Entertainment
- Health
- Home
- HVAC
- MGMT
- Security

volume

- 50
- 100
- 150
- 200



TA2: Insert Autonomous Agents Into Gray Networks

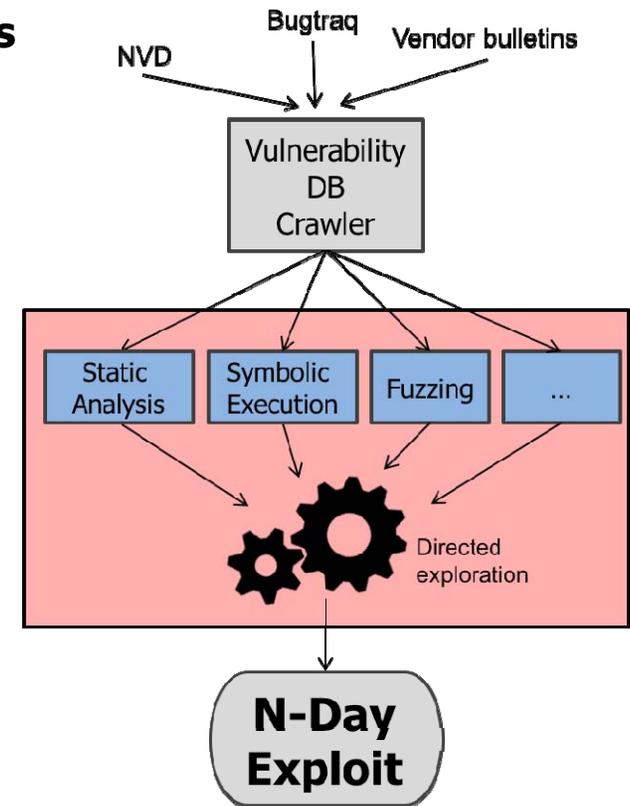
Primary approach: Exploit known (n-day) vulnerabilities

Key Research Challenges

1. Automated generation of n-day exploits for agent insertion
2. Development of IoT- and cloud-specific agent insertion techniques

Possible Approaches

1. Focus Software Reasoning Systems (SRS) analysis on known vulnerable code
 - Example: use Natural Language Processing on unstructured and semi-structured public information to guide software exploration
2. Extend SRS analysis beyond memory corruption vulnerabilities
 - Example classes: web/command injection, authentication bypass, privilege escalation
 - Challenges: symbolic analysis & fuzzing for interpreted languages with different runtime models; determining test conditions; expanding to different types of inputs



Metrics

- Number of exploits
- Vulnerability class coverage
- Stability of exploits



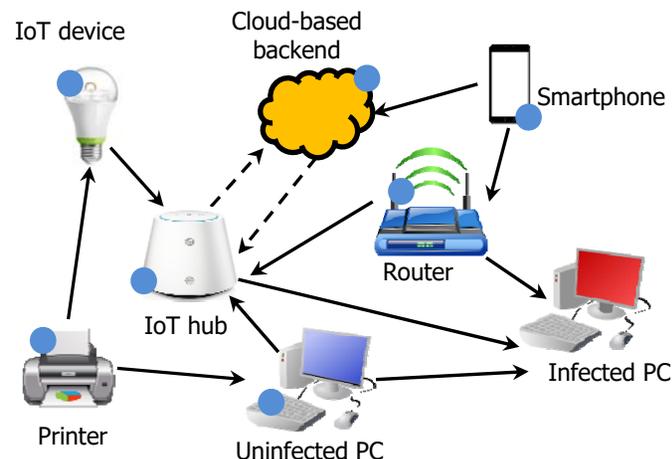
TA3: Identify and Neutralize Botnet Implants

Develop software agents that autonomously navigate within each gray network toward infected devices to safely neutralize the malicious botnet implant

● Potential agent insertion point

Key Research Challenges

1. Autonomous lateral movement in partially known environments
2. Correctness of agent implementation
3. Correctness of rules of operation
 - Understand, encode, and reason about bounding boxes and terminating conditions for the agents



Possible Approaches

1. Learn and generalize from human operators in cyber-exercises, adversary activities, and similar sources
 - Transfer learning for graph traversal
2. Correct-by-construction techniques and tools applied to agent generation
3. Contract-based programming

Metrics

- Success rate and speed in navigating topologies
- Fraction of code proven correct



TA4: Integration

Identify and implement necessary components

- Overall framework (new or existing, e.g., Plan-X)
- Safe anti-implant effects
- Integration of publicly & commercially available sources with performer-provided private/commercial (or Government-only) sources

Conduct full-system testing

Act as Voice-of-the-Offense for the program

Option to act as interface with transition partners if necessary

- Propose optional integration tasks beyond program duration

Key metric: effectiveness in achieving system goals

- Participate in DoD cyber exercises (REDFLAG, CYBERGUARD/CYBERFLAG, etc.)



Program Structure and Schedule

Program duration: 48 months

- Three 16-month program phases

All TAs working in parallel

- Increasing realism and scale in evaluation

Conduct on-demand testing in real conditions as opportunities arise, working with operational/transition partners

	Phase 1	Phase 2	Phase 3
TA1	Characterize 5% of the global IP address space with 80% accuracy of botnet detection and network fingerprinting	Characterize 25% of IP address space, 90% accuracy	Characterize 80% of IP address space, 95% accuracy
TA2	10 n-day exploit instances 1 additional vulnerability class	100 n-day exploit instances 2 additional vulnerability classes	1,000 n-day exploit instances 2 additional vulnerability classes
TA3	Demonstrate lateral movement and effect in 10 computer-simulated topologies 30% of autonomous agent code verified	1,000 computer-simulated topologies 75% of autonomous agent code verified Formally specified Rules of Operation	10,000 computer-simulated topologies 95% of autonomous code verified Formally verified Rules of Operation
TA4	Voice of the Offense	Design and implement integration framework	Demonstrate system in DoD exercises



Evaluation Details

- Each performer conducts their own evaluation for each phase
 - Provide data and prototypes to DARPA and AFRL to conduct an independent validation
 - Government reserves the right to engage third parties to independently validate the results
- DARPA will pursue access to UNCLASSIFIED data sets
 - Proposers strongly encouraged to pursue their own data sets that will facilitate initial development



Program Classification and Clearance Requirements

- The program will be conducted at the UNCLASSIFIED level
 - Technical development
 - Performer-internal testing
- TA4 teams required to include personnel with TS clearance and eligible for SCI
 - Adequate number to allow for extensive T&E in the Washington, DC area
 - Not all team personnel need to be cleared
 - For multi-organization teams, not all participating organizations must have cleared personnel
 - No requirement for SCIF access
- TA1, TA2, & TA3 teams encouraged to include personnel with similar clearances



Programmatic Details

- Proposals due on October 1, 2017 (estimated)
- Anticipated program start date: 1 April 2018
- One proposal per organization as Prime
- Procurement Contract (no Grants)
- To expedite award contracting, proposers are encouraged to have sub-award agreements in place ahead of award notification

- Anticipated number of awards:

TA1	TA2	TA3	TA4
Multiple	Multiple	Multiple	One or more

- Proposals may address any combination of TAs
 - Technical work and cost must be separable to enable partial selection
- The same organization cannot be selected as Prime for efforts under TA4 and TA1, TA2, TA3
- TA4 performers must be prepared to work with all TA1, TA2, & TA3 teams



Meetings and Reporting Requirements

- Two Annual Principal Investigator (PI) Meetings
- Quarterly Technical Reviews between PI Meetings
- Monthly Progress Reports
 - Technical Report describing progress, resources expended and issues requiring Government attention, provided 10 days after the end of each month
- Financial/Technical Progress Reporting to the DARPA Contract Execution Reporting Service (CERS)
- Final Technical Report
- See BAA for full details

- Anticipate high frequency interactions with DARPA technical team

- Agent: DARPA CMO

Harnessing Autonomy for Countering Cyberadversary Systems (HACCS)

Mark Jones
Contracting Officer
Contracts Management Office (CMO)
DARPA

July 31, 2017





DISCLAIMER

If DARPA publishes the HACCS 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 and possibly Grants.gov at 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



HACCS Proposers Day

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.
- May include optional Appendix B – would 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”
- 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 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
- 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

- **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 justify why approach will work
- **Page Limits** – Depth better than breadth
 - Focus on most critical/beneficial aspects
 - 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

Discontinued usage of T-FIMS

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

Submissions must be written in English



PROPOSAL SUBMISSION

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 submit proposals – the submission deadlines 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.

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 (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.



Break

- The HACCS Program Q&A session will begin at 3:55pm.



www.darpa.mil

Harnessing Autonomy for Countering Cyberadversary Systems (HACCS)

Angelos D. Keromytis
Program Manager
Information Innovation Office (I2O)
DARPA

July 31, 2017





Audience Q&A

- Q: Do we care how “stealthy” the agents are when they are deployed? Is this incorporated into “correctness of agent implementation”? Or into the rules of operation?
- A: Stealth of the agents is not a primary concern of the program.



Audience Q&A

- Q: Is precision of agents an important metric? Or are “kitchen sink” approaches to neutralization in scope?
- A: Yes, precision of agent affects is an important aspect of safety and reliability.



Audience Q&A

- Q: Are any impacts to infected networks allowed? E.g. cutting off access of non-botnet comms; E.g. denying access to DNS
- A: It is preferred that side effects are minimized. Understanding and quantifying any unavoidable side effects is required when minimization is impossible.



Audience Q&A

- Q: Are you seeking robust measures of effectiveness integrated as part of the TA4 framework against the stated metrics?
- A: Yes



Audience Q&A

- Q: Will the 'botnet' environments be static or dynamic – that is, will the botnet spread during an experimental run?
- A: Yes



Audience Q&A

- Q: Are you open to a large scale virtualized environment to support enabling parameterized experiment runs as part of the TA4 framework?
- A: DARPA does not seek to fund the creation of such an environment, but if one already exists, its use will be viewed as a strength of the proposal.



Audience Q&A

- Q: Who controls intellectual property?
- A: We desire, at a minimum, unlimited duration GPRs for any technology developed under this program.



Audience Q&A

- Q: TA2: Is it fine looking for zero – days or just restricted to n-days?
- A: Just n-days.



Audience Q&A

- Q: For TA2, if an agent obtains access, can or should it remain persistent to mitigate future bots?
- A: Persistence may be part of the rules of operation. Said persistence is to be a limited time duration.



Audience Q&A

- Q: Are FFRDC's eligible?
- A: Yes



Audience Q&A

- Q: What is the budget for the program?
- A: The budget for this program will not be disclosed.



Audience Q&A

- Q: Can we build vulnerabilities related to any device (IoT, Android)?
- A: Vulnerabilities, in scope, are for any internet connected device.



Audience Q&A

- Q: Can we build vulnerabilities related to any device (IoT, Android)?
- A: Vulnerabilities, in scope, are for any internet connected device.



Audience Q&A

- Q: What kind of data we can expect to have from DARPA?
- A: The proposer should determine the type of data require to support their technical approach.



Audience Q&A

- Q: How will the 5% of IP with 80% accuracy be validated? (Phase 1 evaluation)
- A: Strong proposals will have convincing evaluation plan. DARPA will pursue validation using complimentary data sources.



Audience Q&A

- Q: Does the scope of grey networks include critical infrastructure (electrical grid, manufacturing)?
- A: Yes. The identification of critical infrastructure is of interest and whether and how to act in these networks or on these computing devices is part of the rules of operation.



Audience Q&A

- Q: Clarify relationship of "target" network owner and "GRAY" network owner.
- A: For the purposes of this effort there is no meaningful difference.



Audience Q&A

- Q: What is the outcome of the program?
 - How are the success factors measured?
 - Detecting known or O-day?
- A: The outcome of the program will be technology that will be transitioned to operational partners with the appropriate legal authorities to use them.
 - The success of individual components will be evaluated as delineated in the BAA.
 - To the extent that the question refers to vulnerabilities the program is looking to generate exploits only for known vulnerabilities.



Audience Q&A

- Q: One of the biggest hurdles to fingerprinting a “hack” is knowing where it originated. A lot of times effective botnets & hacks mask their locations and intents. With rules of engagement in mind, and noting your requirement to “insert an agent” into the grey network – are you suggesting that to have true cyber defense, you in actuality have to be authorized to execute offensive cyber?
- A: The program is developing technologies that address a specific threat in a specific manner. Doctrine, operational authorities, and legal framework are outside the technical scope of the effort.



Audience Q&A

- Q: An extensive test environment will be needed & created for this – is the GOV funding?
- A: DARPA is looking to leverage existing test environments and facilities to the greatest extent possible.



www.darpa.mil



National Security Archive,
Suite 701, Gelman Library, The George Washington University,
2130 H Street, NW, Washington, D.C., 20037,
Phone: 202/994-7000, Fax: 202/994-7005, nsarchiv@gwu.edu