Enhanced Attribution

Angelos Keromytis
Information Innovation Office (I2O)

Proposers Day Briefing

April 25, 2016

Program Goal

Develop mechanisms and algorithms to:

1. **Collect and fuse all-sources information to enable at-scale attribution of individual malicious cyber operators**

2. **Validate and enrich the collected corpus with additional information to enable convincing public attribution of malicious cyber activity**

Impact

Entity	Sources	Accuracy	Timeliness	Scale	Shareability
Commercial	Blue space	low	low	low	high
Government	Red space	high	low	low	low
Enhanced Attribution	Red & Blue	high	high	high	high

- If successful, we will enable a range of response options, including many that require public disclosure of evidence

Examples of Response Options
Sanctions
Name and Shame
Selective Revelation
Mass Disclosure
Law Enforcement
Cyber Effects

Attribution is Accomplished Poorly or at Small Scale

The same campaign attributed to 4 different intrusion sets by 3 commercial cybersecurity providers, based on different observables

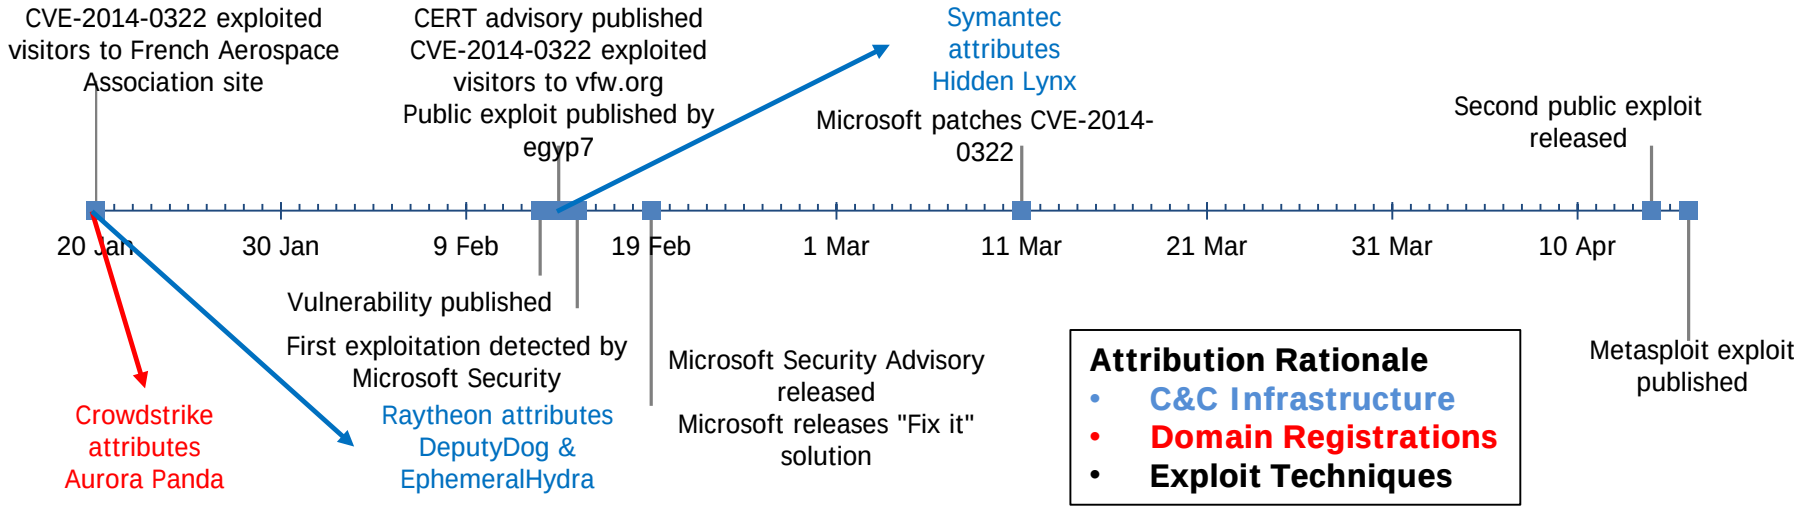

Q: Who is UglyGorrilla?

A: Wang Dong

How do we know?

DNS registrations
DNS use
PLA alumni website
Binary metadata
Passwords
Social media

"Attribution is really really hard … we're using the totality of the sources and methods we have to help inform that. [But] because those advanced persistent threats aren't going away … we can't bring all that information to the fore and be fully transparent about everything we know and how we know it."

- Rob Joyce, Chief NSA/TAO at USENIX Enigma 2016

Monitoring and Understanding of Adversary Cyber Operator Actions Through Collection in Red Space

TA1
Ground Truth via All-Sources Monitoring

TA2
From Data To Information

TA3
Find Adversary Mistakes

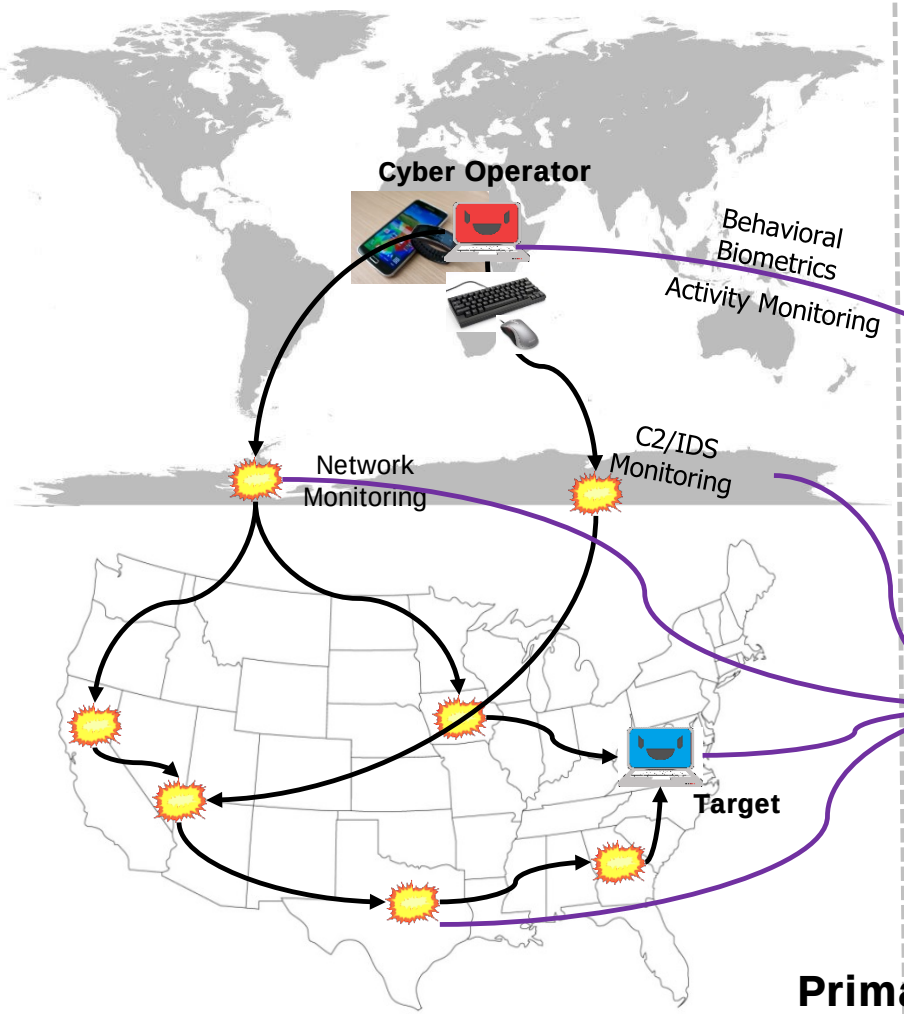

External Information Sources

- Commercial Threat Feed
- End Host Logs (e.g. HBSS, TC)
- OSINT (e.g., social networks)
- SIGINT
- Network IDS

Fuse & Predict

Validation & Enrichment

Shareable Information

Primary challenge: multi-resolution fusion at scale

TA1: Activity Tracking and Summarization

- Create software modules for tracking adversary operators from various vantage points
  - Assume initial access and communications channel is provided
  - Ops desktop, mobile phone, IoT, captured C2 nodes, network infrastructure, …
  - Track personas, extract relevant cyber actions, and link personas to identities
- Possible approaches
  - Behavioral biometrics techniques
  - Decompose attack software toolchain into units of action (e.g., port scan, send spear-phishing email); extract and emit associated metadata (which represents ground truth)
- Challenges
  - Maintaining a low computational-resource profile requires new algorithmic approaches to behavioral biometrics
  - Software partitioning and whole-system analysis at low footprint

Keyboard/Mouse/Touch behavioral biometrics

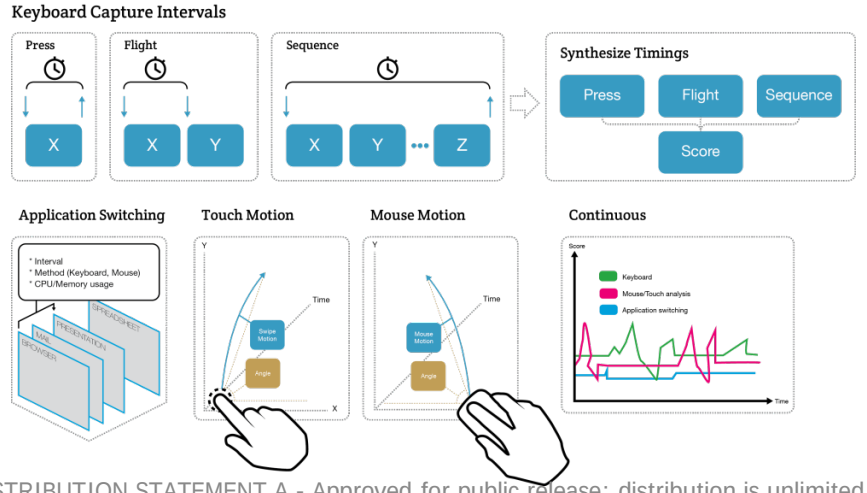

TA2: Data Fusion and Activity Prediction

- Develop multi-INT fusion and tracking for the cyber domain
  - Key enabler: ground truth generated by TA1
- Develop predictive profiles using downstream observables
- Capture ambiguous data associations across diverse data set
- Possible approach: heterogeneous information networks
  - Leverage, reason over, and contribute to structure
  - Scalable and extensible to support wide range and types of data

Challenges
Association Ambiguity
Out of Sequence Data
Large Data Volume

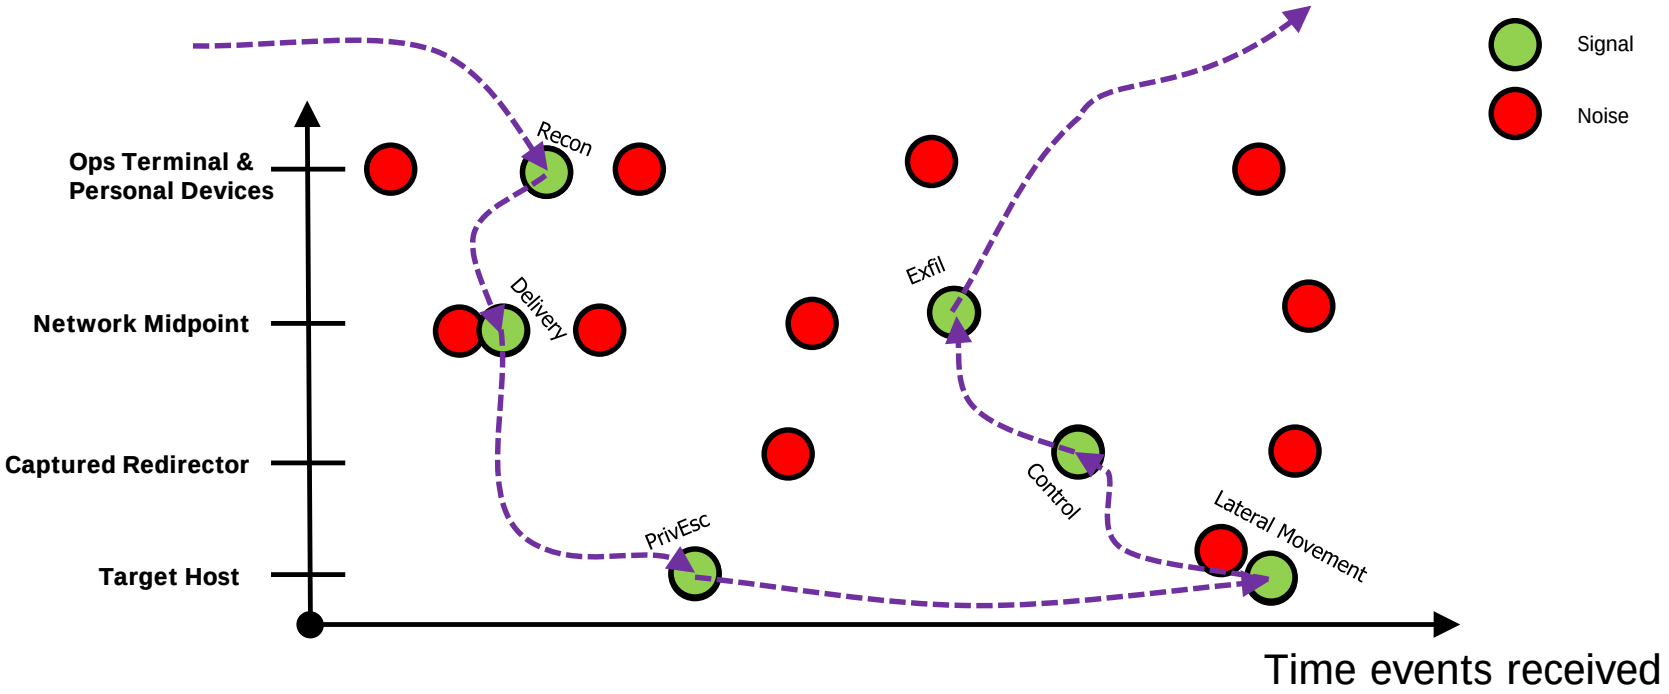

TA3: Validation & Enrichment

- Identify adversary mistakes and externally observable indicators
  - Pattern-matching over an activity graph
  - Fuse and integrate external information sources
  - Compose disclosable explanatory narrative
- Use analytic techniques to expose known but hidden structure
- Create convincing explanation using external data
  - Operate fast enough to make use of rich but transient data

Adversary Mistakes
Weak passwords
Vulnerable software
Weak cryptography
Infrastructure reuse

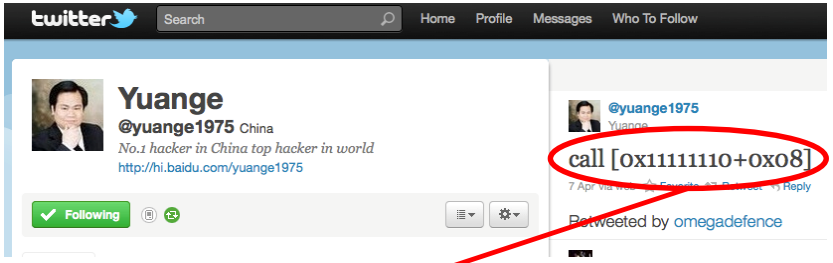

"the Chinese developers ... left clues aplenty about their identities and locations, ... even Tweeting about ... vulnerability days in advance of its use in the wild."

External Indicators
OSINT
Commercial Threat Feeds
Network IDS/analytics
End Host (e.g. HBSS)

Program Structure and Schedule

- Program duration: 54 months
  - Three 18-month program phases
- All TAs working in parallel
- Frequent informal evaluations using various datasets
  - Increasing realism/environments
- Be in position to conduct on-demand testing in real conditions as opportunities arise, possibly starting early in the program

	Year 1	Year 2	Year 3	Year 4	Year 5
	Phase 1		Phase 2	Phase 3	
TA1 Profile/Activity Precision	60%		75%	90%	
TA2 Reconstruction Precision	35%		50%	75%	
TA3 Enrichment Completeness	20%		35%	50%	

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 anticipates integrated evaluations early in the program

- 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

- TA2 and TA3 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 teams encouraged to include personnel with similar clearances

Human Subjects Research (HSR)

- Depending on technical approach, TA1 teams may need to consider HSR implications

- DARPA encourages TA1 teams to consult their IRB and address this matter in their proposals
  - Ideal scenario: proposal includes letter with IRB determination
  - Second best scenario: proposal includes submission to IRB

Programmatic Details

- Proposals due on June 7, 2016
- Anticipated program start date: 1 November 2016
- One proposal per organization as Prime
- Procurement Contract or Other Transaction (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: multiple
  - TA2: one or more
  - TA3: one or more

- Proposals may address any combination of TAs
  - Technical work and cost must be separable to enable partial selection

- TA2 performers must be prepared to work with all TA1 and 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 Technology Financial Information Management System (TFIMS)
- Final Technical Report
- See BAA for full details

- Anticipate high frequency interactions with DARPA technical team

- Agent: AFRL/RIGB