U.S. NAVAL RESEARCH LABORATORY

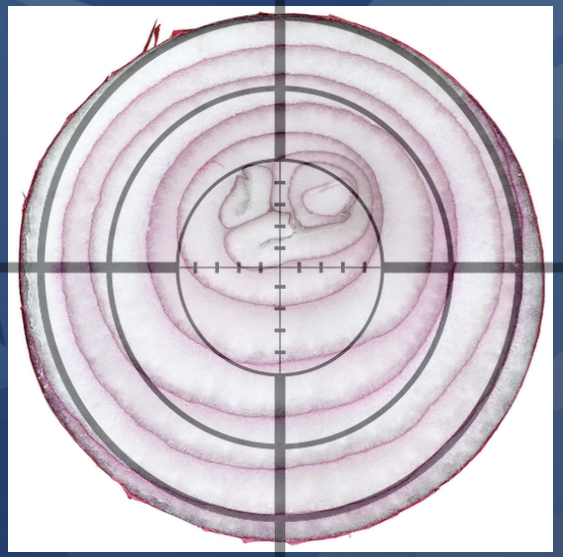

Oft Target

Paul Syverson (joint work with Aaron Jaggard)
Center for High Assurance Computer Systems (CHACS)
U.S. Naval Research Laboratory
Washington DC

HotPETs
2017 PET Symposium
University of Minnesota
July 21, 2017

U.S. NAVAL RESEARCH LABORATORY

Oft Target

(*Tor adversaries that don't miss the mark*)

Paul Syverson (joint work with Aaron Jaggard)
Center for High Assurance Computer Systems (CHACS)
U.S. Naval Research Laboratory
Washington DC

HotPETs
2017 PET Symposium
University of Minnesota
July 21, 2017

U.S. NAVAL RESEARCH LABORATORY

Main take-away

Adversary models in Tor research and design have ignored goals and strategies of realistic adversaries for primary Tor users.

Evaluation and design have thus importantly misconstrued likely risk.

- Will show orders-of-magnitude more efficient attacks than comparably-endowed adversary using hoovering strategies

Going forward, Tor should be designed & evaluated with consideration of targeting adversaries.

U.S. NAVAL RESEARCH LABORATORY

Traditional Tor adversaries

- Primary vulnerability: end-to-end correlation

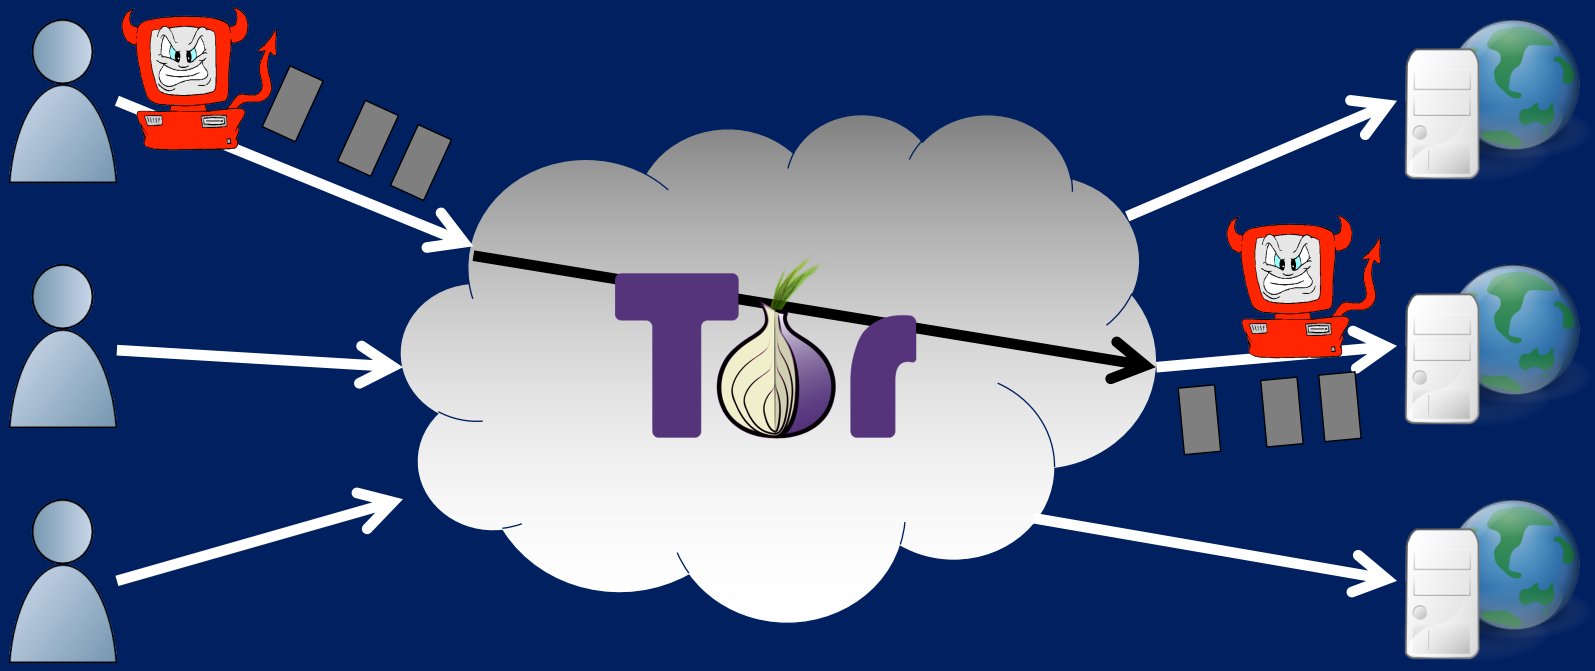

U.S. NAVAL RESEARCH LABORATORY

Traditional Tor adversaries

- Relay-adversary analysis has focused on guards and exits

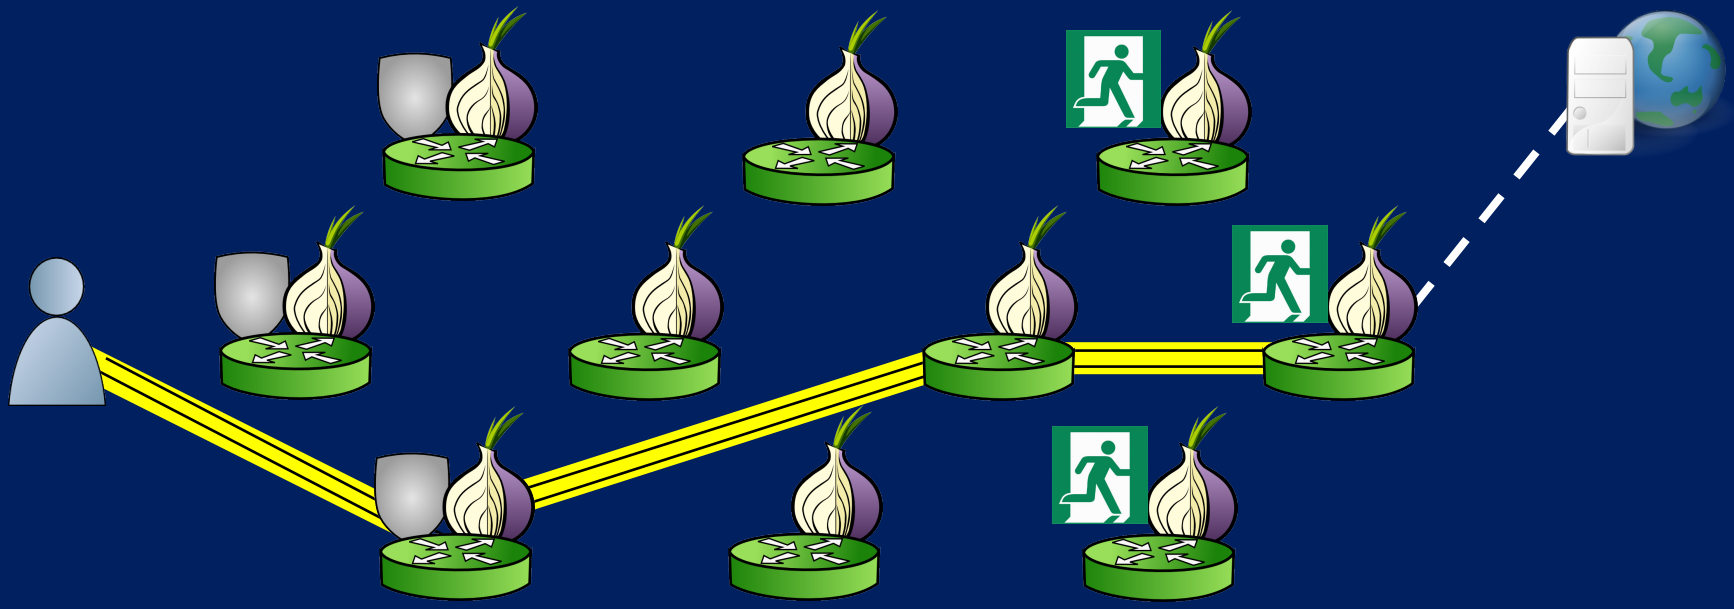

U.S. NAVAL RESEARCH LABORATORY

Traditional Tor adversaries

Prior work generally uses hoovering-strategy adversary

- Adversary wants to suck up ***all*** Tor traffic

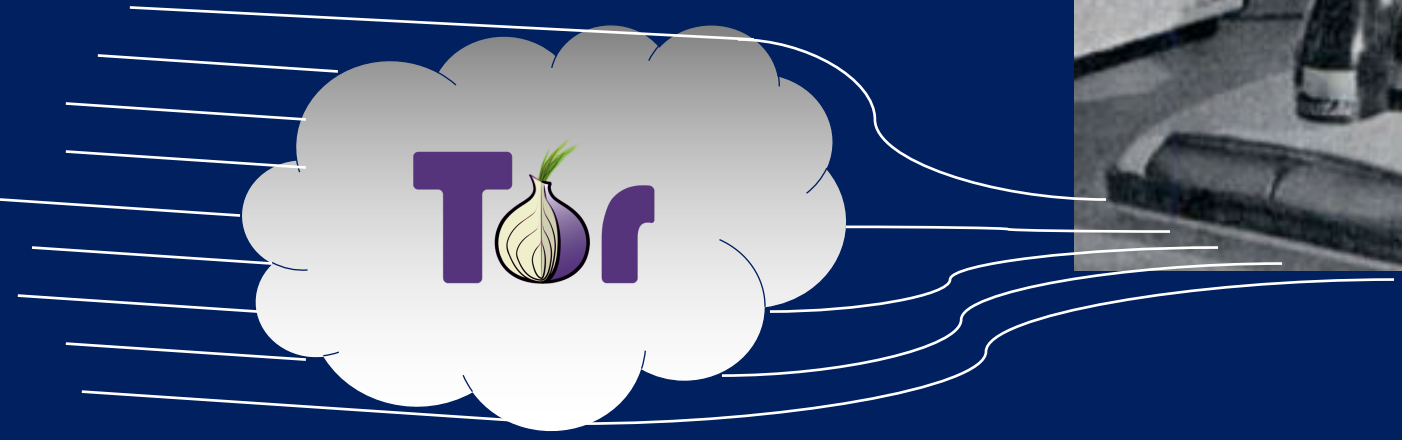

U.S. NAVAL RESEARCH LABORATORY

Who Uses Tor? Who are their adversaries?

- Businesses, Human Rights Advocates, Journalists, Law Enforcement, Military, Normal People, etc.
- Most Tor users likely only need worry about their ISP and potentially hostile/insensitive/incompetent destinations
- Sensitive users may face nation-states or well resourced criminal organizations
  - the people we invented onion routing for, but…
- Such adversaries may employ strategies to target a journalist/human rights advocate/law enforcement agent/ of interest regardless of usefulness against generic users
- What are some examples of potentially targeted users?
- What would such targeting attacks look like?

U.S. NAVAL RESEARCH LABORATORY

Who Uses Tor? Who are their adversaries?

- Businesses, Human Rights Advocates, Journalists, Law Enforcement, Military, Normal People, etc.
- Most Tor users likely only need worry about their ISP and potentially hostile/insensitive/incompetent destinations
- Sensitive users may face nation-states or well resourced criminal organizations
  - the people we invented onion routing for, but…
- Such adversaries may employ strategies to target a journalist/human rights advocate/law enforcement agent/ of interest regardless of usefulness against generic users
- What are some examples of potentially targeted users?
- What would such targeting attacks look like?

U.S. NAVAL RESEARCH LABORATORY

Presented work

- Example of cabal meeting on private IRC channel
- Targeted attacks to learn about that cabal and its members
- Comparison to hoovering attack on IRC cabal based on comparable adversary and usage
- Targeted attacks to learn about particular onionsites
  - Popularity and activity
  - Usage distribution and location of heavy users
- Countermeasure suggestions

See Arxiv paper “Onions in the Crosshairs” for

- Abstract model of targeting adversaries
- Analysis of targeted attacks on cabal meeting via MTor multicast protocol

U.S. NAVAL RESEARCH LABORATORY

Cabal meeting on IRC

- Group meeting regularly on private IRC channel
  - All cabal members access IRC server only via Tor
  - All cabal members make a new circuit for each meeting

U.S. NAVAL RESEARCH LABORATORY

Possible Targeting Adversary Goals

- Learn size of cabal
- Learn guards of (important) cabal members
- Compromise/bridge-across guards of
  - of most important cabal members
  - all cabal members

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

1. Own as many middle relays as possible

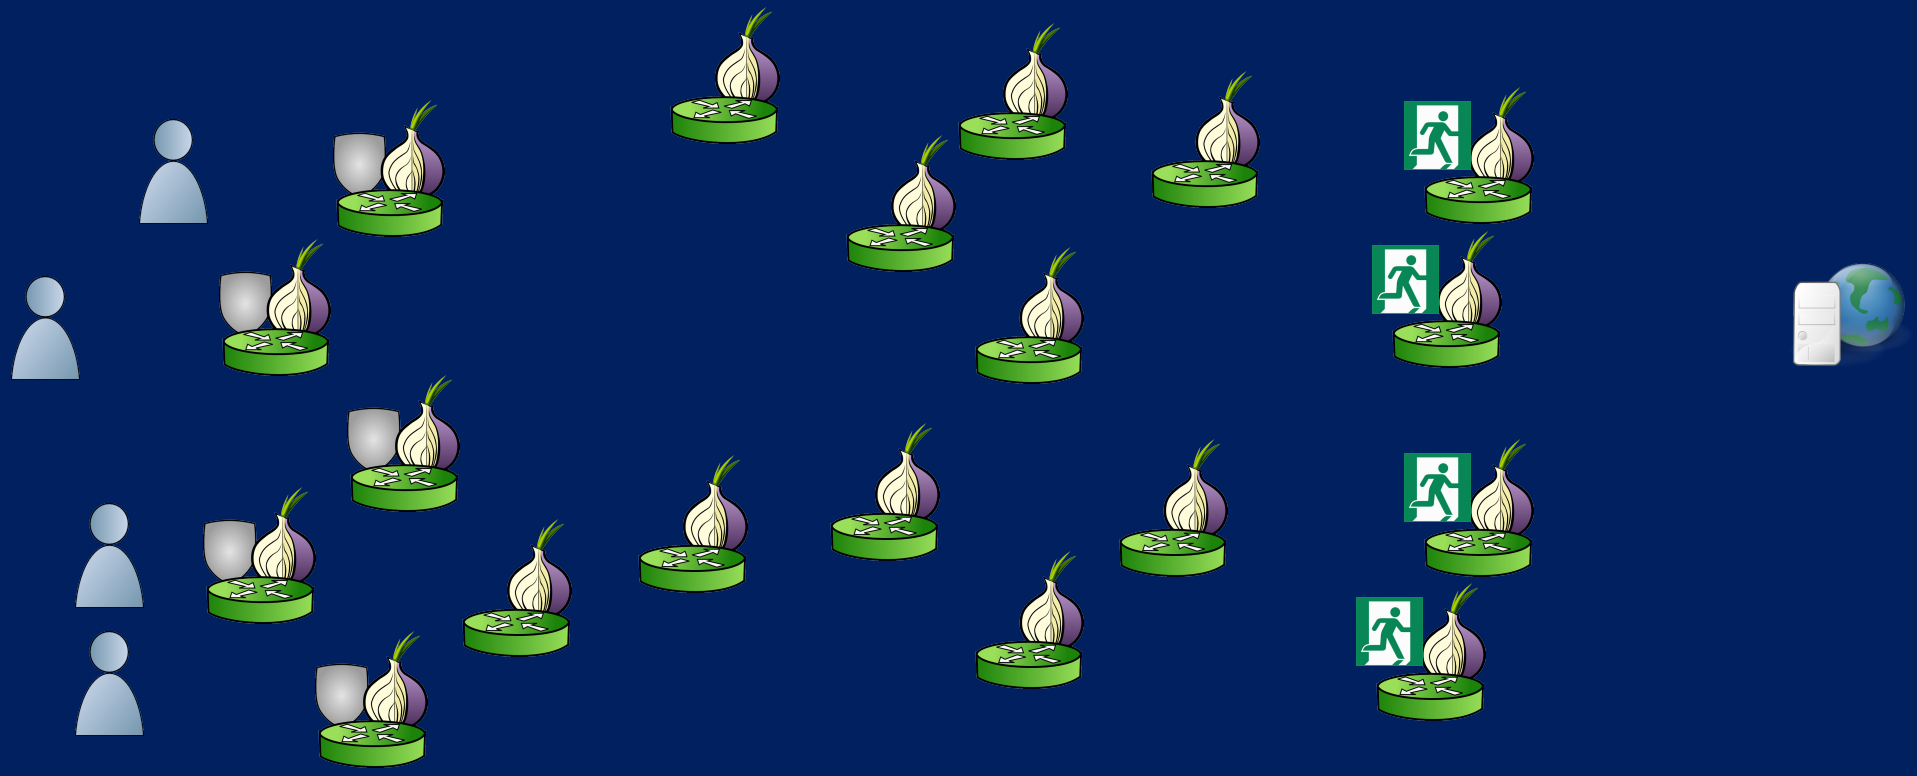

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

1. Own as many middle relays as possible

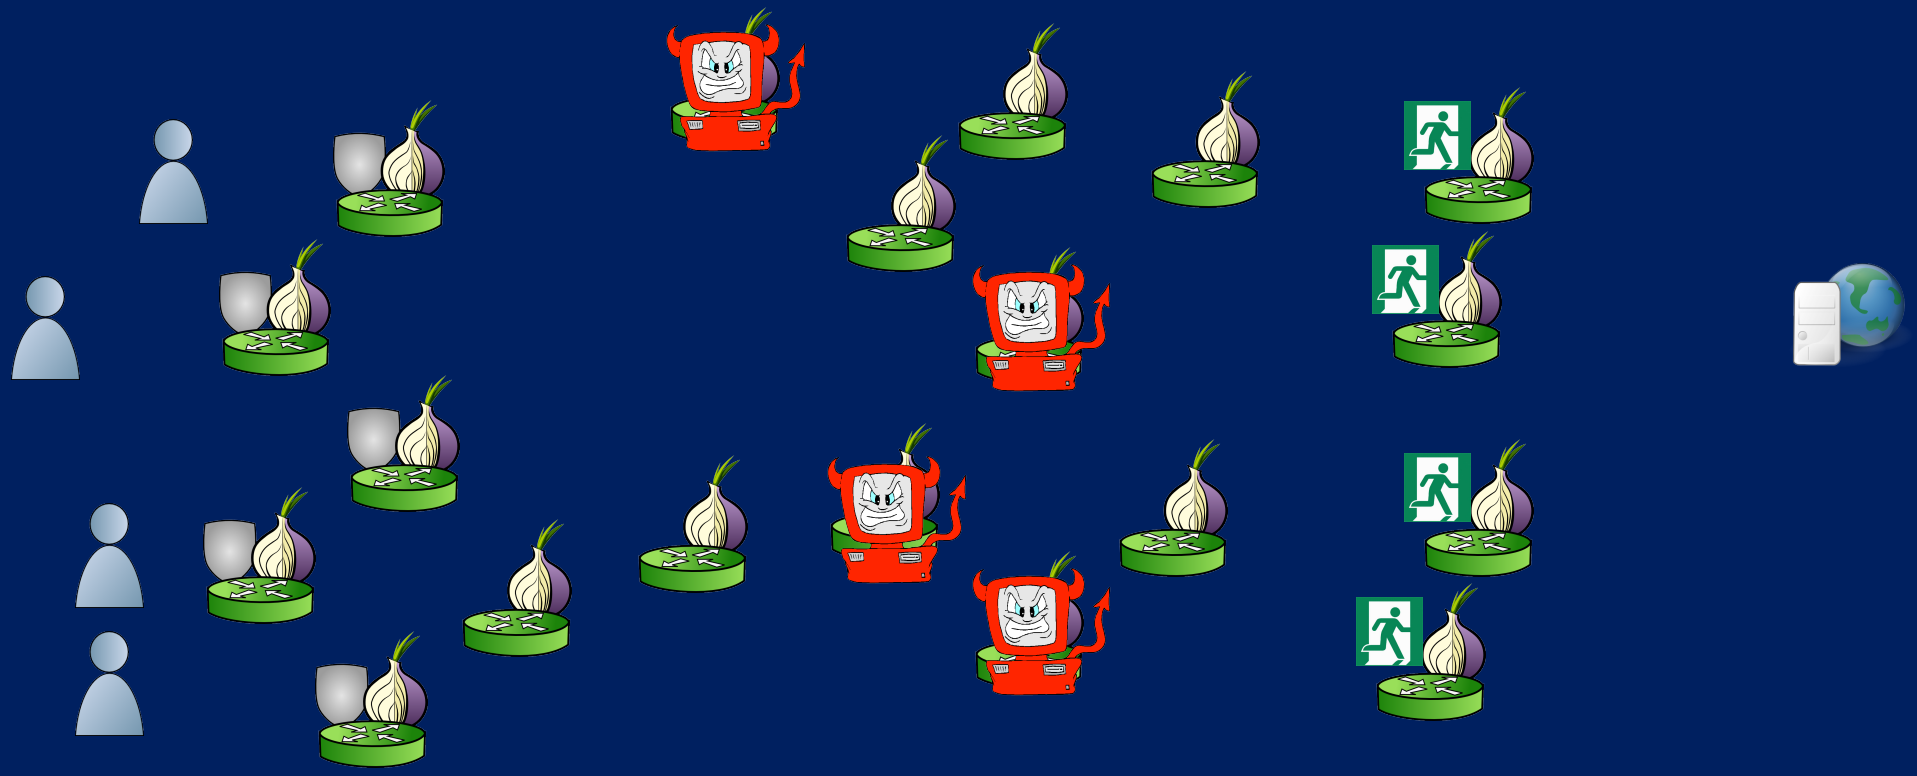

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

1. Own as many middle relays as possible
2. Discover a cabal guard whenever circuit goes through you

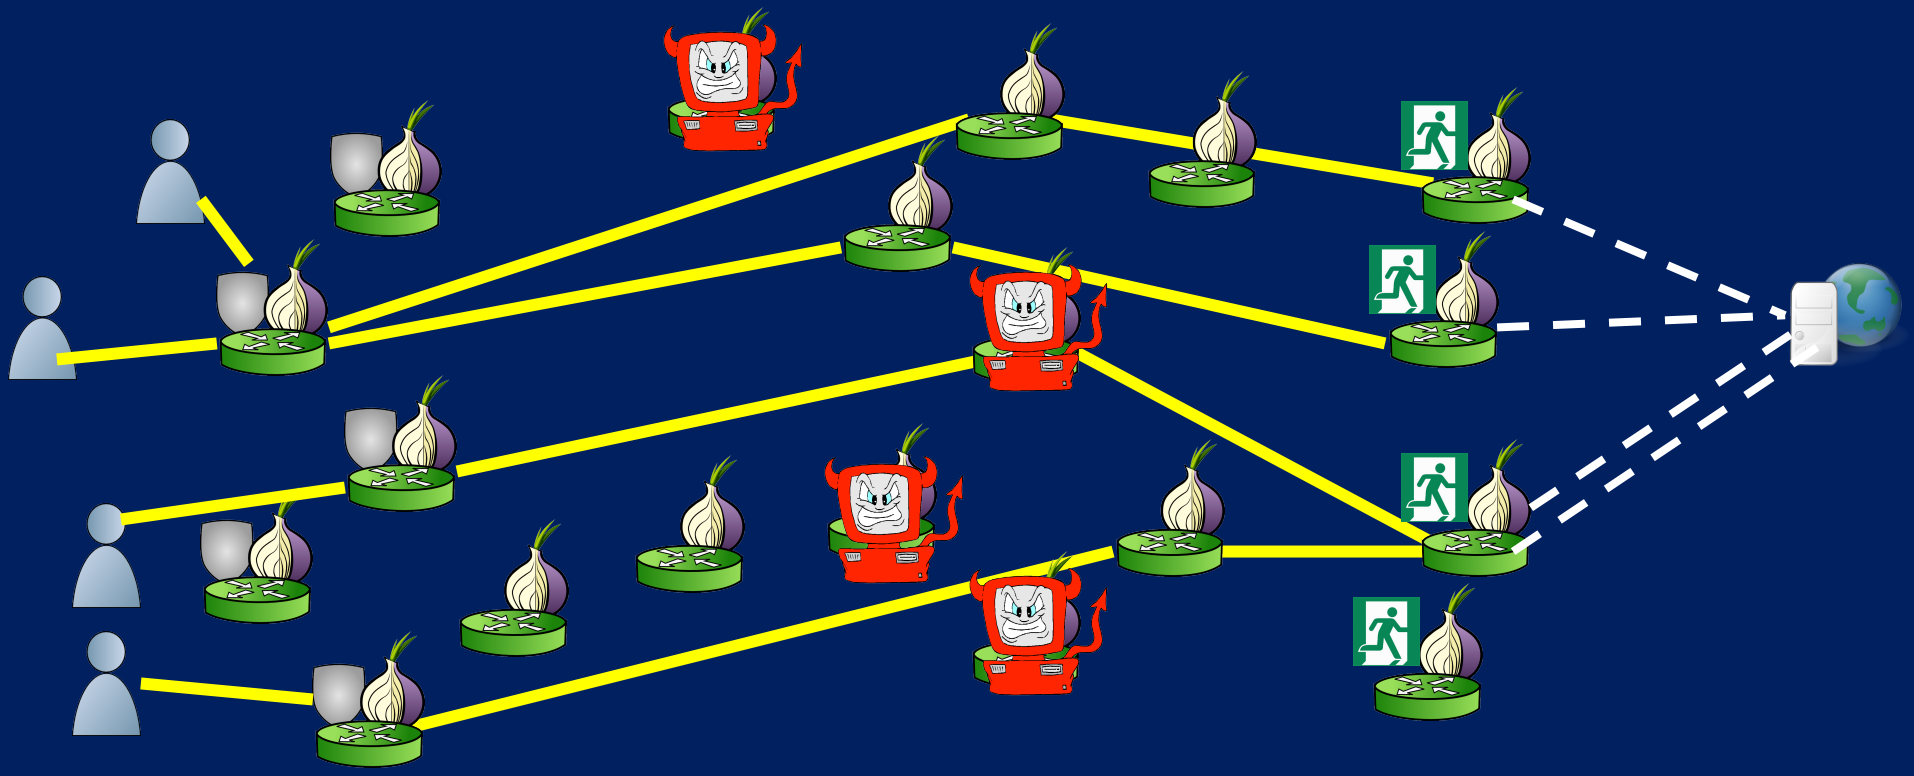

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

1. Own as many middle relays as possible
2. Discover a cabal guard whenever circuit goes through you

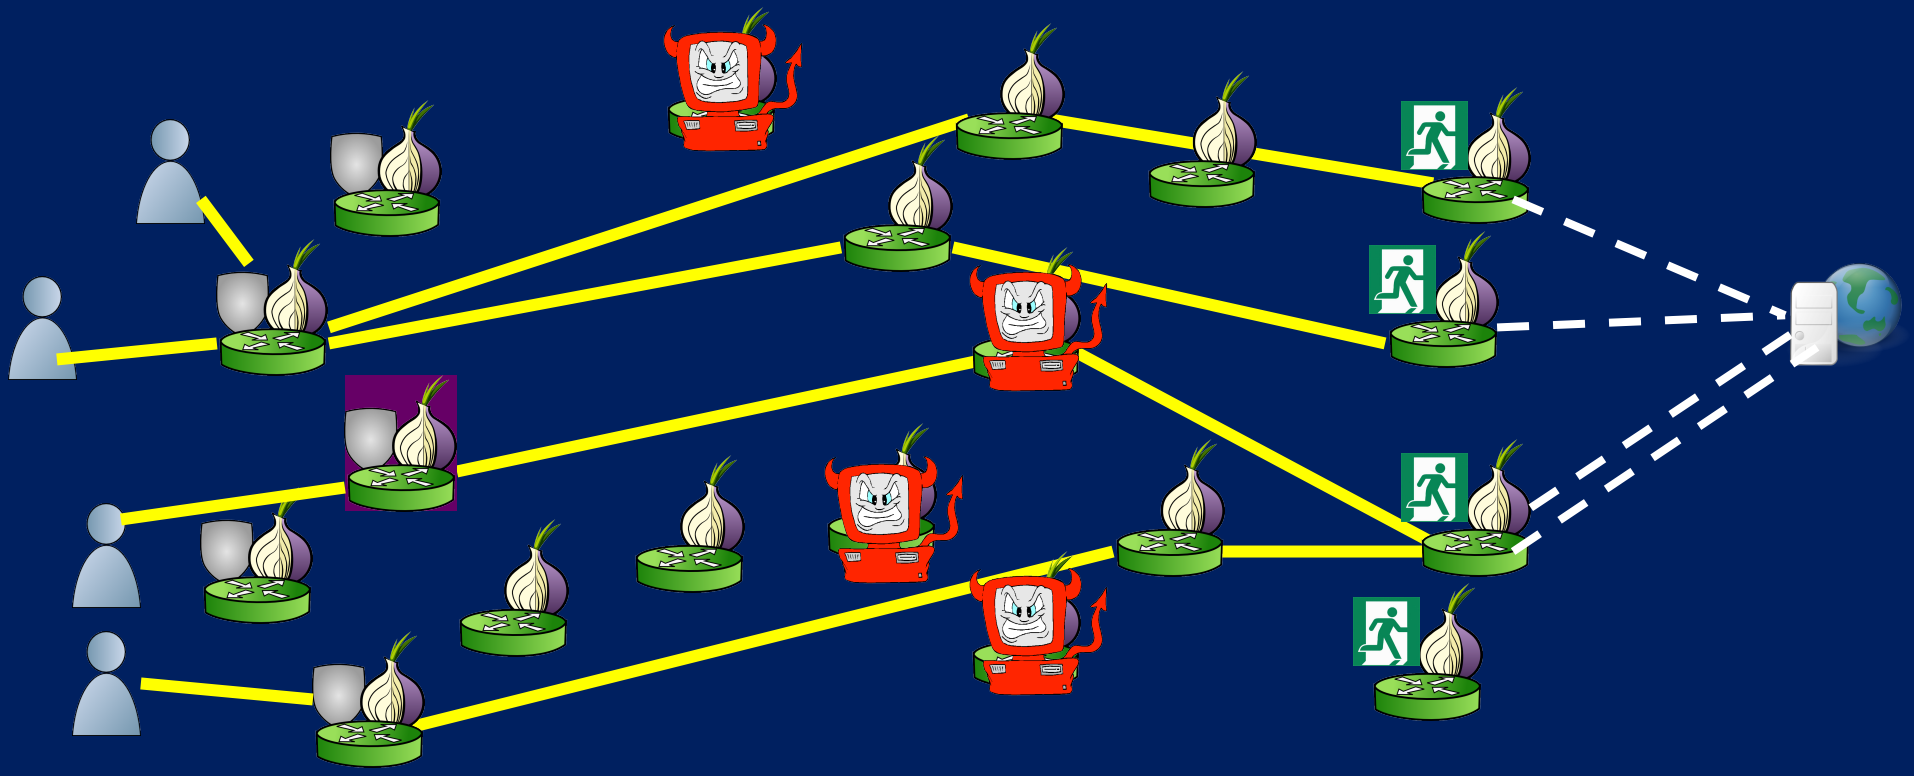

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

1. Own as many middle relays as possible
2. Discover a cabal guard whenever circuit goes through you

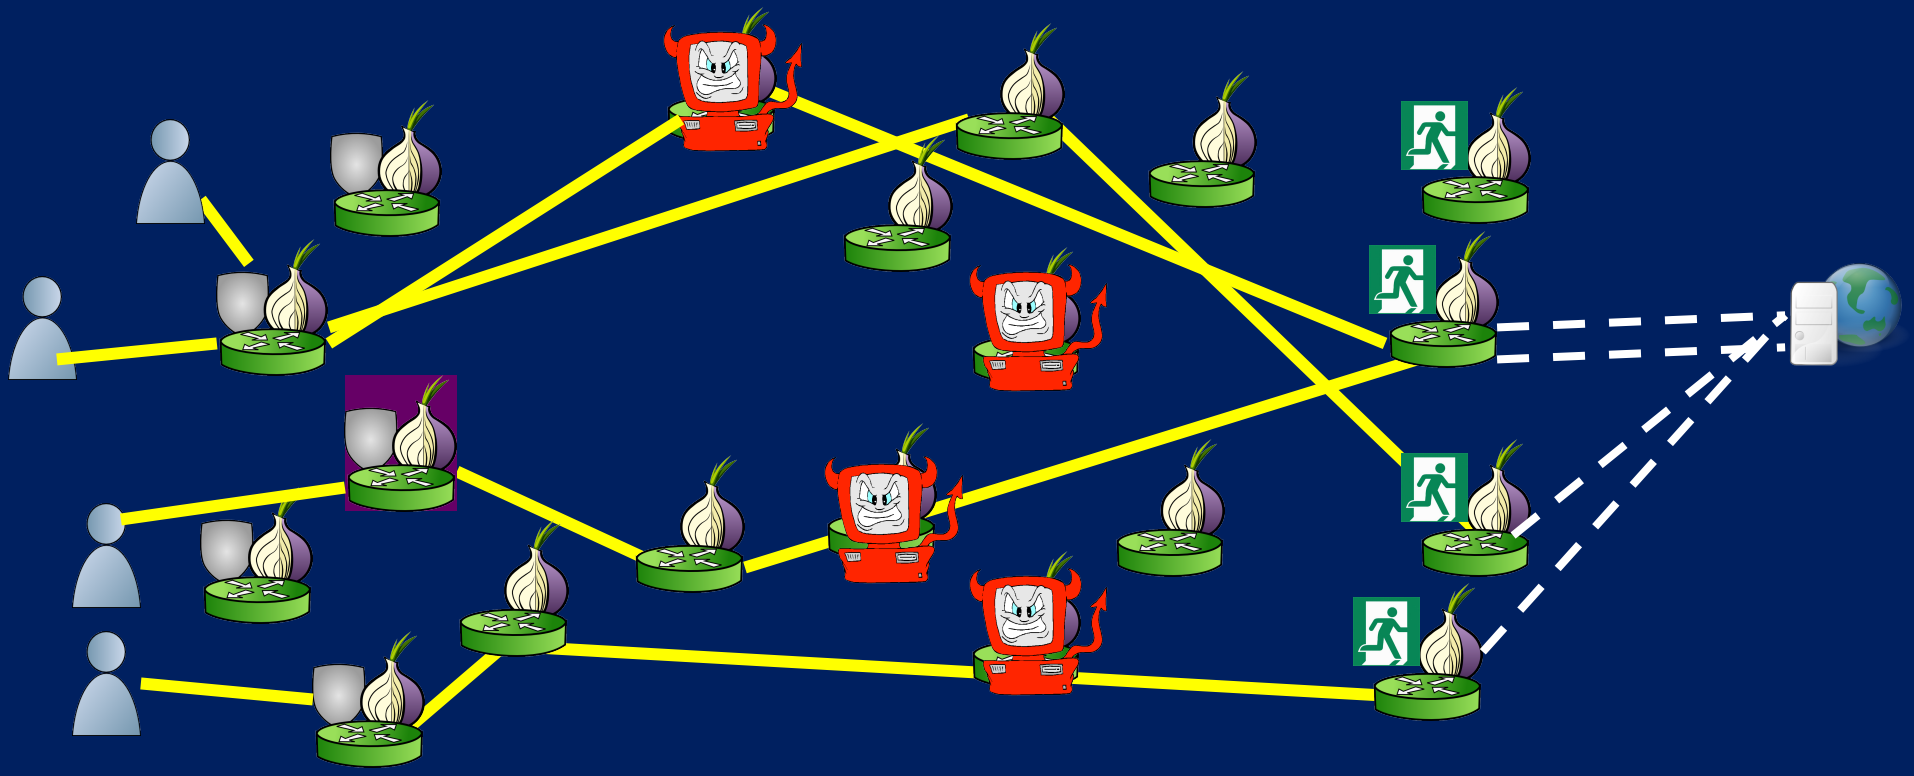

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

1. Own as many middle relays as possible
2. Discover a cabal guard whenever circuit goes through you

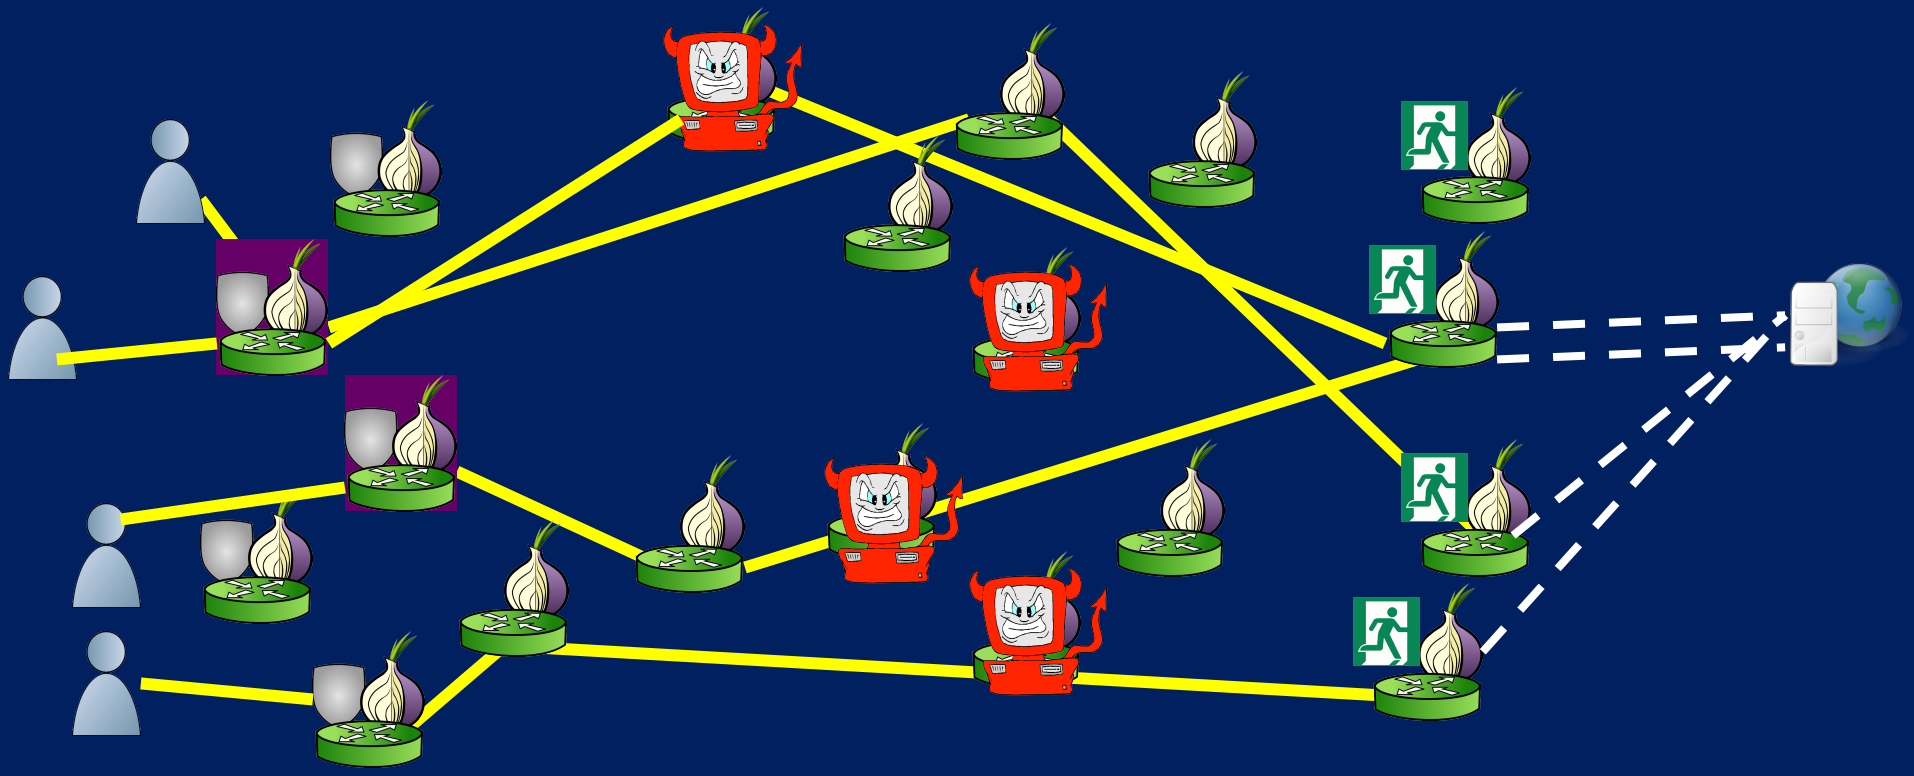

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

1. Own as many middle relays as possible
2. Discover a cabal guard whenever circuit goes through you

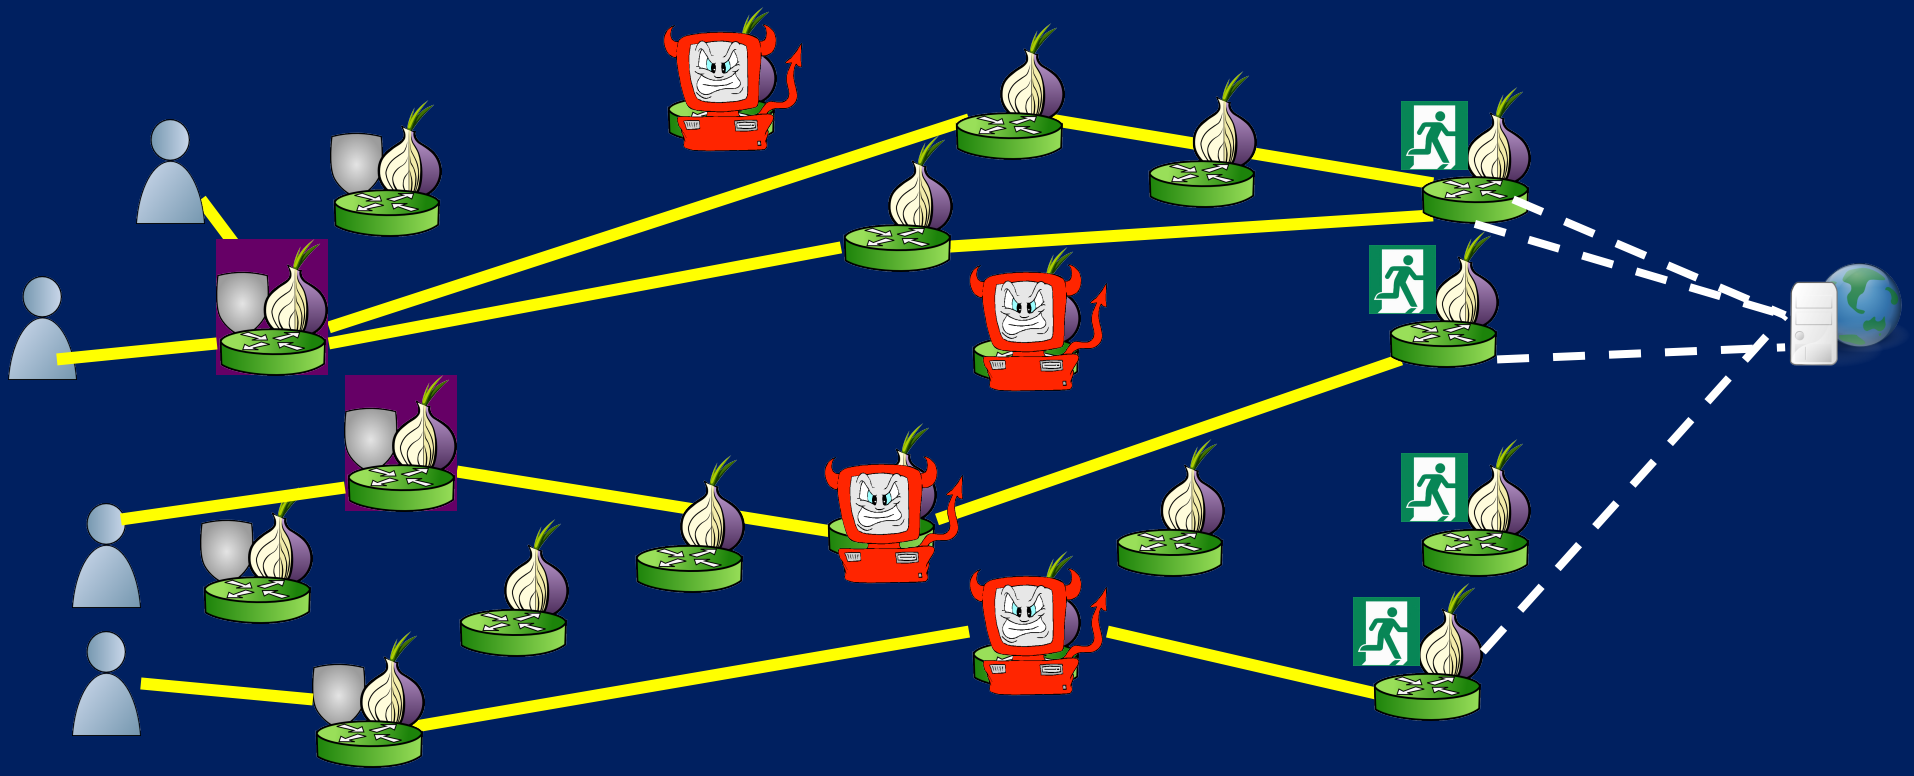

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

1. Own as many middle relays as possible
2. Discover a cabal guard whenever circuit goes through you

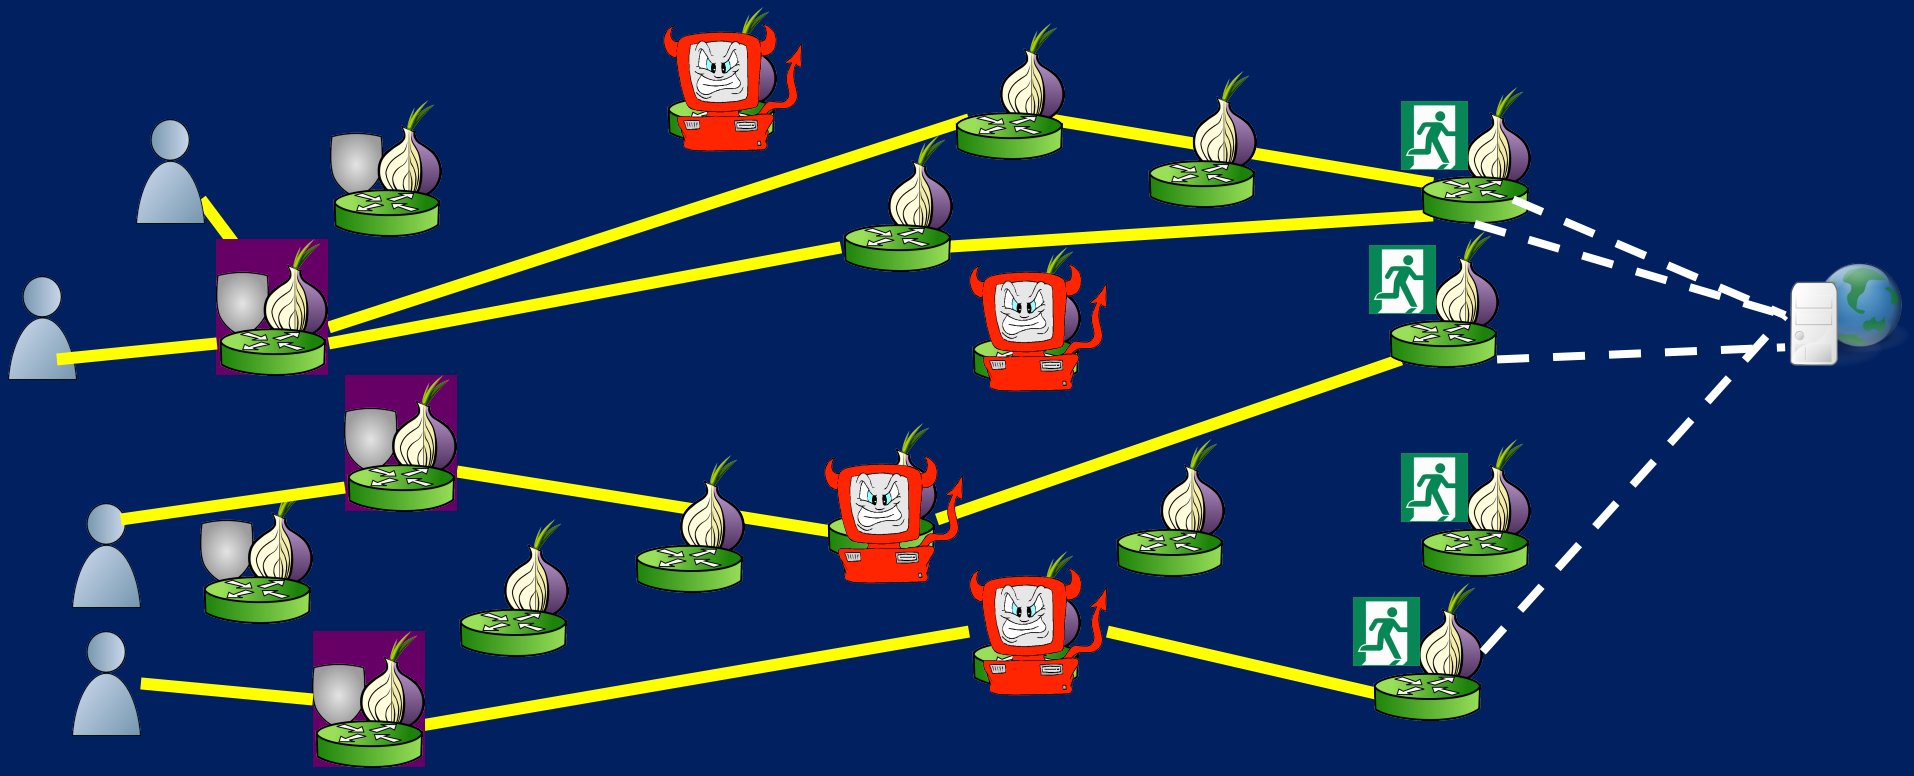

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

1. Own as many middle relays as possible
2. Discover a cabal guard whenever circuit goes through you
3. Bridge or compromise guards of all/interesting cabal members

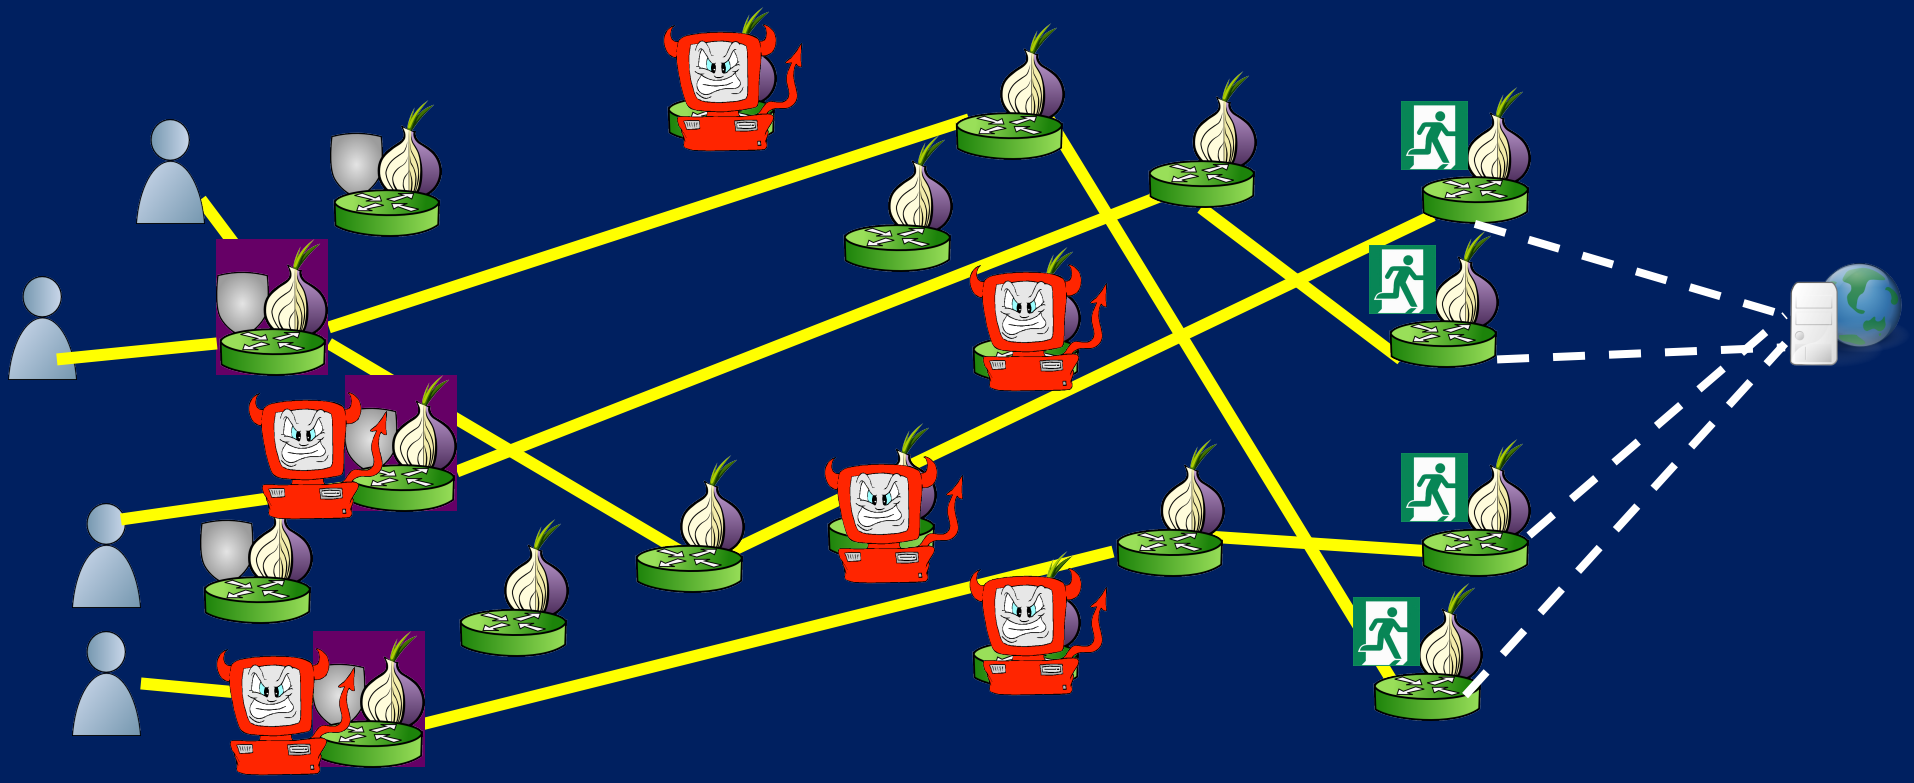

U.S. NAVAL RESEARCH LABORATORY

A Targeting Strategy to Attack IRC Cabal

Adversary assumptions:

- Compromises middle relays independently with prob. B (fraction of owned bandwidth)
- Is targeting the relevant cabal
- Owns either
  - One cabal member
  - ISP or guard of one cabal member
- Meetings are long/varied-in-traffic enough to identify a cabal-meeting circuit

U.S. NAVAL RESEARCH LABORATORY

Identifying at least one guard per cabal member

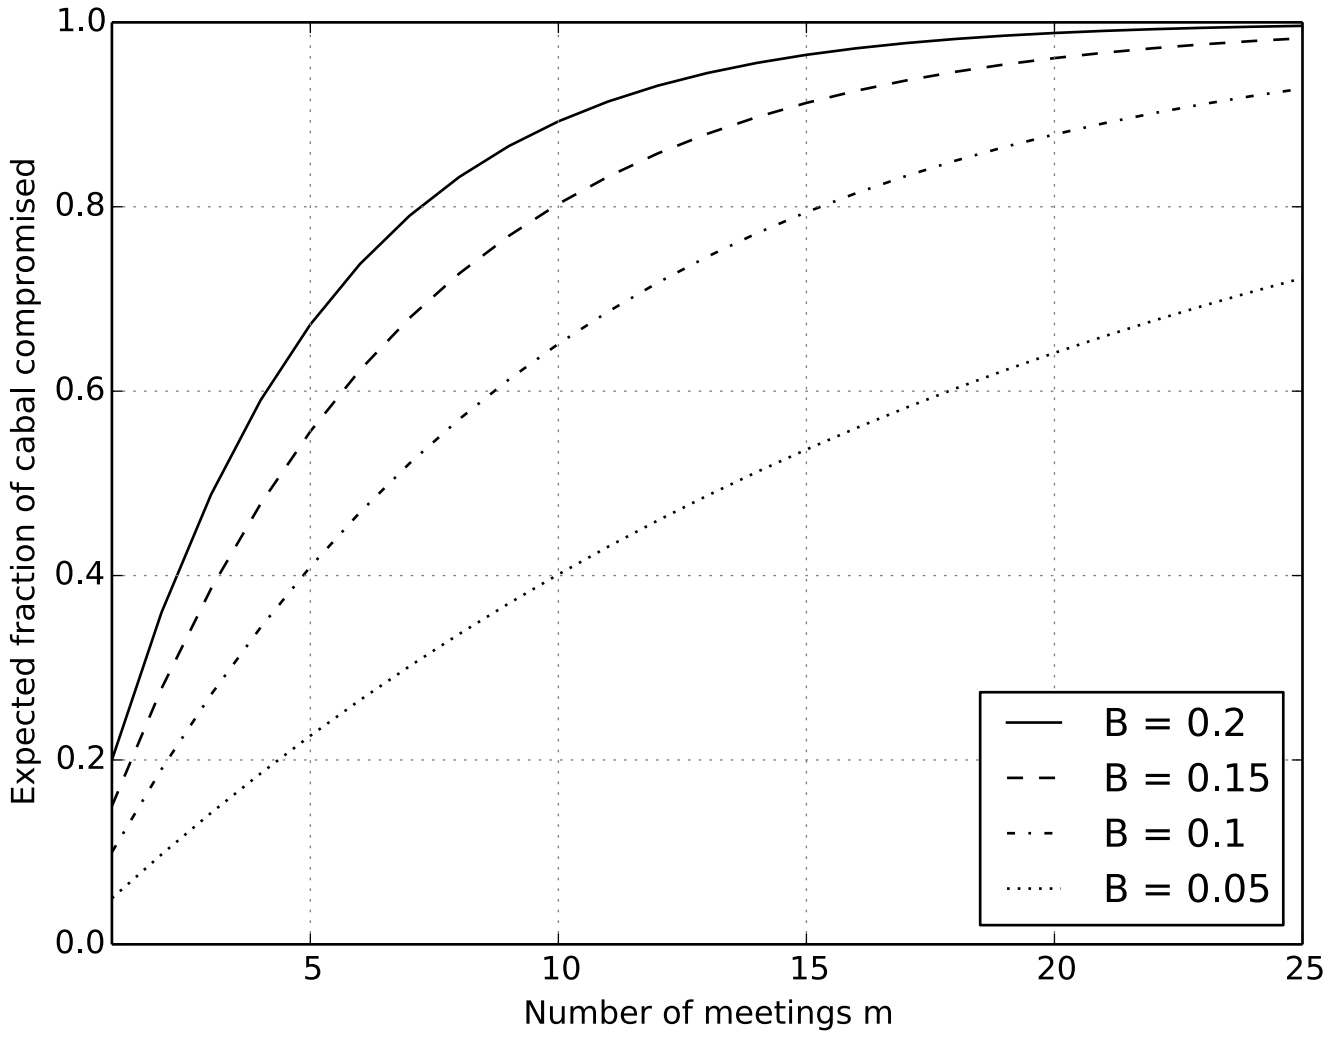

U.S. NAVAL RESEARCH LABORATORY

Estimating cabal size: MLE after five meetings

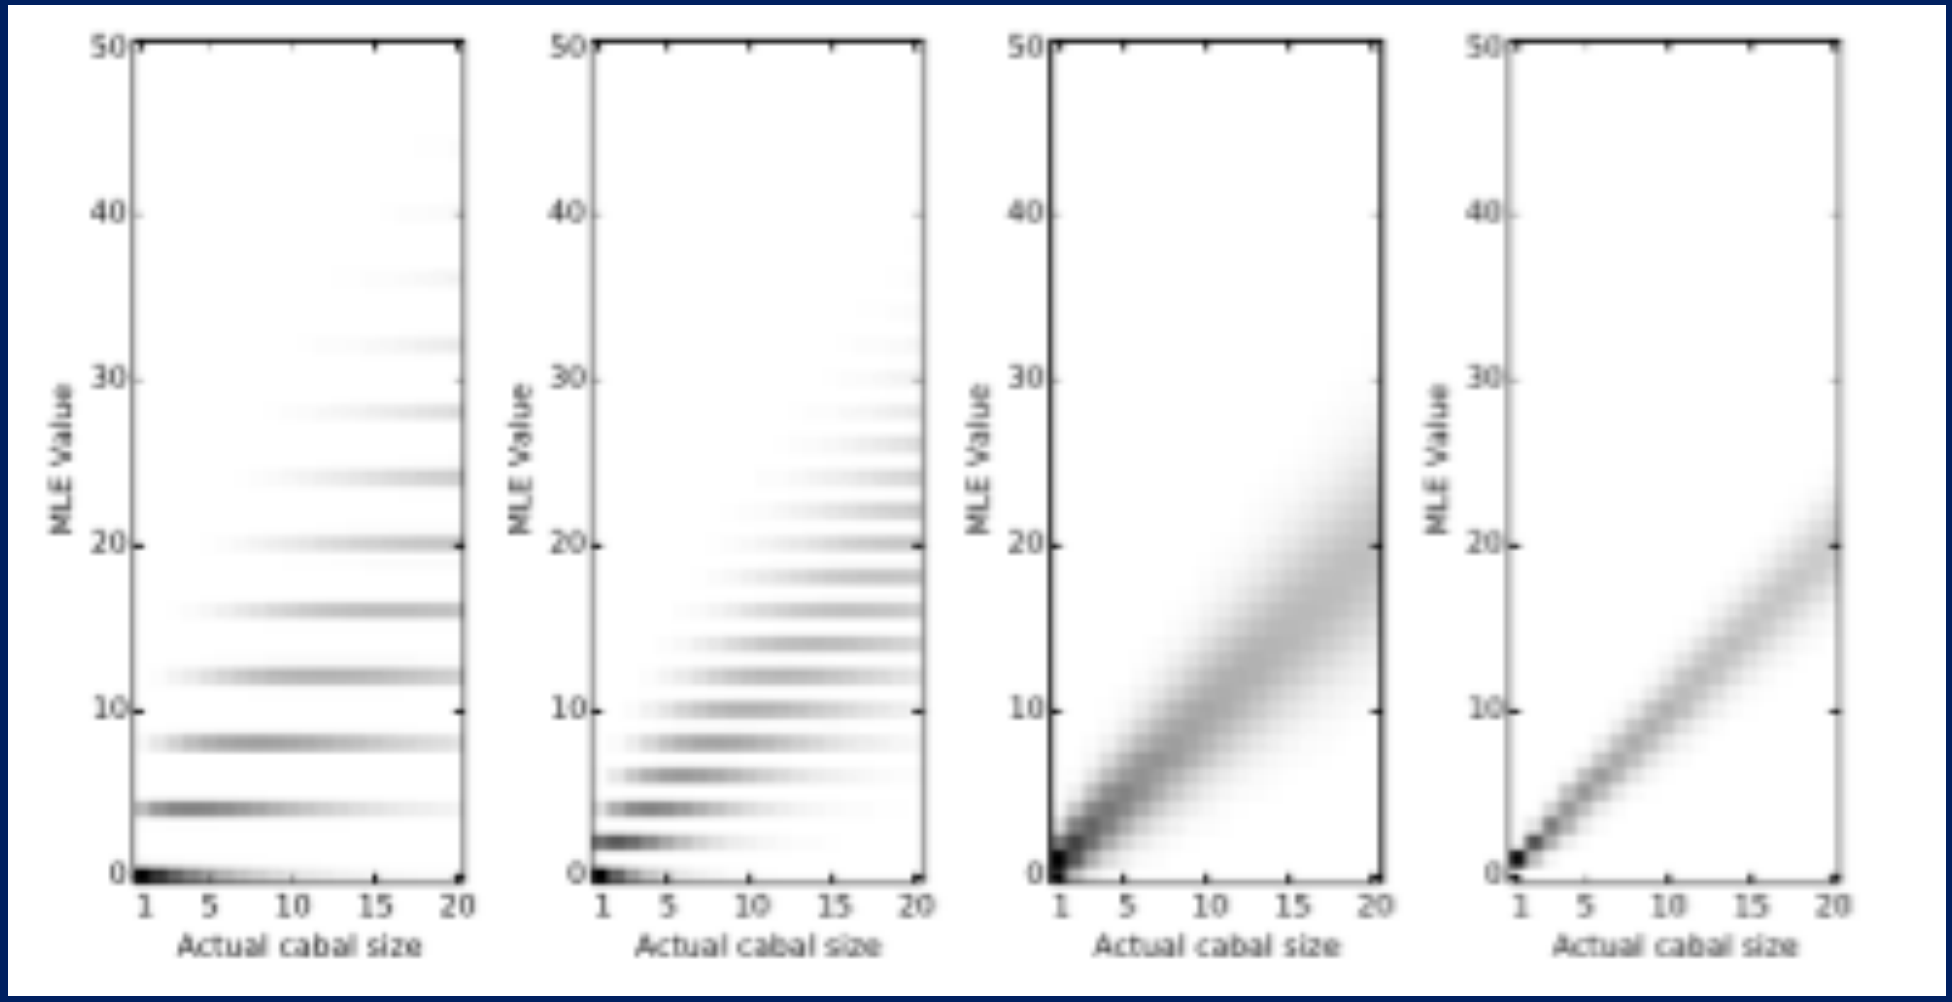

B= .05 B= .1 B= .2 B= .5

U.S. NAVAL RESEARCH LABORATORY

Chances of bridging to cabal leader

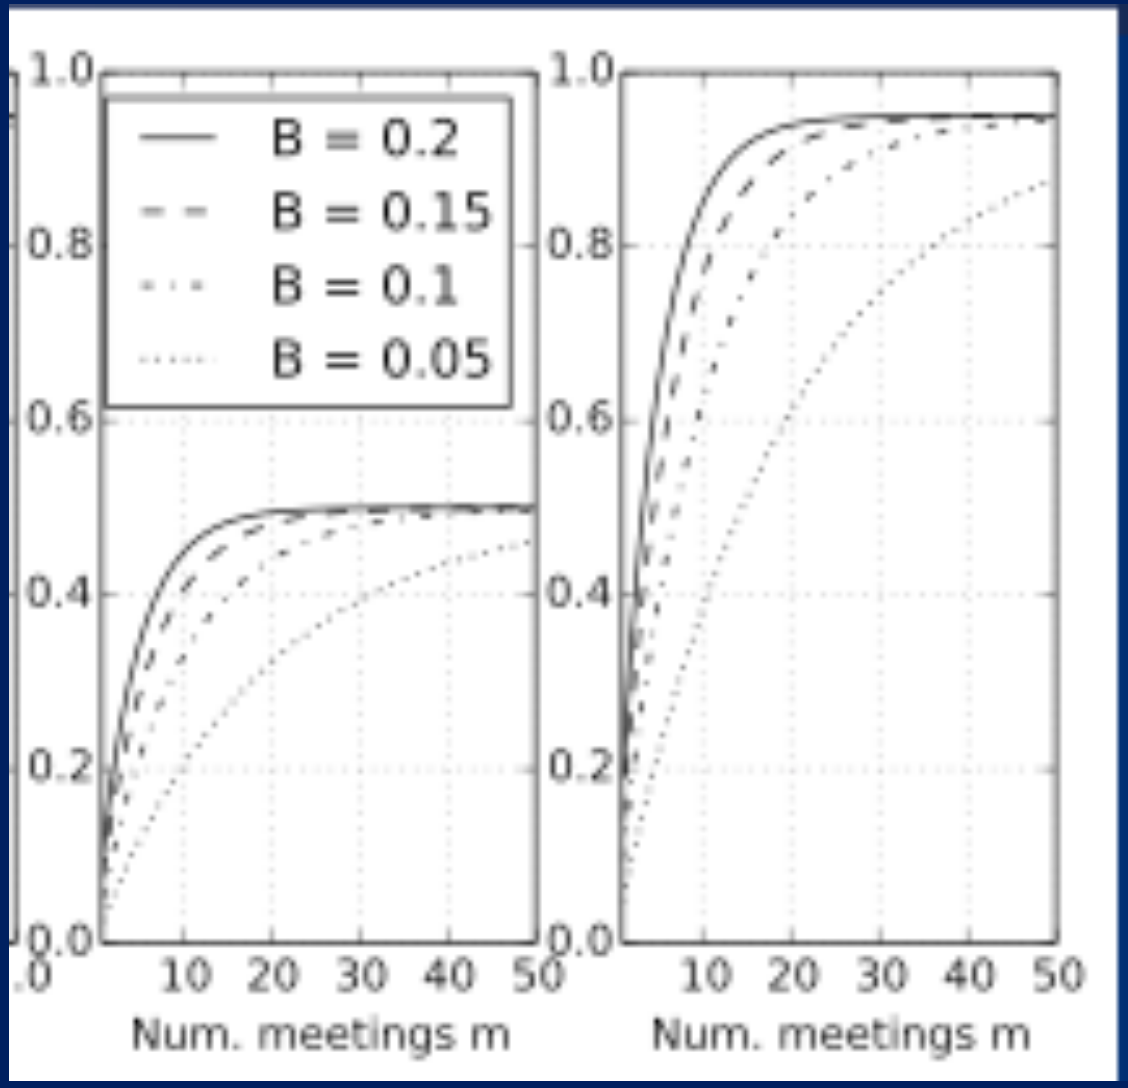

If one guard per client

Probability of Bridging Guard P_b= 0.5 P_b= 0.9

U.S. NAVAL RESEARCH LABORATORY

Chances of bridging to cabal leader

Bridging a guard, e.g., via

- Compromise guard
- Compromise ISP
- Coerce or extort owner/ operator of guard or ISP
- Network attack:
  90% of Tor relays subject to BGP prefix hijack
  -"RAPTOR" (USENIX Sec 15)

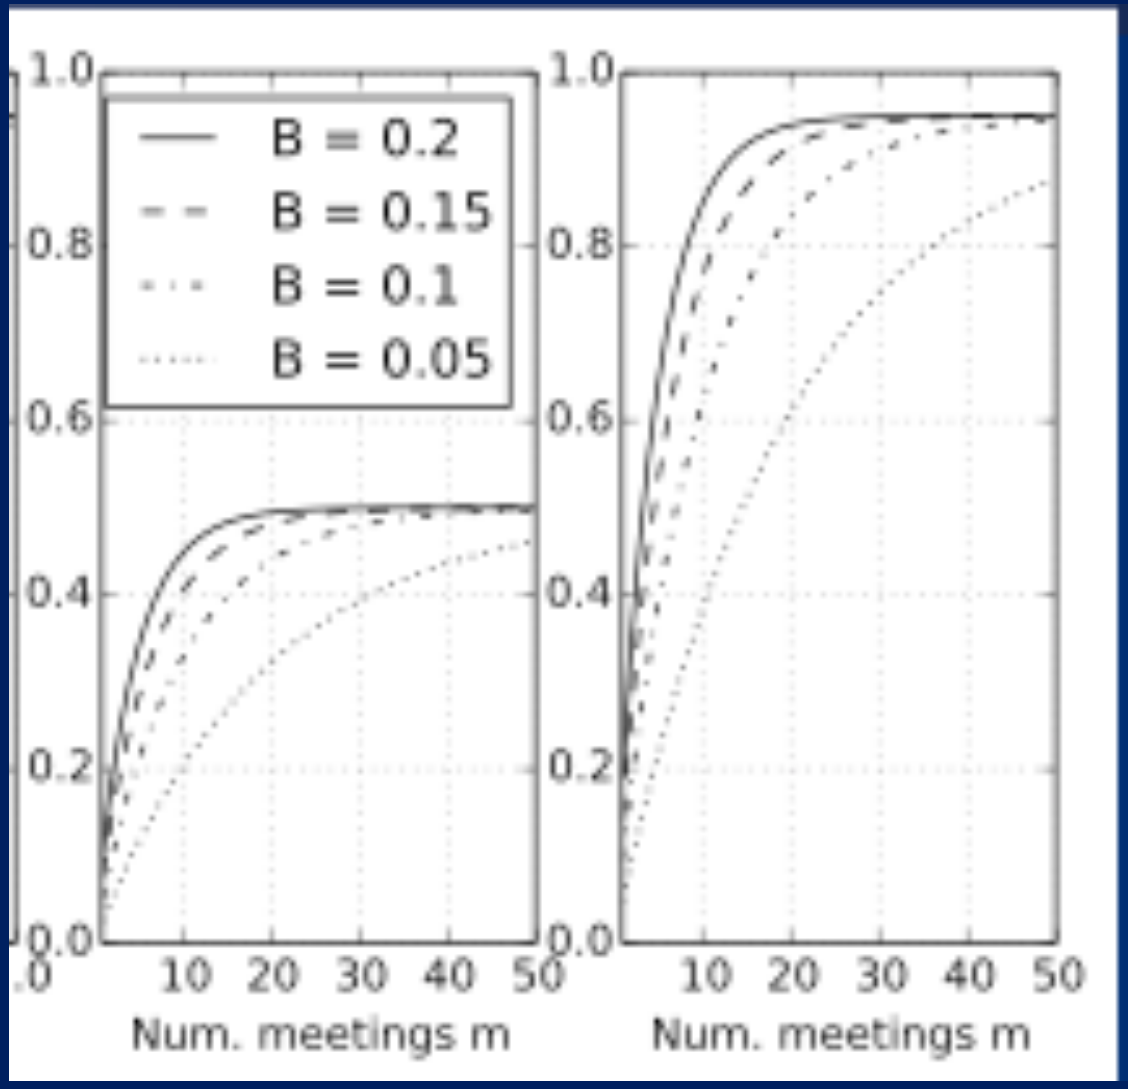

If one guard per client

Probability of Bridging Guard P_b= 0.5 P_b= 0.9

U.S. NAVAL RESEARCH LABORATORY

Targeting vs. Hoovering IRC Cabals

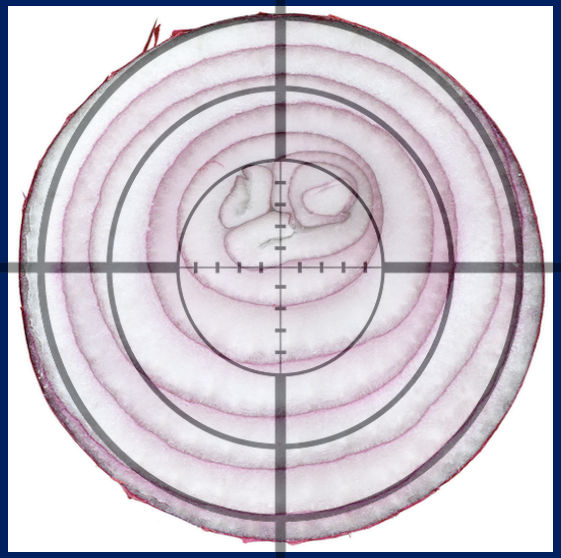

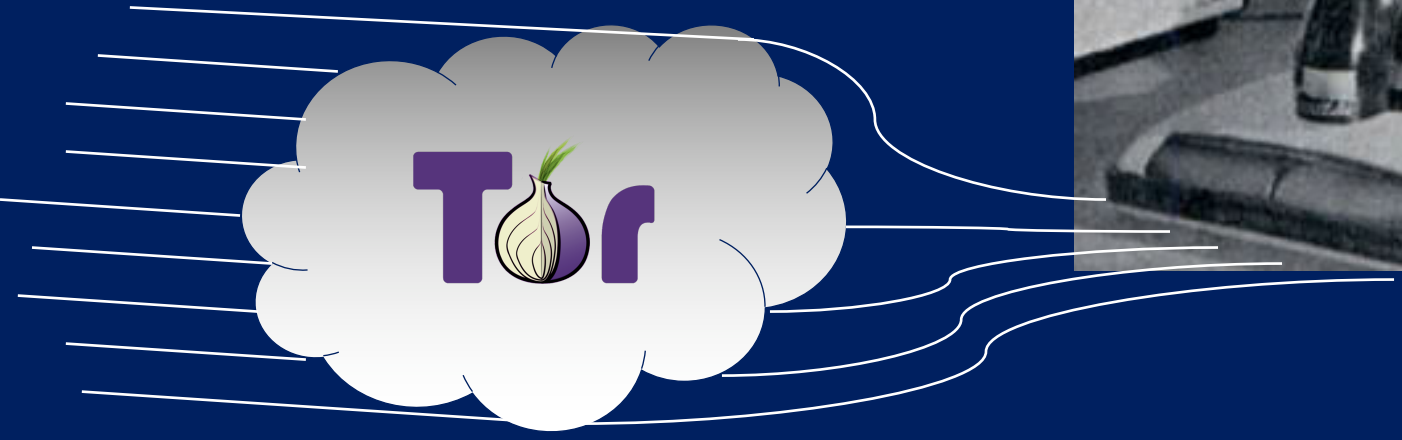

U.S. NAVAL RESEARCH LABORATORY

Targeting vs. Hoovering IRC Cabals

Hoovering adversary (Johnson et al. CCS 2013)

- Examined users making same IRC connection 27 times/day
- Adversary owned c. 4% of relay bandwidth (optimally distributed)
- For cabal with 10-20 members
  - 150-200 days to have identified almost all cabal guards

Targeting adversary (as above)

- After c. 4 days identifies guards of almost all members
  - Good idea of cabal size
  - Good chance of knowing leaders' guards
  - Much faster to steady state (1 day vs. 1 week)
- Has decision points and feedback while conducting attack
  - Spin up more/less relays based on daily reports
  - Do BGP hijack, contact ISP, zero-day a guard, etc.

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

Like IRC cabal example, uses only client-side of connections

Deployed counters to onionsite directory mining do nothing against these attacks

Contemplated onion-service-side protections do nothing to counter these attacks

- layered guards for hidden onion services
- link or multi-hop padding

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

- Capture-recapture basics

Population, N = ?

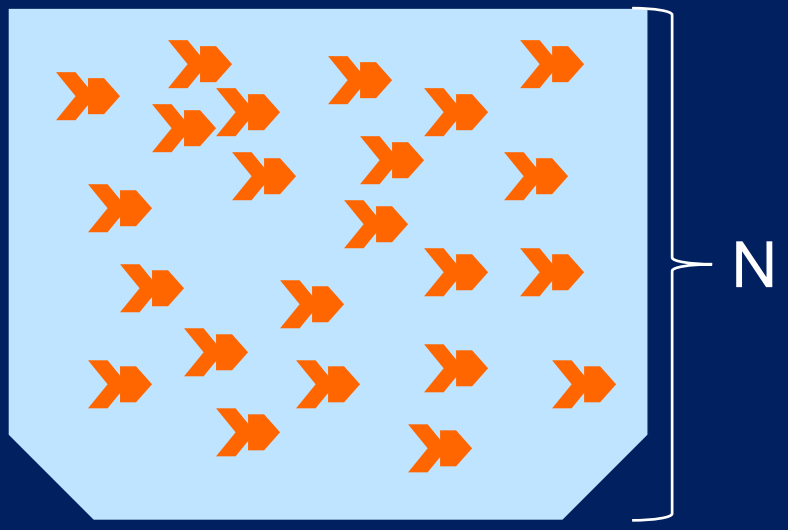

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

- Capture-recapture basics

Population, N = ?

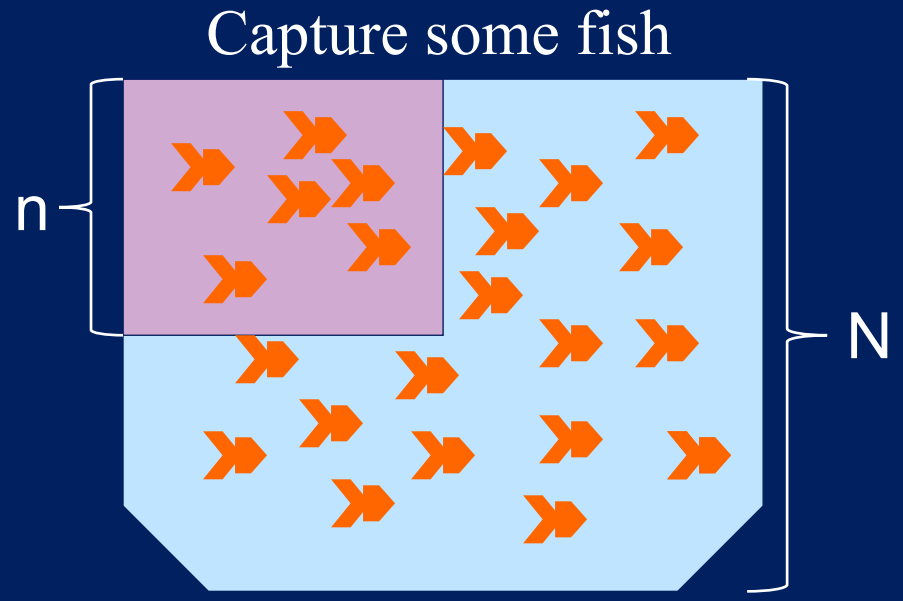

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

- Capture-recapture basics

Population, N = ?

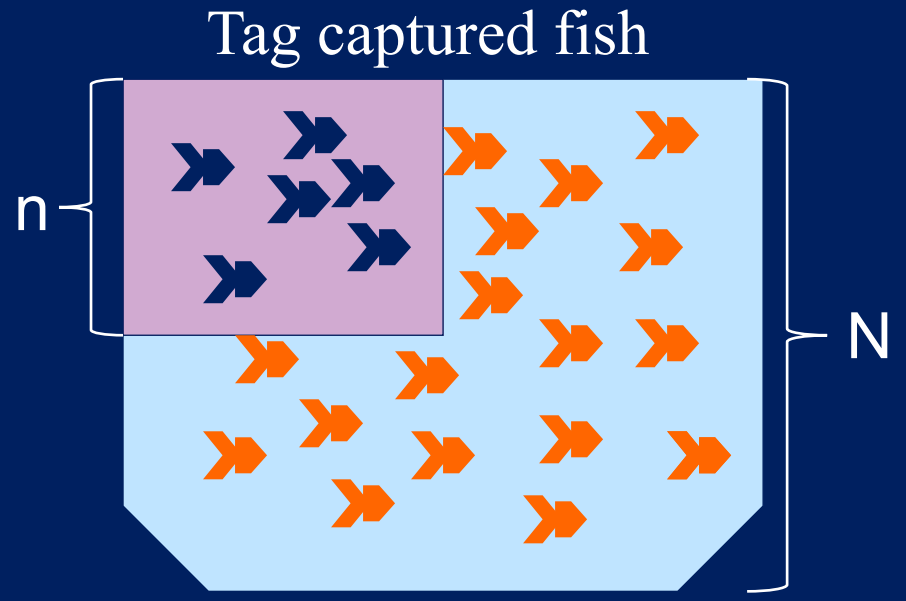

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

- Capture-recapture basics

Population, N = ?

n/N = ?

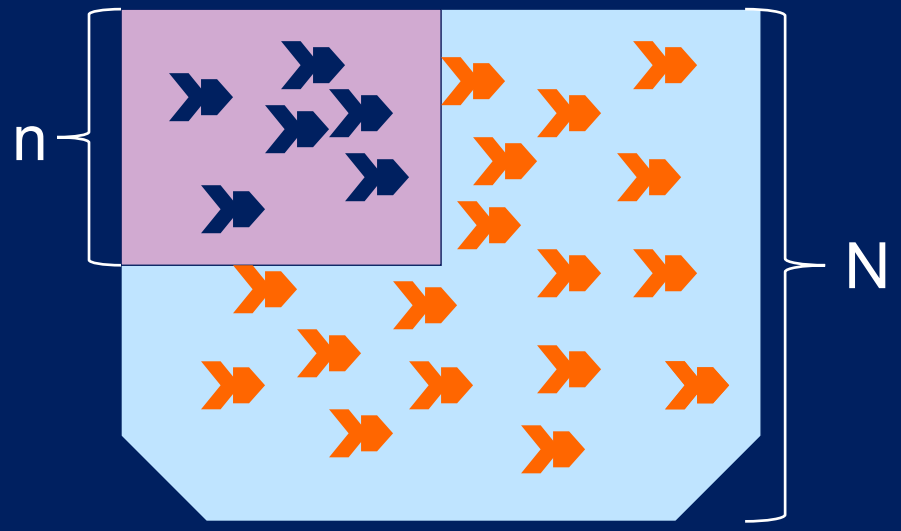

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

- Capture-recapture basics

Population, N = ?

n/N = ?

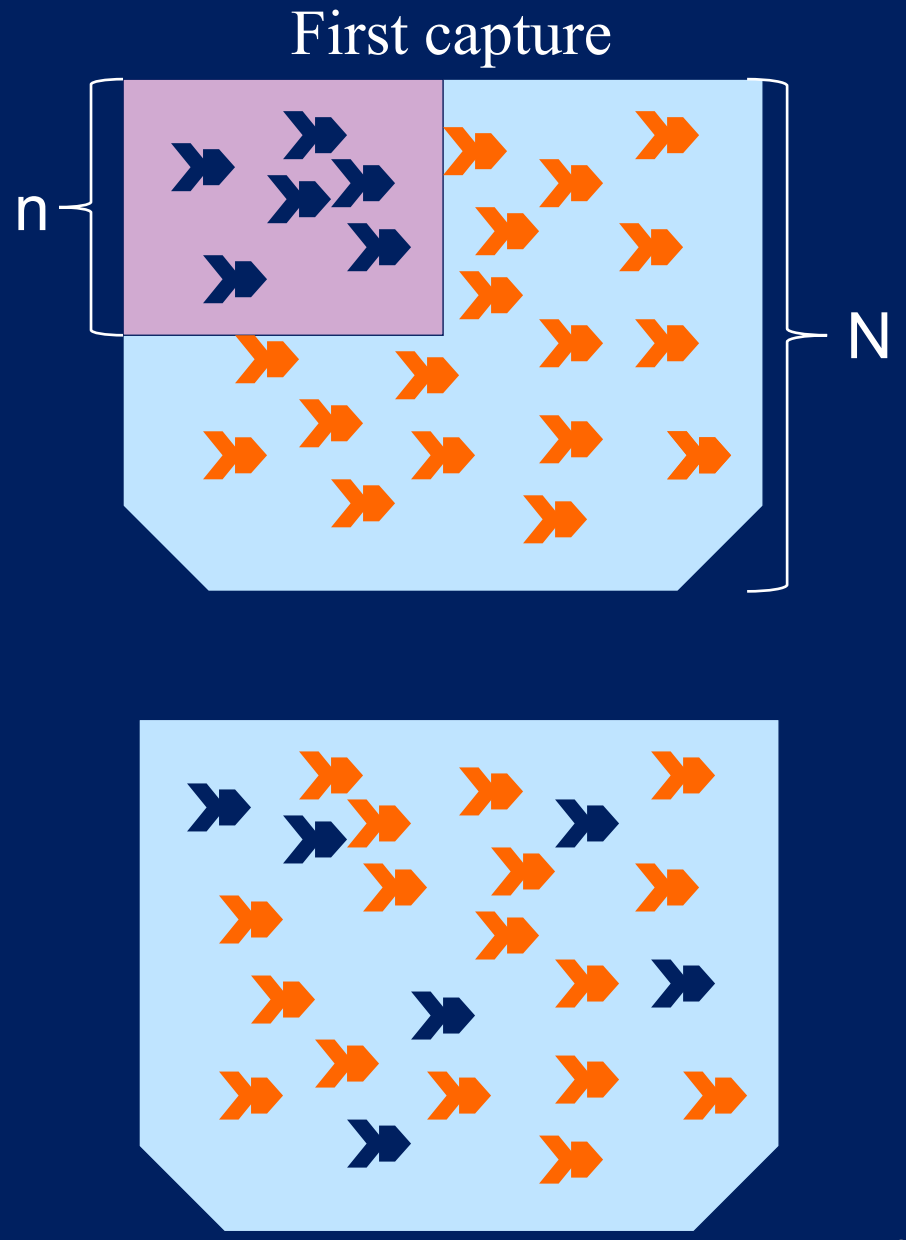

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

- Capture-recapture basics

Population, N = ?

n/N = k/K

N = nK/k

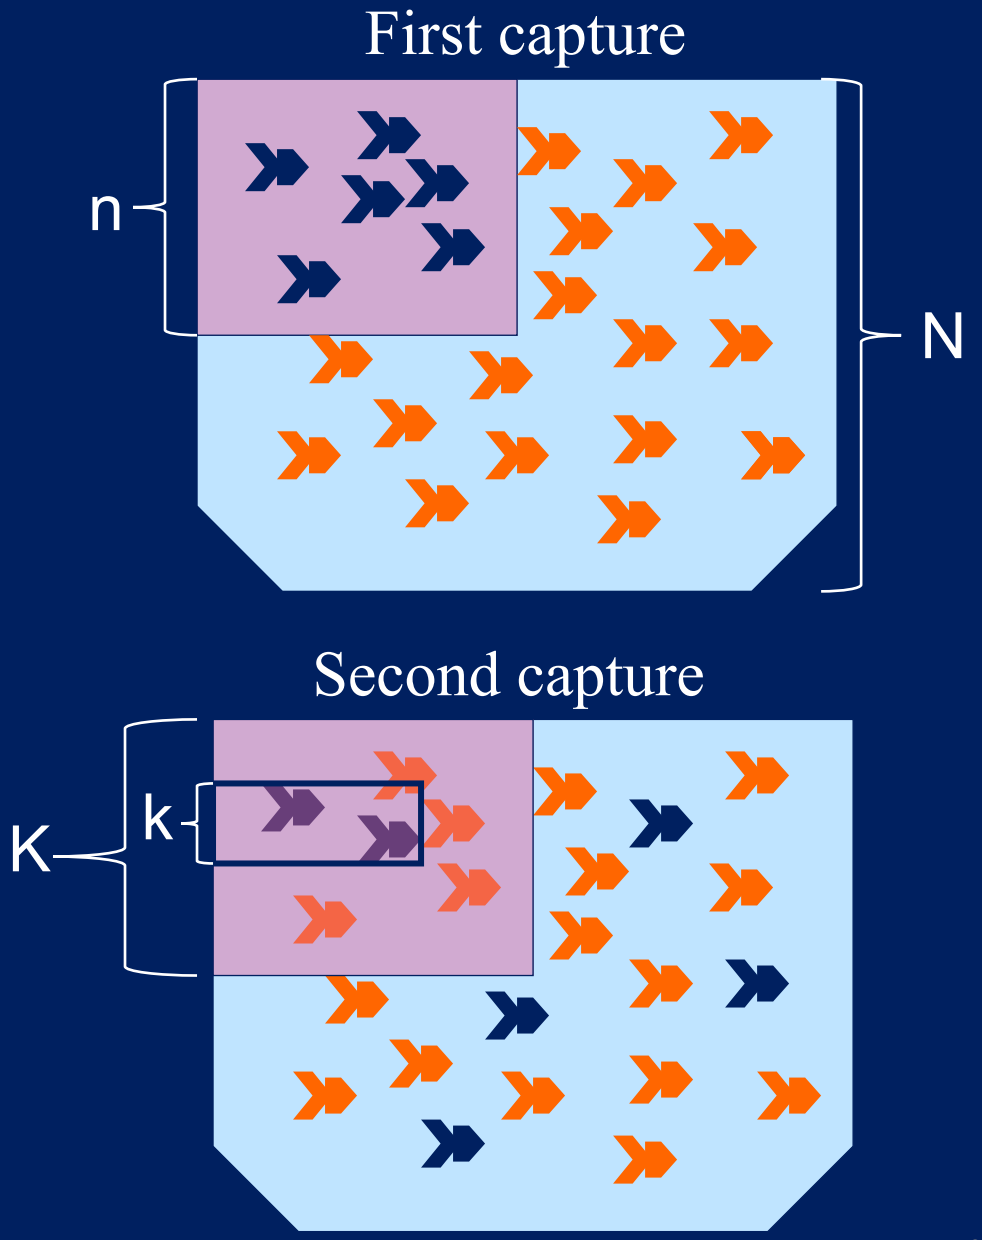

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

What's all this talk of catching fish?

How does it it help us catch targeted onionsite users?

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

- Capture-recapture basics

Population, N = ?

n/N = k/K

N = nK/k

- Onionsite *t.onion* popularity

N = # users visiting > m times/day

n = # guards seen originating more than m/day *t.on* circuits

K = # guards seen originating > m/day *t.on* circuits in 2nd interval

k = # tagged guards seen originating more than m/day *t.on* circuits in 2nd interval

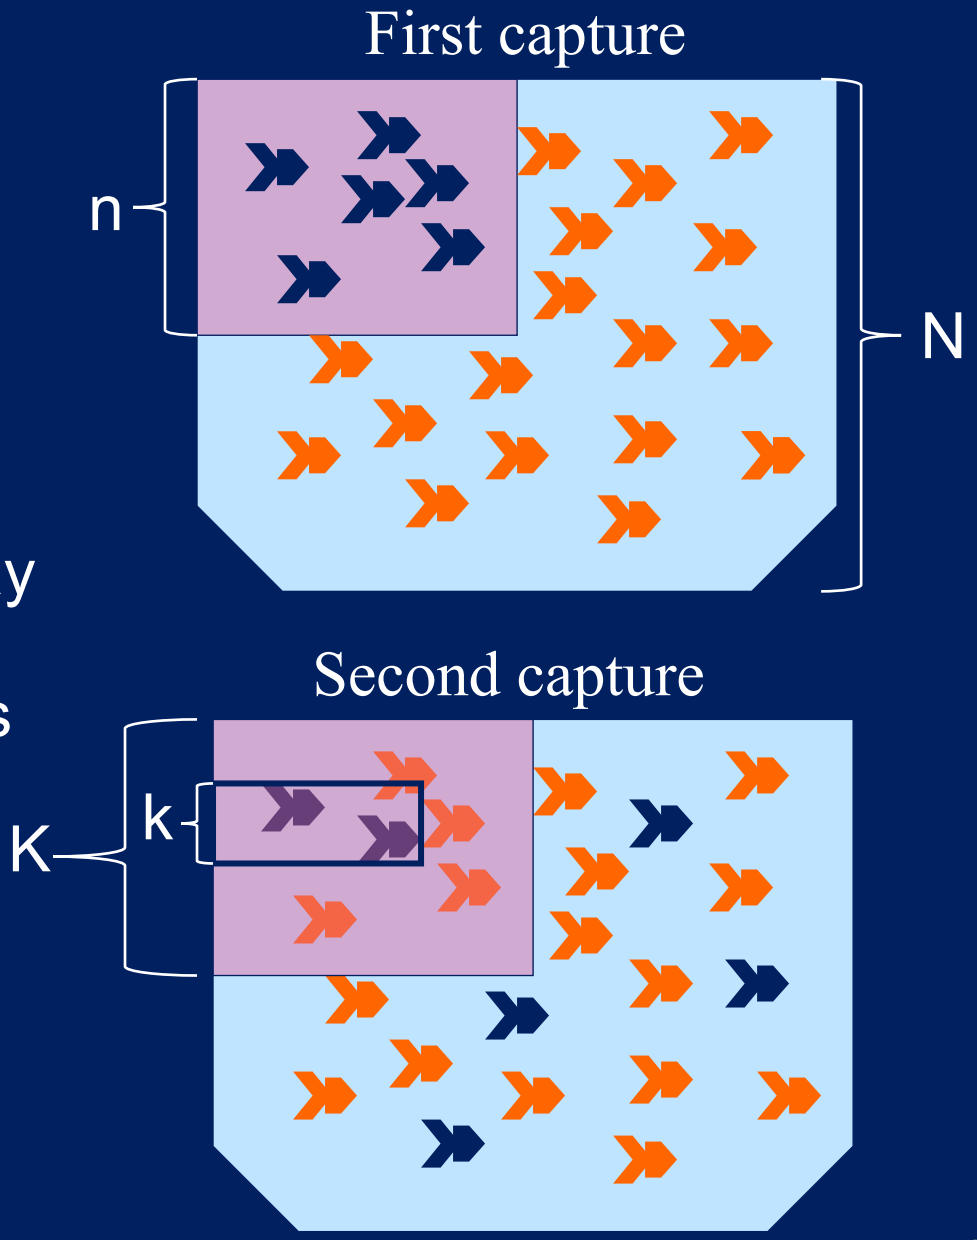

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

Important Assumption:

- Website fingerprinting from middle relay is effective
  - Set of persistent onionsites is small, and target likely unique
  - "Fingerprinting Hidden Service Circuits from a Tor Middle Relay" by Juarez et al. IEEE S&P 2017 Poster showed 99.98% accuracy

U.S. NAVAL RESEARCH LABORATORY

Targeting an onionsite's popularity and its visitors' activity

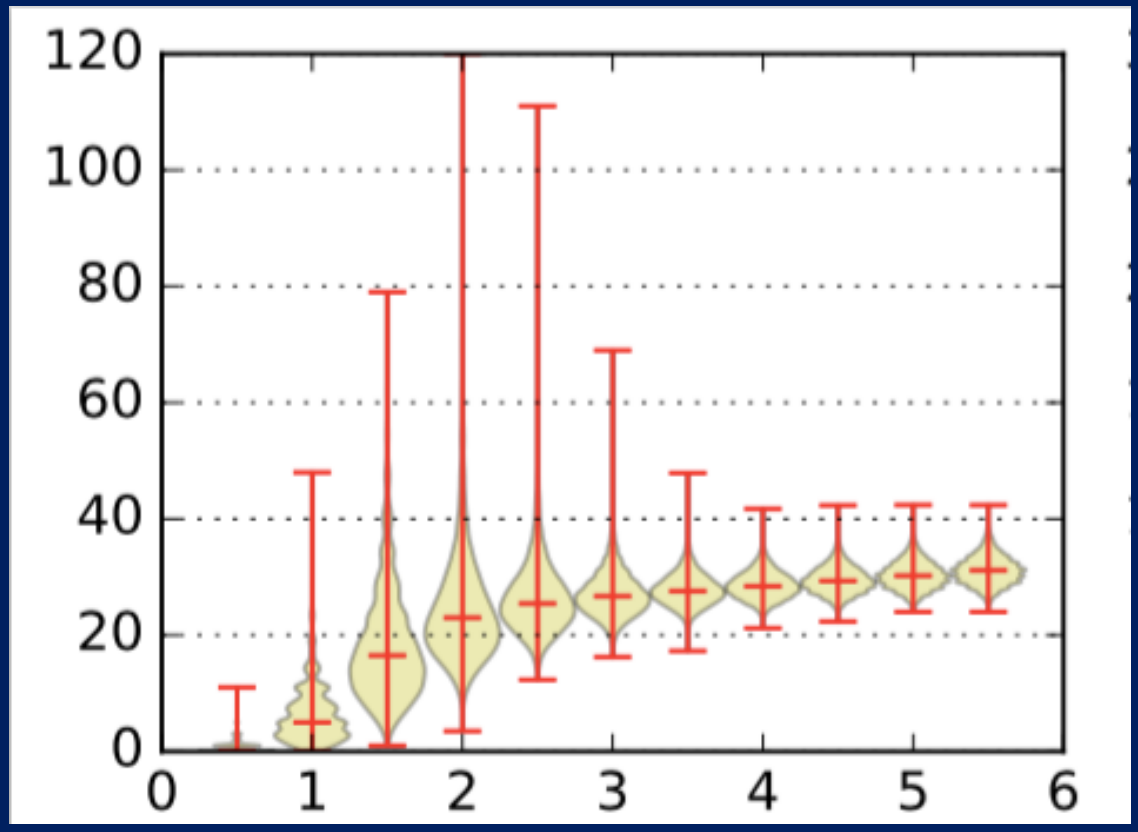

Clients:

- Regular
  visits target 2 x/day
- Interesting
  Visits target 10 x/day

Experiment: 25 interesting clients, 225 regular clients

- 2500 guards, 5000 middles (uniform)
- 10000 runs (capture guards, label, recapture guards, count)
- Capture/Label threshold = 3

U.S. NAVAL RESEARCH LABORATORY

Possible Countermeasures and Future Work

- Layered Guards (Vanguards)
  - First proposed 2006 in same paper that introduced guards
  - Now being finalized to protect hidden services from guard discovery

  Attack strategy still reveals onionsite activity if any relay is chosen per circuit

U.S. NAVAL RESEARCH LABORATORY

Possible Countermeasures and Future Work

- Layered Guards (Vanguards)
  - First proposed 2006 in same paper that introduced guards
  - Now being finalized to protect hidden services from guard discovery

  Attack strategy still reveals onionsite activity if any relay is chosen per circuit
- Randomizing selection of guard-set size & guard duration

U.S. NAVAL RESEARCH LABORATORY

Possible Countermeasures and Future Work

- Layered Guards (Vanguards)
  - First proposed 2006 in same paper that introduced guards
  - Now being finalized to protect hidden services from guard discovery

  Attack strategy still reveals onionsite activity if any relay is chosen per circuit
- Randomizing selection of guard-set size & guard duration
- Standardized traffic templates for sensitive onion services

U.S. NAVAL RESEARCH LABORATORY

Possible Countermeasures and Future Work

- Layered Guards (Vanguards)
  - First proposed 2006 in same paper that introduced guards
  - Now being finalized to protect hidden services from guard discovery

  Attack strategy still reveals onionsite activity if any relay is chosen per circuit
- Randomizing selection of guard-set size & guard duration
- Standardized traffic templates for sensitive onion services

- Adding network links to adversary endowment
- Adding dynamics (of relays, of Internet, of client behavior)
- Example attacks targeting behaviors of a particular client

U.S. NAVAL RESEARCH LABORATORY

Questions?

- Be aware of targeting adversaries
- Move them off target