Resilient Anonymous Communication for Everyone (RACE)

Dr. Joshua Baron
RACE@darpa.mil

Proposers Day

24 July 2018

RACE Goal

Use cryptography and obfuscated communications to build an anonymous, attack-resilient mobile communication system that can reside completely within a contested network environment.

RACE Overview

- 48 month effort with three phases (18 months/12 months/18 months)
- Four TAs to bid on: TA1, TA2, TA3 and TA3.1
- Multiple awards anticipated for TA1 and for TA2, single award anticipated for TA3, single award anticipated for TA3.1
- Total anticipated award amount: ~$44 Million
- Abstract deadline: 12 noon ET, August 14, 2018 (HIGHLY recommended)
- Full proposal deadline: 12 noon ET, September 18, 2018
- One abstract/proposal per TA (can submit multiple abstracts/proposals)
- Email: RACE@darpa.mil

Mobile Communications in Contested Environments

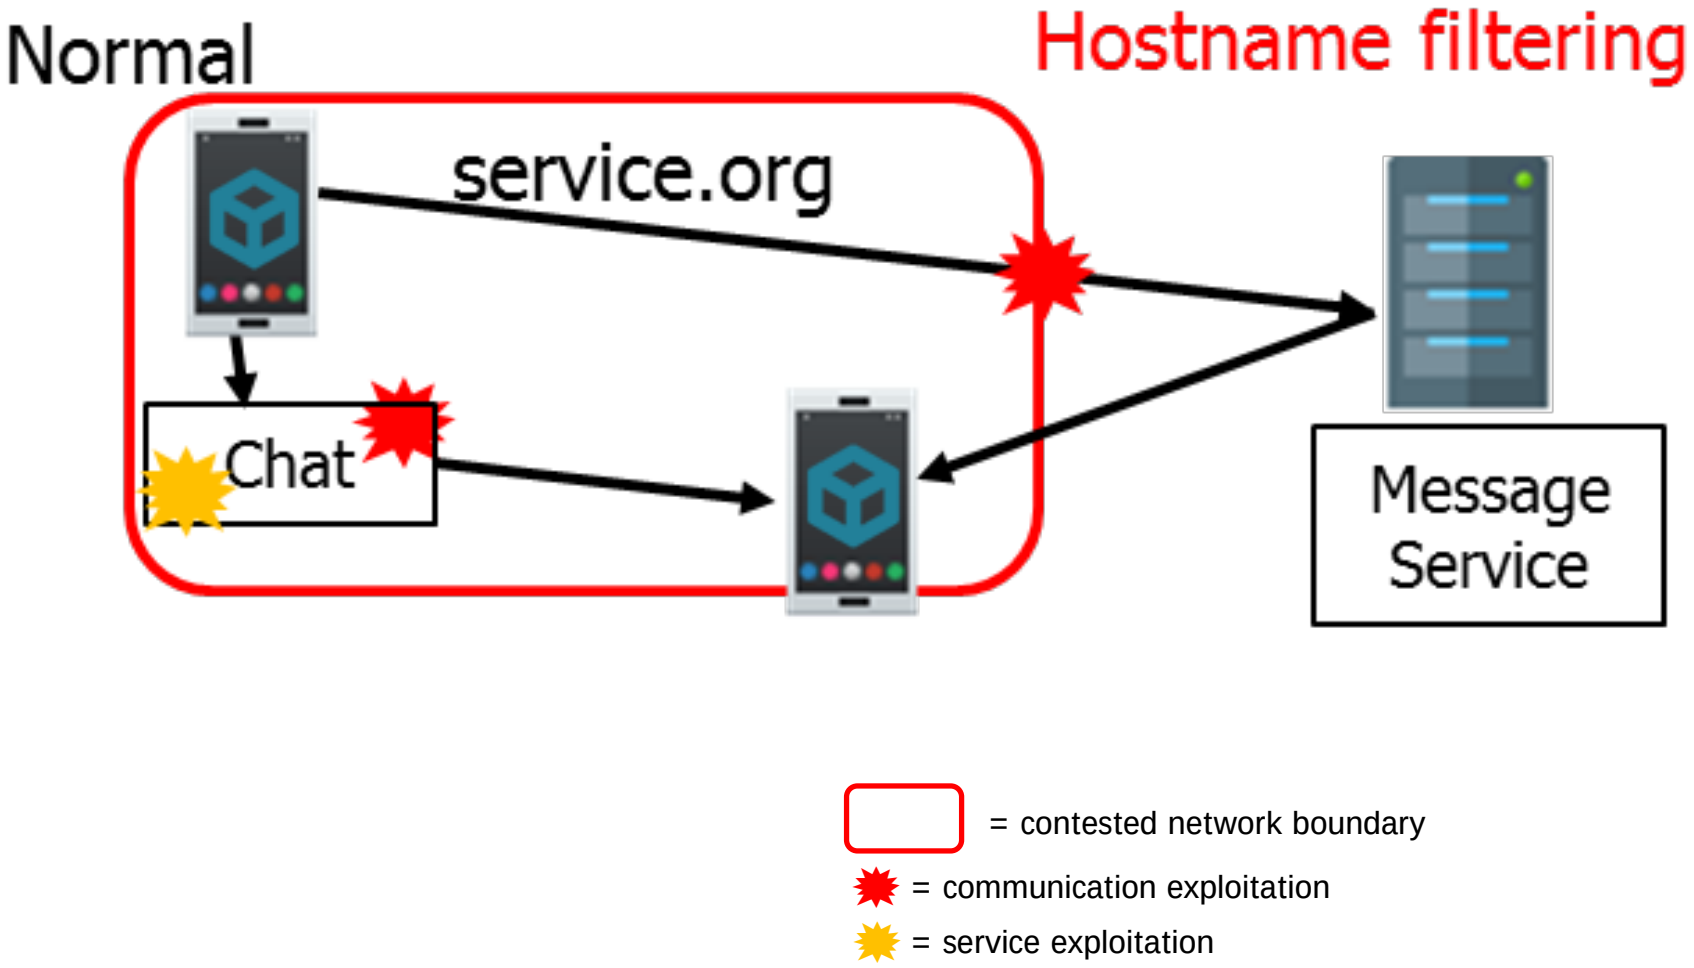

Mobile Communications in Contested Environments

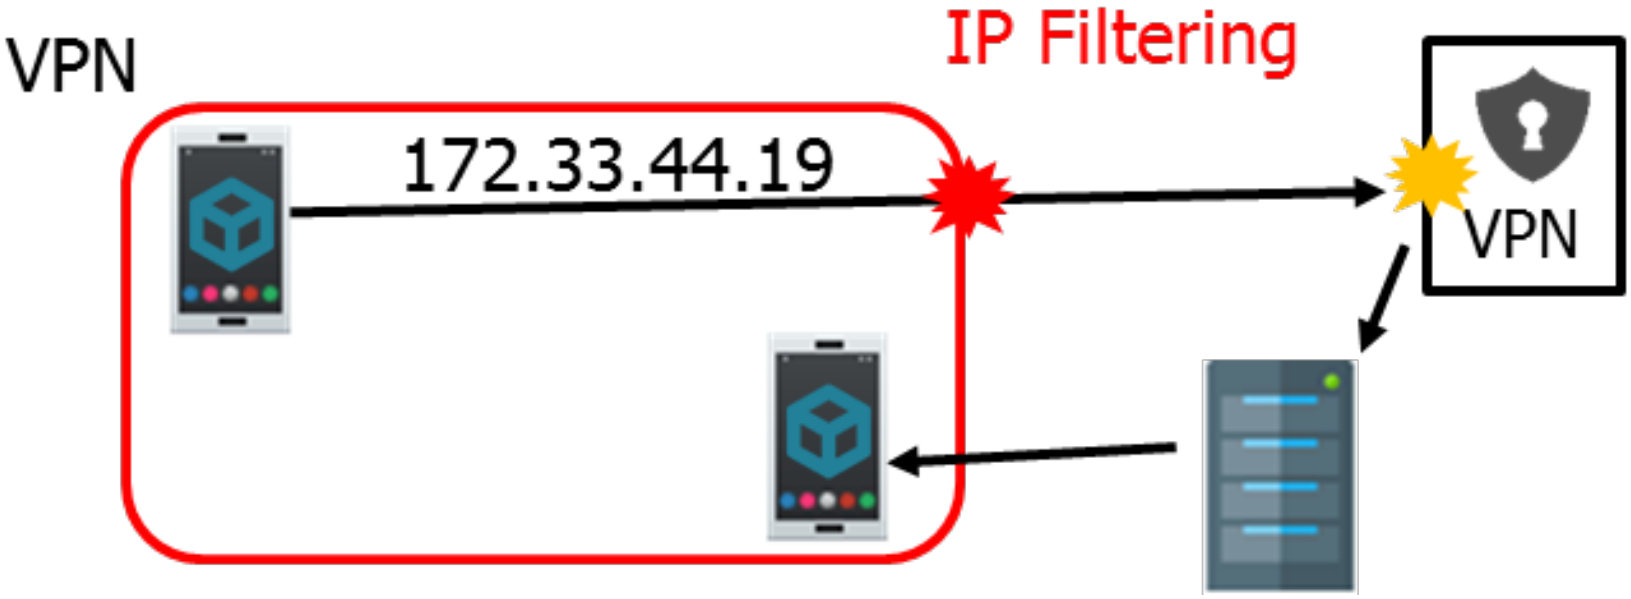

= contested network boundary

= communication exploitation

= service exploitation

Mobile Communications in Contested Environments

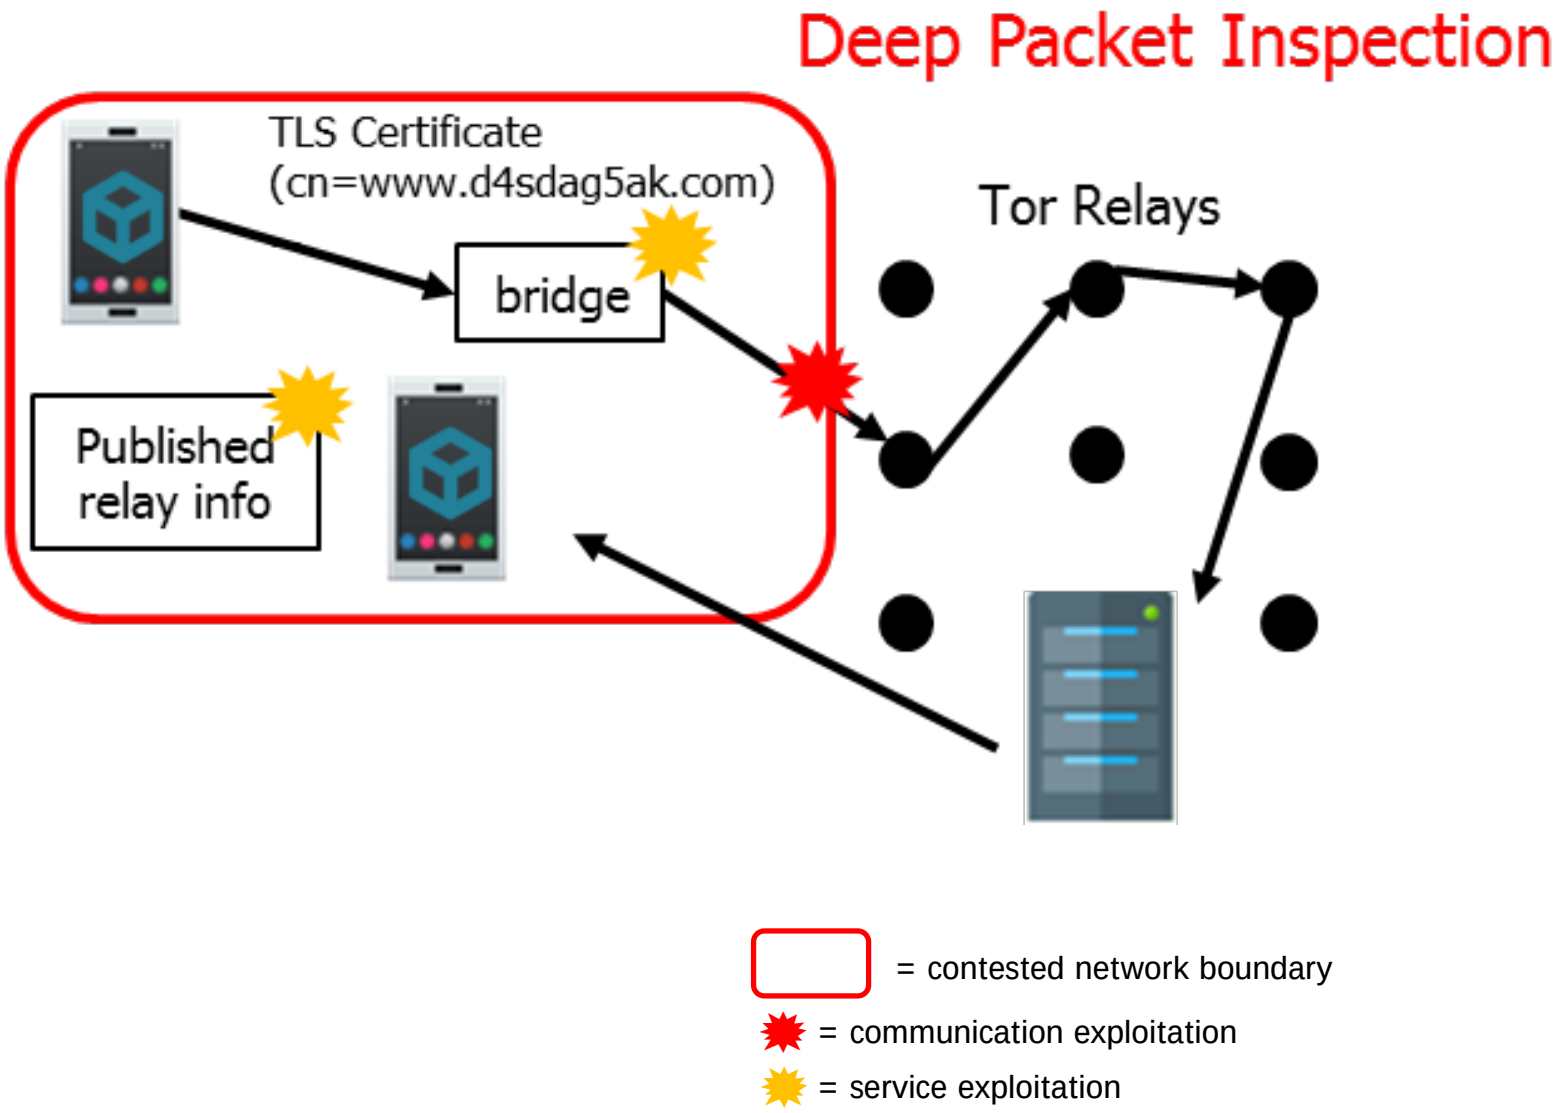

RACE Approach: Avoid Large-scale Targeting

1) Cryptography: Counter service exploitation via computing on encrypted data
2) Obfuscation: Counter communication exploitation via protocol embedding

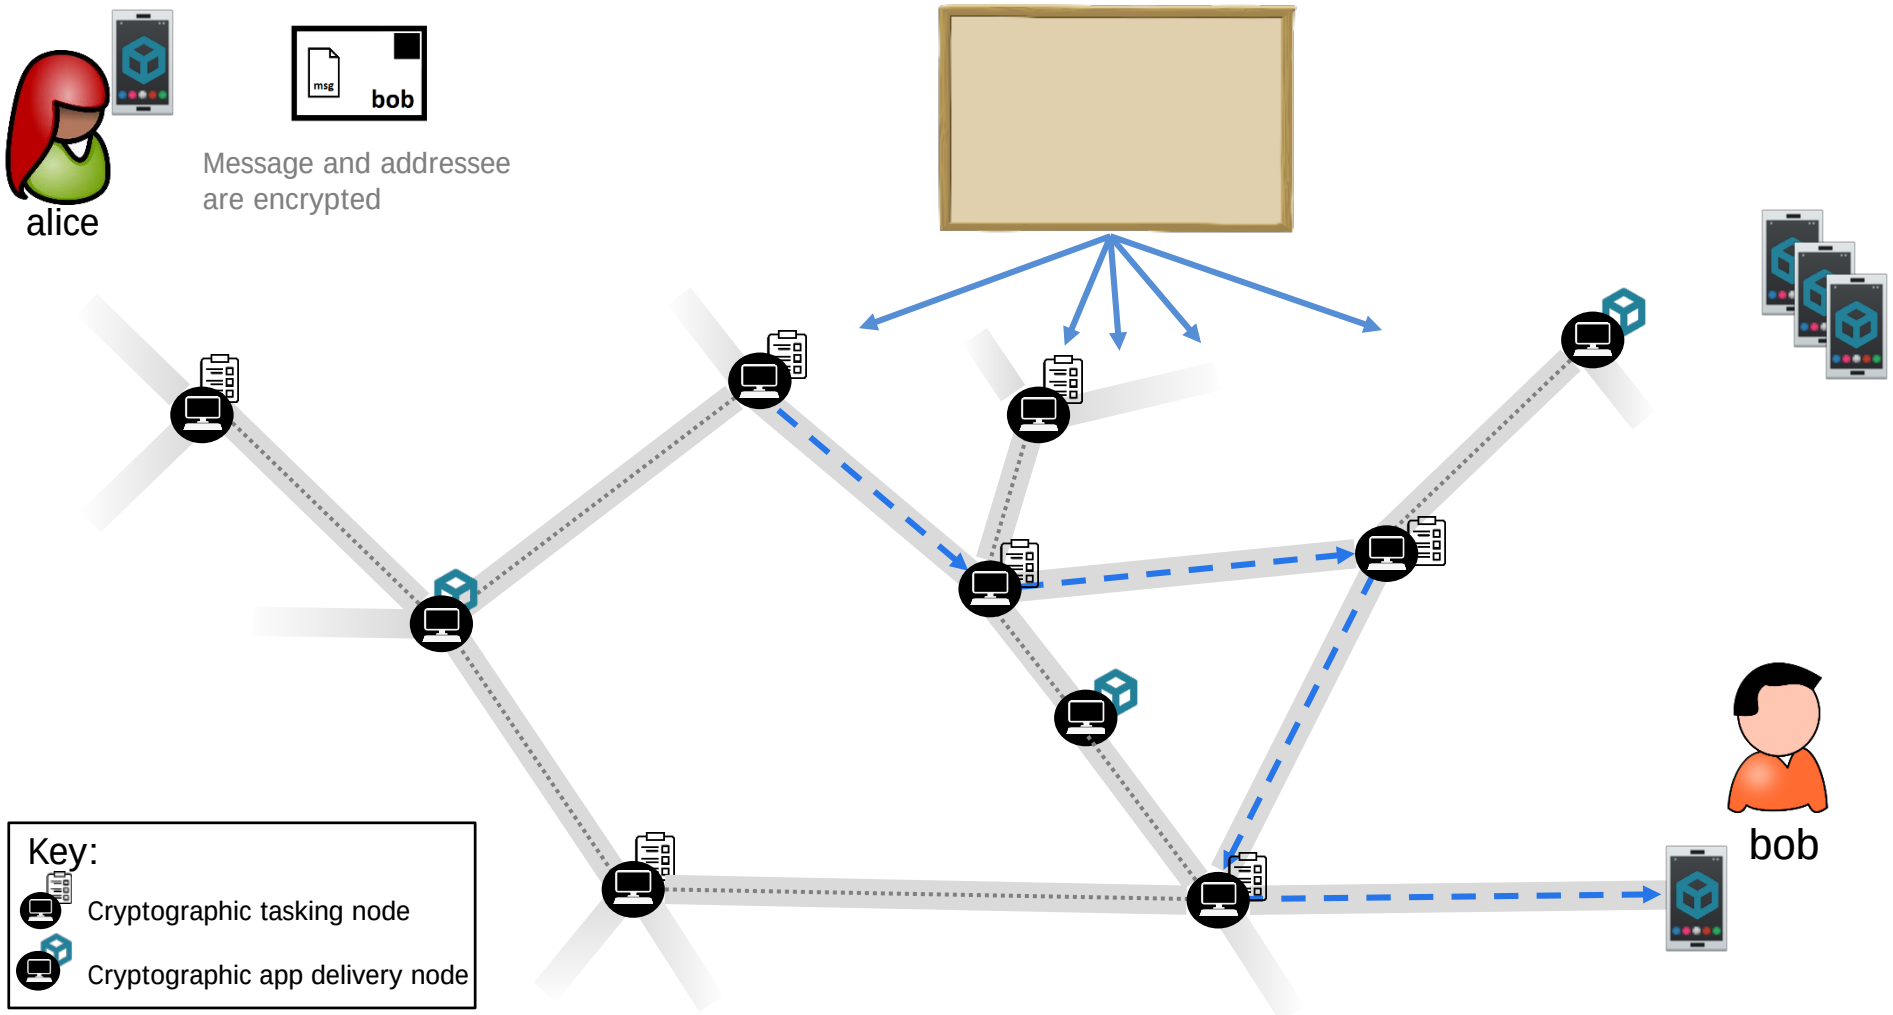

RACE Approach: Avoid Large-scale Targeting

1) Cryptography: Counter service exploitation via computing on encrypted data
2) Obfuscation: Counter communication exploitation via protocol embedding

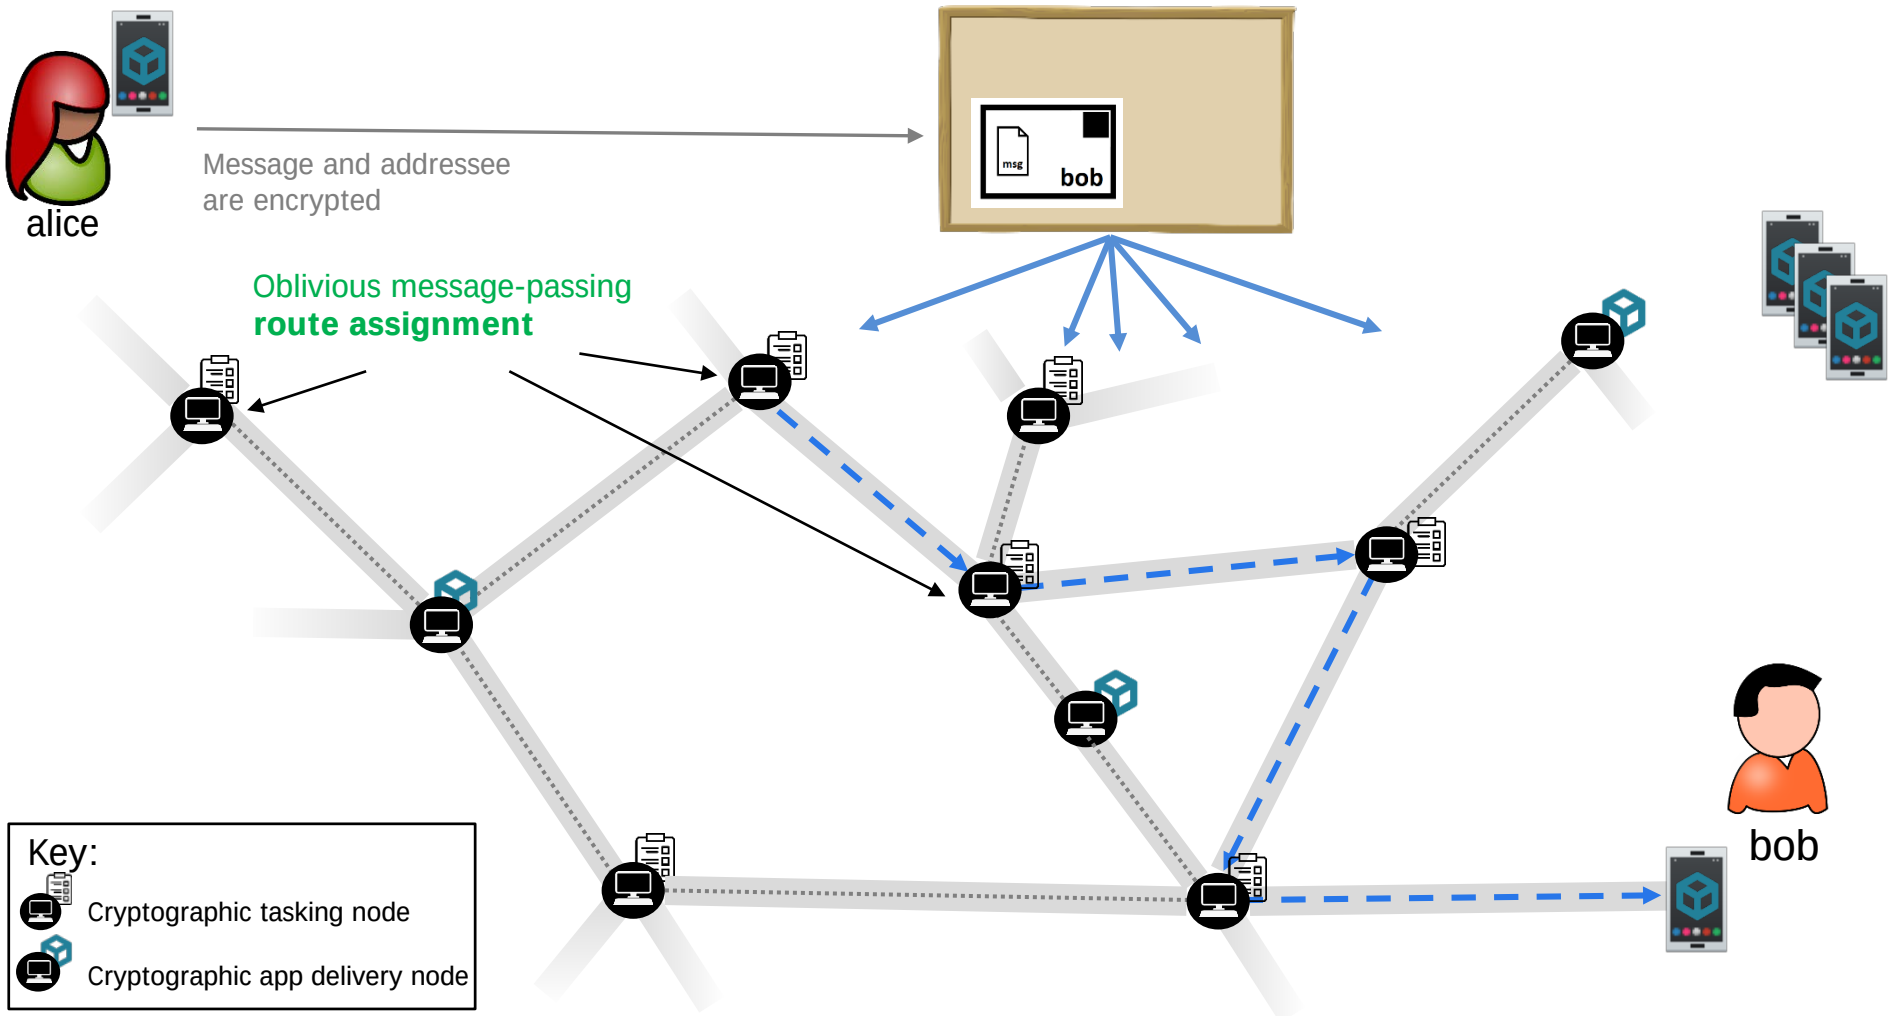

RACE Approach: Avoid Large-scale Targeting

1) Cryptography: Counter service exploitation via computing on encrypted data
2) Obfuscation: Counter communication exploitation via protocol embedding

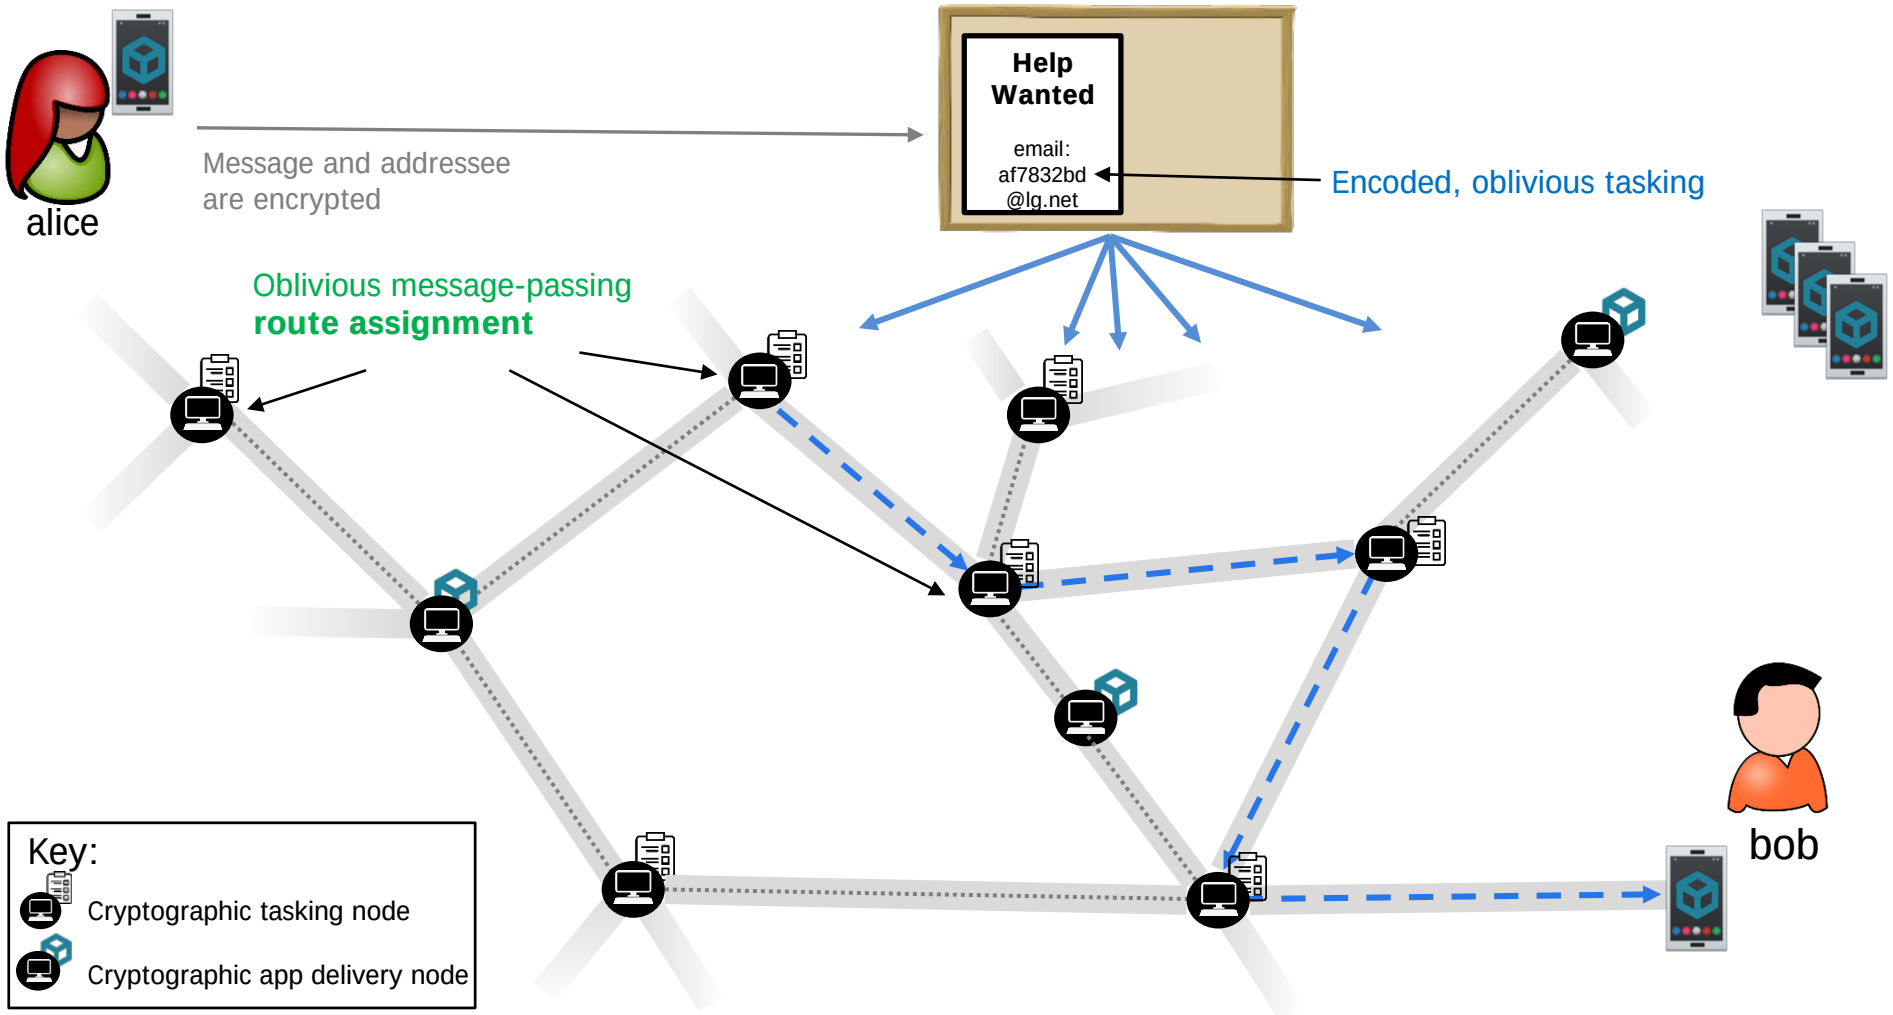

RACE Approach: Avoid Large-scale Targeting

1) Cryptography: Counter service exploitation via computing on encrypted data
2) Obfuscation: Counter communication exploitation via protocol embedding

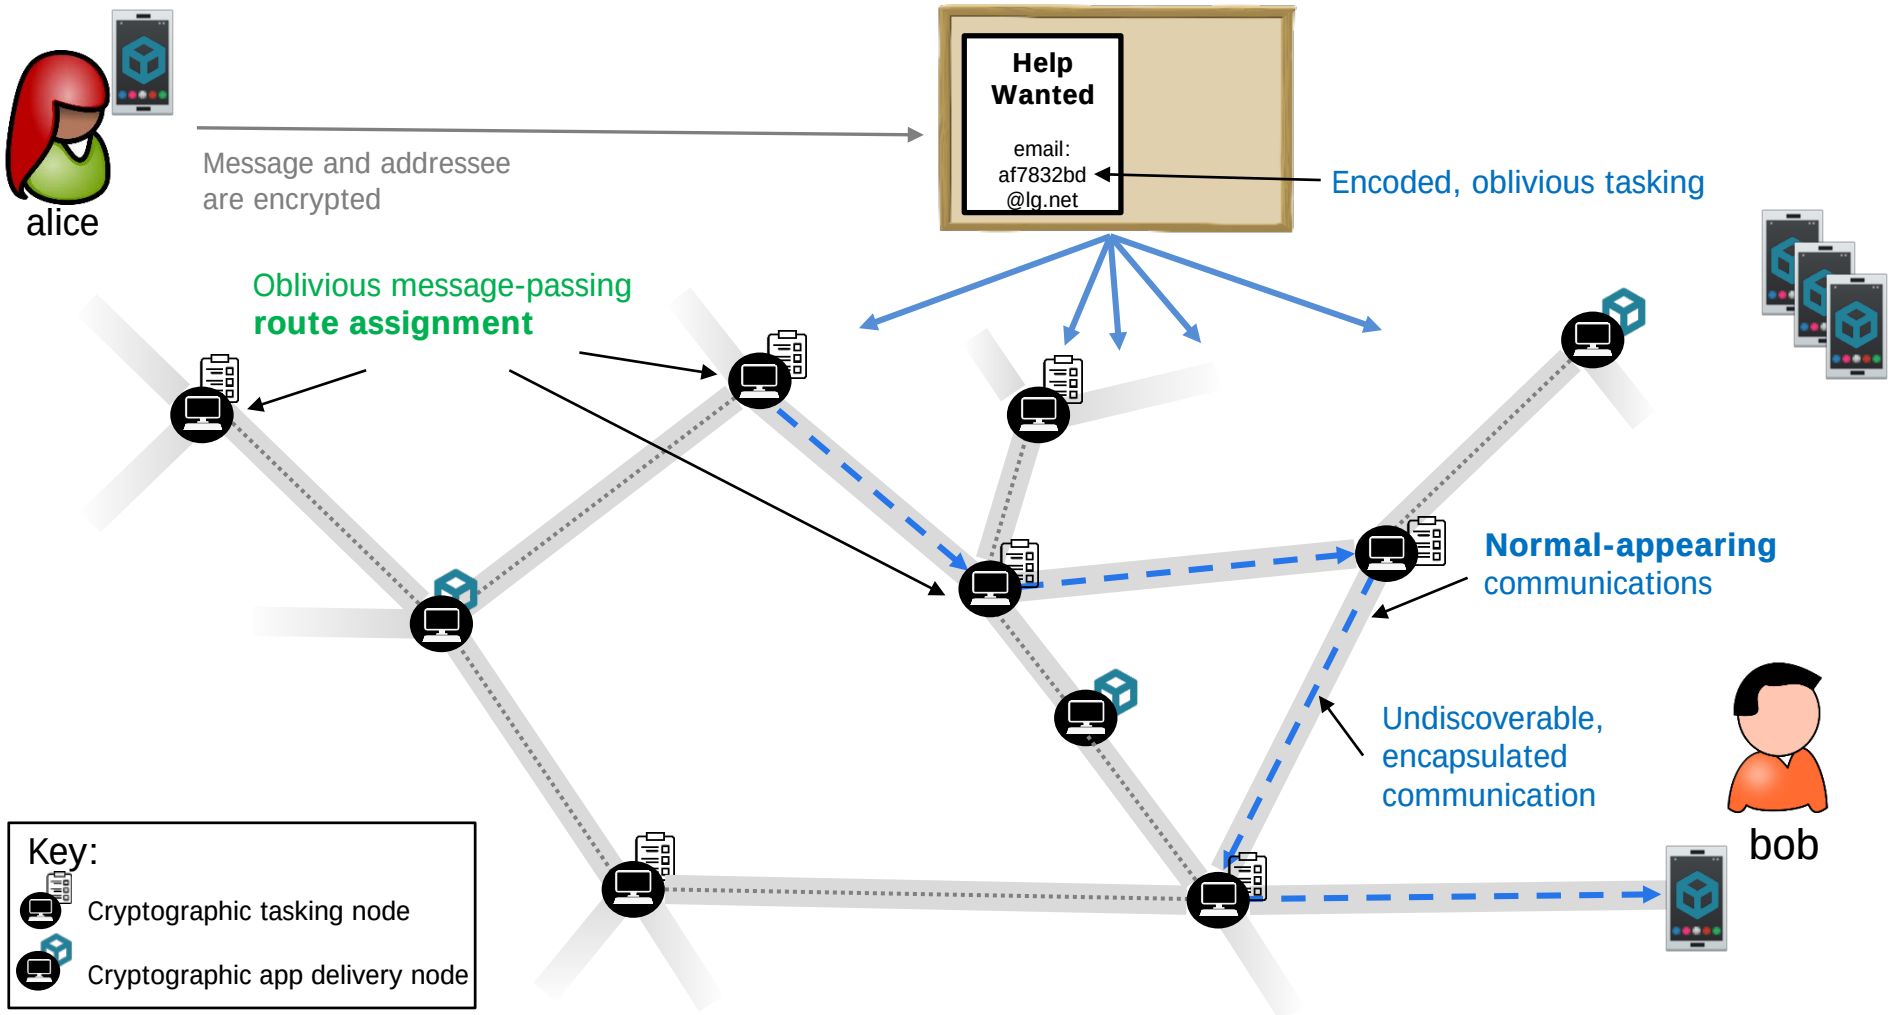

RACE Approach: Avoid Large-scale Targeting

1) Cryptography: Counter service exploitation via computing on encrypted data
2) Obfuscation: Counter communication exploitation via protocol embedding

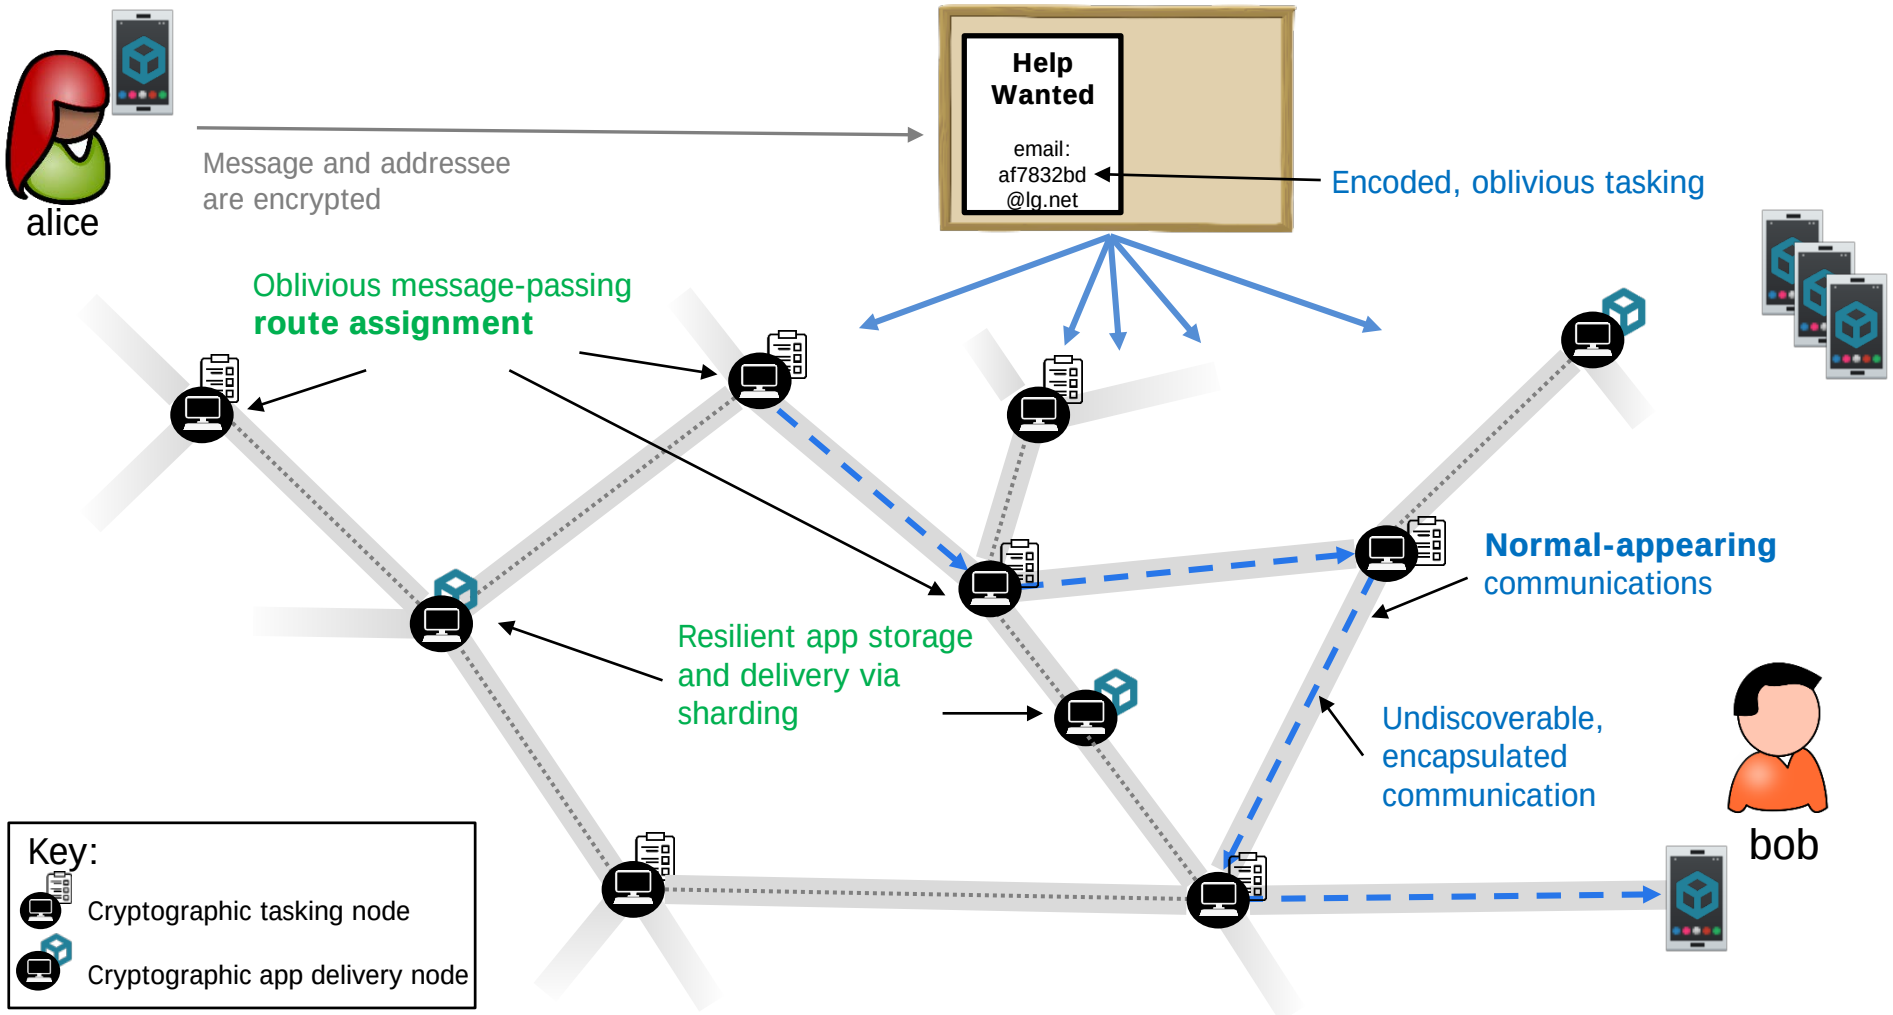

RACE Security Properties

Type	Attribute	Property
Confidentiality	user messages	Only the sender and receiver of a message can see it
	user message metadata	Confidentiality of who talks to whom and when
	unobservable communication	The fact that Alice possesses and uses the mobile application should not be inferable unless Alice's mobile device is compromised
	unobservable service node participation	The fact that Bob is running software to execute service node functionality should not be inferable unless Bob's system is compromised
Integrity	user messages	User messages cannot be changed in transit
Availability	user messages	End-to-end communication time should be one minute

Adversary model

Network level

- Ability to monitor traffic at in order to identify nodes/clients
- Ability to filter/block traffic
- Ability to manipulate/inject traffic

System level

- Ability to shutdown/block identified nodes
- Ability to perform low-level monitoring of a subset of nodes
- Ability to corrupt a subset of nodes (perform arbitrary code execution)
  - Can manipulate/inject communications via corrupted nodes

RACE Program Structure

TA3: System
Distributed Messaging Operations

TA 3.1: Resilient App Distribution

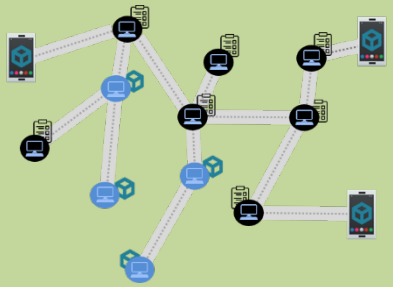

Prototype distributed tasking code

Communications toolbox

TA1: Cryptography
Resilient Distributed System Tasking

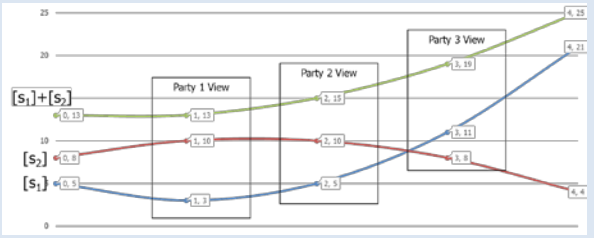

TA2: Obfuscated Communication
Countering Automated Protocol Identification

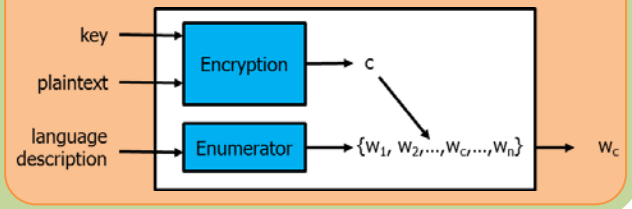

Prototype releases

Testbed and Voice of the Adversary
- Realistic network environments
- Realistic adversarial approaches
- Assess prototype systems

TA1: Cryptography

Objective: Distributed computation protocols for resilient system tasking

	Functionality	Confidentiality	Integrity
Users	Send/receive messages w/system (Rendezvous)	Messages and metadata (who talks with whom and when)	Message content and intended recipient
System	Dynamic, oblivious task allocation for messaging	Local node data cannot enable targeting of other nodes	Task allocation

Challenges:

- Efficient architectures and protocols that maintain resilience and user/network confidentiality
- Possibly incomplete network topologies
- 10Mbps point-to-point bandwidth
- Account for node refresh (e.g., system nodes may need to be removed or be added over time)

TA1 Metrics

Metric	Phase 1 (18 mo)	Phase 2 (12 mo)	Phase 3 (18 mo)
Nodes: users/server	10 / 100	100 / 1k	10k / 1k
Crypto adversary /corruption level	Passive / 20%	Active / 10%	Active / 20%
Crypto key infrastructure	Assumed	Not assumed	Not assumed
msg/day / size / delay	500 / 140B / 5 min latency	5k / 140B / 1 min latency	500k / 1MB / 1 min latency
Node refresh	Demonstrate	1/month	1/week

Strong TA1 Proposals Will:

- Identify either general or specific classes of computations that the RACE system will require and will explain why such computations can be implemented via MPC within the scale and efficiency required

- Account for how their system will eventually scale to 10,000 users who each send, on average, 50 messages per day to each other
  - For architectures where the receiver must retrieve a message from the server, the distributed, resilient storage issues must be addressed

Out of Scope for TA1

- Reliance on specialized or secure hardware
- Cryptanalysis (other than security proofs for provided systems)
- Cryptographic protocols based on non-standard or not commonly accepted cryptographic hardness assumptions

TA2: Obfuscated Communication

Objective: Develop a communications toolbox to counter automated protocol identification

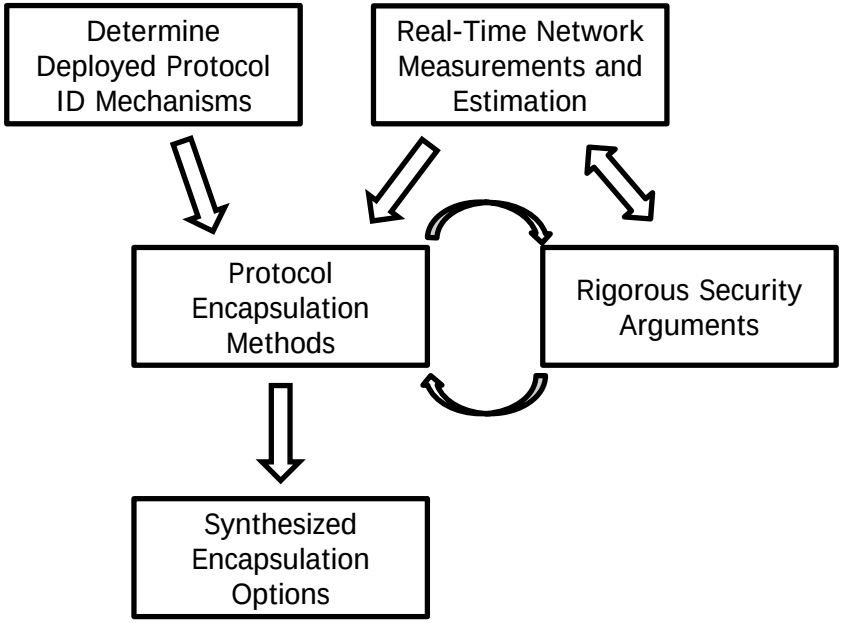

Challenges:

- Two modes of communication: client-server and server-server
  - Can propose one or both modes
- Create protocol encapsulation methods for specified target channels
- Accurately model target channels with statistical distributions
- Rigorously assess statistical closeness of encapsulated traffic to normal
- Adversary may try to actively manipulate links (and nodes that are hacked)

Testability is Critical for TA2 Proposals

- Proposals must discuss the elements required to test their technologies within a simulated network environment

- Proposers should assume that the test environment will be implemented within Amazon Web Services
  - Solutions that require special hardware or software must explain how such systems could be either implemented or realistically simulated within such a commercial cloud environment
  - Proposals should make clear their assumptions about the test environment and propose mitigations in case their assumptions do not hold

- <u>TA2 proposers who cannot fully explain the testability of their solutions will be deemed not selectable</u>

- In addition, TA2 performers should plan on holding regular discussions with the Test and Evaluation (T&E) team in order to help increase the realism of the contested network environment

- The relationship between TA2 and T&E will be collaborative rather than adversarial

TA2 Metrics

Metric	Phase 1 (18 mo)	Phase 2 (12 mo)	Phase 3 (18 mo)
Nodes: users/server	10 / 100	100 / 1k	10k / 1k
Crypto key infrastructure	Assumed	Not assumed	Not assumed
Security	Quantitative/ simulated evaluation	Statistical distance proof sketch	Statistical distance full proof
Adversary	Passive	Active link inject	Link+node inject
Logical bandwidth (server-server)*	5 Mbps	10 Mbps	10 Mbps
Logical bandwidth (client-server broadcast)*	100 kbps outgoing	500 kbps outgoing	500 kbps outgoing
Channel Model	Simulation evaluation	Proof (passive adversary)	Proof (active adversary)

* Proposers must conspicuously state whether they are proposing client-server and/or server-server modes of communication

Strong TA2 Proposals Will:

- Justify their targeted channels within the objectives of the RACE program and TA2. Choosing too-rich channels within which to encapsulate can be nearly impossible to realistically model
- Discuss how possibly unrelated network effects and traffic may affect proposed TA2 statistical models
- Explain why the targeted channel is such that the adversary will not simply filter out the entire channel in order to broadly filter the RACE system as well
- Discuss how TA2 solutions can resist adversarial manipulation of the contested network environment
  - For instance, if a timing side-channel solution is proposed, why will this be resilient against an adversary that introduces random timing jitter across the network?
- Discuss how TA2 solutions can retain the RACE TA2 security and functionality objectives even as adversaries are able to successfully exploit server nodes, clients, and their respective systems

Out of Scope for TA2:

- RF-based communication (other than IP-based mobile communication)
- Techniques that are untestable other than by deployment in an actual contested network environment

TA3: System

Objective:

- Integrate RACE technologies to build the prototype RACE system
  - Android mobile application for communications
  - Software application for the system nodes
  - Integration of application distribution approach from TA3.1

- TA3 proposals should discuss what additional capabilities they think need to be developed, potentially including:
  - Networking technologies on top of TA2 communications links
  - Technologies to counter denial of service attacks
  - Technical mechanisms, to include mechanisms for trusted introduction, to help enable the introduction of new server nodes
  - Technical mechanisms to leverage TA3.1 technologies to create (possibly trusted) introductions to help spread the mobile application to potential users

Challenges:

- Ensure that the overall RACE system that TA3 develops does not introduce degrade security capabilities of TA1, TA2, and TA3.1 technologies

TA3 Integration Tasks

- The TA3 performer will lead bi-weekly phone discussions amongst the various other RACE performers

- The TA3 team will host a shared, program-only software repository for TA1, TA2, and TA3.1 teams to deliver modules and for the T&E team to obtain artifacts to run their version of the messaging system infrastructure
  - This repository should support a secure shared workspace area (e.g. wiki, docs)
  - The TA3 team should also maintain the ability to validate their system (though the most extensive testing, especially from an adversarial perspective, will likely only be doable within the T&E team's environment)

TA3 Metrics

Metric	Phase 1 (18 mo)	Phase 2 (12 mo)	Phase 3 (18 mo)
Nodes: users/server	10 / 100	100 / 1k	10k / 1k
System	Architecture	Full prototype integration	Full demo system
Adversarial exploitation	Passive	Active node exploitation	Full spectrum exploitation
Communications channels	Mock channel	Single TA 2 (server-server and client-server) channel	Switch between channels

Strong TA3 Proposals Will:

- Address how they will work to integrate TA1, TA2, and TA3.1 technologies from a software development perspective
  - Proposals should mention previous experience, if any, in managing diverse teams of academics and companies to develop complex systems
  - Proposals should discuss their approach to the design of the RACE system architecture and integration API/plugin system

- Describe what additional capabilities will be needed on top of TA1, TA2 and TA3.1 technologies to build a useful, prototype RACE system and how the TA3 team will develop those capabilities
  - In particular, proposals will discuss how their integration will maintain the security delivered by individual TA technologies

TA3.1: Resilient Application Distribution

Objective:

- Develop techniques as well as associated software to enable a distributed storage and reconstruction functionality for the RACE mobile app
  - Create three functionalities:
    1. take as input an application and split it into shards;
    2. store and maintain these shards across service nodes
    3. upon appropriate command (which will be determined by TA3.1 performers), reconstruct the application at a desired service node and/or mobile device.

Challenges:

- Enable true application sharding and reconstruction that enables application execution on the output

Any proposer can bid on this TA, not just TA3 proposers. It is anticipated that there will be a single performer for this TA

TA3.1 Metrics

Metric	Phase 1	Phase 2	Phase 3
Crypto adversary/ corruption level	Passive / 20%	Active / 10%	Active / 20%
Crypto key infrastructure	Assumed	Not assumed	Not assumed
Node refresh	Demonstrate	1/year	1/month
Logical sharding	Demonstrate	Atomic functionality	Innocuous “gadgets”
Nodes: total w/ shards/need to reconstruct	50/10	250/30	1000/50
App reconstruction	10 min	5 min	5min
App size	1MB	10MB	50MB

(Selected) RACE Metrics

	Metric	Phase 1 (18 mo)	Phase 2 (12 mo)	Phase 3 (18 mo)
Common	Nodes: users/tasking	10 / 100	100 / 1k	10k / 1k
	Crypto adversary /corruption level	Passive / 20%	Active / 10%	Active / 20%
	Crypto key infrastructure	Assumed	Not assumed	Not assumed
TA 1	msg/day / size / delay	500 / 140B / 5 min latency	5k / 140B / 1 min latency	500k / 1MB / 1 min latency
	Node refresh	Demonstrate	1/month	1/week
TA 2	Security	Quantitative/ simulated evaluation	Statistical distance proof sketch	Statistical distance full proof
	Adversary	Passive	Active link inject	Link+node inject
	Bandwidth (c-s/s-s)	100 kbps / 5 Mbps	500 kbps / 10 Mbps	500 kbps / 10 Mbps
	Channel Model	Simulation eval	Proof (passive adversary)	Proof (active adversary)
TA 3	System	Architecture	Full prototype integration	Full demo system
	Adversarial exploitation	Passive	Active node exploitation	Full spectrum exploitation
	Comm channels	Mock channel	TA 2 channel	Switch b/t channels
TA 3.1	Logical sharding	<5	Atomic functionalities	Innocuous “gadgets”
	Nodes: total/reconstruct	50/10	250/30	1000/50
	App reconstruction	10 min	5 min	5 min

RACE Test and Evaluation

- T&E team is anticipated to consist of an FFRDC/UARC team
- The T&E team will create a realistic environment within which to deploy and evaluate technologies built by the other RACE teams
  - General development and functional testing is expected to be performed separately by individual TA teams
  - TA2 performers are in particular expected to work closely with the T&E team in order to help ensure network environment realism
- The T&E team will also act as the voice of the adversary
  - Evaluate communications traffic
  - Evaluate how data resident on systems can be exploited for further targeting

Schedule:

- Technical exchange meets every three months
  - Meetings will alternate between PI meeting locations and at T+E site

System Evaluation

- Two scrimmages in Phase 1, one in Phase 2, two in Phase 3
- End of phase formal measurement of individual teams plus system

RACE Classification and Technology Transition

- RACE is an UNCLASSIFIED program
  - Only UNCLASSIFIED submissions will be accepted
  - For TA3 performers: clearances are desired, but not required, for some end user discussions

- Performers will be expected to open-source their technologies by the end of the program
  - Intellectual property rights asserted unless otherwise directed by proposers are strongly encouraged to be aligned with open source regimes
  - Teams that do not intend to open-source their technologies developed in this program will be judged less favorably

www.darpa.mil

Email: RACE@darpa.mil

People/phone/computer icons from https://openclipart.org/ and https://github.com/KDE/breeze-icons