

## A Privacy-First Platform for Data Backhaul

Jean-Luc Watson, Tess Despres, Alvin Tan, Shishir Patil Prabal Dutta, Raluca Ada Popa



**UC Berkeley** 



## Goal: enable sensor deployments at large scale



Air quality measurement **on every block** Public infrastructure use **throughout a city** Wildlife counting **over square miles** Package tracking **at every stop** 



Air quality measurement **on every block** Public infrastructure use **throughout a city** Wildlife counting **over square miles** Package tracking **at every stop** 



Public infrastructure use **throughout a city** Wildlife counting **over square miles** Package tracking **at every stop** 



Air quality measurement **on every block** Public infrastructure use **throughout a city** Wildlife counting **over square miles** Package tracking **at every stop** 



Existing backhaul options are:

- Not scalable (manual data collection)
- **Expensive** (app-specific base stations)



Air quality measurement **on every block** Public infrastructure use **throughout a city** Wildlife counting **over square miles** Package tracking **at every stop** 



Air quality measurement **on every block** Public infrastructure use **throughout a city** Wildlife counting **over square miles** Package tracking **at every stop** 

## Instead, 'mules' in close proximity could gather data





## Instead, 'mules' in close proximity could gather data



**App Server** 

## Mules send data to a backhaul provider



## The provider forwards to paying servers



## The provider forwards to paying servers



## Centralized backhaul threatens privacy



## Centralized backhaul threatens privacy



## Centralized backhaul threatens privacy



**Platform** Provider

General-purpose

Scalability

**Preventing misuse** 

Payment

Privacy

Provider can run a general-purpose service that is not private



Provider can run a **private** service if it limits application scope



Provider can run a private and general-purpose service but not at scale



Nebula accomplishes each of these at the same time



A backhaul system that *preserves mule privacy* while:

1. Charging for backhaul use and incentivizing mule participation

A backhaul system that *preserves mule privacy* while:

1. Charging for backhaul use and incentivizing mule participation

Out-of-band accounting: app servers buy untraceable *tokens* from provider and exchange them for data.

A backhaul system that *preserves mule privacy* while:

1. Charging for backhaul use and incentivizing mule participation

Out-of-band accounting: app servers buy untraceable *tokens* from provider and exchange them for data.

2. Detecting mule and application server misbehavior

A backhaul system that *preserves mule privacy* while:

1. Charging for backhaul use and incentivizing mule participation

Out-of-band accounting: app servers buy untraceable *tokens* from provider and exchange them for data.

2. Detecting mule and application server misbehavior

Allow mules to submit complaints with *proof* of misbehavior to the provider.

A backhaul system that *preserves mule privacy* while:

1. Charging for backhaul use and incentivizing mule participation

Out-of-band accounting: app servers buy untraceable *tokens* from provider and exchange them for data.

2. Detecting mule and application server misbehavior

Allow mules to submit anonymous complaints with *proof* of misbehavior to the provider.

3. Supporting millions of mules

Nebula can process ~445k tokens/second, or 250M tokens/dollar

A backhaul system that *preserves mule privacy* while:

1. Charging for backhaul use and incentivizing mule participation

Out-of-band accounting: app servers buy untraceable *tokens* from provider and exchange them for data.

2. Detecting mule and application server misbehavior

Allow mules to submit anonymous complaints with *proof* of misbehavior to the provider.

3. Supporting millions of mules

Nebula can process ~445k tokens/second, or 250M tokens/dollar

#### Nebula's approach: decentralize protocol onto mules



#### Nebula's approach: decentralize protocol onto mules



- System usage is tracked with PrivacyPass [DGS+18] tokens
  - Unlinkable: when tokens are redeemed to a malicious server, cannot be linked to the client(s) that generated them

- System usage is tracked with PrivacyPass [DGS+18] tokens
  - Unlinkable: when tokens are redeemed to a malicious server, cannot be linked to the client(s) that generated them

Provider controls access to system resources without knowing the token owner

- System usage is tracked with PrivacyPass [DGS+18] tokens (S)
  - Unlinkable: when tokens are redeemed to a malicious server, cannot be linked to the client(s) that generated them

Provider controls access to system resources without knowing the token owner

• One-more-token security: even knowing many valid tokens, a malicious client cannot forge more

- System usage is tracked with PrivacyPass [DGS+18] tokens
  - Unlinkable: when tokens are redeemed to a malicious server, cannot be linked to the client(s) that generated them

Provider controls access to system resources without knowing the token owner

• One-more-token security: even knowing many valid tokens, a malicious client cannot forge more

Only way to acquire more tokens is to participate in backhaul

### (1) Token Pre-purchase

• At the beginning of an epoch, app servers *pre-purchase* tokens from the provider



## (2) Payload Delivery

• Mules send data directly to application servers over anonymous connections and receive tokens in exchange



### (3) Token Redemption

• At the end of an epoch, mules redeem tokens with the platform provider in exchange for compensation



#### A backhaul system that *preserves mule privacy* while:

1. Charging for backhaul use and incentivizing mule participation

Out-of-band accounting: app servers buy untraceable *tokens* from provider and exchange them for data.

#### 2. Detecting mule and application server misbehavior

Allow mules to submit complaints with *proof* of misbehavior to the provider.

3. Supporting millions of mules

Nebula can process ~445k tokens/second, or 250M tokens/dollar
#### Big problem: delivery misbehavior



#### Big problem: delivery misbehavior



#### Big problem: delivery misbehavior



# (4) Complaint

• After an epoch, if a mule notices misbehavior (e.g. invalid token), it can *complain* to the provider for a new one



#### Complaints are based on app server commitments

Before receiving data, app servers commit to token they will use if payload is uploaded



#### Complaints are based on app server commitments

Before receiving data, app servers commit to token they will use if payload is uploaded

A new token is granted for proof of misbehavior





#### Overall, the provider learns:

1. Count of prepaid tokens purchased by each app server



#### Overall, the provider learns:

- 1. Count of prepaid tokens purchased by each app server
- 2. Count of payloads a mule uploaded, across all app servers



#### Overall, the provider learns:

- 1. Count of prepaid tokens purchased by each app server
- 2. Count of payloads a mule uploaded, across all app servers
- 3. Set of anonymous complaints against app server misbehavior



#### Formal Proof Sketch in Paper

#### Overall, the provider learns:

- 1. Count of prepaid tokens purchased by each app server
- 2. Count of payloads a mule uploaded, across all app servers
- 3. Set of anonymous complaints against app server misbehavior

(1) Mule encounters and authenticates a sensor, receiving a E2E-encrypted payload and a sensor-signed payload hash



(2) Mule sends signed hash to desired application server. Server checks hash and decides whether they want the payload.



(2) Mule sends signed hash to desired application server. Server checks hash and decides whether they want the payload.



(3) If the server would like the payload, they sign a **commitment** to the hash and an encrypted token they are willing to spend in exchange



(3) If the server would like the payload, they sign a **commitment** to the hash and an encrypted token they are willing to spend in exchange



(4) Commitment in hand, the mule sends the actual payload to the app server.



(5) The server checks the data against the hash it previously got; if successful, it signs a response containing the unencrypted token



What happens if the mule gets a bad token, or no token at all?



Provider will notify mule of duplicate/invalid tokens it redeems at end of epoch

What happens if the mule gets a bad token, or no token at all?



(1) At end of epoch, mule will send a complaint token and commitment to provider



(2) Provider will check complaint token, then decrypt and verify the delivery token, invalidating it



(2) Provider will check complaint token, then decrypt and verify the delivery token, invalidating it



(3) If the token was rejected by the platform provider (wrong or a duplicate), the mule sends over the signed token it got from the app server



(3) If the token was rejected by the platform provider (wrong or a duplicate), the mule sends over the signed token it got from the app server



(3b) If no token was received by the mule, it has the option of sending the original data payload matching the committed hash



(3b) If the data matches the committed hash, it's forwarded to the app server



(3b) If the data matches the committed hash, it's forwarded to the app server



(4) Platform provider and mule work together to blindly sign a new token!



(4) Platform provider and mule work together to blindly sign a new token!



For duplicate tokens: first-come first-served to avoid generating more tokens than were used during delivery while incentivizing complaints

# **NEBULA**

#### A backhaul system that *preserves mule privacy* while:

1. Charging for backhaul use and incentivizing mule participation

Out-of-band accounting: app servers buy untraceable *tokens* from provider and exchange them for data.

2. Detecting mule and application server misbehavior

Allow mules to submit anonymous complaints with *proof* of misbehavior to the provider.

#### 3. Supporting millions of mules

Nebula can process ~445k tokens/second, or 250M tokens/dollar

We measured likely *frequency* of sensor-mule interaction



We measured likely frequency of sensor-mule interaction, and duration of data transfer



We measured likely frequency of sensor-mule interaction, and duration of data transfer



*Example*: in a park, we can expect a mule every few minutes

We measured likely *frequency* of sensor-mule interaction, and *duration* of data transfer



*Example*: in a park, we can expect a mule every few minutes,

each with a transmission window of 5 to 10 seconds,

We measured likely *frequency* of sensor-mule interaction, and *duration* of data transfer



*Example*: in a park, we can expect a mule every few minutes,

each with a transmission window of 5 to 10 seconds, supporting 2 to 16kB payloads

We measured likely *frequency* of sensor-mule interaction, and *duration* of data transfer



*Example*: in a park, we can expect a mule every few minutes,

each with a transmission window of 5 to 10 seconds, supporting 2 to 16kB payloads

Analytical Energy/Memory models in paper
## Provider performance

High-throughput database to check token validity and detect duplicates.



## Provider performance

High-throughput database to check token validity and detect duplicates.



## Find us at our poster!

