WISDOM

Rethinking Security in the *Era of Cloud Computing*

Thomas Ristenpart

Public IaaS clouds (Infrastructure-as-a-Service)

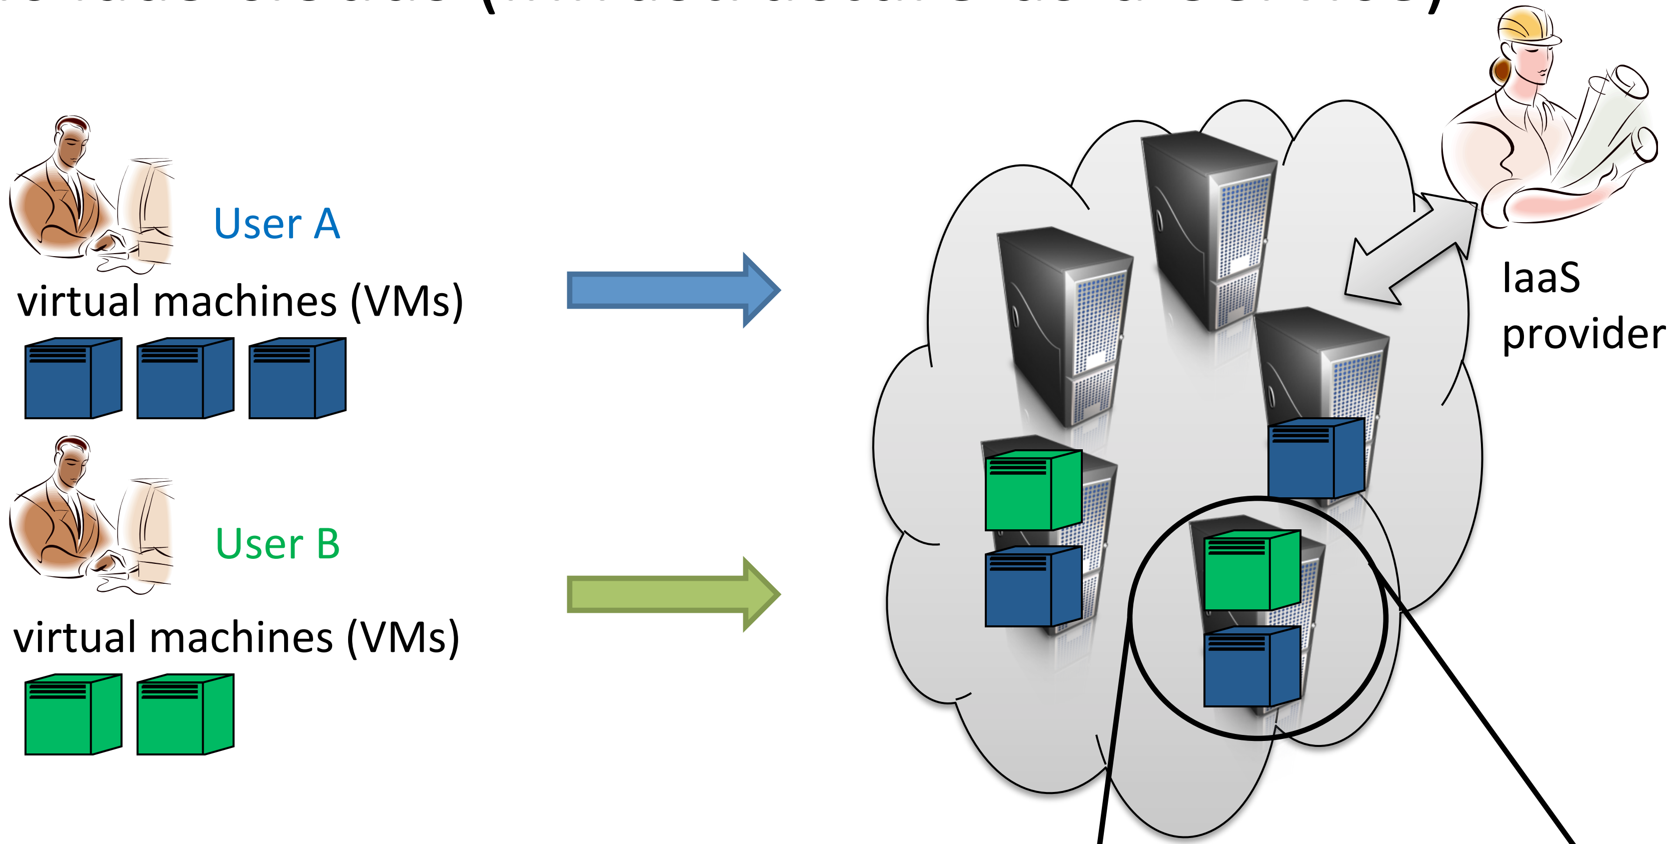

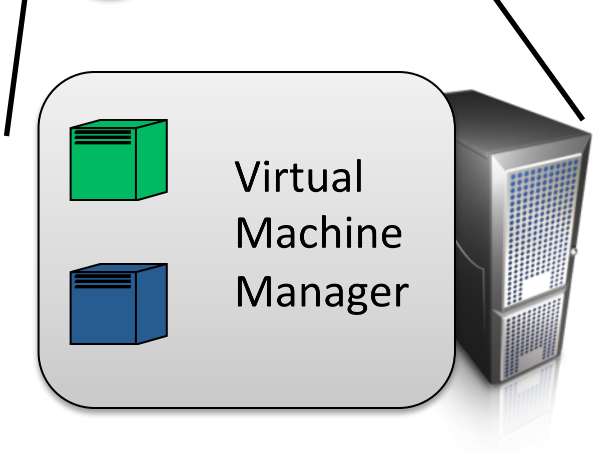

User's pay for resources used

Multitenancy (users share physical resources)

Virtualization of resources (compute, networking) as managed by provider

Threat models and cloud research

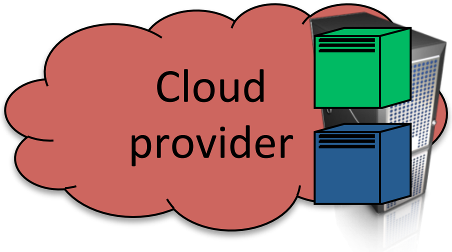

(1) Cloud-as-adversary

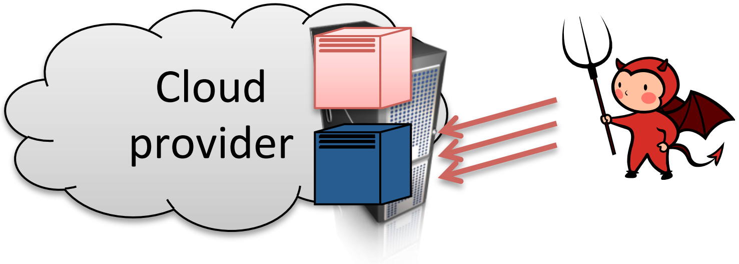

(2) Adversarial tenants and outsiders

<u>Real-world examples:</u>

Insiders

Compromise of control plane

Government surveillance

<u>Real-world examples:</u>

Co-location attacks / side-channel attacks

Compromised VMs

External attackers (SQL injection, DoS, etc.)

Threat models and cloud research

<u>(2a) *New threats in public clouds*</u>

Focuses on intersection of resource sharing and adversarial tenants; new technologies used

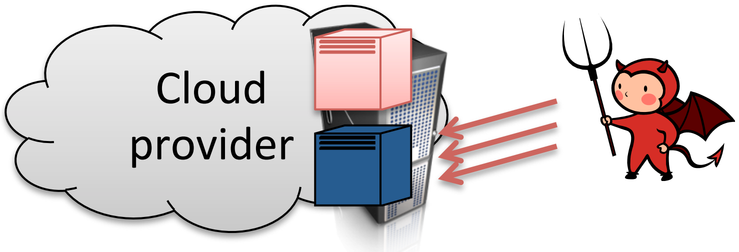

(2) Adversarial tenants and outsiders

Side-channel attacks and defenses
(See Venkat's talk)

Pricing and resource abuse
(Resource-freeing attacks,
placement gaming,
billing measurements – See Rob's talk)

Technology issues
(RNGs in virtualized environments – See Adam's talk)

Threat models and cloud research

(2a) *New threats in public clouds*

Focuses on intersection of resource sharing and adversarial tenants; new technologies used

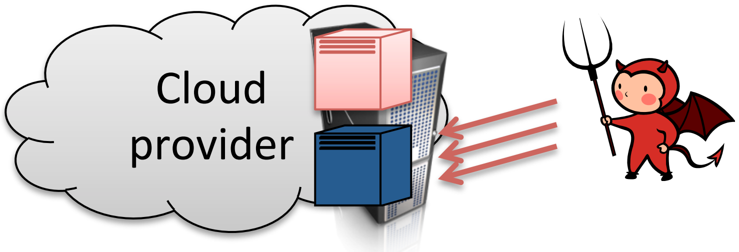

(2) Adversarial tenants and outsiders

(2b) *Dealing with old threats, better*

Focuses on leveraging provider & control plane to help tenant security

Project Silver

Broad research agenda on how cloud providers can help improve security for the tenant ecosystem

The goal: It is *safer* to run in the cloud

The opportunity

The migration to cloud services:

- 4% of Alexa Top Million websites using EC2/Azure

  (See Keqiang's talk)
- Centralization of hosting into fewer large providers
- Cloud providers (or third-parties) adding features

Multi-institution effort

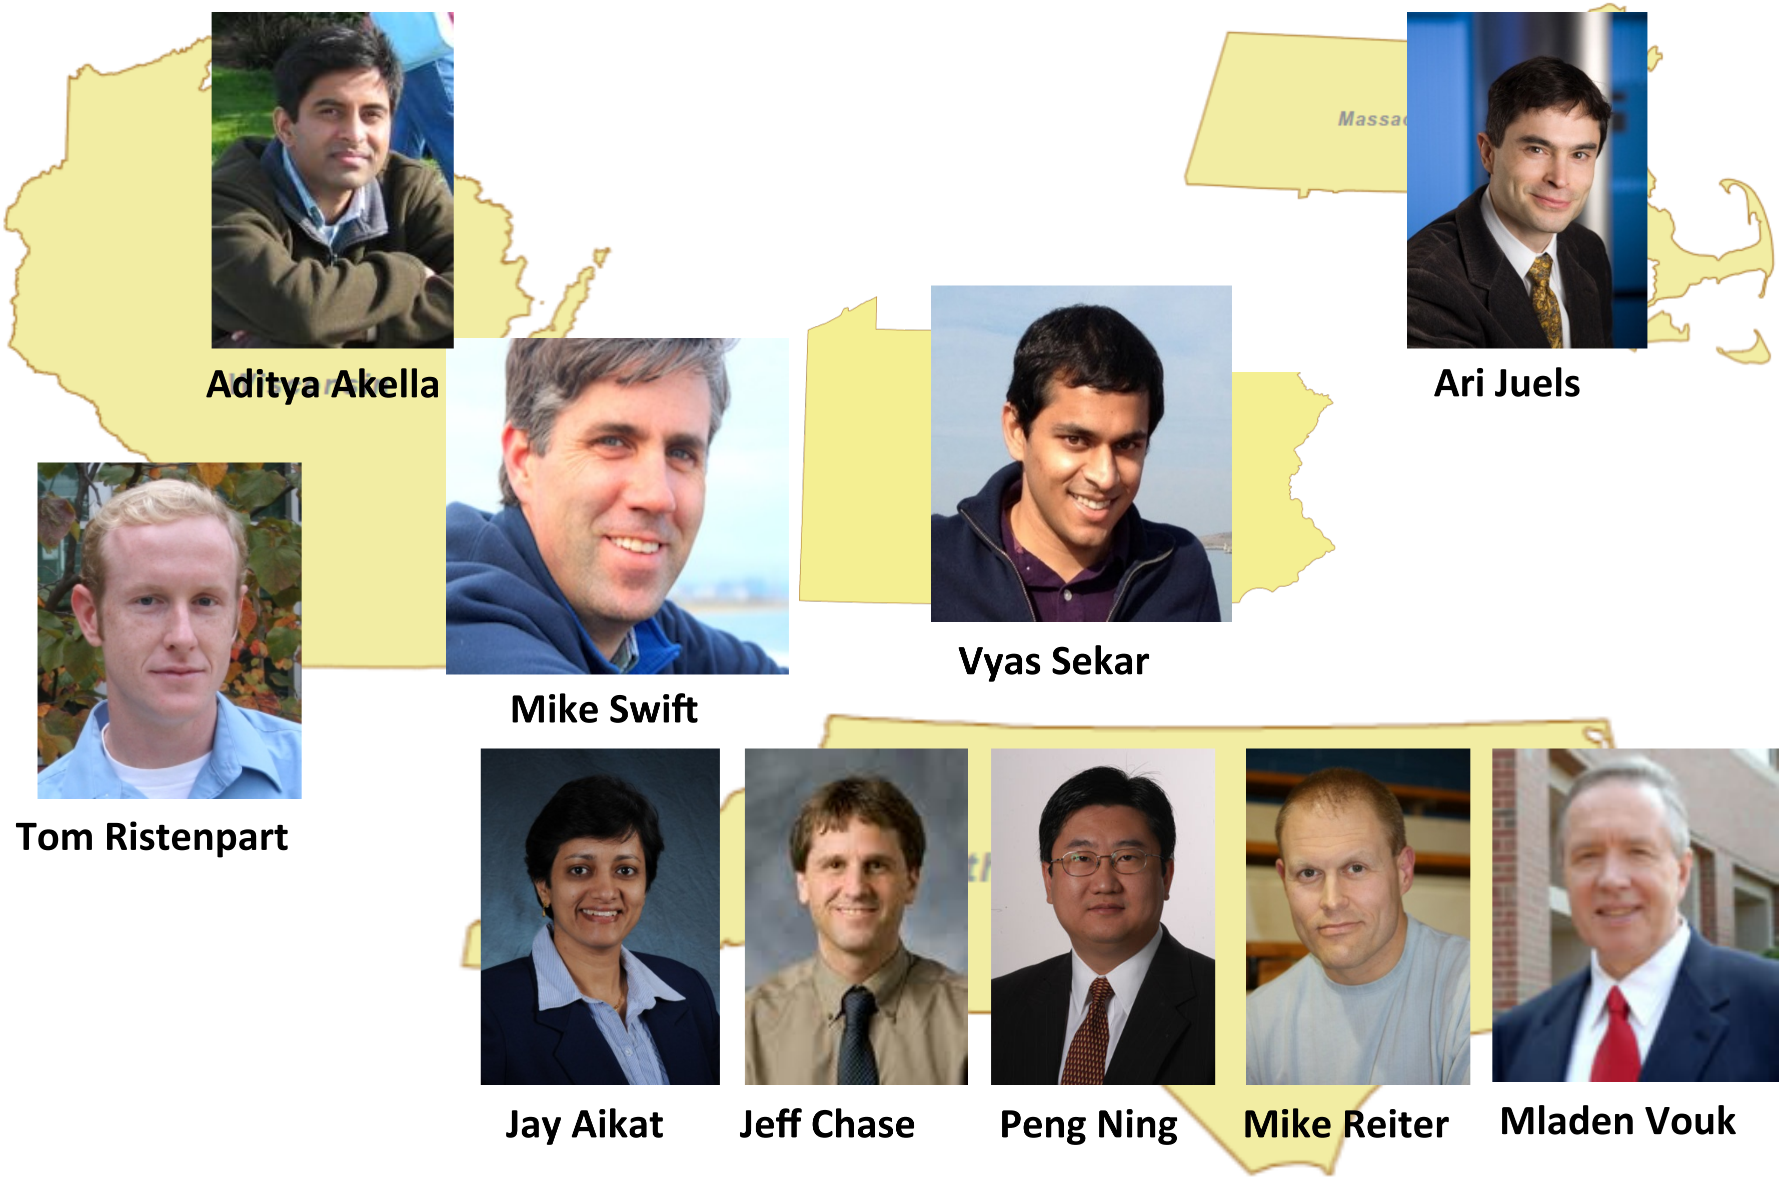

Multi-institution effort

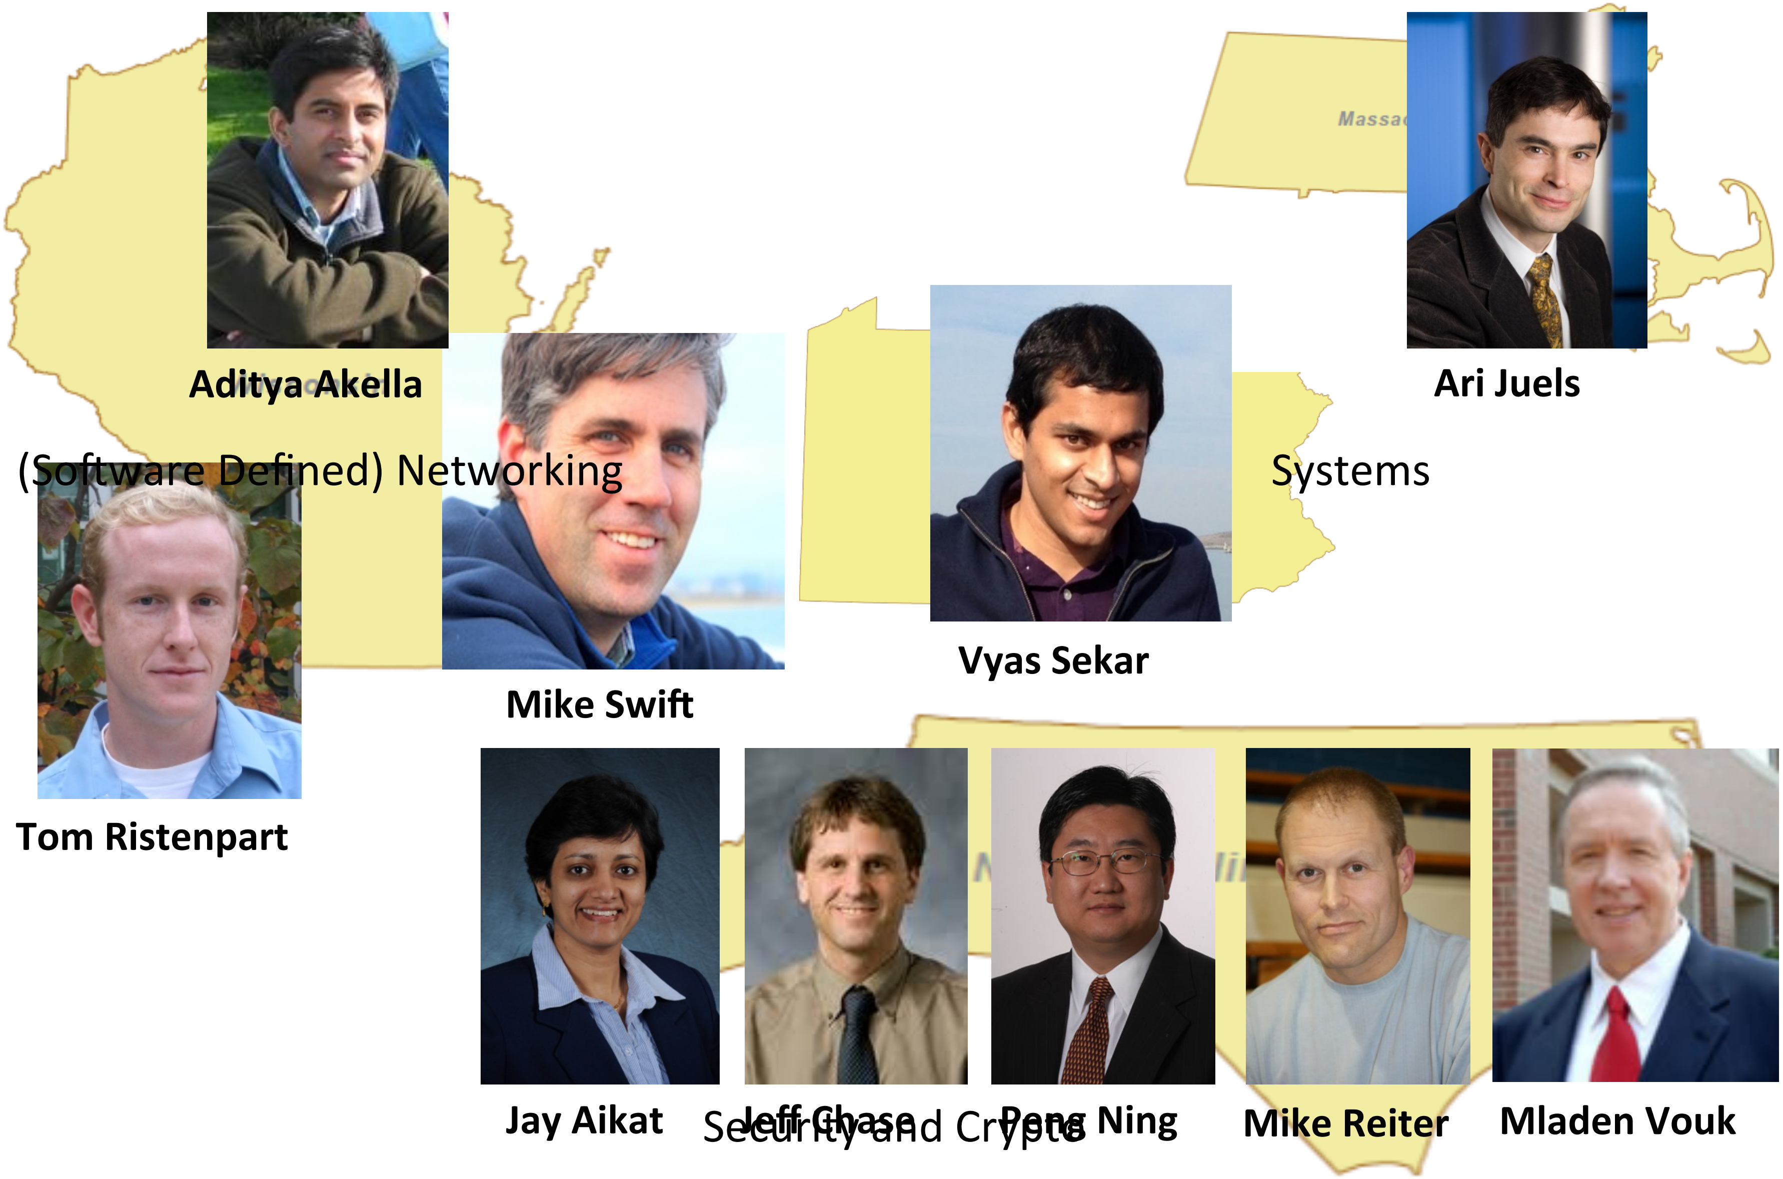

Aditya Akella

(Software Defined) Networking

Mike Swift

Tom Ristenpart

Vyas Sekar

Systems

Ari Juels

Jay Aikat

Jeff Chase

Peng Ning

Mike Reiter

Mladen Vouk

Security and Crypto

Project Overview

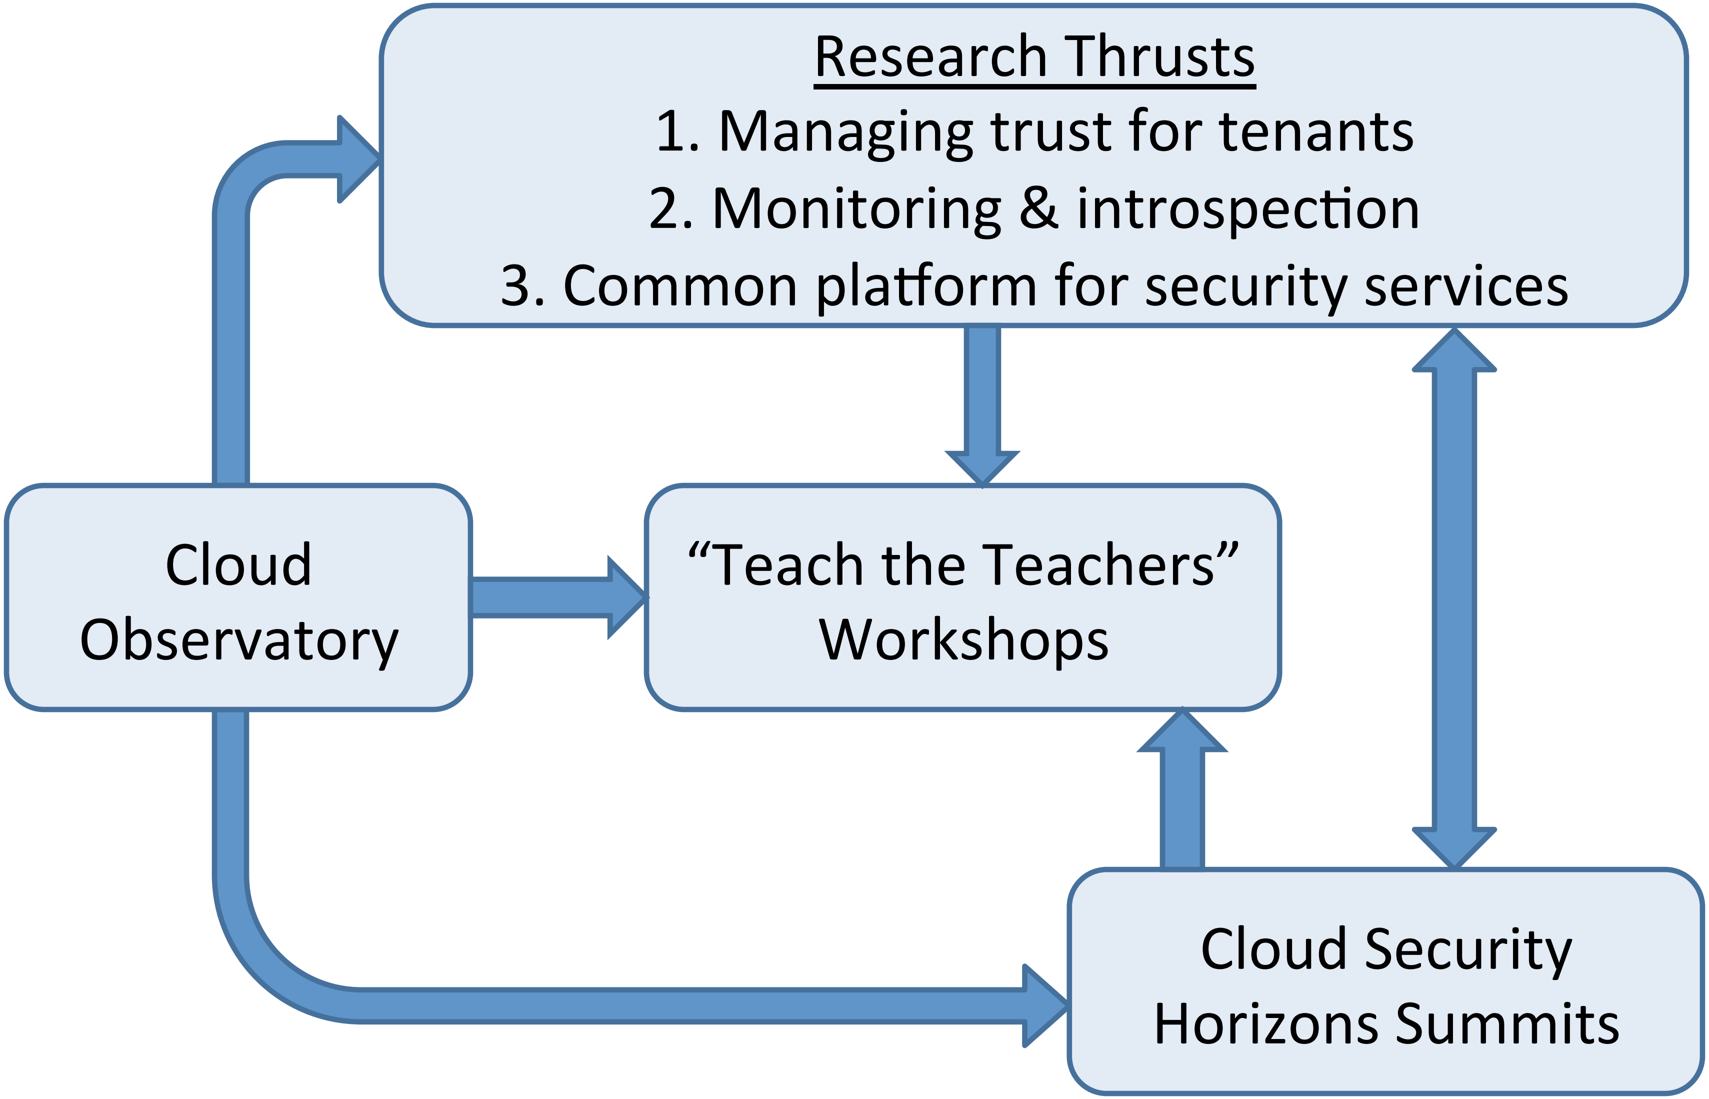

Research Thrusts

1. Managing trust for tenants
2. Monitoring & introspection
3. Common platform for security services

Cloud Observatory

“Teach the Teachers” Workshops

Cloud Security Horizons Summits

Project Overview

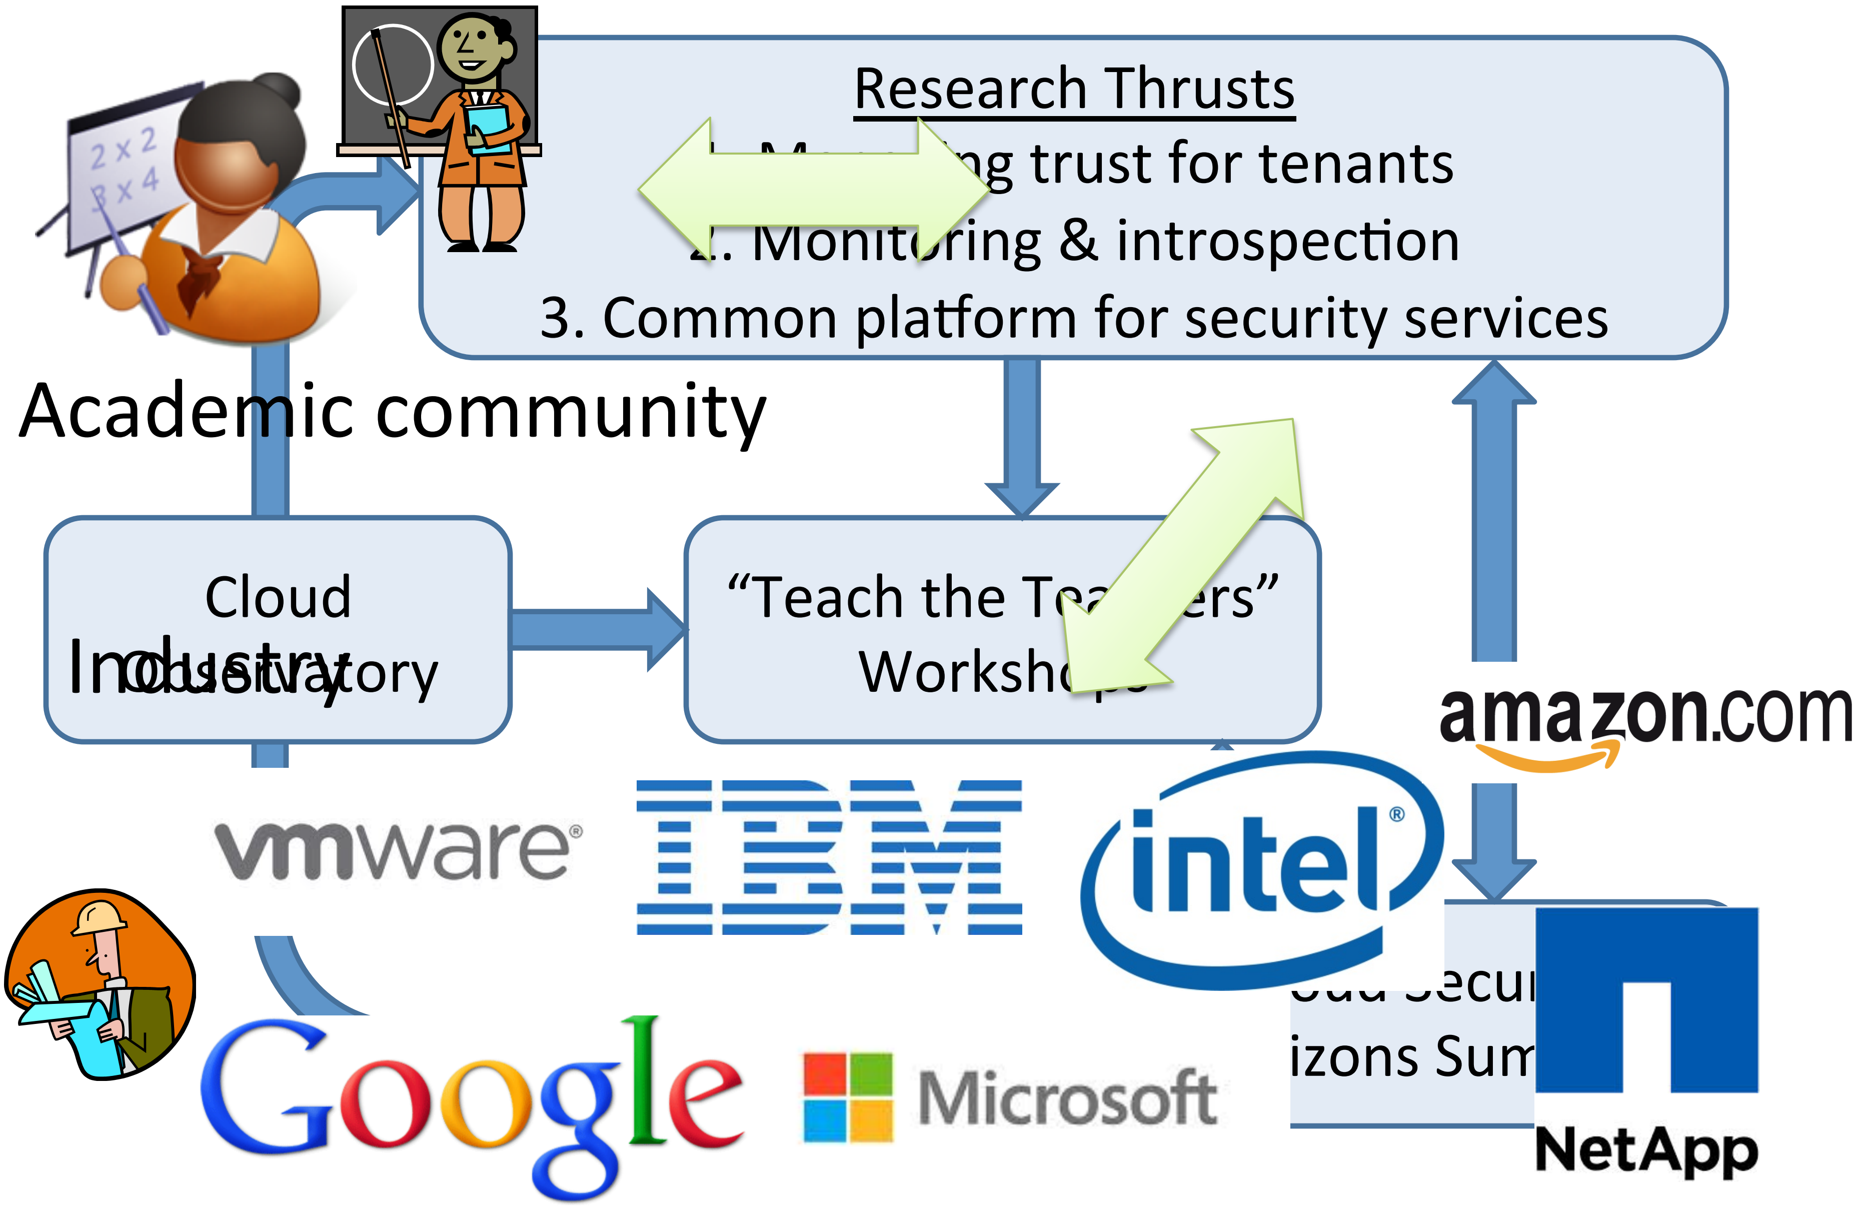

Research Thrusts

1. M[illegible]ng trust for tenants
2. Monitoring & introspection
3. Common platform for security services

Academic community

Cloud Observatory

Industry

“Teach the Teachers” Workshops

oud Secur

izons Sum

Today:
Ongoing projects involving WISDOM

- Cloud observatory
  - Provide data sets and methodologies for understanding how cloud usage evolves
- New IaaS Security Services
  - Security-posture audit tools (SPATs)
  - Other projects

Cloud observatory

- Measure usage, security posture of cloud tenants
  - Generating several rich datasets
  - Analysis and opportunity finding

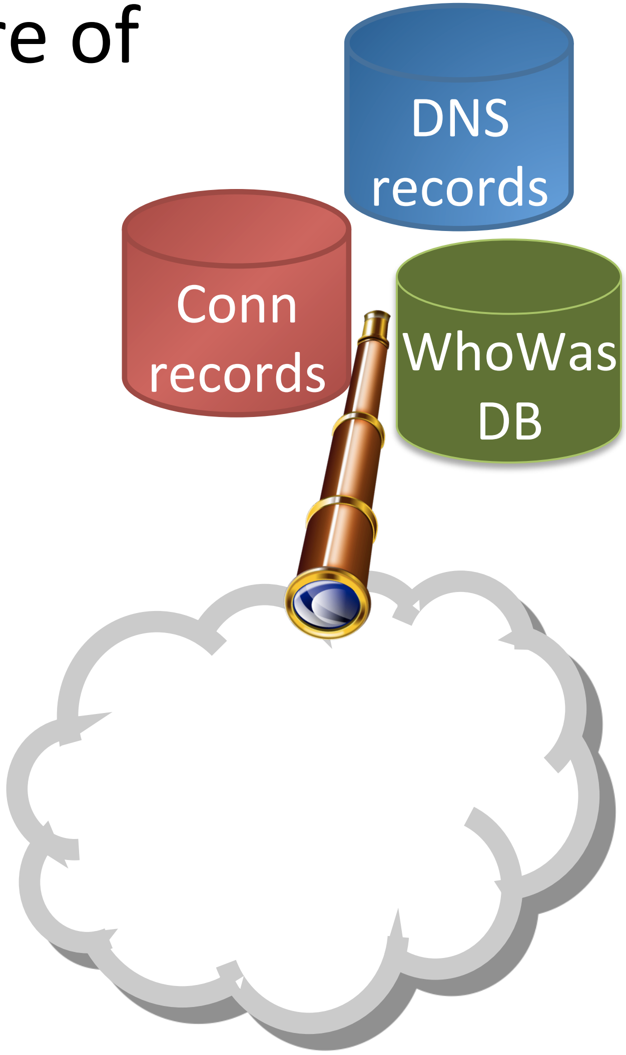

Example questions to answer

- What is distribution of deployment types?
- How much churn is there? (Turnover rate per IP address)
- Are software updates reaching cloud tenants quickly?
- What kinds of malicious activity arise? Are IP-based blacklists working well for IaaS clouds?

Cloud observatory data sets

Alexa subdomains DNS records

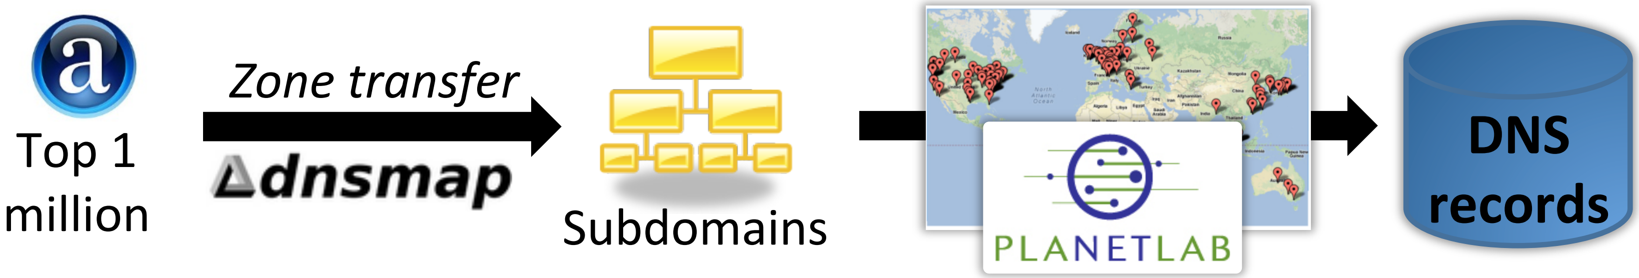

Top 1 million

Zone transfer

Subdomains

DNS records

University packet capture

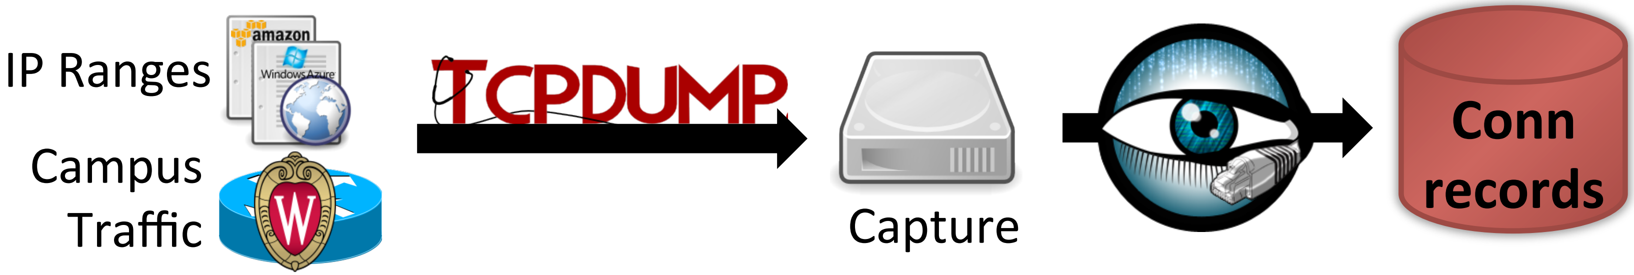

IP Ranges

Campus Traffic

Capture

Conn records

IP crawl dataset

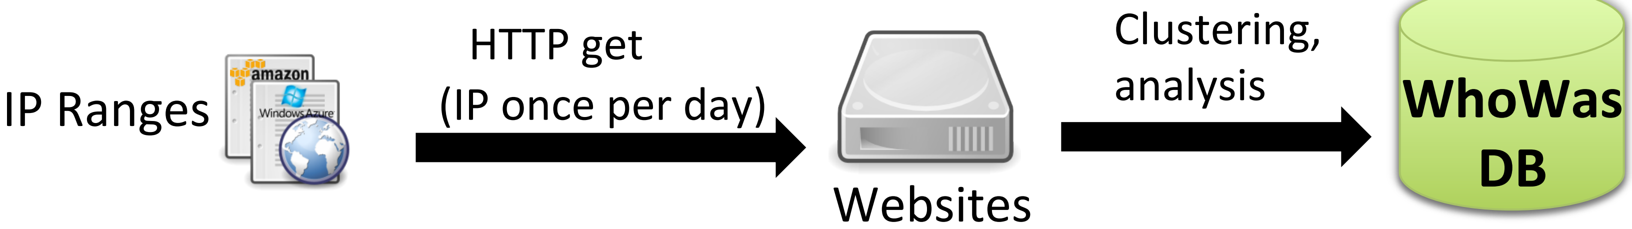

IP Ranges

HTTP get (IP once per day)

Websites

Clustering, analysis

WhoWas DB

Cloud observatory data sets

Fetch HTML content of web pages ~ every 3 days (using IP address)

Extract features to cluster IP addresses for same web page

MySQL database with front-end for running analyses

EC2: 3 months (Oct, Nov, Dec 2013)
Azure: 2 months (Nov, Dec 2013)

IP crawl dataset

900 GB of data

EC2: ~1.4 million unique IPs respond. ~300K unique clusters

Azure: ~150K unique IPs respond. ~40K unique clusters

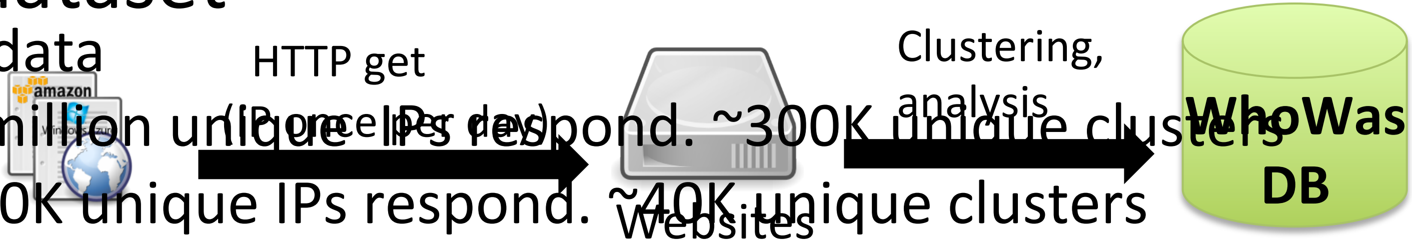

Cluster-based analysis

- Six-tuple to ***fingerprint*** an IP during a measurement round
  - <title> </title>   content
  - Keywords
  - Server software and version
  - Generator tags (e.g., PHP vs. Ruby backend)
  - Google Analytics ID number
  - SimHash of HTML textual content
- Use unsurpervised clustering. Parameters chosen using gap analysis

IP address responses over time (Azure)

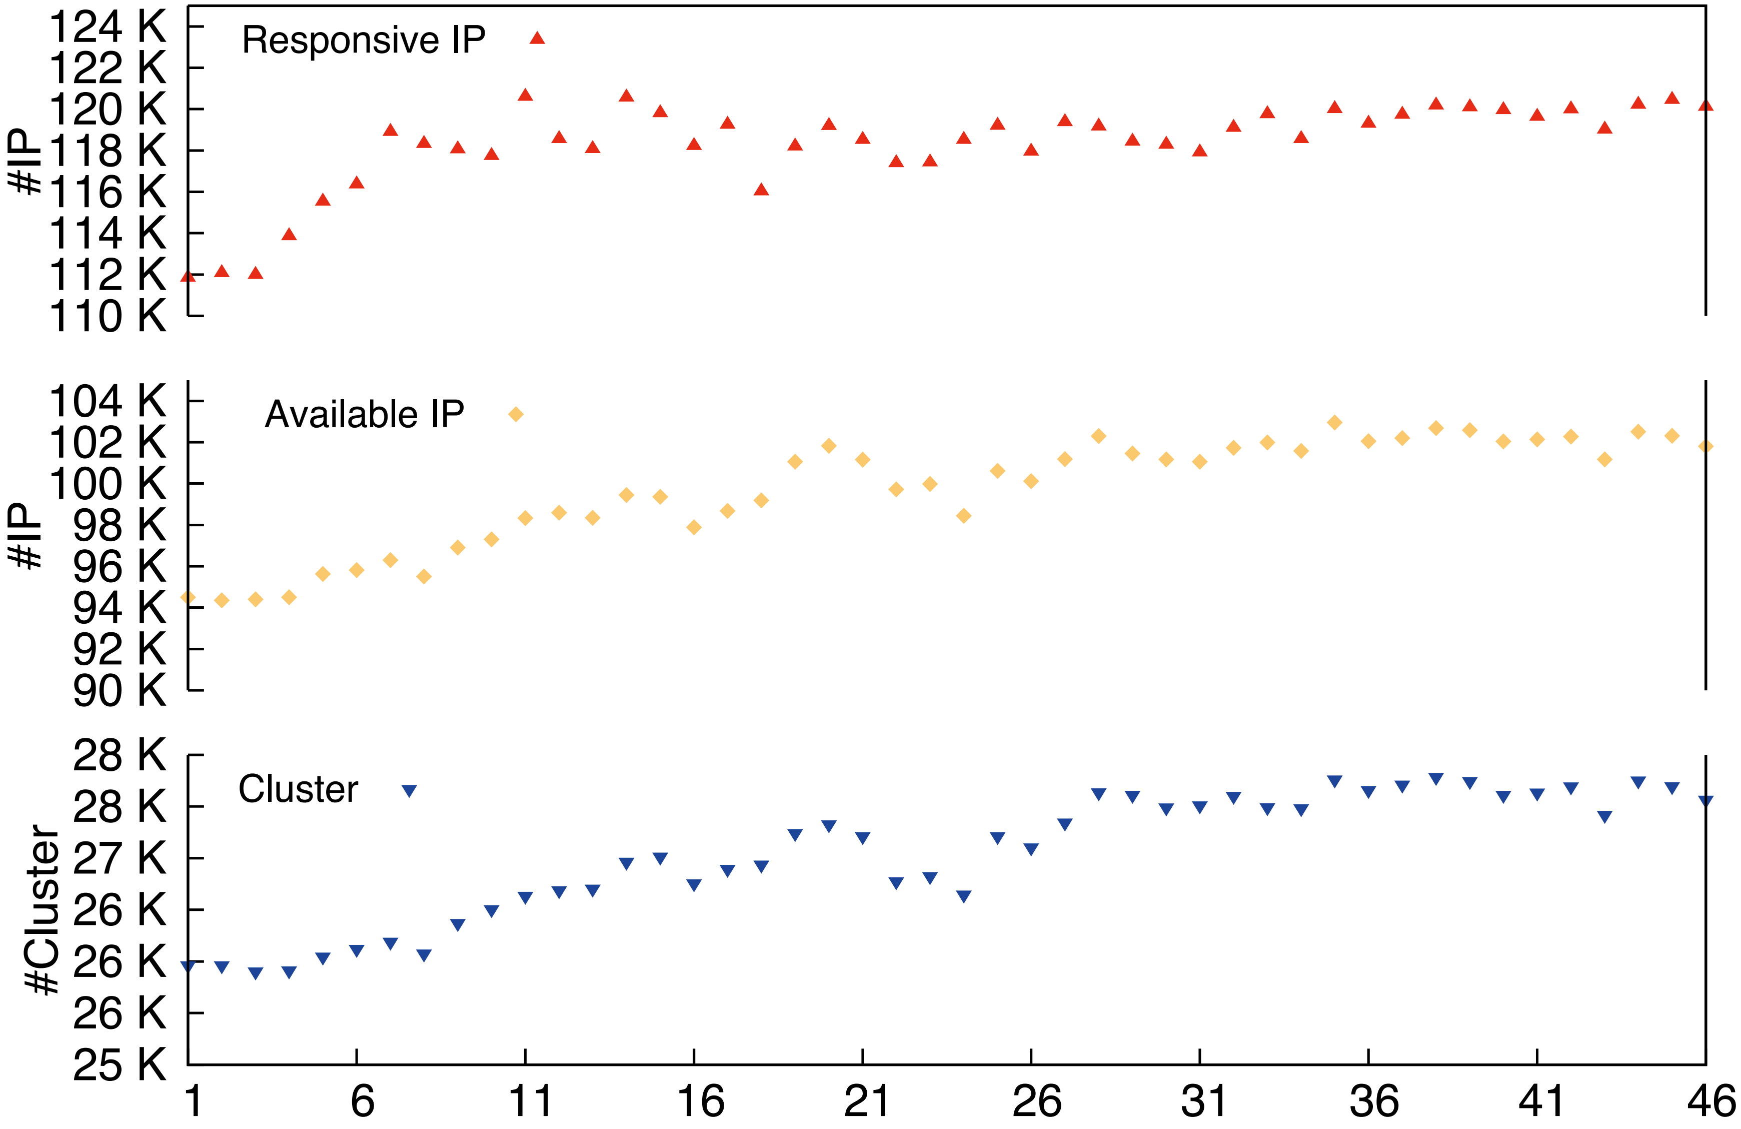

IP address responses over time (EC2)

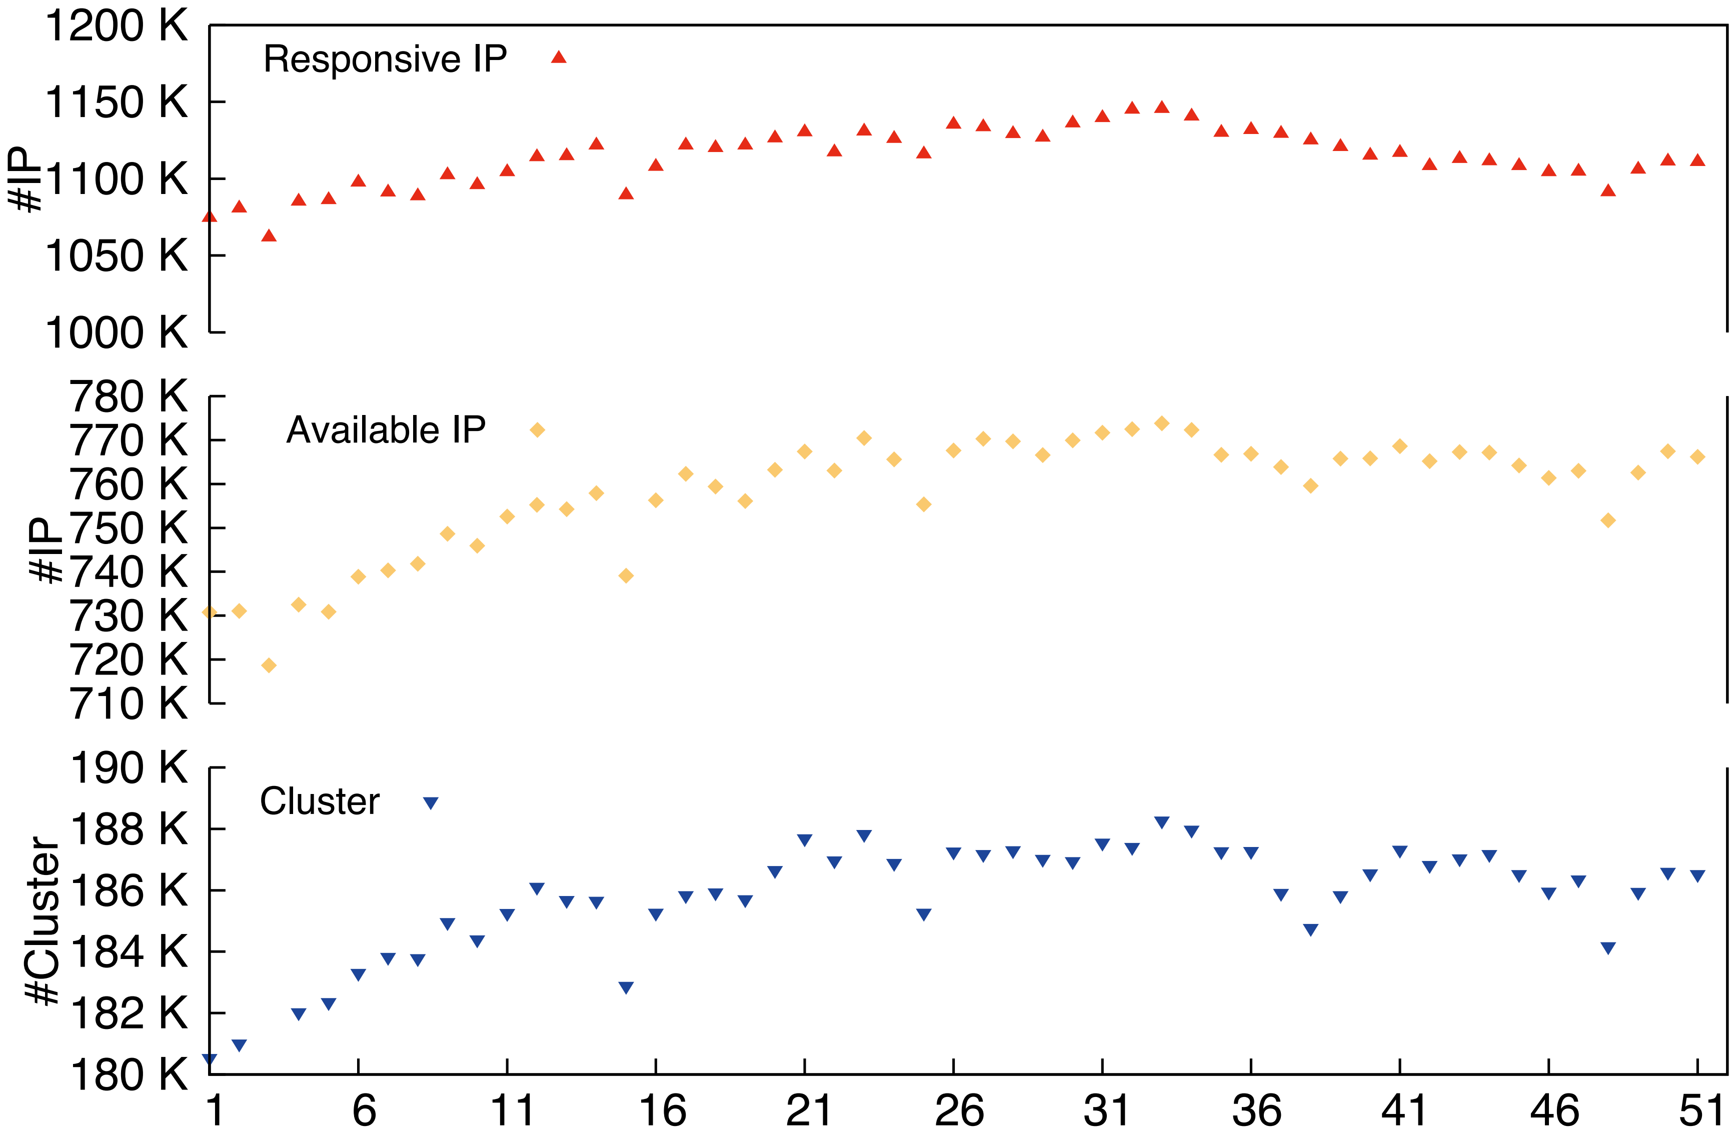

Number of page clusters over time

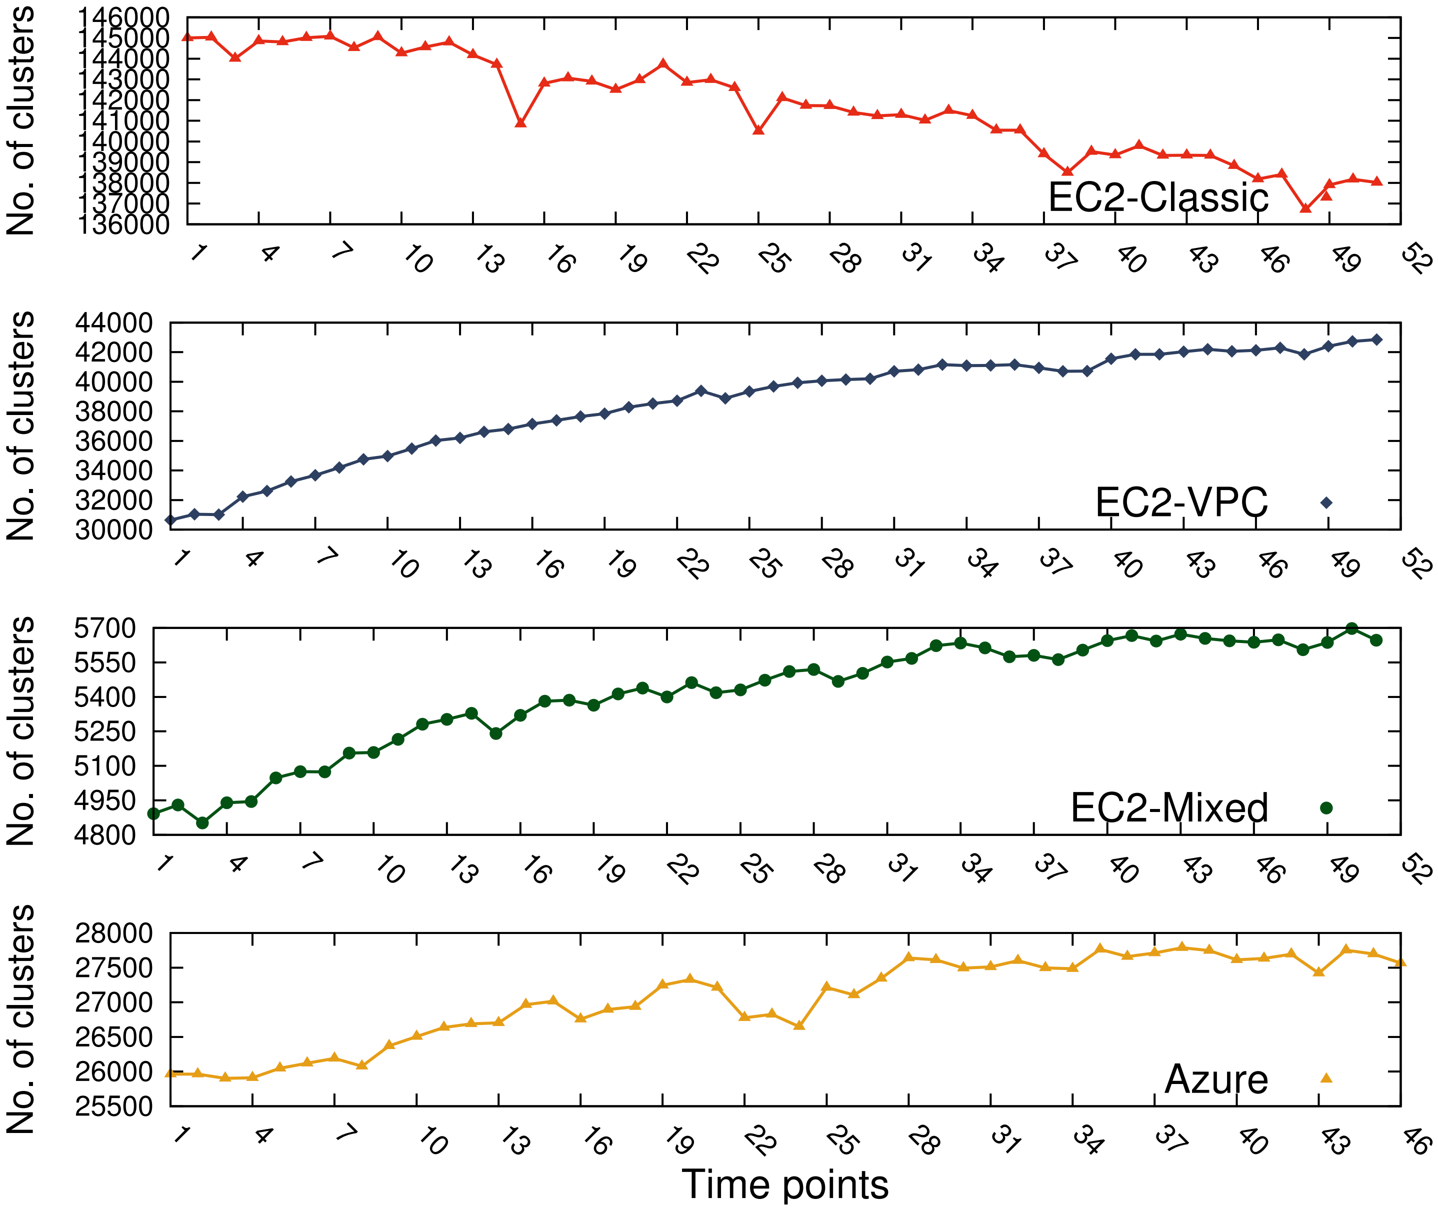

Number of IPs in EC2 reporting different PHP versions over time

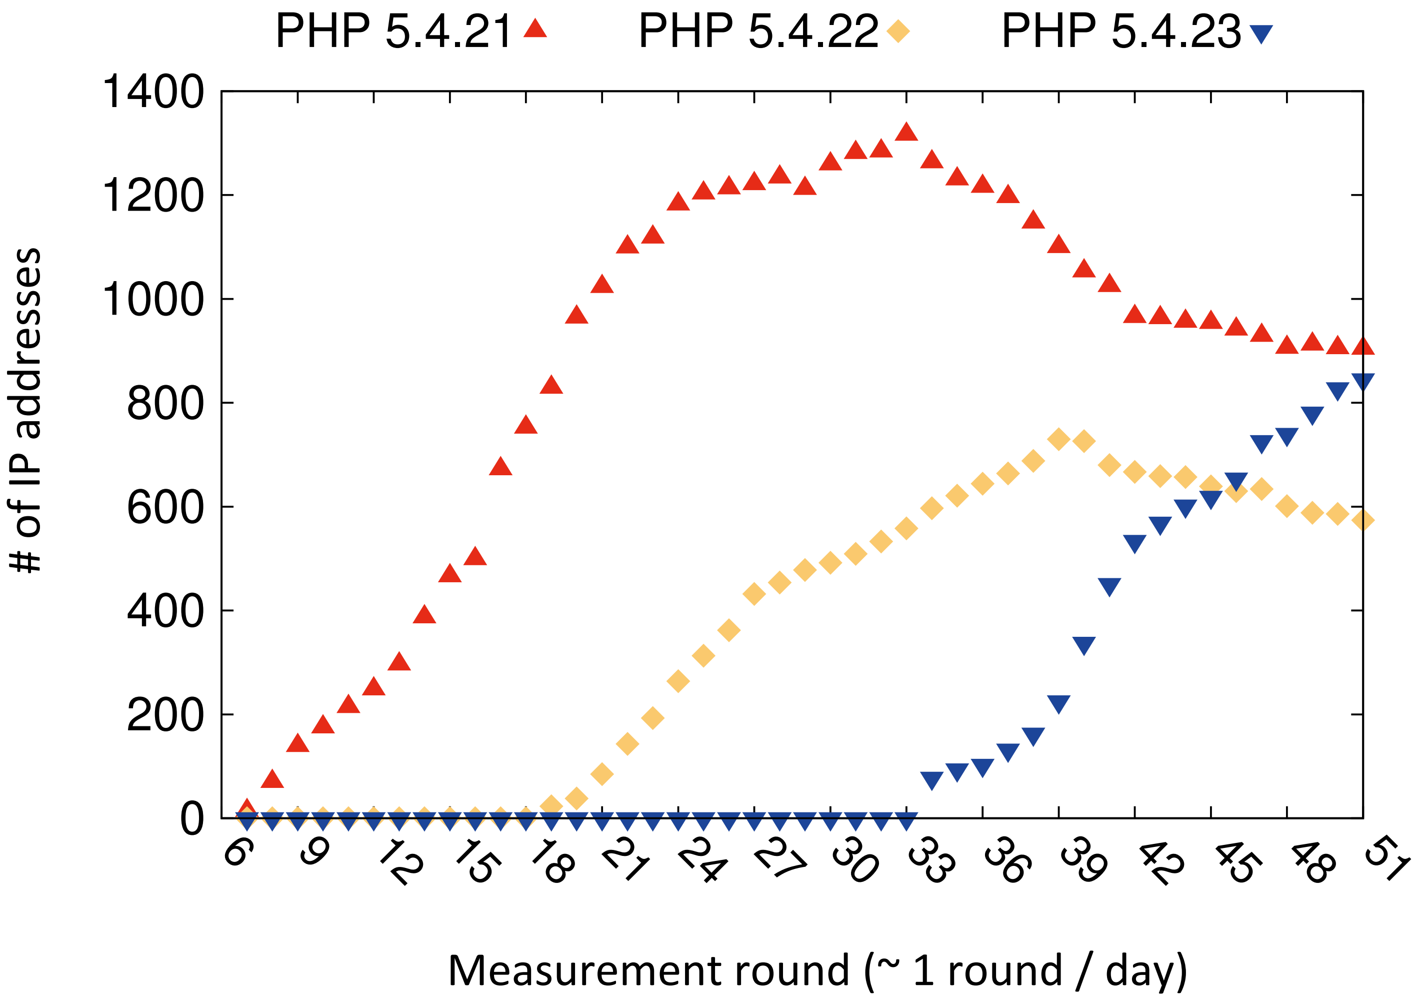

Malicious Activity?

- 3.2 million URLs collected in average 3-day period. Ran through Google SafeBrowsing.
  - 197 unique EC2 IPs contained >=1 malicious URL
  - 13 unique Azure IPs contained >=1 malicious URL
- VirusTotal (Feb 2014): 3,840 unique EC2 IPs
  - Most associated with URLs (typical keyword in domains: “download”)
  - Investigated 98 in depth:
    - use clustering to find further IPs (199 extra IPs found)
- Either case: Average uptime is ~7 days (outliers: 90+ days)

VirusTotal blacklist uptime for 98 malicious webpages

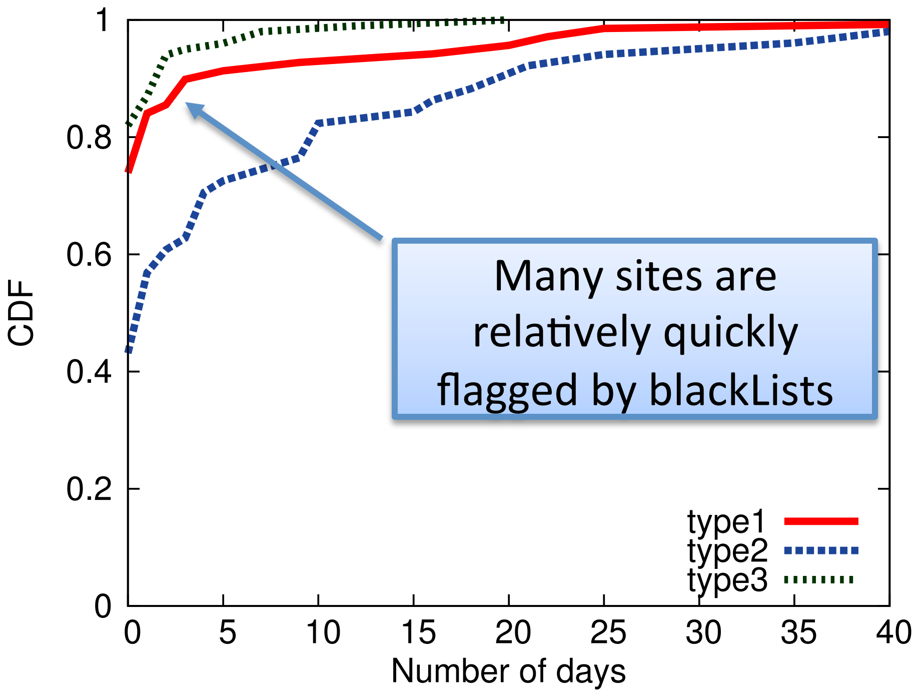

Many sites are relatively quickly flagged by blackLists

# of days website available ***before*** appearing on blacklist

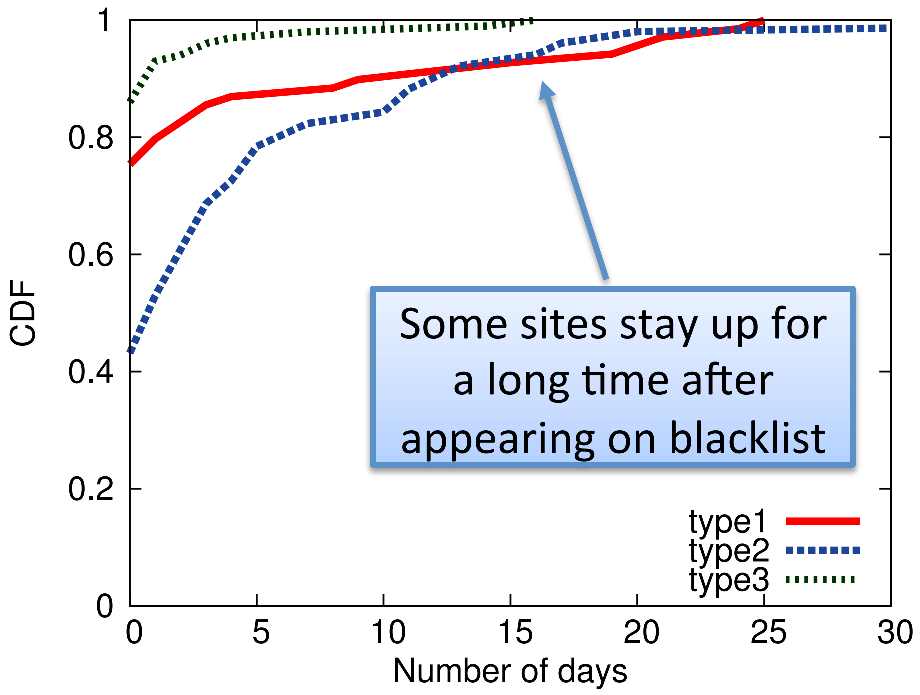

Some sites stay up for a long time after appearing on blacklist

# of days website available ***after*** last appearing on blacklist

Type 1, 2, 3 refer to different patterns of malicious deployments

Cloud Observatory is Ongoing Work

- Measure usage, security posture of cloud tenants
  - Generating several rich datasets
  - Analysis and opportunity finding

- ***Questions for you:***
  - Other questions to ask?
  - Other ideas for methodologies?
  - Further data sets?

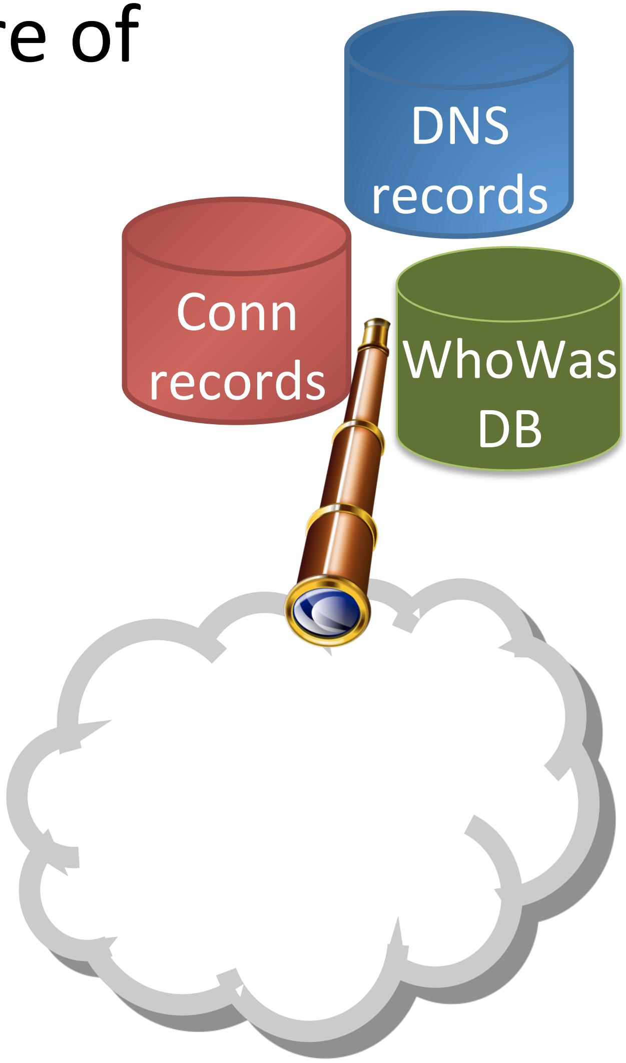

Today:
Ongoing projects involving WISDOM

- Cloud observatory
  - Provide data sets and methodologies for understanding how cloud usage evolves
- New IaaS Security Services
  - Security-posture audit tools (SPATs)
  - Other projects

Security Services for Tenants

IaaS control plane as trusted third-party for tenants

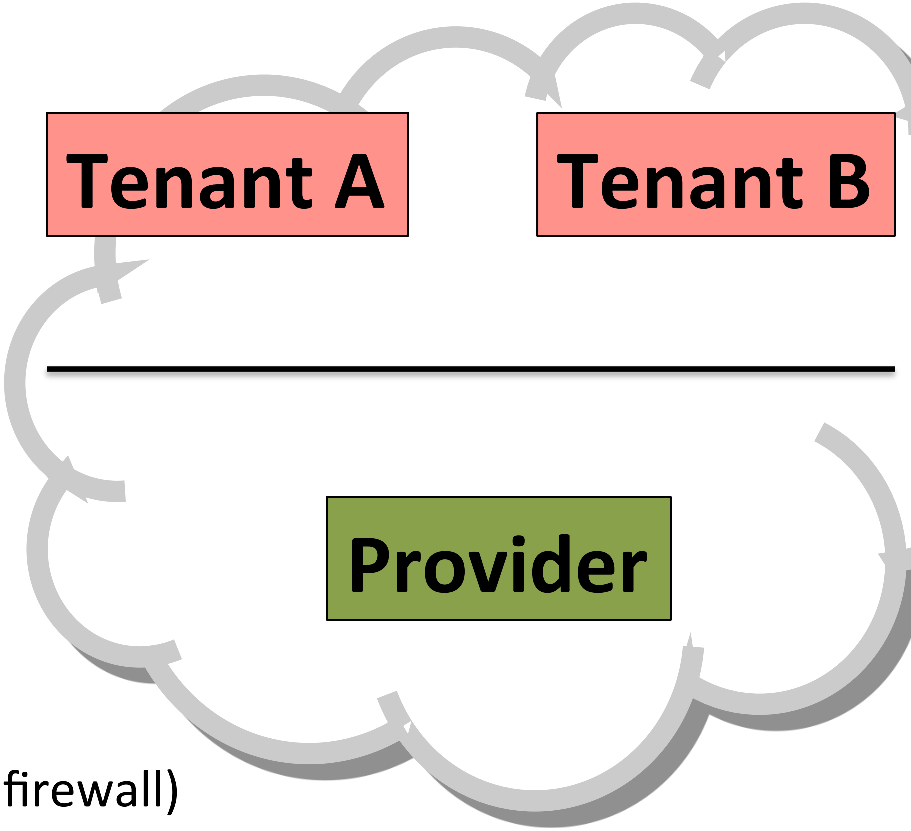

Somewhat analogous to kernel/userland interface

What can be done with this viewpoint?

Currently:

Security group settings (firewall)
Logging / billing records
HSM
PaaS/SaaS value-added services

A motivating example: Confidentiality-preserving data mining

Analysis by Alice

Alice wants to run her computations over Bob's data, but doesn't want to give Bob her code

Data owned by Bob

Bob's wants to allow this, but needs guarantees about the use of his confidential/private data

<u>Examples:</u>
Clinical outcomes data
Demographic information
Advertising data sets
Survey responses
Network security logs
…

Unsatisfying approach #1

Alice sets up IaaS VM(s)

Bob gives Alice access to data

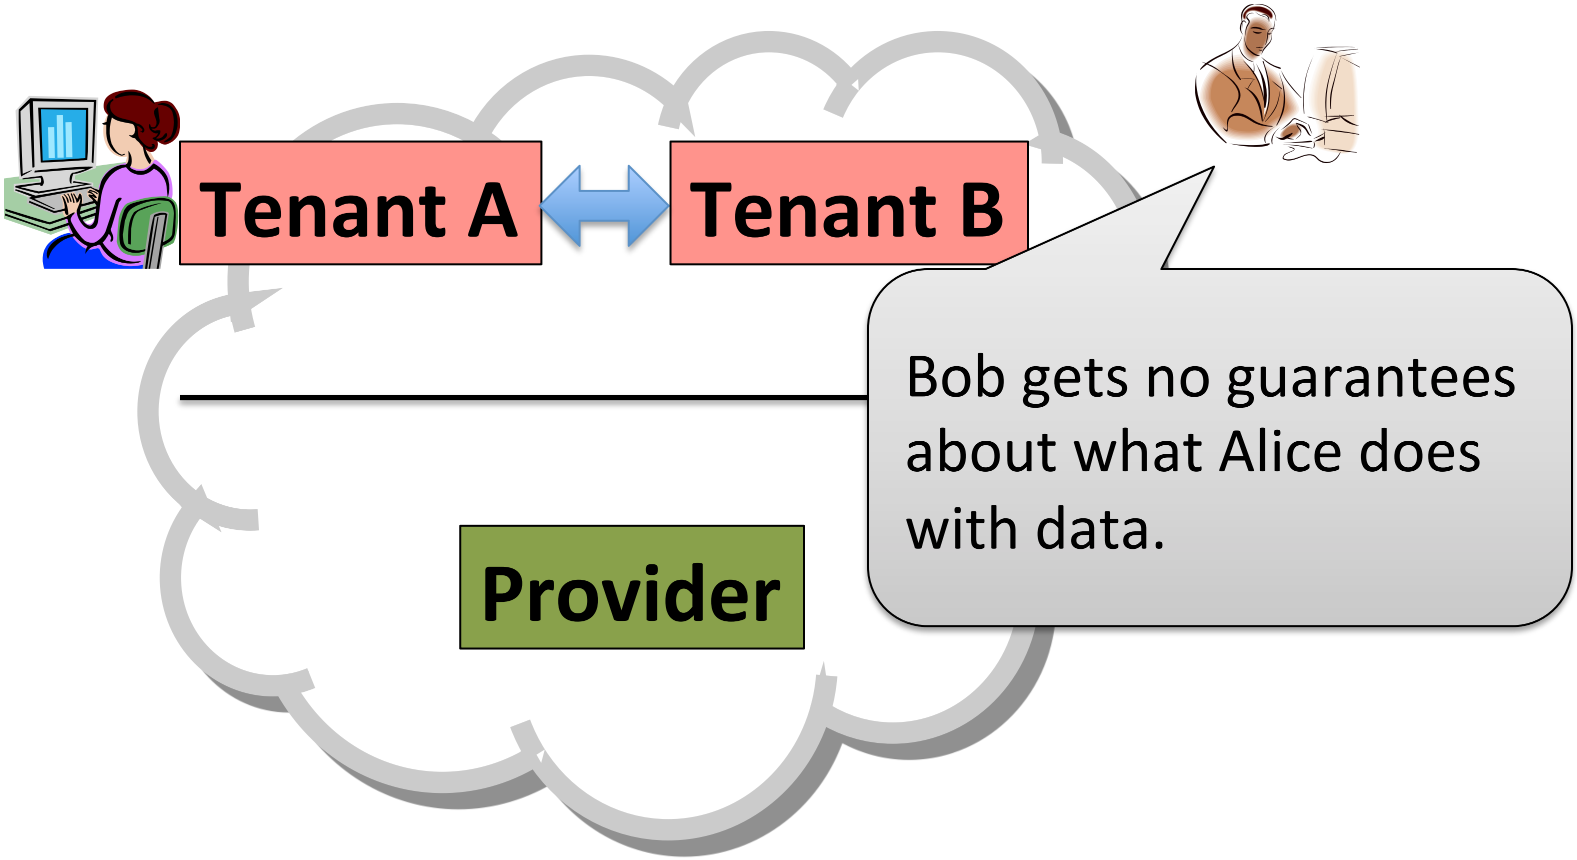

Tenant A

Tenant B

Provider

Bob gets no guarantees about what Alice does with data.

Unsatisfying approach #2

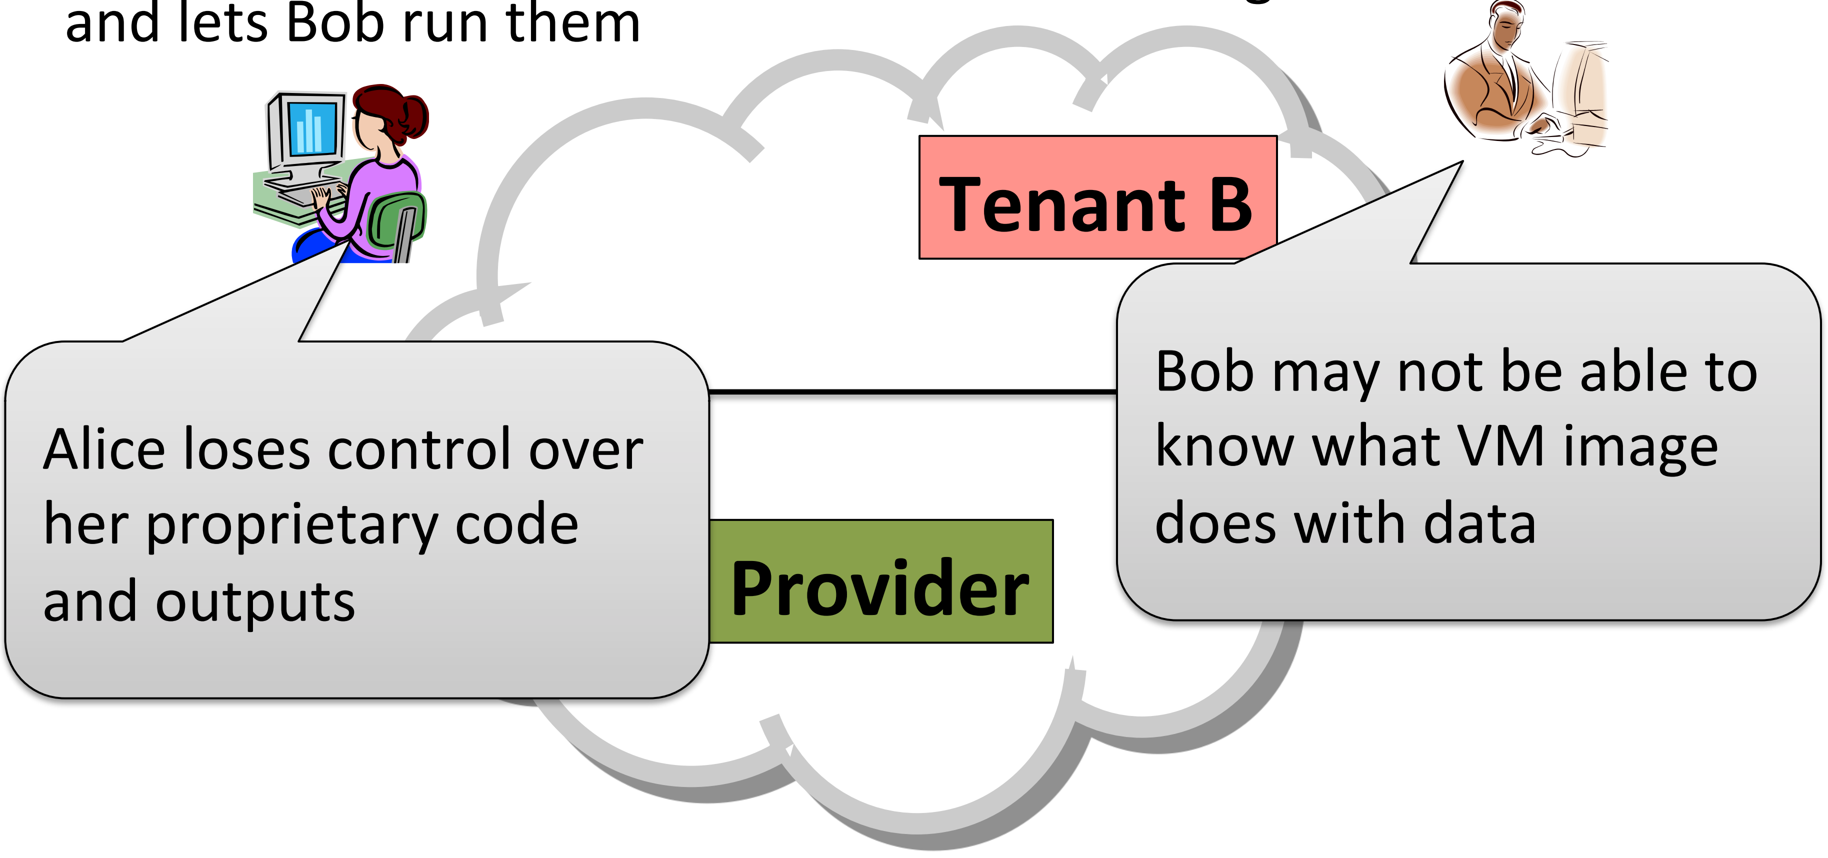

Alice setups up IaaS VM images and lets Bob run them

Bob runs image and gives it access to data

Tenant B

Provider

Alice loses control over her proprietary code and outputs

Bob may not be able to know what VM image does with data

Unsatisfying approach #3

Alice and Bob use cryptography (multiparty computation)

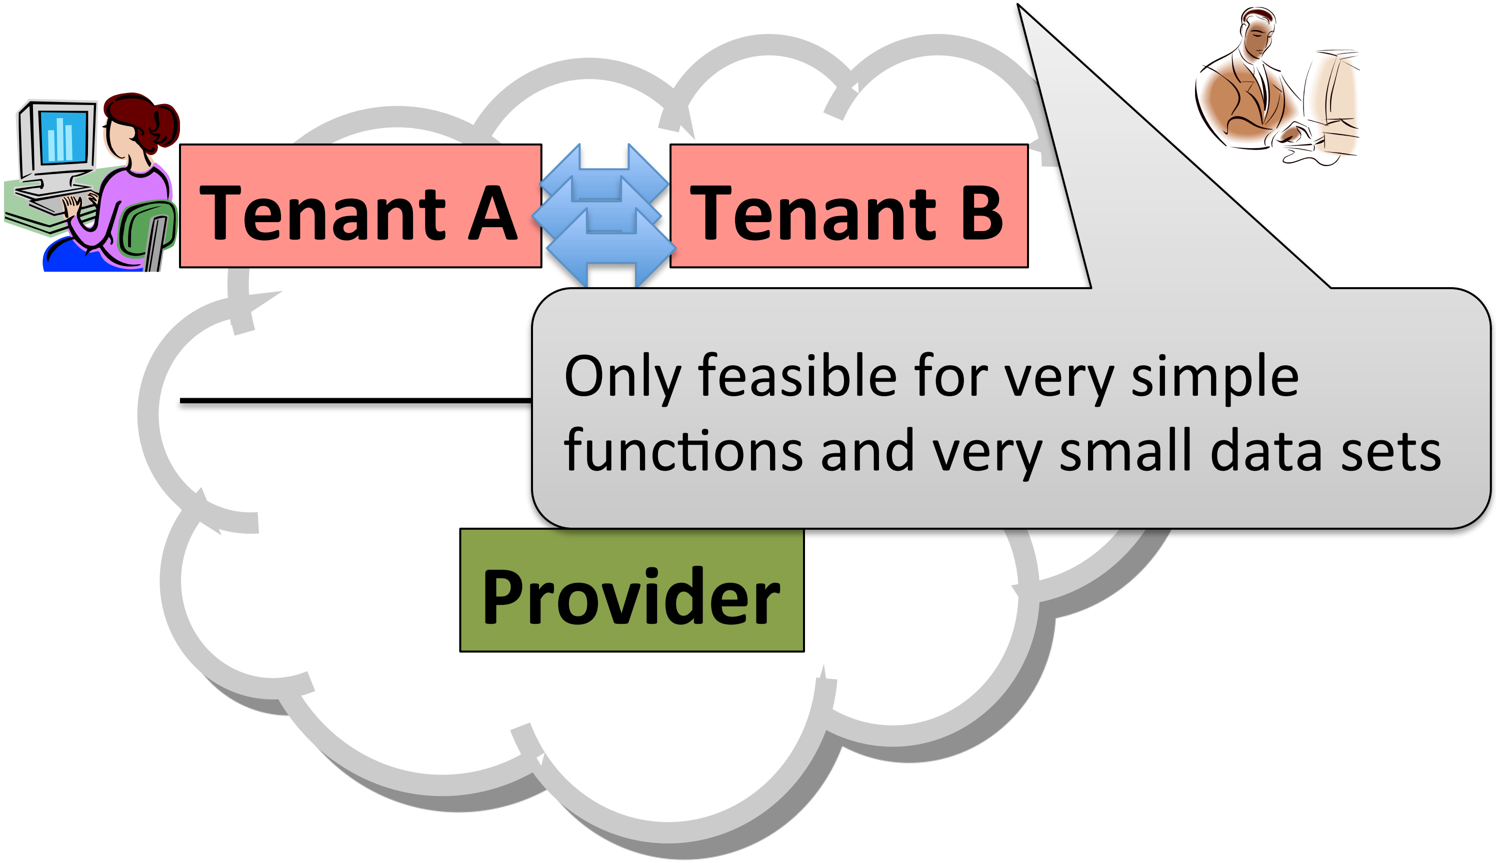

Only feasible for very simple functions and very small data sets

Instead: Leverage the provider

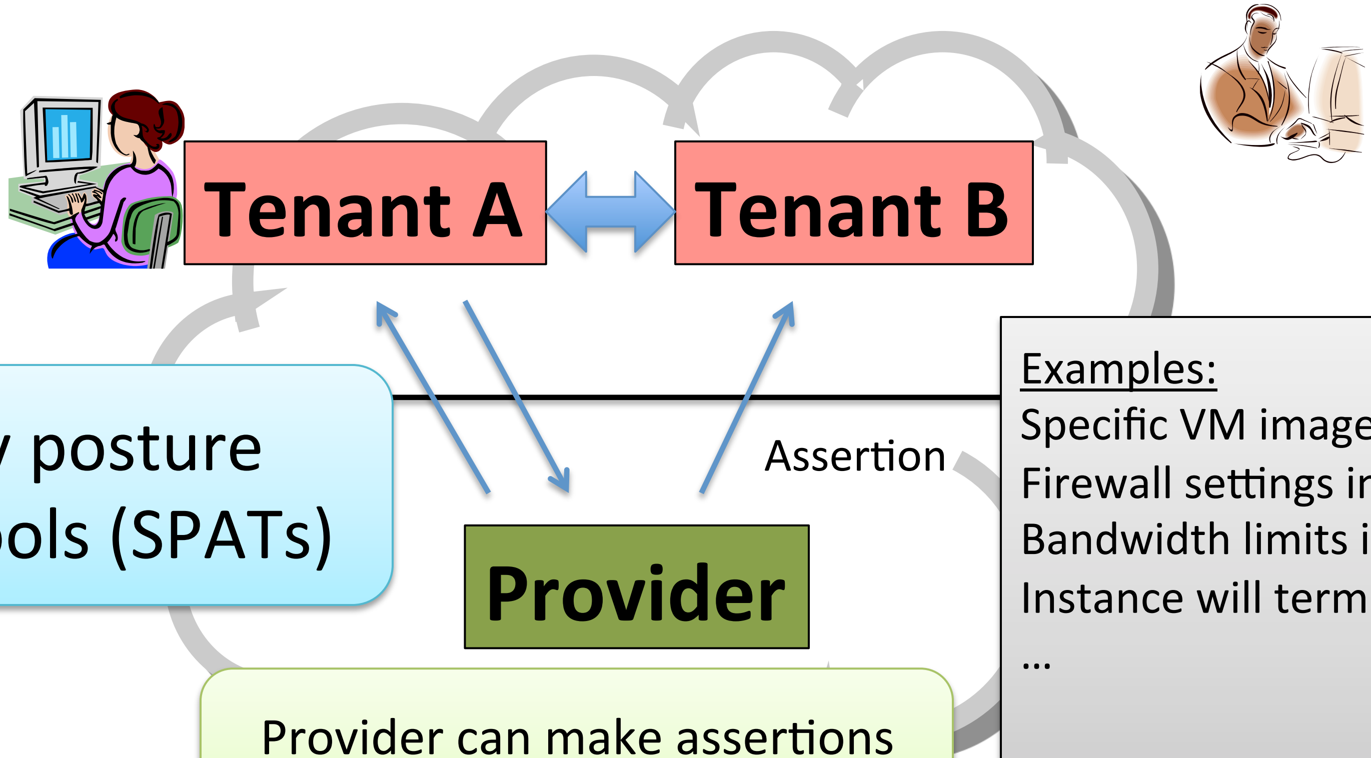

Security posture audit tools (SPATs)

Provider can make assertions about Alice's VM to Bob

Examples:
Specific VM image booted
Firewall settings in order
Bandwidth limits in place
Instance will terminate soon
…

Security-Posture Audit Tools

What SPATs are useful?

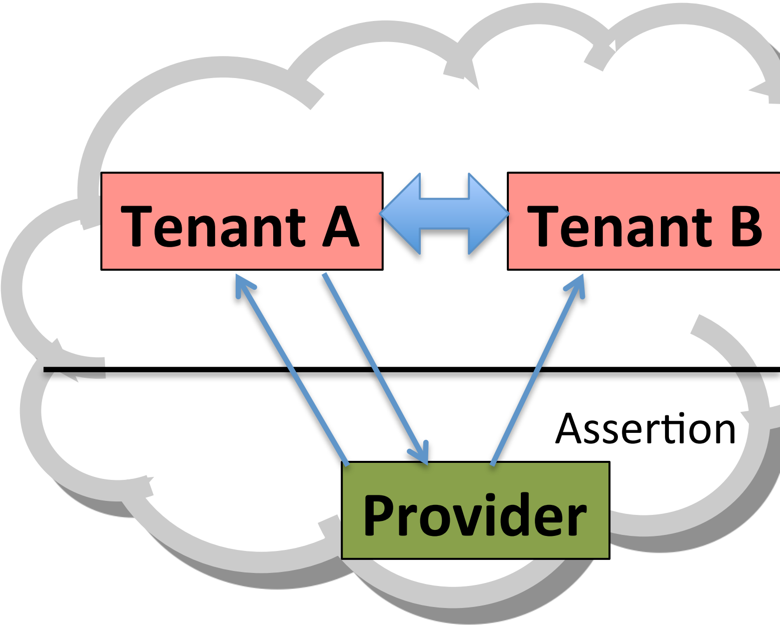

How does Alice opt-in to let Bob use SPATs on her VM instances?

How do tenants identify audited instances?

Can we make this transparent to users? SPAT-audited PaaS platform?

Ongoing work: SPATs on OpenStack

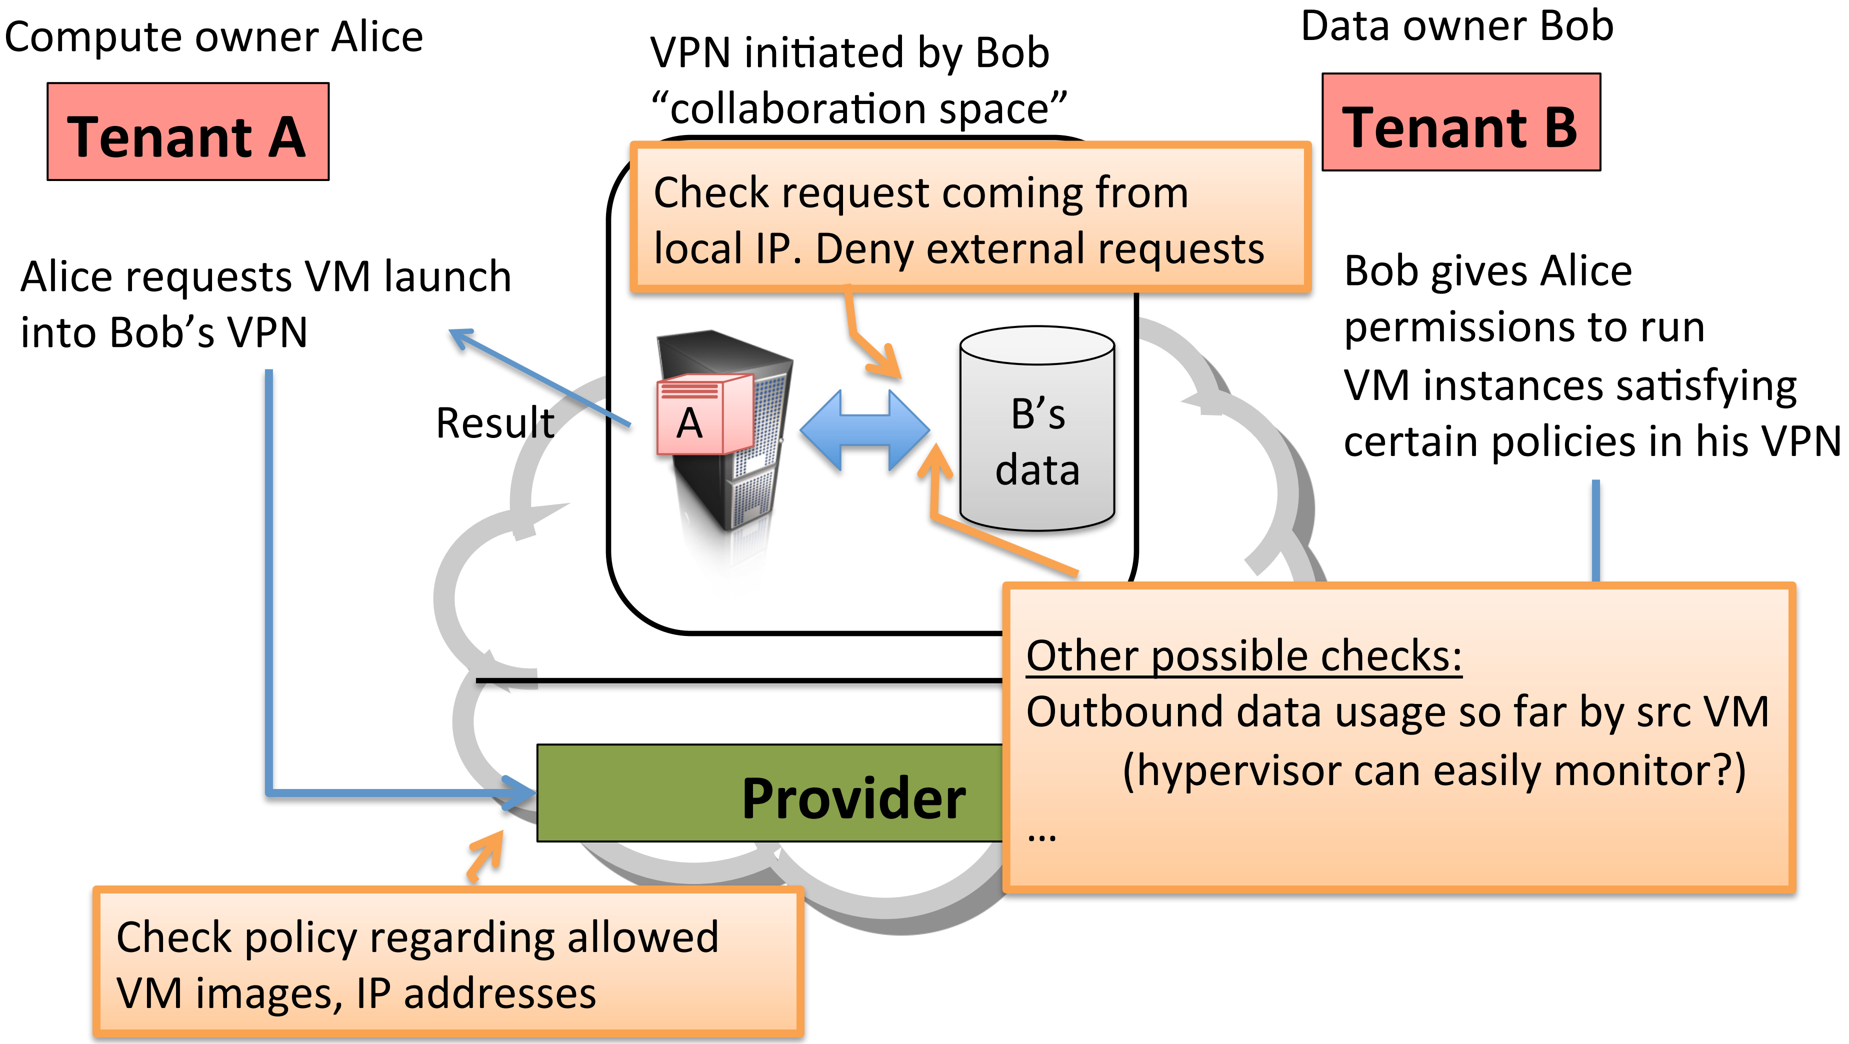

Alice requests VM launch into Bob’s VPN

Check request coming from local IP. Deny external requests

Bob gives Alice permissions to run VM instances satisfying certain policies in his VPN

Other possible checks:
Outbound data usage so far by src VM
(hypervisor can easily monitor?)
…

Check policy regarding allowed VM images, IP addresses

Security-Posture Audit Tools

- The future?
  - PaaS confidentiality-preserving data-mining platform with IaaS provider SPAT-based root-of-trust

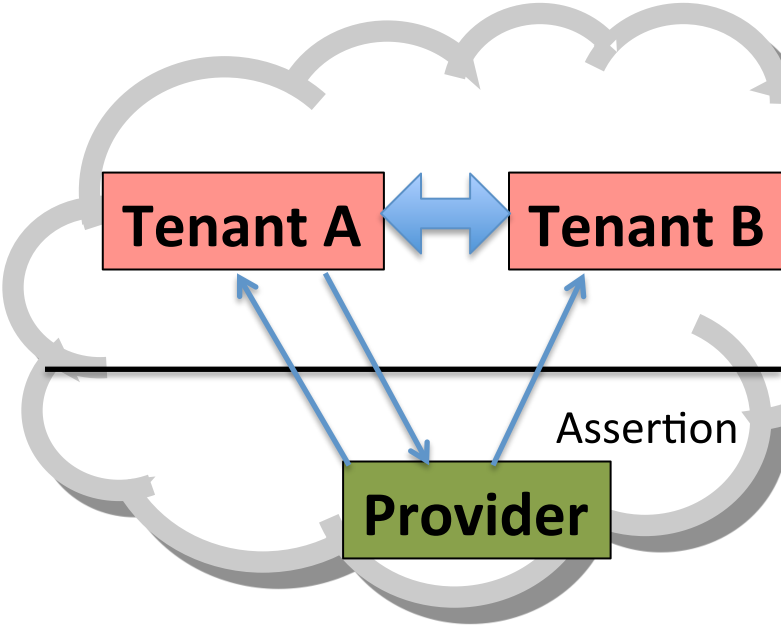

- ***Questions for you:***
  - Ideas for SPATs?
  - SPAT use cases and requirements?
  - Integration thoughts?

Some Other Silver Projects

- Unifying approach to authorization with trust logics
  - SAFE (safeclouds.org)
- Policy management
  - SDAC (Software-defined access control), user-facing interfaces, tools to aid policy configuration
- Infrastructure
  - SDN, middleboxes, hypervisors
- Encryption services
  - >90% of EC2 web connections are HTTP (circa 2012)
  - Can we change that to HTTPS (or something even better)?

New encryption primitives

Format-transforming encryption

- Encryption whose ciphertexts guaranteed to match against input regex

[Dyer et al., CCS 2013]

Message-locked encryption

Dropbox

Google Drive

- Encryption for which outsourced storage can dedup given just ciphertexts

[Bellare et al., Eurocrypt 2013], [Bellare et al., USENIX 2013]

Honey encryption

- Password-based encryption for which decrypting with wrong password leads to plausible plaintext [Juels and Ristenpart, Eurocrypt 2014]

Rethinking Security in the Era of Cloud Computing

- Cloud observatory
- SPATs and IaaS root-of-trust primitives
- Other Silver Projects

- Feedback please!