# What we really do with NSF facilities

Susan E. Sons
Executive Director, OmniSOC & ResearchSOC
sesons@iu.edu

# Current ResearchSOC Members:

**NRAO**: the National Radio Astronomy Observatory

**NOIRlab**

**GAGE**: Geodetic facility for the Advancement of GEoscience

**NSO**: National Solar Observatory

**ARF**: Academic Research Fleet

**ACCESS** *coming soon*

*-- OmniSOC also serves 11 higher education institutions and 2 regional R&E networks. --*

# Always watching...

- OmniSOC Core

  shared 24x7x365 security operations center for research & higher ed.

- Project Liaisons

  Onboarding support, single point of contact for emerging issues, interpretation, coordination of mass security events

- SOC-plus

  Honeypots, vulnerability scanning, tabletop exercises, and more

- Virtual Cybesecurity Services (staffing)

  CISO, CISO advisory, virtual cybersecurity teams, partial FTE security analysts/engineers, and specialized incident response teams

# When incidents happen

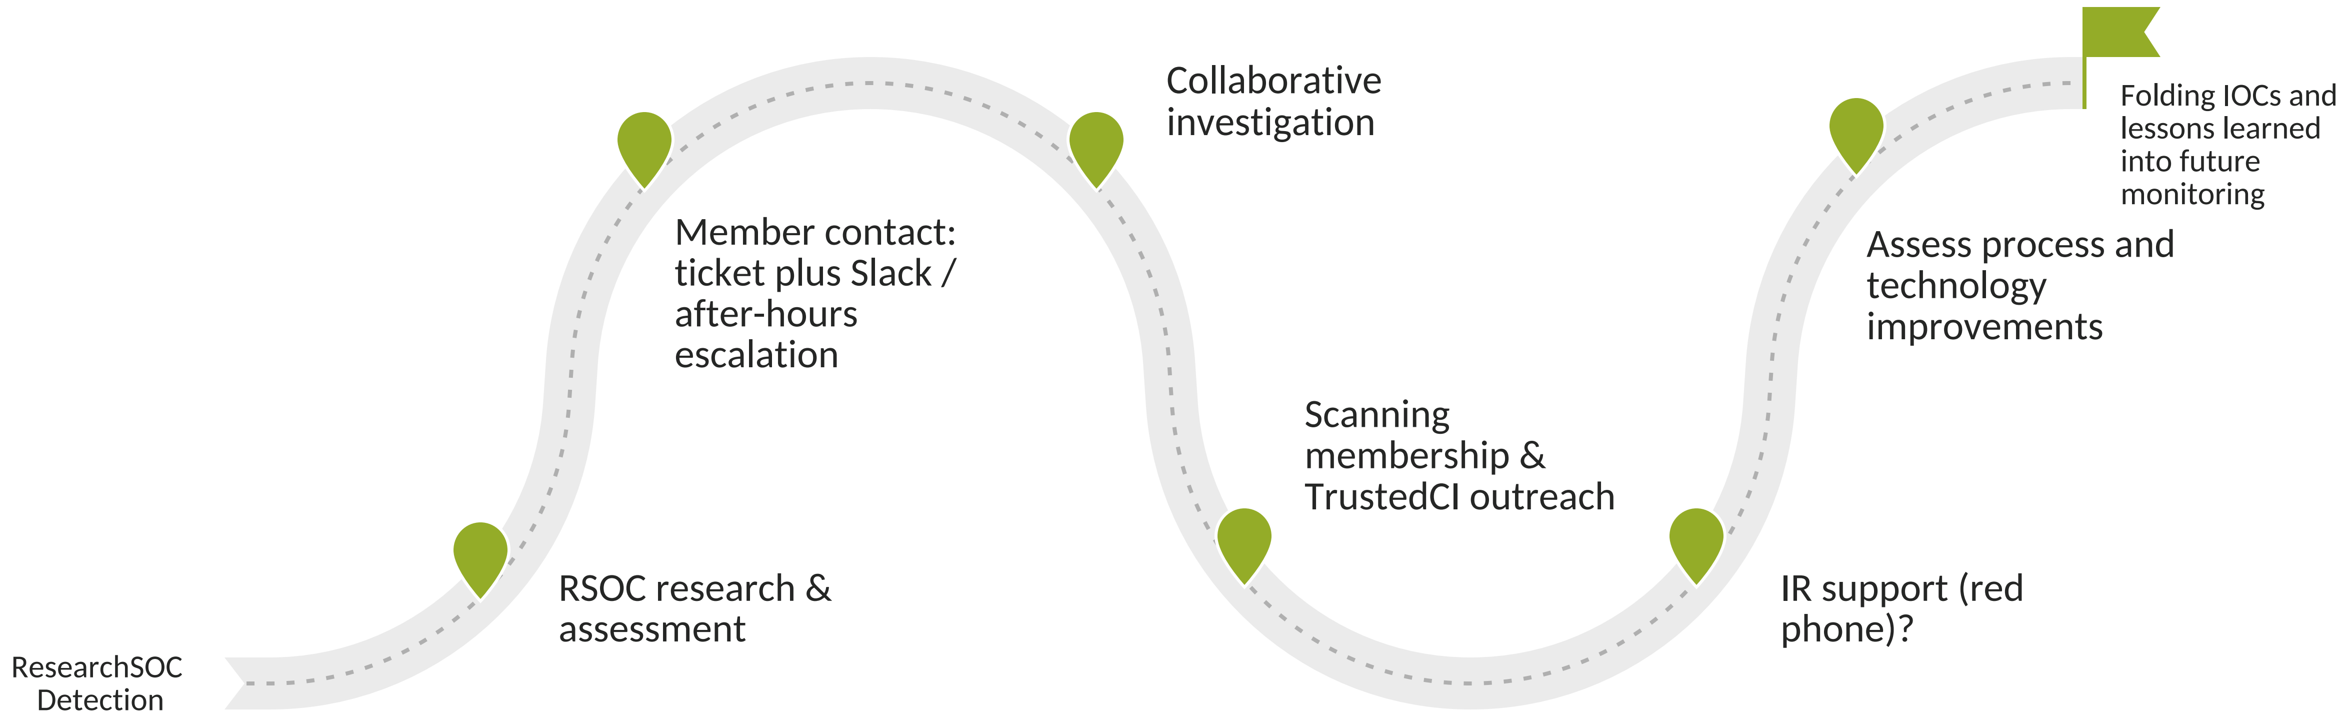

ResearchSOC Detection

RSOC research & assessment

Member contact: ticket plus Slack / after-hours escalation

Collaborative investigation

Scanning membership & TrustedCI outreach

IR support (red phone)?

Assess process and technology improvements

Folding IOCs and lessons learned into future monitoring

# A Community Approach to Cybersecurity

# Advantages over security silos:

## Reducing costs:

There's a minimum level of overhead to stand up a SOC…regardless of how big or small the SOC is.  We enable members to share those expenses rather than bearing them individually.

## Understanding threats in our vertical:

Emerging threats are immediately scanned for across all of our members.  TrustedCI is notified of threats which are likely to impact the NSF community as a whole.

## Evolving with CI:

Meeting each member facility where they are, and helping them to up their cybersecurity game over time.

## Career paths for SMEs:

Reduced turnover of cybersecurity SMEs due to career paths within ResearchSOC and OmniSOC

## Learning Together:

Facilitating the development of shared tools, information sharing, and cybersecurity research within our community.

# Bringing in new facilities…

1 **Scoping and Contracts**

Choose the right services, establish the contractual relationship and billing.

2 **Deeper Discovery**

Set up communication channels, better understand your infrastructure, create onboarding plan

3 **Onboarding: Platform**

Set up monitoring appliances and aggregators, assess feed quality, begin normalization and enrichment workflows, deploy endpoint solution

4 **Onboarding: Sec. Engineering**

Establish monitoring rules and workflows, begin threat hunting,

5 **Onboarding*: Virtual Services**

- Get to know your Project Liaison
- Establish vulnerability scanning
- Set up honeypot management platform
- Introduce VCS personnel

6 **Move to production**

# Member Ops...

Quarterly leadership meeting

Biweekly ops meeting

Ad hoc comms via Slack

Project Liaison

Reporting and Dashboards

Shadow Sessions

VCS Ops

Continuous improvement

# The Dirt

Walk first, then run.

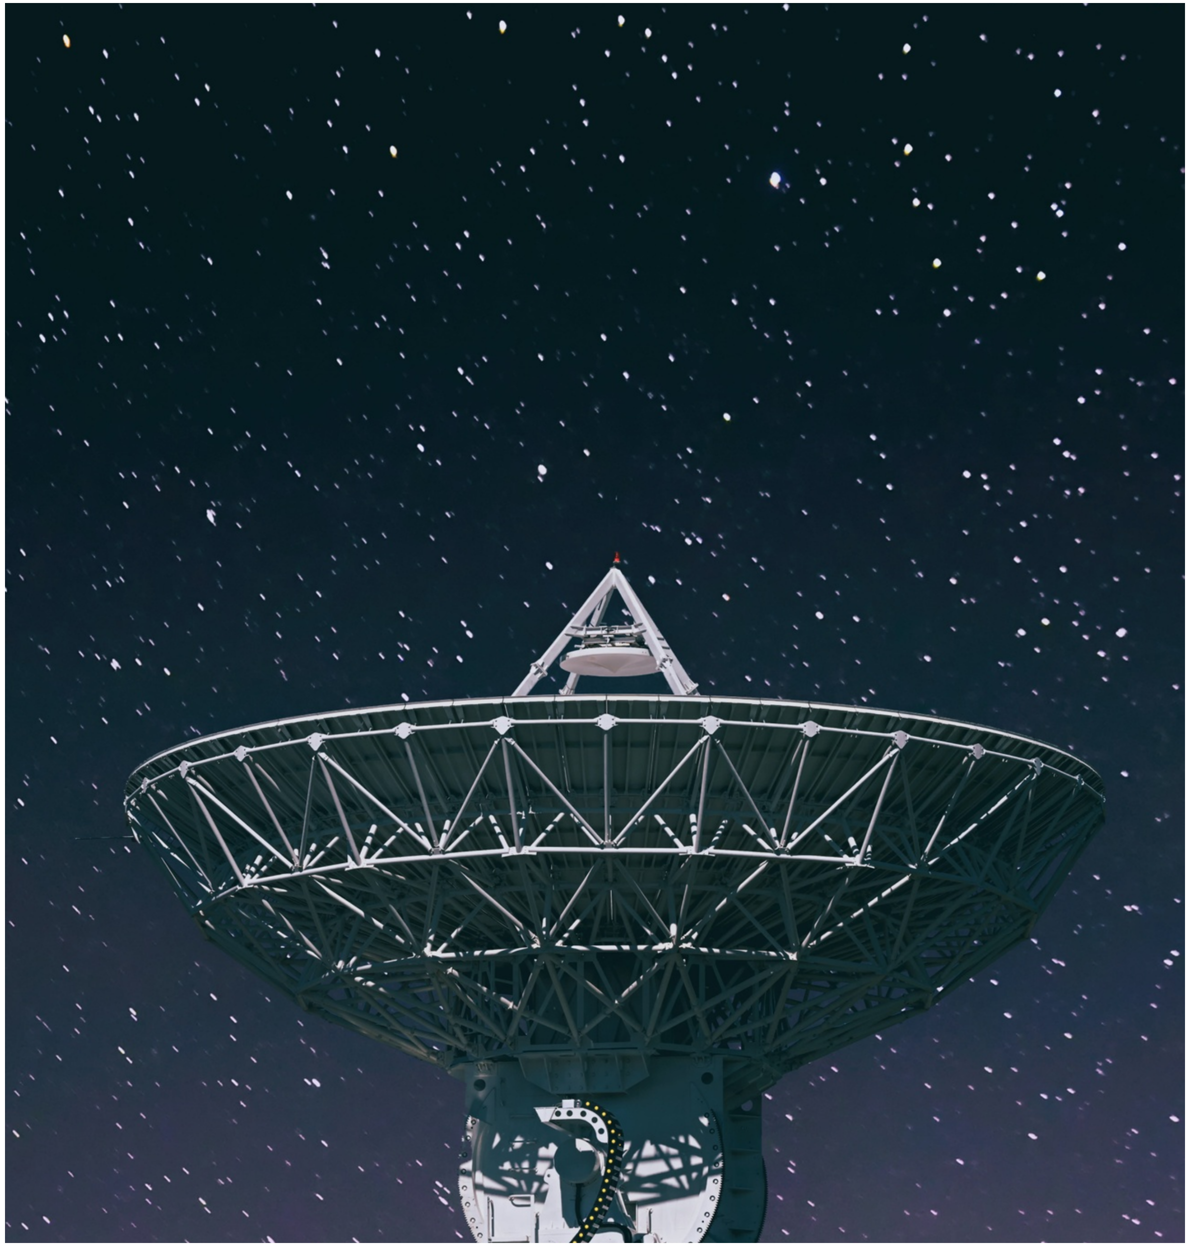

# OT is special IT

- Scientific instruments which are computer controlled and/or network connected
- Supporting ICS/SCADA: telescope domes, UAVs, power regulation, temperature regulation, dish aiming, etc.
- Facility systems (e.g. fire, HVAC)
- Safety threats

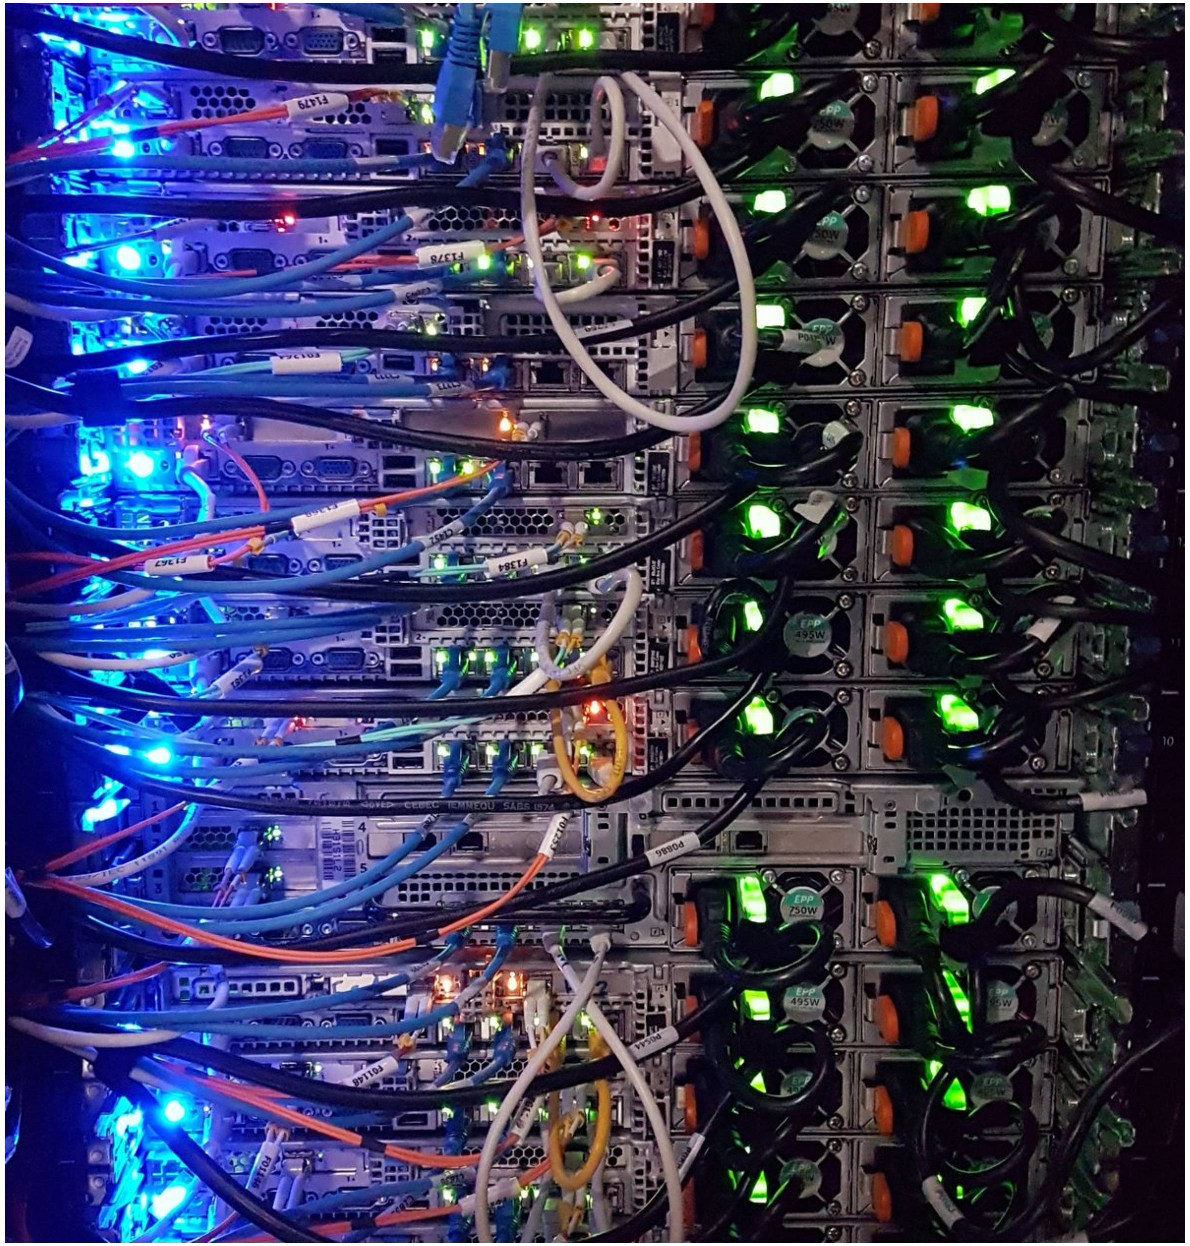

# Blinky Box Obsession

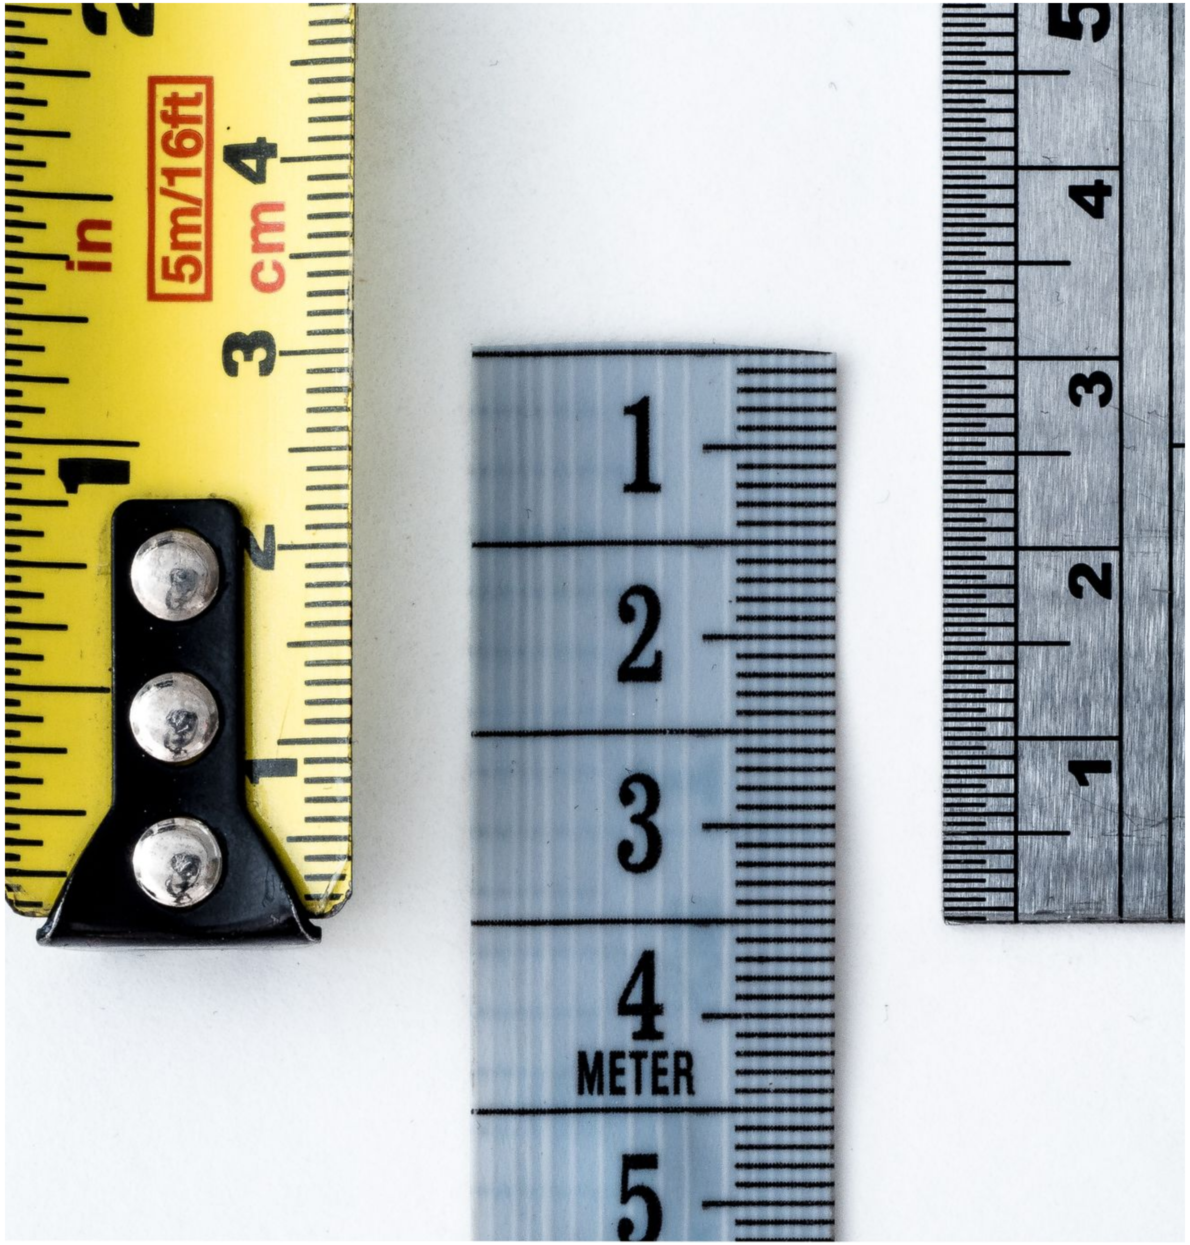

Measurement is hard.

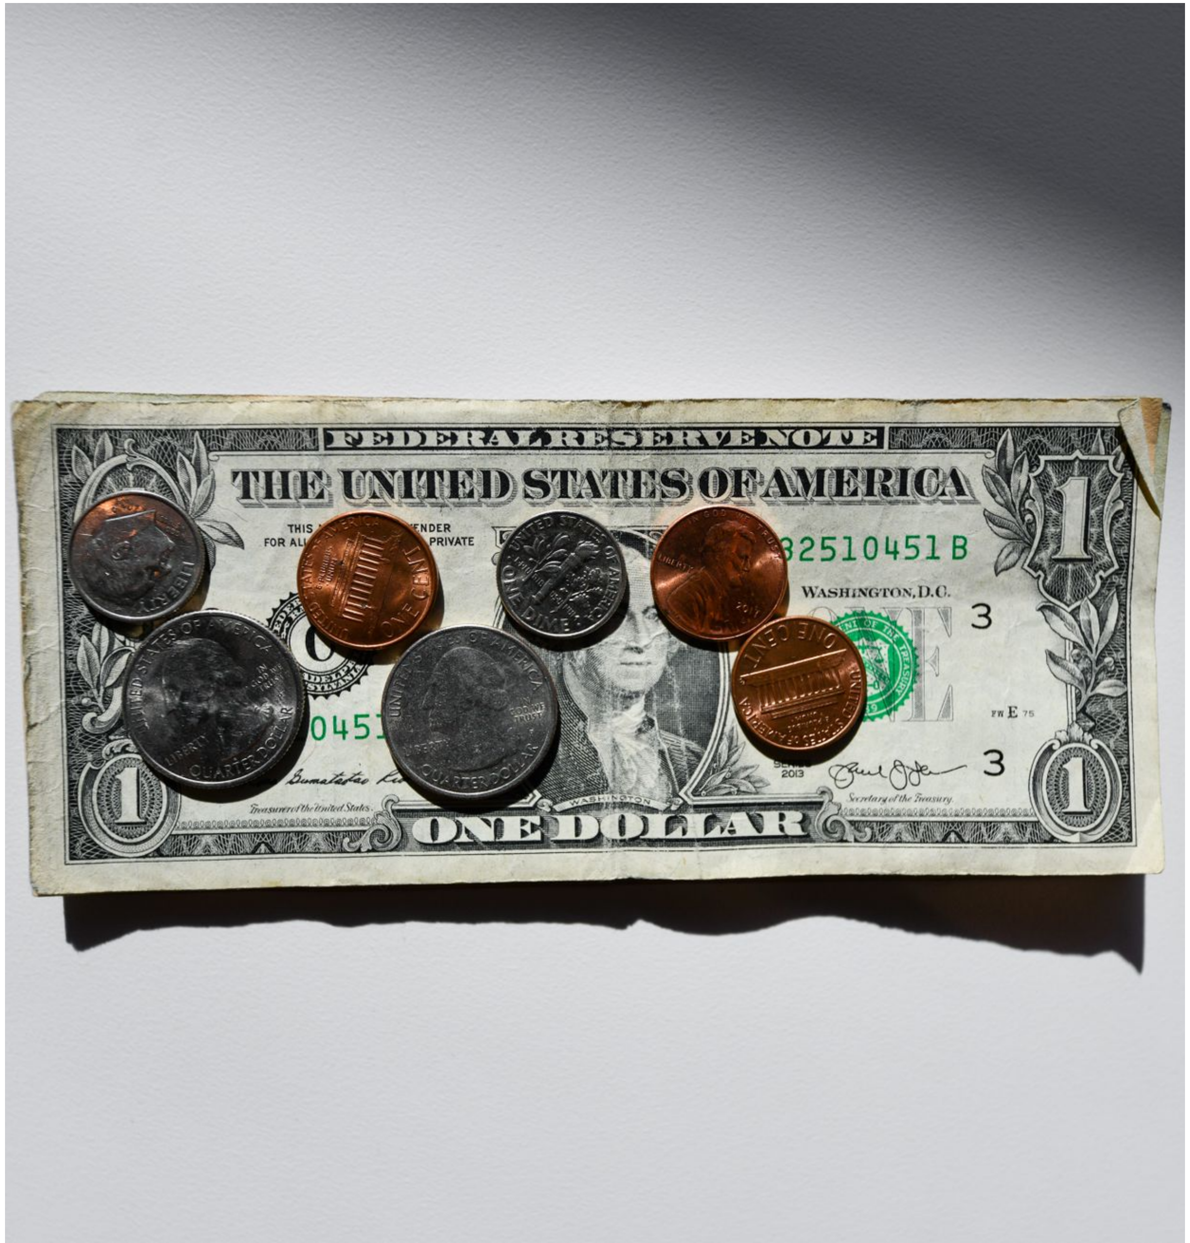

# Funding Challenges:

- Security as an unfunded mandate
- Stone soup & bolt-on Security
- IT budget distortion
- Security budget uncertainty

# Cheap Wins:

1. NSF PO Engagement
2. Start during facility planning phases
3. Security advocates on IT & OT teams
4. TrustedCI resources (Framework, engagements, fellowships, and more)
5. OmniSOC & TrustedCI webinars
6. NSF Cybersecurity Summit
7. CIS Critical Security Controls (use v8)
8. Security Ops Assessment (ResearchSOC)
9. CISO Advisory (ResearchSOC)
10. Security SME Advisor (ResearchSOC)

# Hidden OmniSOC Activities

SOC Labs

Member Community

Threat Intelligence

# Questions?

omnisoc.iu.edu

sesons@iu.edu

# How OmniSOC Works

# Example Log Requirements

### Must have (bare minimum):

- Traffic Session - N/S (Examples: Netflow, Zeek Conn logs)
- NIDS (Examples: Suricata, Palo Alto Threat logs)
- DNS query logs (Examples: Bind, Zeek DNS)
- Endpoint logs for critical systems (Crowdstrike, Elastic Endpoint, MS Defender)

### Adds significant additional value:

- Application layer protocol analyzers (e.g. Zeek HTTP, SMTP, TLS, Weird)
- Centralized Authentication (e.g. AD, kerberos, o365)
- Malware (e.g. MS Defender)
- Web Proxy (e.g. Palo Alto, Squid)

### Helpful but not required:

- NAT logs for public/private IP mapping
- DHCP (examples: Zeek DHCP)
- Traffic Session E/W - (e.g. Netflow, Zeek Conn logs)
- Endpoint logs for high value (less than critical) systems
- Wireless / VPN Authentication
- Vulnerability scanner results (e.g. OpenVAS, Nessus, Qualys)
- Honeypot (e.g. Duke STINGAR)
- Service specific access logs (e.g. MSSQL, MySQL, Apache)

# Typical Architecture

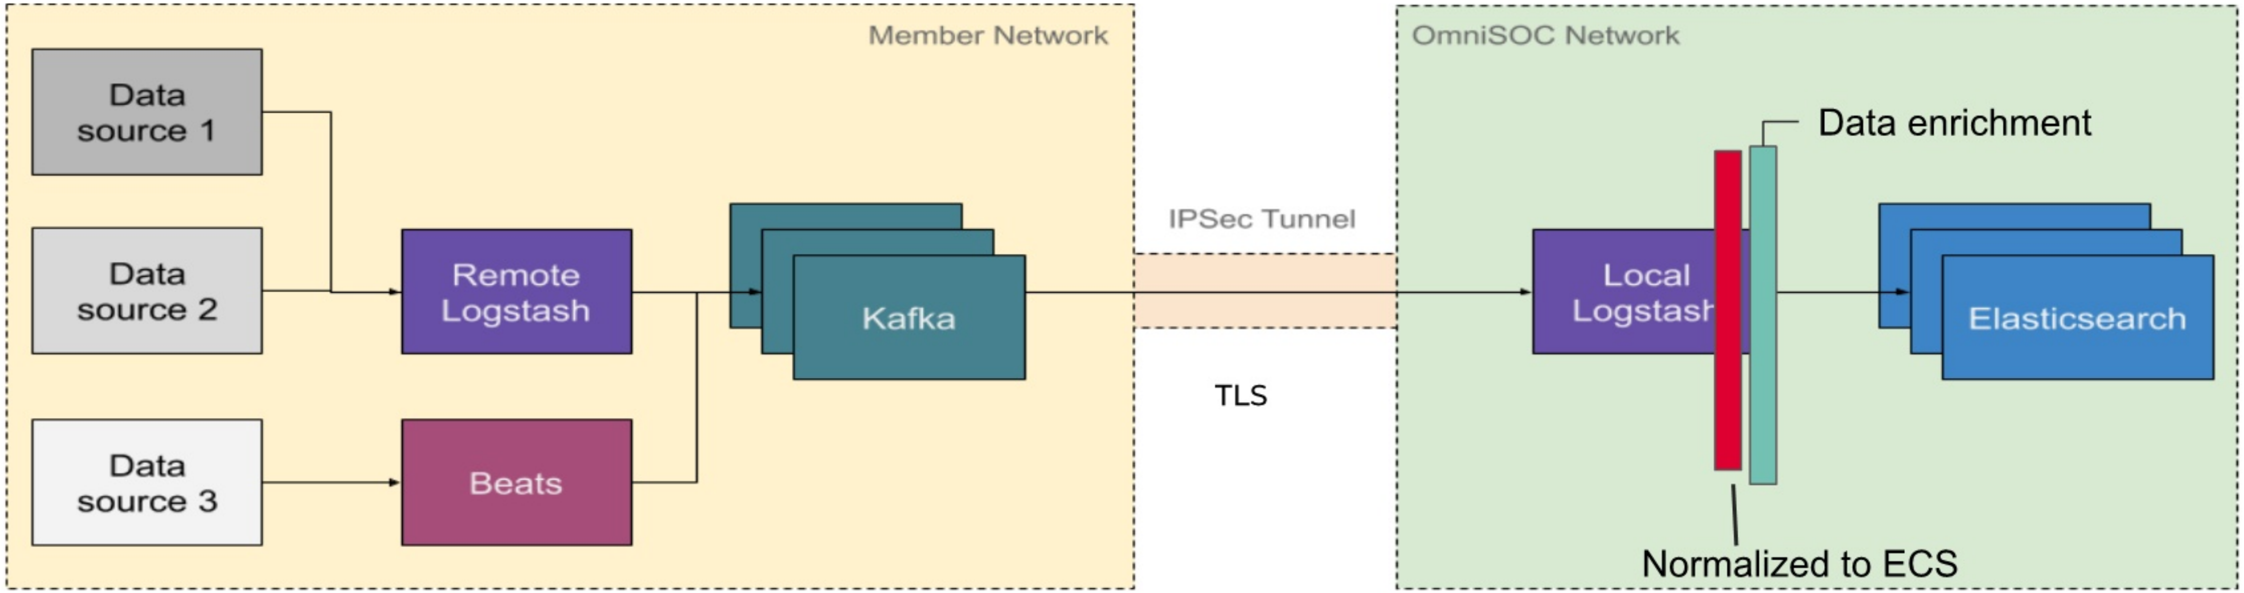

- ➔ Raw data collected by Logstash on aggregators
- ➔ Inserted into Kafka as-is
- ➔ Picked up by IU-based Logstash instances
    - ◆ “Store and forward”, no data loss due to network interruption

- ➔ Data normalized to ensure consistency between different sources
- ➔ Data augmented with additional metadata
- ➔ Finally indexed into Elasticsearch

# Enrich and then Analyze

- GeoIP tagging
- ASN information
- Critical IPs, netblocks
- Threat feed matching
- Reverse DNS
- Whois lookups, domain age
- Standardization & Normalization
- Use of ECS

- Central Data Collection/Enrichment
  - Elastic
- Automated Alerting
  - ElastAlert
  - Elastic SIEM Detection Engine
- Big Data Analysis
  - C2 Beacon Detection: Flare
  - Anomaly Detection: Elastic Machine Learning
- Incident Case Management
  - TheHive (including Cortex OSINT modules)
- Threat Intelligence Sharing

  CIF

  REN-ISAC

# Visualize, search, alert, rinse, repeat…

- Elastalert (we’ll likely replace with Elastic’s Signal Detection rule engine)
  - Alerting engine used for creating custom, high fidelity, alerts.
- Elastic Common Schema gets us:
  - SIEM app+workflows
  - Pre-built Machine Learning jobs.
  - Pre-built Signal Detection rules
  - and more…
- Endgame (Elastic Endpoint Security)
  - Elastic acquired Endgame in 2019. We are working on a PoC now, as well as ‘dog fooding it’ ourselves. Best of breed EDR.

# Visualize, search, alert, rinse, repeat…

- What’s new?
  - Alert, ASN, DNS, User-Agent.
    - Looking toward Elastic ML outlier detection
- What’s high/low volume?
- What’s in the news (threat intel)?
  - Cross cluster search is awesome.
- What’s seen at other members?
- Can we detect beaconing?
  - Flare.
  - Elastic ML?

Maintain hunting cadence independent of calendar.

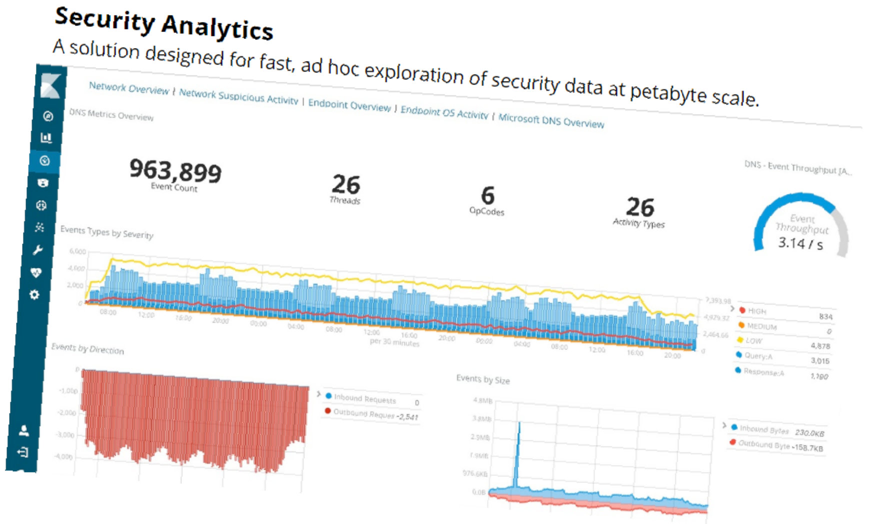

**Security Analytics**

A solution designed for fast, ad hoc exploration of security data at petabyte scale.

# Handling events, finding evil

- ➔ 350K events per second peak
  - ◆ 5 IDS/IPS alerts per second, bursty
  - ◆ 200 distinct “new” signature IDs per day

- ➔ Don’t assign tickets (busy work/false positives); improve UI (case management integration) and assign “beats”.

- ➔ Stay in Kibana (as much as possible), zoom in/zoom out, filter, filter, filter.
  - ◆ Looking toward Elastic SIEM and ECS dependent tooling.