FORENSIC AND MEASUREMENT NETFLOW
BRIDGING A GAP WE DIDN’T NOTICE

Michael Collins (mcollins@isi.edu)

Flocon 2024

January 9, 2024

Talk Outline

1. Biography, Thesis, Front Matter
2. A Brief History of the Two NetFlows
3. ISI’s work in Forensic NetFlow Analysis
4. Conclusions

INTRODUCTORY MATERIAL

Biography

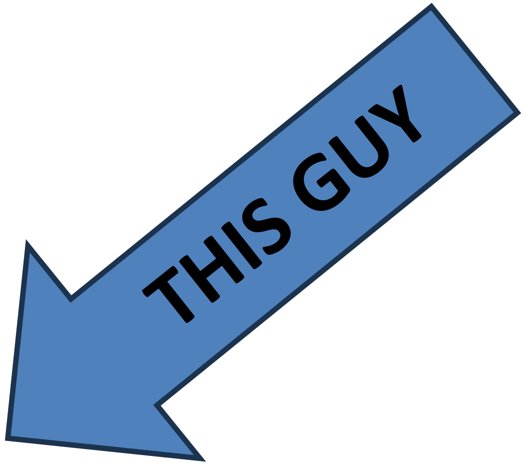

- Dr. Michael Collins
- Senior Computer Scientist, University of Southern California
  - Located in Arlington
- Started work in information security in 2000 working at the CERT, CMU
  - Developed core data collection/analysis systems for DoD and DHS
  - Investigated multiple insider threat, hacking and data breach incidents
- Focus areas: network monitoring, security operations, testbeds and experimentation, satellite security

Thesis of the Work

- There are two NetFlows and we haven’t really come to grips with that
  - *Measurement Flow:* a sampled measurement standard for traffic summarization located a router, and a low-priority monitoring process
  - *Forensic Flow:* an unsampled forensic tool used for security investigations, to complement EDR and because we don’t have better options
- These two flows have different needs which are increasingly diverging, and the pressure on the security side is getting stronger
- In the majority of this talk, I am going to discuss what future work in forensic flow may look like
  - I will discuss things USC-ISI is doing here as well

WHY THERE ARE TWO NETFLOWS

NetFlow Bifurcated on December 8, 2000

- Core flow measurements described in Claffy's "A Parameterizable Method for Internet Traffic Flow Profiling" (IEEE J. Sel. Comm., 13(8), Mar. 1995)
  - Claffy built on previous work in packet trains
- CISCO Implements NetFlow in 1996 as an internal switching technology
  - NetFlow as we think of it was originally the debugging and monitoring data for that switching
- Romig, Fullmer and Luman write a set of analysis tools published in "The OSU Flow-tools Package and CISCO NetFlow Logs" (Proc. 2000 LISA Conference, December 2000)
- Suresh Konda builds on flow-tools to build the SiLK Suite.

Measurement Work Since 2010

- Research based on SDN has opened up new forms of *software-defined traffic measurement.* These techniques use languages such as P4 to implement *ad hoc* measurements at the switch.
- Canary Testing: continually pushing samples of live traffic through modified network appliances to test impact of design changes
- Sketches: Diverse stateful traffic summarization techniques implemented at the switch.
- These techniques are, in general
  - Switch-resident
  - Highly flexible
  - Header-focused
  - Statistical or summary-oriented

The Fundamental Difference

SAMPLING

Non-Sampled Flow and its Implications

- Forensic flow analysis requires *reconstructing events;* given conventional sampling rates it is possible the event will not be present in the record.
- (Non-)Sampling impacts
  - Data storage sizes
  - Regulatory record retention requirements
  - Response time

The Encrypted Elephant In The Room

- DPI is on its deathbed
  - IAB recently held a workshop, M-TEN, focusing on network management in an ell-encrypted world
- Increasingly, the SOCS I talk to use a combination of EDR and Flow collection.
  - EDR is the smoking gun evidence
  - Flow covers EDR blind spots
- Inventory management between Flow and EDR is very challenging
  - Network traffic is often *confounded* by middleboxes such as NATs and VPN concentrators

What We Need

- Deep network visibility that can correlate with EDR
- Non-redundant data collection with a reduced footprint
- Prioritization and partitioning mechanisms

ISI’S WORK ON THIS PROBLEM

Motivation: Merge Testbeds

- ISI's NCD provides network testbed services for security and networking research.
  - Historically, ISI ran the DETER testbed for DHS
  - Merge is our current technology
- To provide increasingly valid experiments, we are running larger and larger Merge testbeds
  - We have run virtual CDN's with 5000+ individual VM's within them

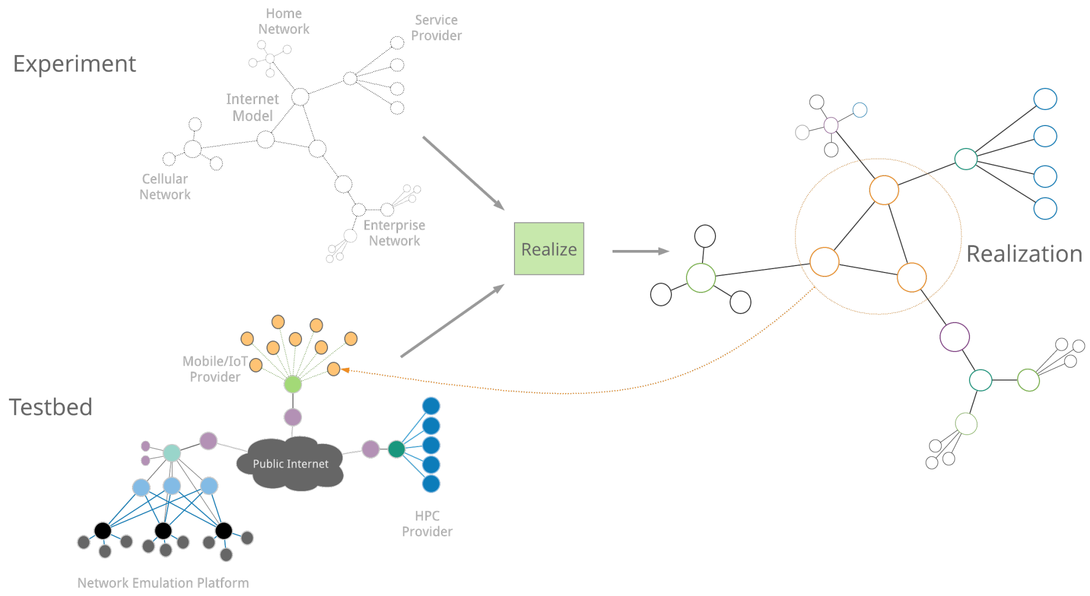

Merge Flow Collector

- The Merge Flow Collector is a network sensor for Merge testbeds based around these ideas
- The sensor can currently output NetFlow and pcaps based on simple filtering rules including protocol, port, source or destination netblocks
- Sensor also includes initial packet sizes, packet flags
- MFC is intended as a software platform to build in different collection and analysis ideas
  - Initial ideas are superflows and sensor placement optimization

SuperFlows

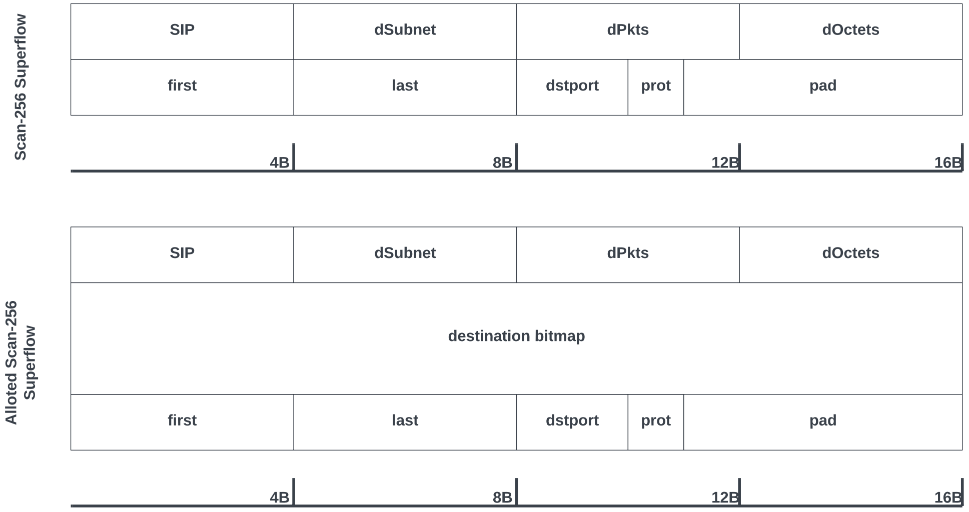

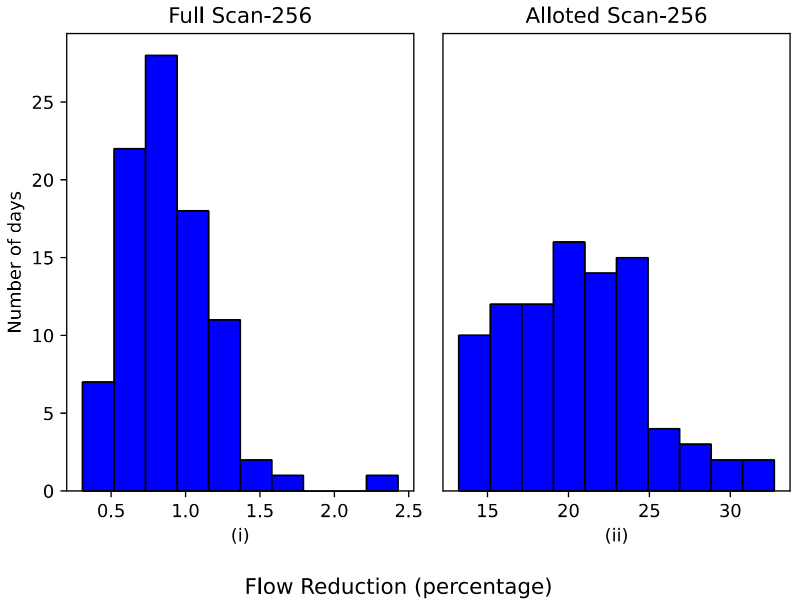

- A SuperFlow is a multiple-address traffic summary which groups together flows which are part of a common phenomenon. Examples we are exploring include web page fetches and scanning.
- The figures above show the footprint reduction for a “rigid” SuperFlow for scanning, and a SuperFlow with a more flexible hypothesis
- SuperFlows are about *data representation* not *detection;* our hypothesis is that multiple overlapping SuperFlows will still have a smaller footprint than the original flows with acceptable data losses

Superflow Test Cases

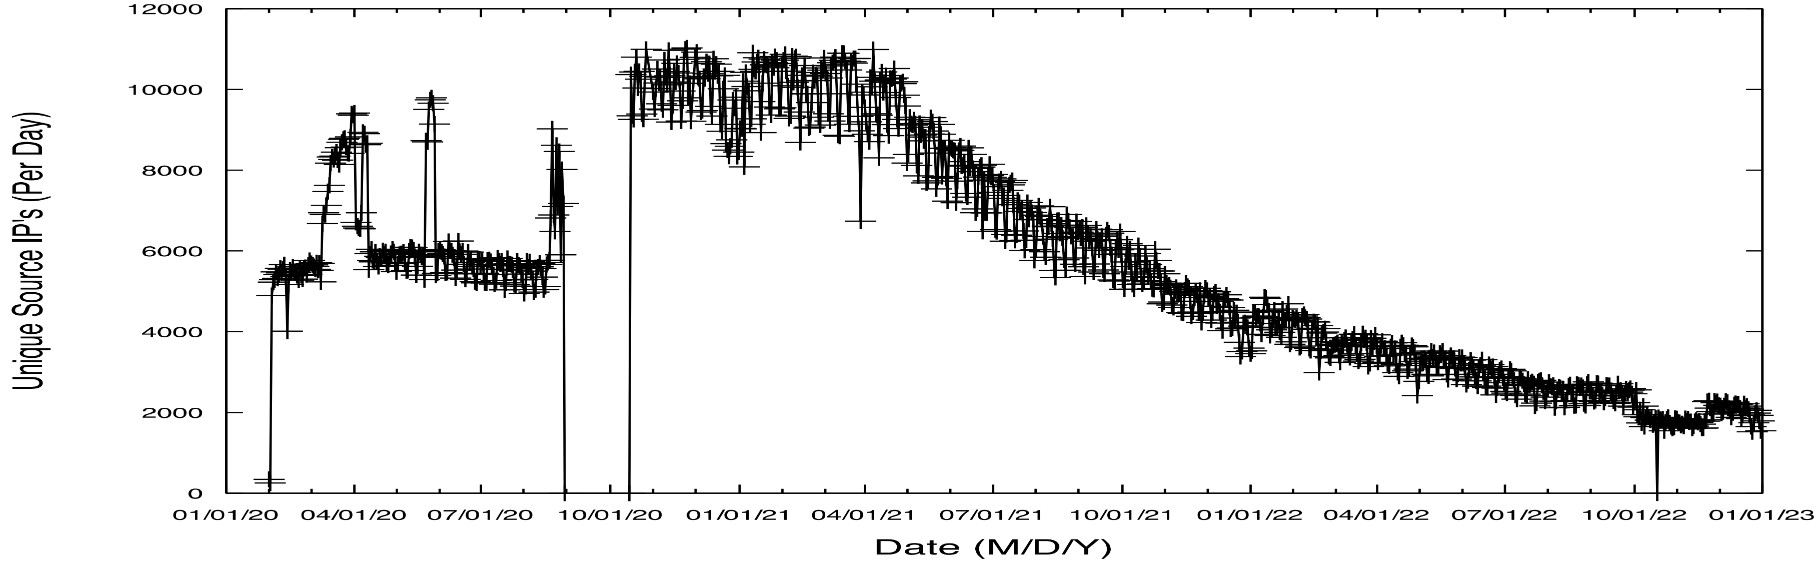

- Basic test cases: scanning, web browsing, DoS
- Try to figure out how to create superflows describing more exotic test cases, like the weird UDP phenomenon plotted here
  - Arbitrary high port changes randomly at 0000Z daily

Conclusions

- Forensic flow analysis is distinct from measurement flow analysis, with its own problems
  - These problems are fundamentally driven by a need for unsampled flow
- Recognizing this need, we need to develop new techniques for addressing forensic flows, which we should recognize as distinct from the measurement community's needs
- I have discussed several preliminary efforts, and anyone who wants to discuss further is welcome to hunt me down for discussion