Bro for Real-Time Large-Scale Understanding

Seth Hall
International Computer Science Institute

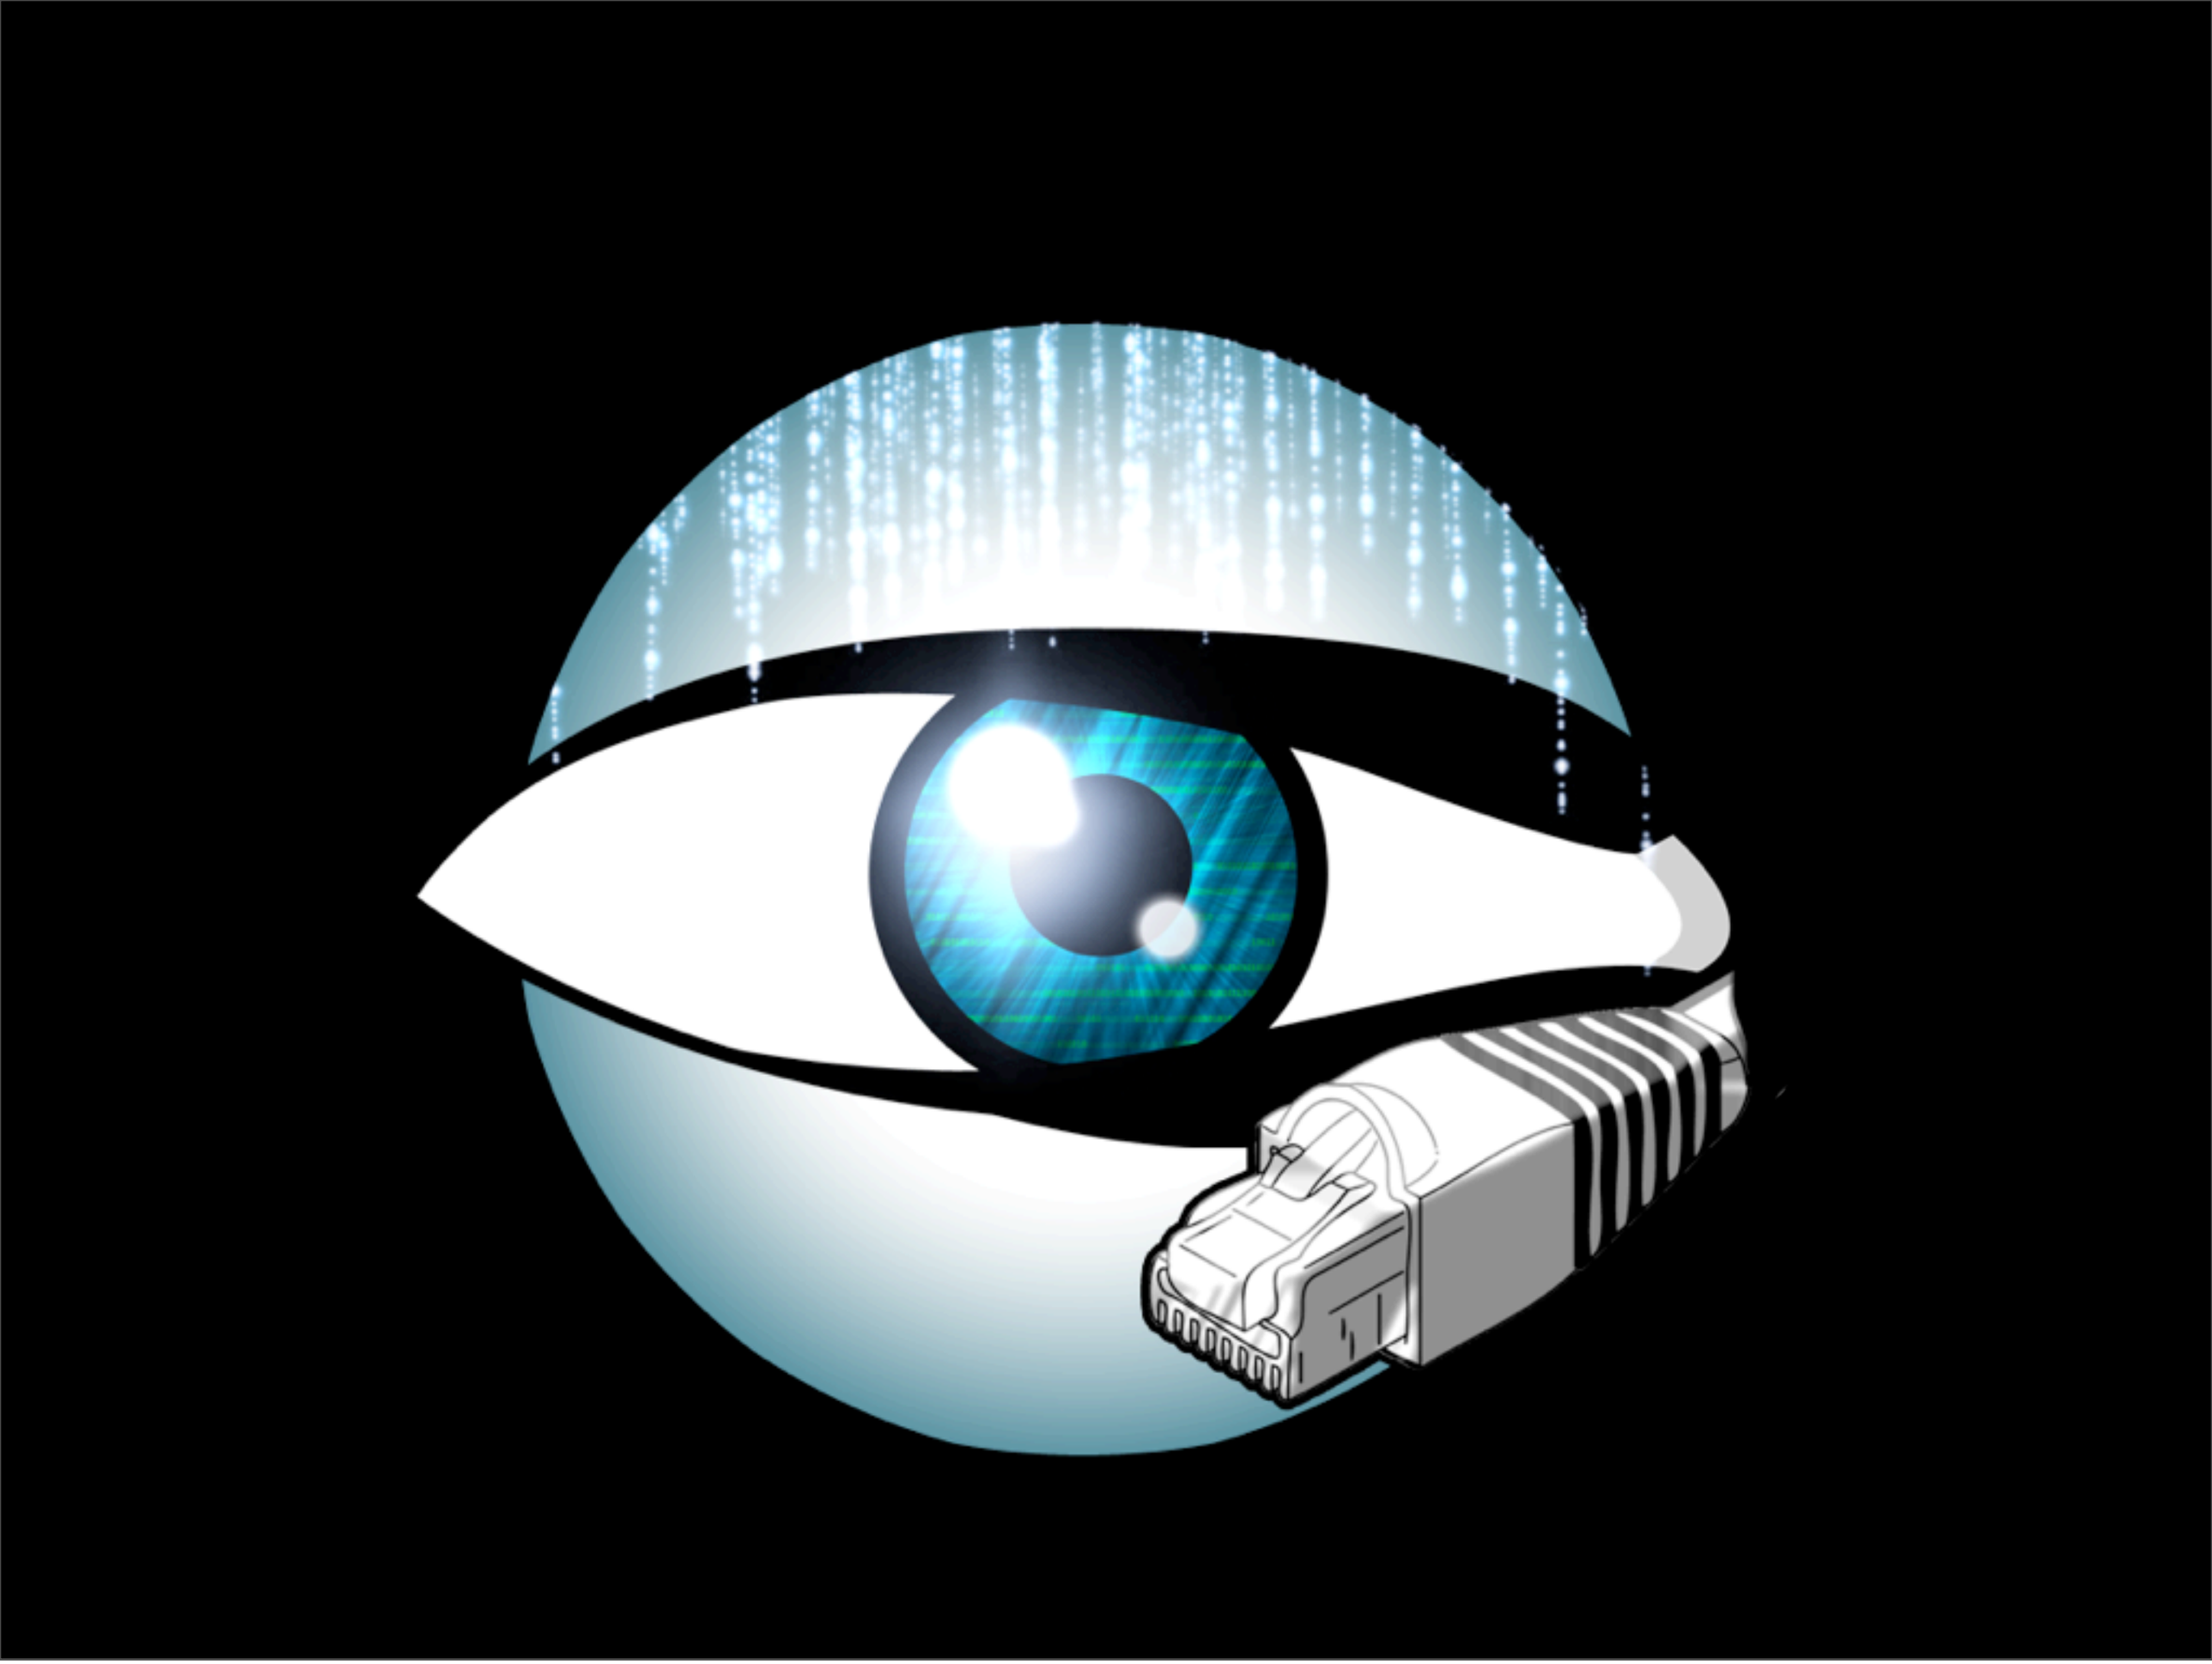

About Me

- Incident response at The Ohio State University for 7 years. 2 extra as a student.
- Spent quite a few years sitting and running flow-tools searches.
- OSU -> GE
- GE -> ICSI (International Computer Science Institue)

No NetFlow?

- NetFlow analysis served well at OSU for many years.
- Year over year degradation of detection capability with NetFlow.
- IRC Botnets were the first decline.

Bro

- Real time event analysis language and platform with protocol analysis.
- $3 million NSF grant for engineering work.
- Strong focus on usability and capability while fixing many, many bugs.
- 84th most popular programming language on GitHub.

Metrics Framework

- Return to measurement roots with new abstractions.
- Programmatic interface for measurement.

Metrics Framework Motivations

- Load balancing made previous techniques all fail. Metrics framework hides cluster abstraction.
- Better and more repeatable interface and approach for measurement and thresholding.
- Give more people the ability to write real world deployable measurement scripts.

Metrics Approach

- Discrete time slices (still investigating sliding window).
- Only streaming algorithms allowed.
- Every measurement must be merge-able for cluster support.
- Probabilistic data structures coming.

Why do any of this?

Measurement is fun!

Lavender Martini Border Gateway Protocol

Facts about Justine's TCP Sessions

I've been doing some self-experimentation this week, running tcpdump persistently and logging the results to a (huge) directory of pcap files. Here's a taste of 24 hours of the data...

Length of time in seconds:

MAX: 62182.238625 <— ssh
MIN: 1.7e-05 <— probably a bug?
MED: 2.046867
AVG: 60.9999384359015
STD. DEV: 736.853599230884

Same numbers for just HTTP/Port 80 traffic:

MAX: 5607.822474 <— seriously?
MIN: 0.003673
MED: 1.12272
AVG: 57.0826988090778
STD. DEV: 184.986012558513

Replicate those results!

```
event bro_init()
    {
    Metrics::add_filter("conns.originated",
                        [$every=1hr,
                         $measure=set(Metrics::AVG, Metrics::MAX, Metrics::MIN,
                                      Metrics::VARIANCE, Metrics::STD_DEV),
                         $period_finished(ts: time, metric_name: string, filter_name: string, data: Metrics::MetricTable) =
                            {
                            for ( index in data )
                                {
                                local val = data[index];
                                print fmt("Connection duration stats for %s", index$host);
                                print fmt("        Period from %s to %s", strftime("%F-%R", val$begin), strftime("%F-%R", val$end));
                                print fmt("        Number of conns: %d", val$num);
                                print fmt("        Max: %.2f", val$max);
                                print fmt("        Min: %.2f", val$min);
                                print fmt("        Std dev: %.2f", val$std_dev);
                                print fmt("        Average %.2f", val$avg);
                                }
                            }
                         ]);
    }

event connection_state_remove(c: connection)
    {
    if ( c$id$orig_h == 192.168.1.105 )
        Metrics::add_data("conns.originated", [$host=c$id$orig_h], [$dbl=interval_to_double(c$duration)]);
    }
```

Connection duration stats for 192.168.1.105
Period from 2009-11-18-13:34 to 2009-11-18-14:26
Number of conns: 77
Max: 3188.27
Min: 0.00
Std dev: 362.48
Average 54.79
Connection duration stats for 192.168.1.105
Period from 2009-11-18-14:34 to 2009-11-18-14:47
Number of conns: 19
Max: 1183.11
Min: 0.00
Std dev: 263.29
Average 67.83
Connection duration stats for 192.168.1.105
Period from 2009-11-18-19:00 to 2009-11-18-19:20
Number of conns: 81
Max: 30.11
Min: 0.00
Std dev: 7.08
Average 3.58

Other Detections using the Metrics Framework

200.29.31.26 had 349 failed logins on 2 FTP servers in 14m47s

92.253.122.14 scanned at least 29 unique hosts on port 445/tcp in 1m4s

88.124.212.10 scanned at least 41 unique hosts on port 445/tcp in 1m13s

212.55.8.177 scanned at least 75 unique hosts on port 5900/tcp in 0m36s

200.30.130.101 scanned at least 66 unique hosts on port 445/tcp in 1m20s

107.22.92.186 scanned at least 64 unique hosts on port 443/tcp in 0m1s

5.254.140.123 scanned at least 29 unique hosts on port 102/tcp in 4m1s

122.211.164.196 scanned 15 unique ports of host 75.89.37.60 in 0m5s

Coming in Bro 2.2

Thanks!

- seth@icir.org
- @remor on Twitter
- http://www.bro-ids.org
- info@bro-ids.org
- @Bro_IDS on Twitter