LOGSTASH: BFD*

- Security Weekly: December 4, 2014
- Phil Hagen
- phil@redcanary.co / phil@lewestech.com
- @PhilHagen / +PhilHagen

**Big Forensic Data*

ALL ABOUT PHIL

- SANS Certified Instructor and Course lead, FOR572: Advanced Network Forensics & Analysis
- Red Canary Managed Threat Detection Service
- Forensic/infosec consultant
- Former DoD/IC/LE contractor, USAF Comm Officer
- USAFA Computer Science
- Likes running and craft beer

PHIL'S TAKE ON THE IR/FORENSICS "HUNT"

- Two primary battlefields: **Endpoint** & **Network**

- Endpoint generates massive data stores

  - Process execution tree, DNS resolver, module loads, parent/child relationships, socket creation, file/registry modifications

- Network generates massive data stores

  - NetFlow, infrastructure logs, full packet capture

ENDPOINT VISIBILITY

- Full disclosure: Red Canary uses Bit9 + Carbon Black to provide fast behavioral detection of realized threats based on endpoint observations
  - 100k/endpoint/hour not uncommon

- Our service freaking rocks, but is not this presentation's focus
- Automatically trawling hundreds of millions of records in a timely manner is **REALLY HARD** (especially to do well)

NETWORK VISIBILITY

- Most scalable solutions still target the NetOps market, not the forensic and IR segments
  - NetFlow, SIEM, asset mapping, availability monitors
- Security-focused tools are INSANELY EXPENSIVE and often incur the "Cable TV effect"
  - If you watch all 700ch of reality TV, we need to talk

DEATH BY DISTRIBUTION

- Network devices tend to have volatile storage
  - Reboot, overflow, corruption? Lost logs
- Distributed log storage = distributed analysis
  - Best case: Collect from far and wide
  - Inefficient use of analysts' time
- Multiple log formats require multiple tools

DEATH BY DISTRIBUTED LOGS

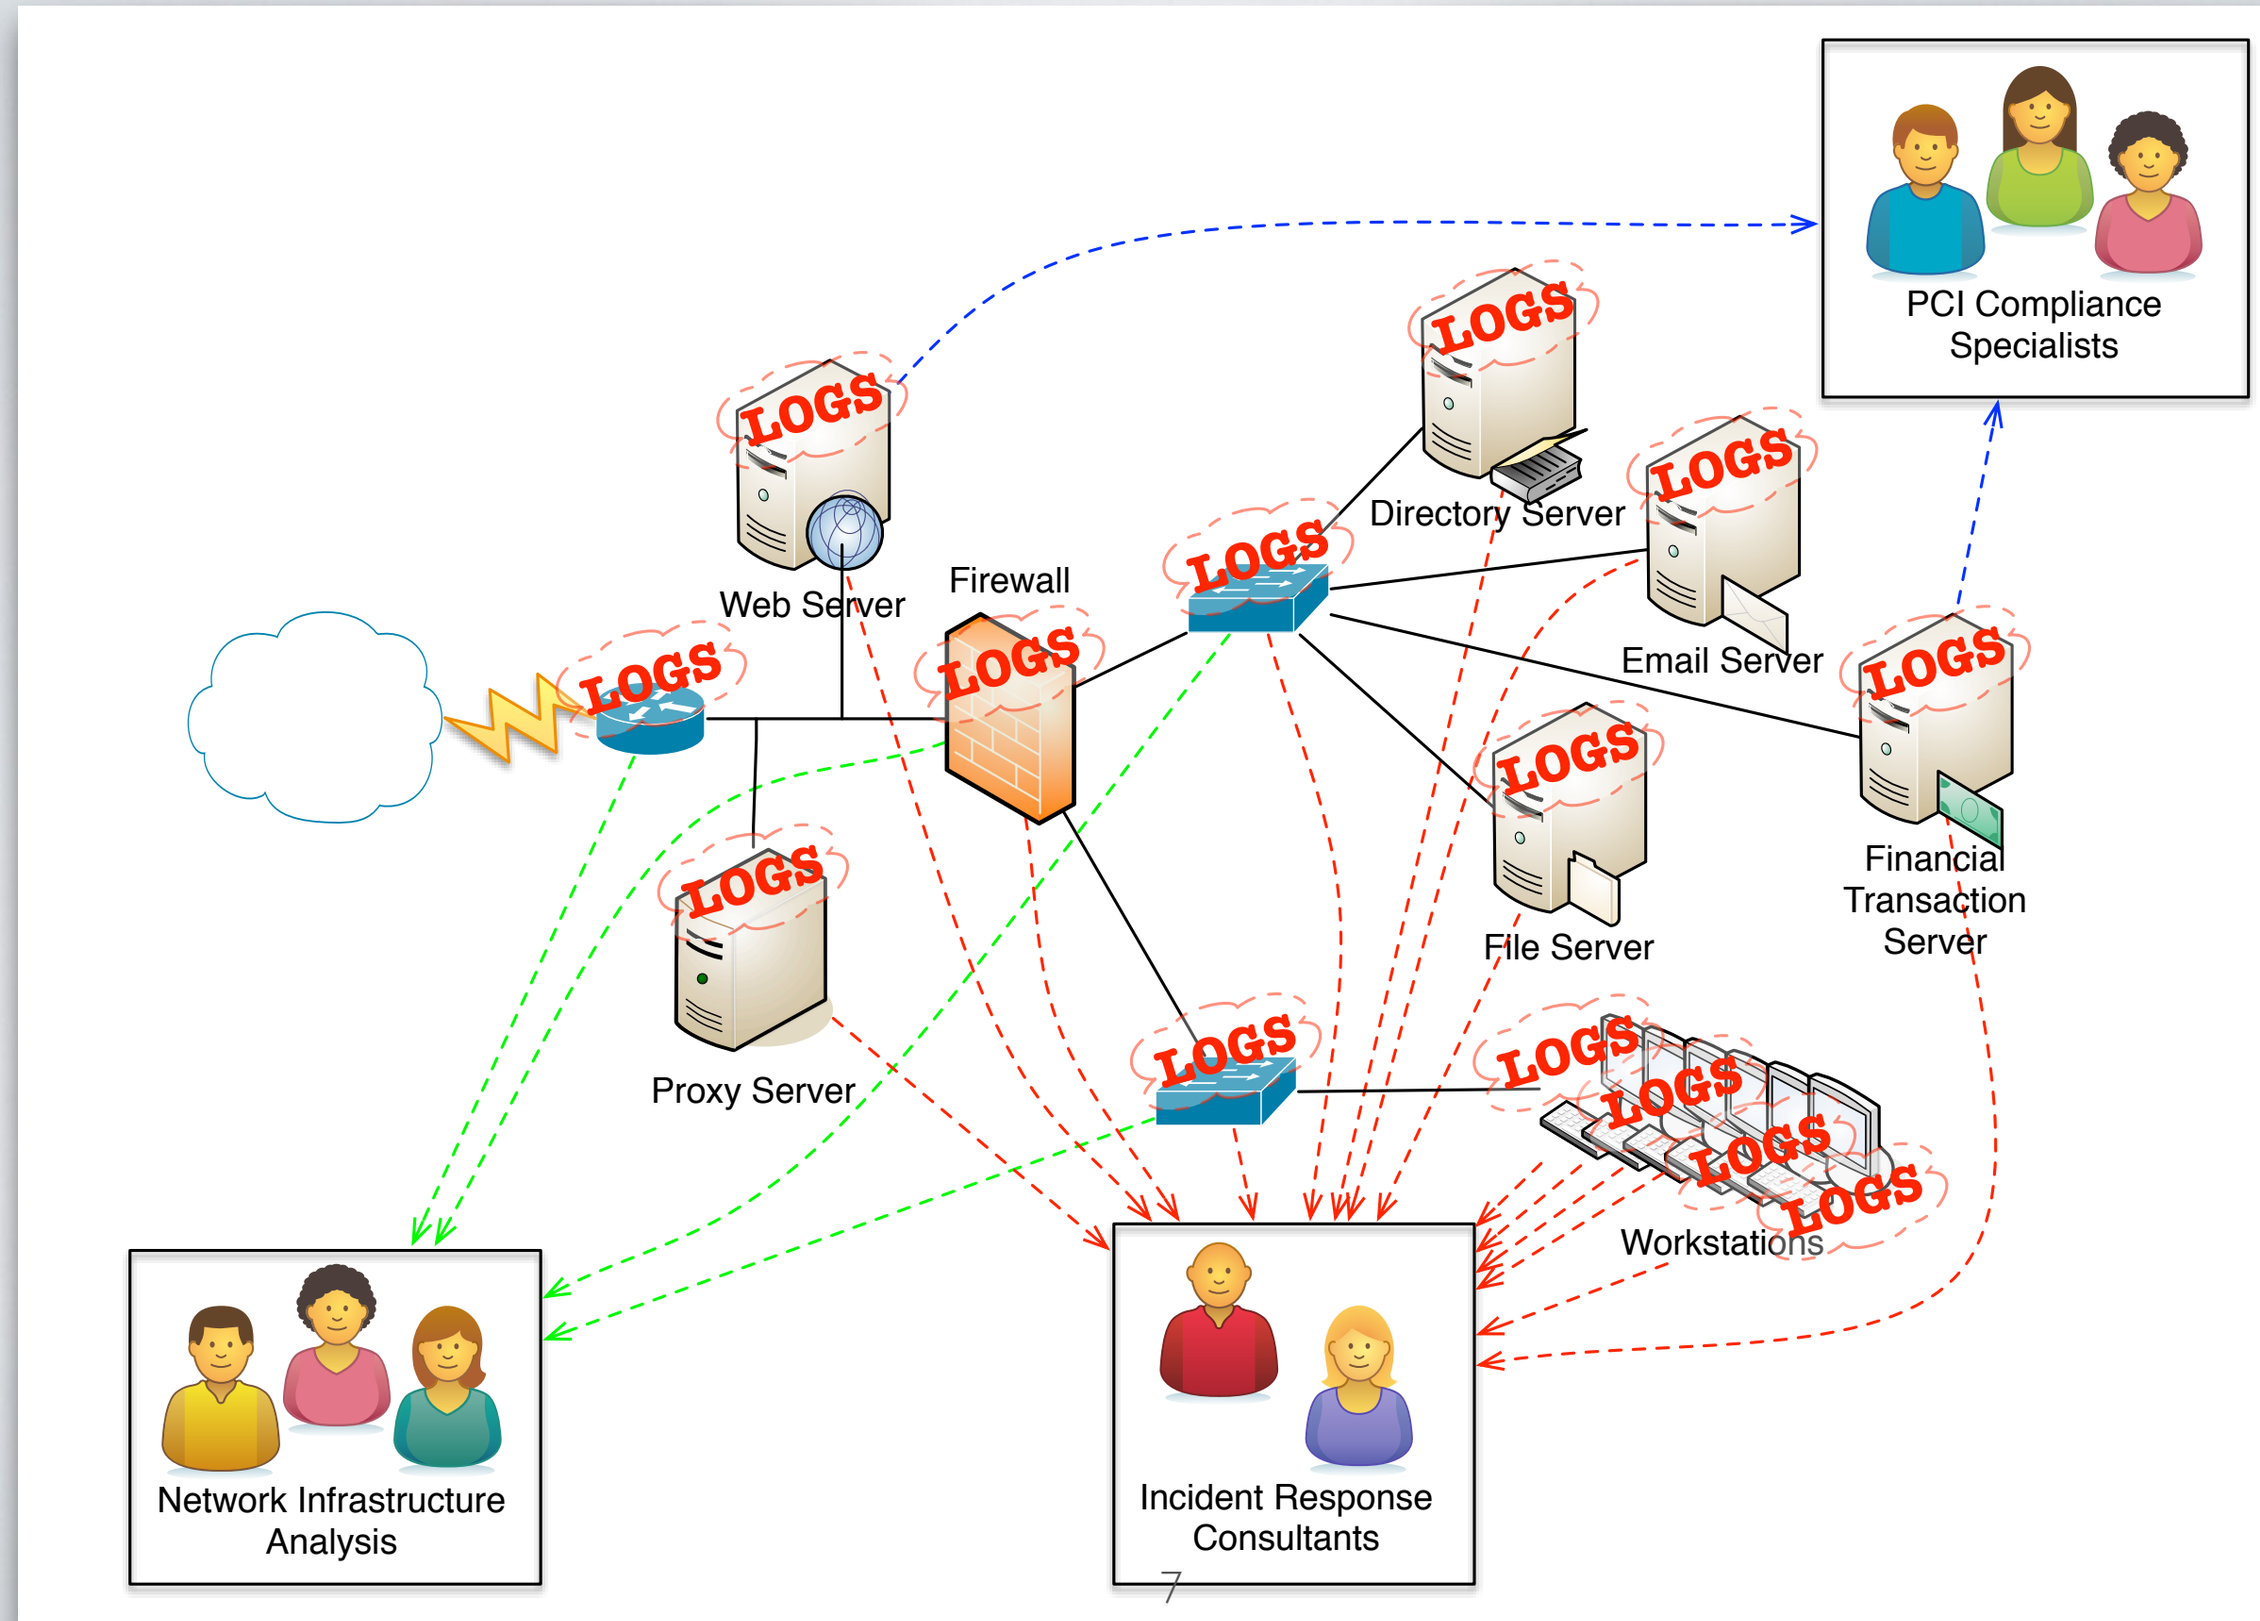

PURSUE THE 80% SOLUTION

- …often at less than half the cost
- A solid, scalable network evidence aggregation solution will save your team **time** and your organization **money**
- Seek Artifacts of Communication in network investigations to hunt, scope, and enrich

PUT THAT ANOTHER WAY

ONLY PAY FOR THIS

SEEK VALUE OF THIS

AGGREGATE TO FORESNICATE

- Built-ins: rsyslog, syslog-ng, MS Eventing
- SIEMs: Arc Sight, Trustwave, Tenable
- Pure aggregators: ELSA, Splunk,

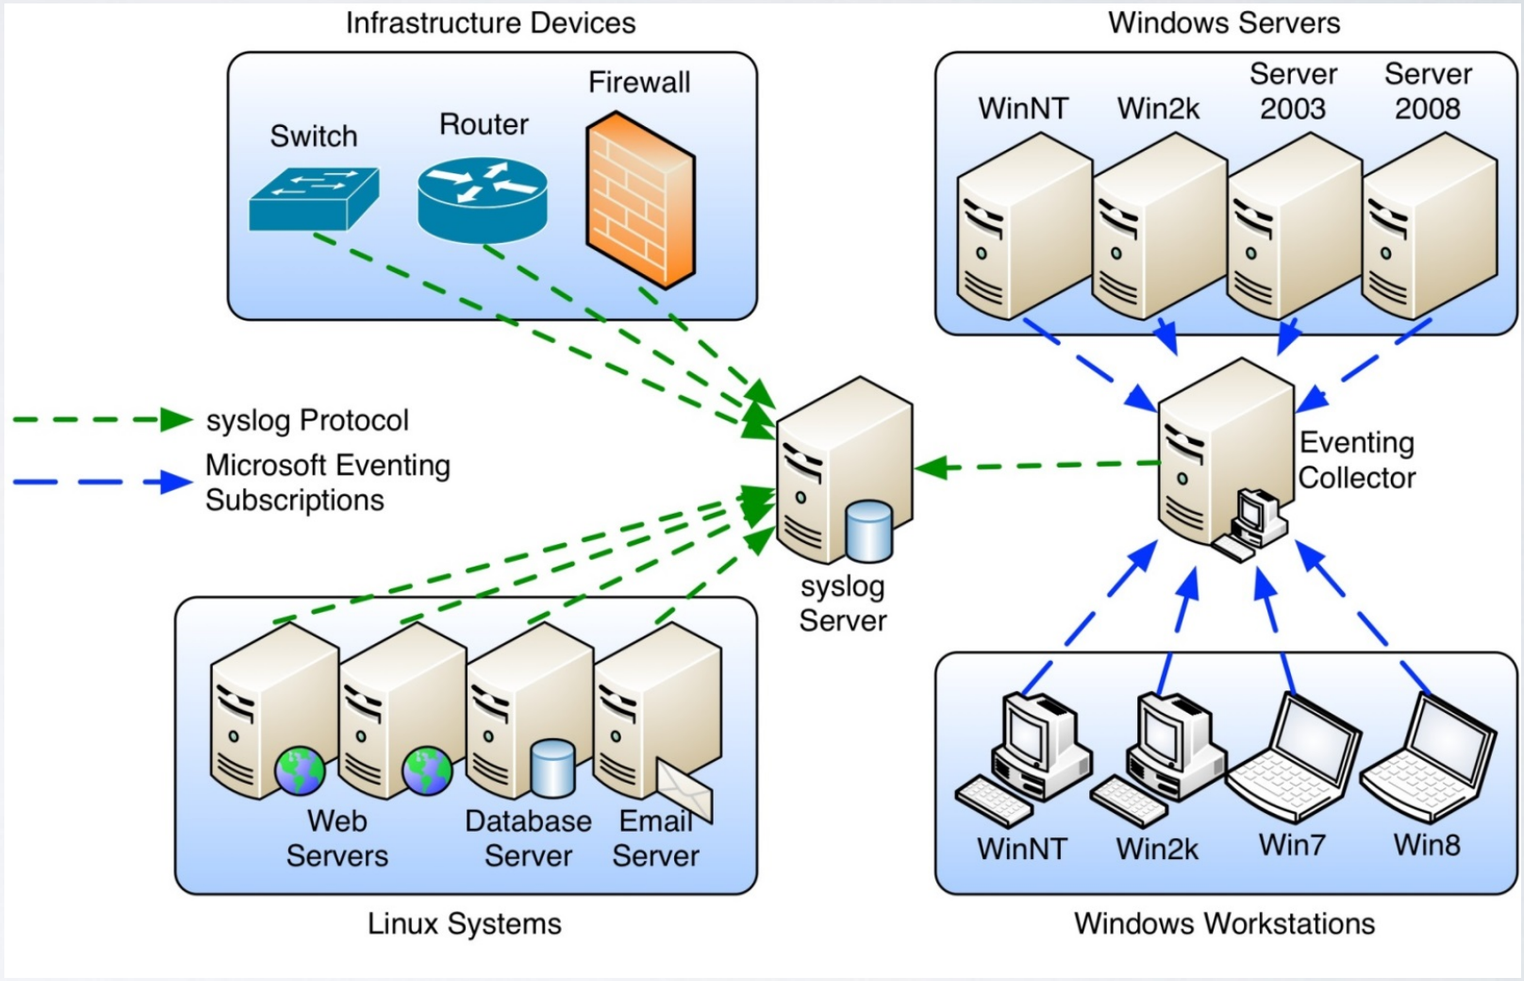

Logstash!

DISCLAIMER

- I'm not a paid Elasticsearch guy, just a huge fan
- All comments are mine, not approved by or specifically endorsed by anyone else
- Do not taunt happy fun ball
- FWIW, YMMV, ASAP, etc.

THE VITALS

- Free, open-source:
  `http://logstash.net`

- Developed at Dreamhost
  - Made to scale **huge**

- Now part of Elasticsearch

- Kibana web frontend

- Great developer and community support

LOGSTASH AT A GLANCE

- Extensive "grok" pattern-matching syntax
  - Test at `http://grokdebug.herokuapp.com`
- Can operate on standalone system or in larger hierarchy
- LOTS of inputs: Files, syslog via network socket, NetFlow, Twitter, IMAP, SNMP, named pipes, more
- Filter inputs to match relevant fields
- Many outputs: E-mail, pagers, file, search database, HTTP JSON

FOR572 LOGSTASH VM

- Created for **SANS FOR572**, Advanced Network Forensics and Analysis: `http://for572.com/course`
- Pre-configured with many data parsing formats, three Kibana dashboards
  - Linux syslog, Apache access logs, live/archives NetFlow, Cisco ASA, iptables, passivedns, dhcpd, bind DNS server
  - Ready to ingest data and provide interactive visualizations!
- **Free VM appliance for the community!**
  - `http://for572.com/logstash-readme`

LOAD REAL-WORLD DATA

- Logs from:
  - Linux syslog
  - Live streaming NetFlow
  - Apache HTTPD access
- LIVE DEMO TIME!

WHAT COULD POSSIBLY GO WRONG?!

DEMO TIME!

This is where we hope Phil's demo goes according to plan, and Murphy's Law doesn't take over

LATEST NEW HOTNESS

- Parse existing NetFlow to text and ingest alongside live
- Log2timeline/Plaso CSV output into Logstash
- **Forensicator FATE (Barry Anderson's SANS Gold Paper)**
  `http://for572.com/uag8v`
- **log2timeline in Logstash**
  `http://for572.com/x2wen`

SUMMARY

- Endpoint and Network are vital for hunting, scoping, and investigation, but need proven solutions and good tech to handle at scale

- Logstash is a great tool in the forensicator's arsenal - especially for network evidence

- README: `http://for572.com/logstash-readme`

Phil Hagen

phil@lewestech.com / phil@redcanary.co
@PhilHagen / +PhilHagen

Slides: http://for572.com/lsbfd-slides