**Aplura, LLC**
5653 Blithaire Garth
Columbia, MD, 21045
301-523-2110 (w)
410-864-8386 (f)

**Focused Information Security**
http://www.aplura.com


# Log Centralization for Security

# CALUG – March 11, 2009



# Agenda

- **Introduction**
- Log Centralization for Security
- Trends (Predicted and Current)
- Enterprise Log Collection
- Splunk
- Examples
- Get Splunkin'

# Introduction

- Log Collection for Security
- Last CALUG logging presentation
  - Before You SIM – October 2007
- Why am I talking about this again?
- What do I know about log centralization?

# Disclaimer

- Last talk focused on design considerations
- This talk focuses on a solution
- I am very pro FOSS
- Log solution is VERY Linux and CLI Friendly
- Log solution is Commercial Software
  - I have come to grips with this, you might too

# Minimap

- *Introduction*
- **Log Centralization for Security**
- Trends (Predicted and Current)
- Enterprise Log Collection
- Splunk
- Examples
- Get Splunkin'

# Log Centralization for Security

- CentralSecLogs: Why
- CentralSecLogs: What
- CentralSecLogs: How
- CentralSecLogs: Example

# CentralSecLogs: Why

- Why is often dictated by policy, regulatory compliance requirement, business need
  - 'Reg compliance' != 'security value'
- Beyond requirement, logs are essential for security analysis, trending, and reporting
- Search across all data types for investigation
- Not necessarily an extension of admin logging
- See “Why” Examples

# CentralSecLogs: Why Ex1. Before

- Situation without Centralized Security Logs:
  - PC-based IPS warns of attack from PC Alice01
    - Phone call to user shows she is oblivious to attack - 3m
    - Validate Alice01 patch levels through central utility - 5m
    - Query central AV server for unusual activity - 5m
    - Login to firewall and pull logs related to Alice01 - 5m
    - Login to web-proxy and pull logs related to Alice01 - 5m
    - Review IDS alerts for Alice01 - 15m
    - Call Network Admin to pull flow-data for Alice01 – 20m
  - This investigation could take nearly an hour

# CentralSecLogs: Why Ex1. During

- All the above steps require individual logins and client (thick, WebUI, Command line)
- All of the steps require manual data parsing to try and find evidence to support the analysis
- These steps can take lots of time

# CentralSecLogs: Why Ex1. After

- With a Central Security Log system
  - PC-based IPS warns of attack from PC Alice01
    - Phone call to user shows she is oblivious to attack - 3m
    - Query Central Log Server for Alice01 (name/IP) – 2m
  - 5 minutes instead of 30
  - What if the investigation found 20 infected hosts?
  - Log Centralization greatly reduces reaction time

# CentralSecLogs: What

- System: Servers, Desktops, Network Devices
- Application: Auth, action, failures
- Activity: Network, Proxy, IDS, Flow
- Ad-hoc: script output, client, debug
- Application **data**
- “What” - Frequently limited by Log vendor
- Note: Might be further defined by regulation

# CentralSecLogs: How

- Precise collection and reporting
- Flexible reporting and queries
- Flexible alerting capabilities
- Distributed collection and queries
- Use agents whenever possible *
- Due-diligence during planning = $ Savings!

# CentralSecLogs: Agents *

- Agent: Application on the data-generator that collects the targeted data and forwards it
- Opposed, then on the fence, now a believer
- Value:
  - Real-time reporting
  - Confidence in data-set: queuing, controlled hand-off
  - Retrieve logs from off-network hosts
  - Dig deep - (registry, file integrity, client app logs)

# CentralSecLogs: Ex2. Situation

- Security admin receives e-mail from Log server of a correlated event
  - System Bob01's IP browsed to a “watched domain”
  - Bob01's AV reported a “Failed to Clean” warning
  - Bob01 browsed to over 10 distinct URLs in < 1min
- Because these logs are centralized, reusable rules can be written to leverage this data

# Minimap

- *Introduction*
- *Log Centralization for Security*
- **Trends (Predicted and Current)**
- Enterprise Log Collection
- Splunk
- Examples
- Get Splunkin'

# Projected Trends

- Predicted in October 2007 (at CALUG)
  - "SIM Type 1 & 2 Space Steadily Growing"
  - "SIM Type 3 Space Shrinking"
  - "Well-developed Security Teams Ditching SIMs"
  - "Analysts Turning Toward Extrusion Detection"
  - "Flow Analysis an Increasing Player"
  - "Regulatory/Auditors Require Central Log Mgmt"

# Current Trends

- Regulatory compliance requirements tighten
- End-user attacks increase exponentially
- Attackers pillage weak inter-Node security
- Increase in network/system investigations
- A lot of work in flow collection and analysis
- Achilles heal is STILL the user-node

# Minimap

- *Introduction*
- *Log Centralization for Security*
- *Trends (Predicted and Current)*
- **Enterprise Log Collection**
- Splunk
- Examples
- Get Splunkin'

# Enterprise Log Collection

- Current State
- Demands
- Challenges

# Enterprise: Current State

- Most don't collect node logs for security or at all
- Collected-logs often aren't analyzed
- Application logs frequently forgotten/discarded
- Few trend or search **DATA**

# Enterprise: Demands

- Regulatory Compliance Reporting
- Requirements
  - Fast/Efficient/Scalable
  - User-level Controls
  - Distributed collectors
  - Diverse data types (I didn't say log-types)
  - Flexible reporting
- Log reduction frequently not appropriate

# Enterprise: Challenges

- Collect logs from internal resources
- Admins with varied responsibility and access
- Many data types
- ROI demands

# Minimap

- *Introduction*
- *Log Centralization for Security*
- *Trends (Predicted and Current)*
- *Enterprise Log Collection*
- **Splunk**
- Examples
- Get Splunkin'

# Full Disclosure

- Aplura staff do not work for Splunk directly
- We do not get reimbursed for promoting Splunk
  - This is my first time to actually do this
- Splunk has hired Aplura for PS engagements
- Goal is to demo features valuable to analysts
- This is not designed to be a Splunk ad: however,
- It might be my favorite COTS product...ever

# Splunk

- Splunk is IT Search. (period)
- Not your grandmother's log manager
  - Yet, simple enough for her to use
- Not a SIM (Self-proclaimed)
- No Custom parsers or connectors
- Federated Search
- Flexible reporting
- Many uses outside of security

# Splunk Works

- Supports: *NIX, OSX, FreeBSD, Win
- Same code regardless of size or intended use
- One set of binaries with adjustable roles
- Roles: Splunk Index, forwarder, Web UI, etc.
- Deployed in seconds useful in minutes
  - SIM admins probably think this a joke
- Reads: Input files, tcp/udp socket, etc.
- Licensed based on Indexed-data/day
- Free license for anyone up to 500mb/day

# Splunk is *NIX Friendly

- Three config paths (WebUI, Conf Files, CLI)
- The WebUI/CLI takes the “|” (pipe) command
- Familiar Terms: (dedup, sort, top, rename)
- Built in man-like help
- Send script output to Splunk as Search input
- Send Splunk search output to scripts as input
- Community: http://www.splunkbase.com

# Splunk is Windows Friendly Too

- Accommodates 32/64bit desktop and server
- Inputs: Registry, WMI, Event Logs
- Msi Install
- Integrates with AD
- Works cleanly with other Splunk systems
- Splunk on Windows is fully-featured

# Splunk Downsides

- Splunk has a few downsides
  - User/Role Administration via UI needs improvement
    - Can be done with CLI/Confs but not UI as well
    - Some have leveraged Splunk's API to write their own
  - Enhanced visibility into current processing
    - This is just an App waiting to happen

# Minimap

- *Introduction*
- *Log Centralization for Security*
- *Trends (Predicted and Current)*
- *Enterprise Log Collection*
- *Splunk*
- **Examples**
- Get Splunkin'

# Examples

- In the last 6 months:
  - Log collection for 5 different US agencies
  - Log collection for commercial entities
  - Used many forwarder techniques: syslog, syslog-ng, copied rolled log files, Splunk
  - Used many agents: Snare, DAD, LogLogic's Lasso, Windows Log Shipping, Splunk

# Example: CYA

- Moderate-sized US Federal Entity
- Vastly mis-understood their sec-data volume
- Day after FWs added, tripled their license
- Exec response “turn off FW feed.”
  - I convinced him to leave it on temporarily
- The next morning Splunk highlighted a 2-hour gap in Enterprise FW data.

# Example: Many Clients

- Descent-sized US Federal Entity
- 6000+ Windows hosts, Linux, Solaris, Router, IDS, Urlsnarf, dns-query
- Only a single server was dedicated to Splunk
  - We encouraged customer to use more resources
  - Deployed the one server on modest hardware
  - Splunk performed fairly well considering load
  - Eventually they took the suggestion to redeploy

# Example: Splunk vs SIM

- World-wide US Federal Agency
- Devices to log to SIM:
  - 140 Enterprise Firewalls
  - 110 NIDS Sensing Interfaces
  - McAfee EPO data
  - ISS Site Secure HIDS data
  - Some syslog data

# Example: vs - SIM1

- Traditional SIM 1
  - 7 Enterprise Systems (HP DL580)
  - HP SAN - 45 15K SCSI Drives in RAID 10 Config
  - $900,000 Initial License Investment
  - ~15 hours/week admin cost to maintain
  - >1 rebuild required/year
    - Cost 120 hours admin time and 40 hours PS
  - Single IP Query across 30 days = 26 - 80min

# Example: vs - SIM2

- Traditional SIM 2
  - 6 Enterprise Systems
  - HP SAN – 30 Drives in RAID 5 Config
  - $400K Initial License Investment
  - ~38 hours/week admin cost
  - >3 Rebuilds in six months (all with PS)
  - Single IP Query across 30 days = Impossible

# Example: vs - Splunk

- Changed Requirements:
  - 140 NIDS and Urlsnarf Sensing Interfaces
  - Syslog from Enterprise Mail Appliances
  - No EPO and NO ISS
- Splunk
  - 2 1-RU Servers
  - $70K Initial License Investment
  - < 5 hours/week admin cost
  - Single IP Query across 30 days = < 1 min

# Example: vs - 1 Year Cost Chart

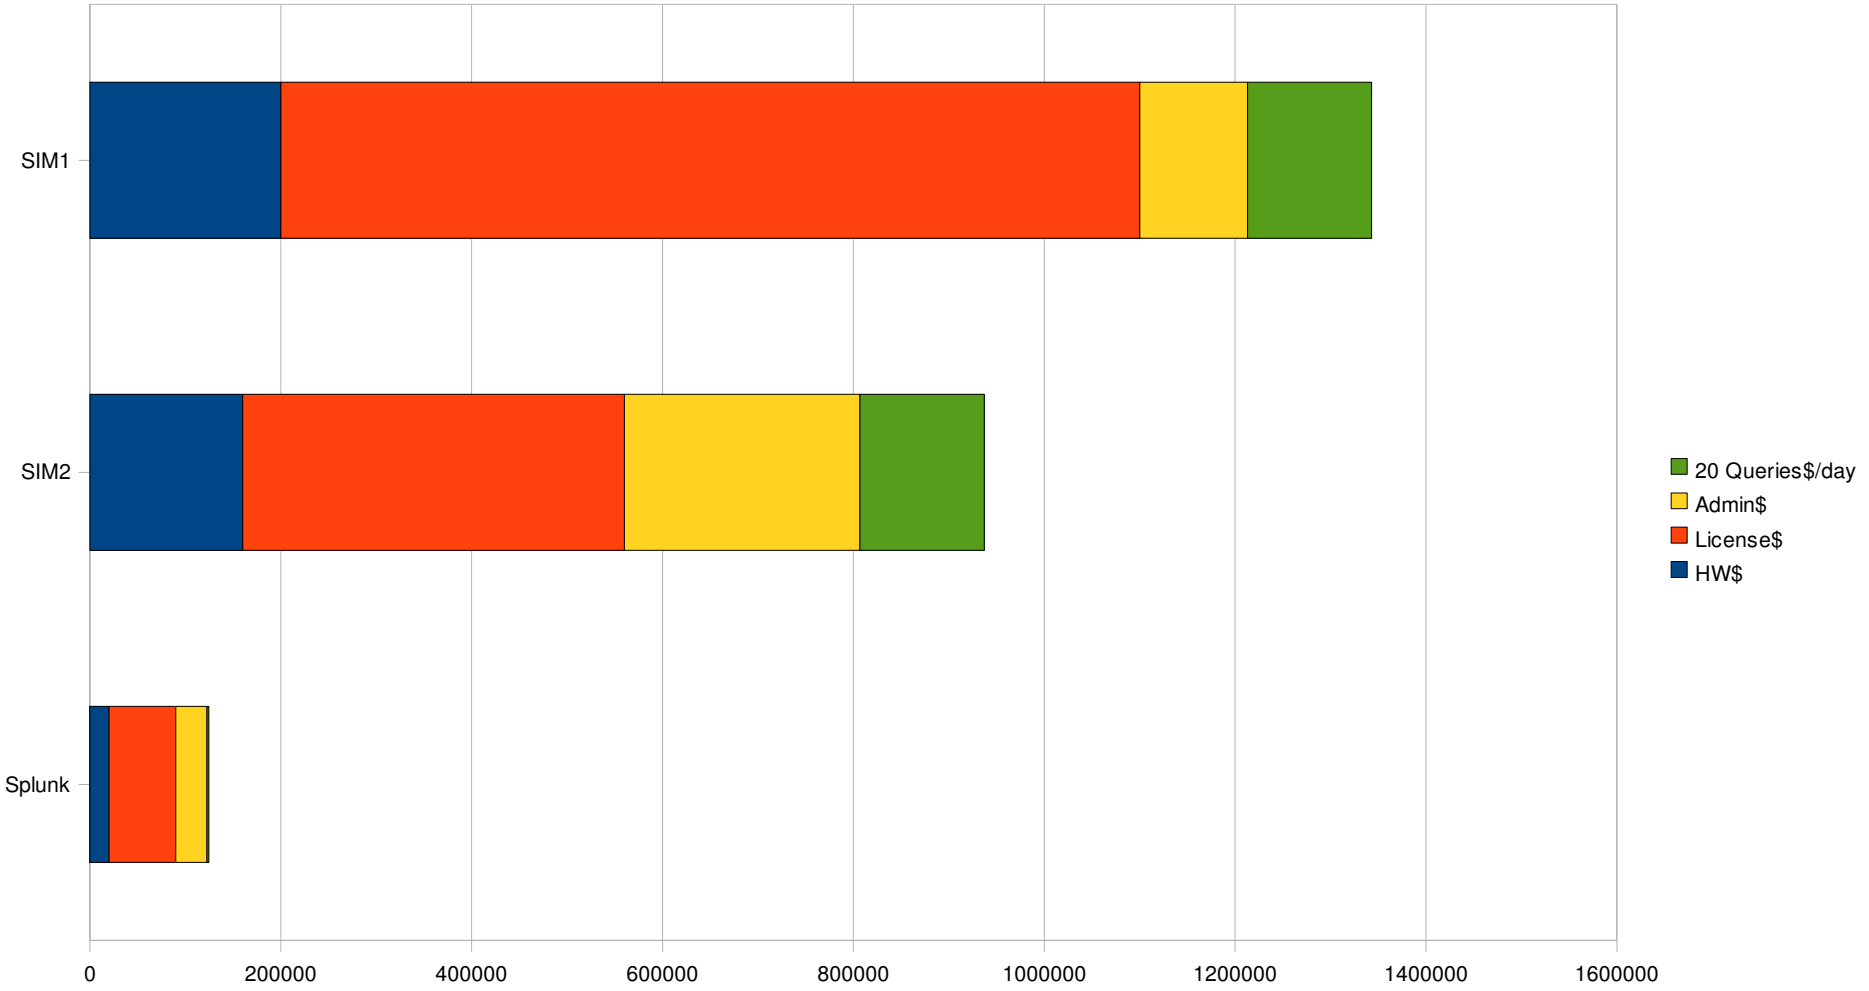

# Minimap

- *Introduction*
- *Log Centralization for Security*
- *Trends (Predicted and Current)*
- *Enterprise Log Collection*
- *Splunk*
- *Examples*
- **Get Splunkin'**

# Get Splunkin'

- Go Splunk, Yourself
- Browse to www.splunk.com
  - Sign up (no hassle)
  - Download splunk for free for your distro
- Browse to www.splunkbase.com
  - Download, tryout and enhance splunk apps
- Create, Explore, Search

Presenter:

Sean Wilkerson

Aplura, LLC

swilkerson@aplura.com