Look Before you SIM

Feb. 17, 2011

- Log Mgmt is My Profession and Passion
- I Have Witnessed Many Companies Jump Head First into Log Mgmt/SIM Before Knowing What They were Getting into
- Hopefully, this presentation will help you to **Look Before you SIM**

My Experience With Log Mgmt

- Over a Dozen Years in Professional IT
- Focus on Security Engineering since 2004
- Many (Projects | US Agencies | Countries)
- Log Expertise Includes
  - Part of the Gartner's SIM eval. Since 2008
  - Spoke at last two SANS Log Mgmt summits
  - Frequent talks on SIM and security for feds
  - Speaks often at cons and community events

Aplura's Recent Log Projects

- Commercial
- NGOs
- State Agencies
- Civilian Federal Agencies
- FFRDCs
- Defense Contractors
- Defense Commands
- Intelligence Community Agencies

Why Do you Centralize?

- “Net Video Tape” - What happened and When
- Regulatory Compliance
- IG (Auditor) Requirement
- Support Your SOC or Daily Security Analysis
- Facilitates Incident Response
- Builds Inter-Group Relationships
- Improve Security Posture

What Do you Centralize?

- System: Servers, Desktops, Network Devices
- Activity: Network, Proxy, IDS, Flow
- Security Application: Auth, Action, Failures
- Ad-hoc: Script Output, Client, Debug
- Application **Data**

- “What” - Frequently Limited by Log Vendor
- Note: Might be Further Defined by Regulation

How Do You Centralize?

- Precise Collection and Reporting
- Flexible Reporting, Queries and Alerts
- Distributed Collection and Queries
- Use Agents Whenever Possible

Due-diligence During Planning = **$** Savings!

SIM, SEM, SIEM... Oh My

- Gartner's Push Toward Uniformity
- SIM – To Support Compliance
  - Lots of data stored un-altered for a long time
- SEM – To Support Analysts
  - Reduced data to show “malicious”
- A SIEM Generally Includes: Event (Collection, Normalization, Aggregation, Classification, Correlation), Real-Time Analysis, State Engine
- SIM Space Is Still Evolving and Not Absolute
- Don't Get Fooled on Terms, Consider Qualities

Data Collection

- Standard SIM Collection Process

Collect->

Normalize->

Classify->

Aggregate->

Store

Normalize

- Process to “Map” Incoming Event Data into a Fixed Structure
  - src_ip, dest_ip, src_port, dest_port...
- Pros:
  - Data is Consistent
- Cons:
  - Expensive
  - Generalized
  - Done Once (a format change is detrimental)

Example: Normalize

- A Juniper NetScreen FW Log Has 56 Fields
- A SIM Schema Might Have 38 Fields
- During Noramlization, Processes Will Run to Sift Through The 56 Fields and Fit Many Of Them Into the Ones in the Schema
- The Fields That Didn't Fit Are Usually Trashed
- The Ones That Did Fit Might Change Context

Classify

- Apply Pre-Developed Taxonomy
- Pros:
  - Events Can be Filtered/Alerted by Type
- Cons:
  - Generalization
  - Outliers are grouped or left behind
  - Your site may make different decisions

Example: Classify

- Your Network May Have/Use:
  - 2000 IDS Signatures
  - 6 Firewall “Actions”
  - 20 Anti-Virus “Actions”
- Classification Is Taking All of these Outputs and Group Them Into Vendor Presets

Aggregate

- Stores “Counts” of Critical Items
- Pros:
  - Good for trends
- Cons:
  - You may not agree with groupings
  - Done Once (a format change is detrimental)

Store (Data)

- The Data is Saved for Collection or Analysis
- Traditional SIM Vendors Use RDBS
- Many Vendors Use Flatfile DB Today

Project Scope

Scope: Target Logs

- Which Devices
  - All Security-Related Devices
  - All Business-Critical Devices
  - Consider All Publicly-Available Systems
  - Physical/Logical Separation of Networks
  - Regional or Device-Type Collector
  - **Whatever You Can Afford**
- Will Agents Be Required

Scope: Reporting

- Most SIMs Include Many Built-in Reports
- Most SIMs Do Correlation
- Determine If the Built-in Reports Will Suffice Your Needs
  - Many may not
- Determine if Ad-hoc Reporting Is Supported
  - Most products DO NOT support this
  - Some non-SIM products excel here

Scope: Network Map

- Physical/Logical Separation of Networks
- Some SIMs Accommodate Regional Collectors
  - This is is a function of their scalability and if they support distributed collection/reporting
- Group Systems By Type or Region – Example:
  - NY Office
  - HR Desktops
  - Server Network

Scope: Users

- User Planning
  - User/Role Permissions
    - Many SIMs allow granular RBAC
      - This is frequently very confusing
  - Technical Vs Decision Maker
  - Security Implication of Providing Access
- User Client
  - Java and Web Clients Are Big In This Market
  - Will the Client Pass Your Software Requirements?
  - Consider Server Resources (i.e. Proc/Query)

Scope: Hardware

- Prepare for Several/Many Systems
- Most Benefit From High CPU/Memory
- Databases Require High Reliability
- Expect to Leverage a SAN or Storage Array
  - Most Vendors Require RAID1+0 15K Disks

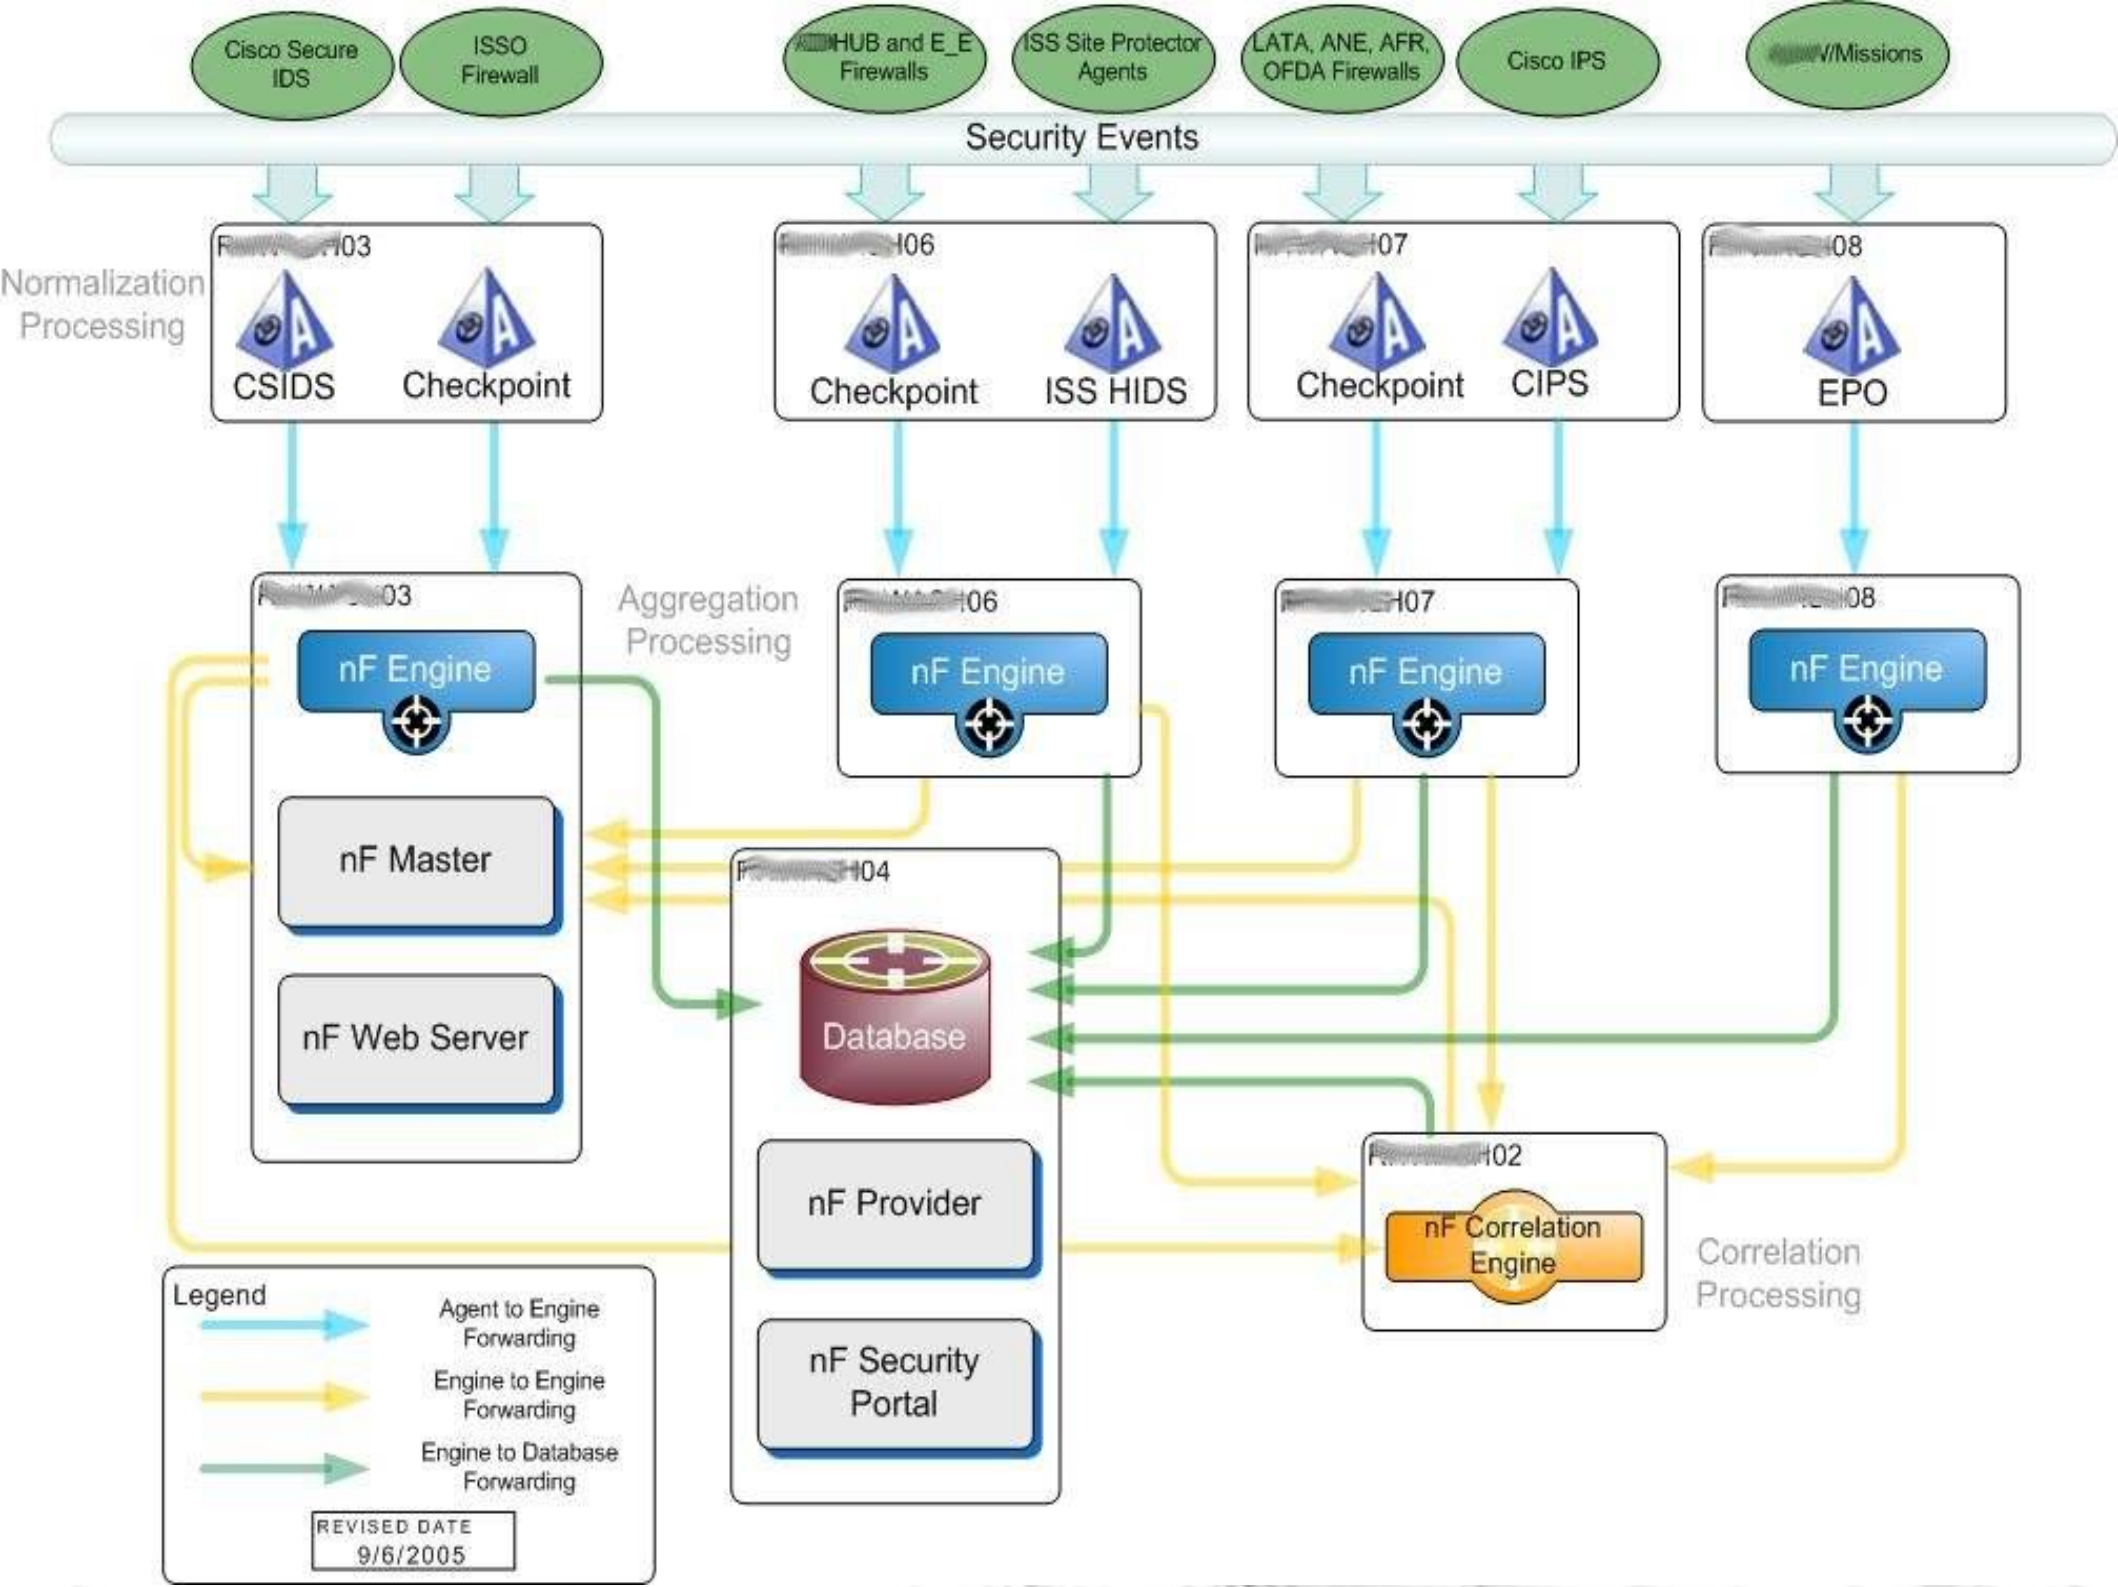
Cisco Secure IDS
ISSO Firewall
HUB and E_E Firewalls
ISS Site Protector Agents
LATA, ANE, AFR, OFDA Firewalls
Cisco IPS
/Missions
Security Events
Normalization Processing
03
CSIDS
Checkpoint
06
Checkpoint
ISS HIDS
07
Checkpoint
CIPS
08
EPO
Aggregation Processing
03
nF Engine
nF Master
nF Web Server
06
nF Engine
07
nF Engine
08
nF Engine
04
Database
nF Provider
nF Security Portal
02
nF Correlation Engine
Correlation Processing
Legend
Agent to Engine Forwarding
Engine to Engine Forwarding
Engine to Database Forwarding
REVISED DATE
9/6/2005

Scope: SIM Management

- Dedicated Admin is Recommended for Most Log Products and for all Traditional SIMs
- For a Traditional SIM, You Might Want to Buy Admin a Headset For The Long Support Calls
- For Most SIM Deployments, Professional Services are a Necessity
- Most SIMs are Unix/Linux-Based

Scope: Trending

- Review Trends From Collected Devices
- Long Term Trends (Years)
- Must Collect Metadata From Stored Logs
  - Few SIMs do this out of the box
- Most SIMs Display Trends Only On Currently-Stored Logs

Scope: Use Case

- Why Are You Deploying Log Management?
- One Product Can't Do All These Well:
  - Trend Reports
  - Regulatory Compliance Requirement Reports
  - To Support a SOC
  - Forensics/Intrusion Analysis
  - Administrator Reporting
  - Availability Monitoring
- However, Some Products Do Many of Them

Scope: Disaster Recovery

- Defer to Management Direction
  - Full System
  - Configs
  - Data (Raw, Parsed, Summaries, etc.)
- This Might Increase Your Cost Considerably
- Vendor Will Provide Method to Backup/Restore Their Application
    - This may not work as well as stated

Consider: The Vendor's “Terms”

- Many SIM-Specific Terms
- Product Pricing is Term-Critical
- Examples:
  - EPS – Events Per Second
  - Distinct Device Types
  - Event Generators
  - Event (Collection|Normalization|Classification |Aggregation|Correlation)

Consider: Implementation Timing

- Experts Agree That a Standard SIM Deployment is **6-18 Months**

- I Have Seen Products:
  - Show Value in Minutes (This is very unusual)
  - Have No Value in 6+ Months

Purchase

Purchase: Vendor Selection

- Market Survey
- Reputation (Gartner and Others)
- Not Apples-to-Apples
- Get Vendor-Referrals
- Demo, demo, demo
  - One Product Failed During Demo
  - Another Should Have

Purchase: Appliance Vs Software

- Appliance is Easier, But at What Cost?
- What is the Analytic Requirement?
- What is the Expected Load/Event Counts?
- Price
- Extensibility:
  - An appliance generally can't
  - Many software SIMs don't do this well either

Purchase: Hidden Costs

- Traditional SIMs Are Outrageously Expensive
- Normal SIMs Still Are Pretty Expensive
- Your Cost Doesn't End With the License
  - Hardware could be considerable
  - Professional services is often mandatory(2)
  - Some Require External Software (e.g. DB Lic.)
- Product Costing Model Might Require Additional License for New Data/Source/User

Purchase: Alternatives to COTS

- MSP (Managed Service Provider)
  - Have Experts Do It (Questionable??)
  - Security Concerns
  - IT Staff Can Lose Motivation and Abilities
- Roll Your Own
  - Central Syslog
  - Open Source SIM - http://ossim.org
  - Prelude IDS http://prelude-ids.org

Prepare/Deploy

Prepare: NTP

NTP NTP NTP
NTP NTP NTP
NTP NTP NTP

Prepare: Vendor Guides

- Steep Learning Curve for SIM
- Additional Learning Curve for Each Vendor
- This Will Not be a Seamless Implementation
  - Should it Be?

Prepare: Things Will Go Wrong

- A Good Vendor Will Overcome the Hurdles
- These Have Happened During Install
  - Reinstall during install
  - Dozens of OOM (Out of Memory) Crashes the First Week
  - Millions of Events Lost in First Few Days
  - Upgrade/Migration Lots Billions of Events

Prepare: Contract

- What determines a successful install?
- How Is This Measured?
- What Do You Do If There Are Problems?
- Deployment Agreement/Statement of Work
- Ensure the SOW Meets Your Requirements

Prepare: SIM Administration

- Backups
- Monitoring
- Configuration Management
- Device Support
- Maintenance
  - Patching, Signature Updates, Etc.
  - Full-Time Admin In Some Organizations
  - Database Pruning/Purging

Prepare: SIM Monitoring

- SIM Systems (resources, network, etc.)
- SIM Processes (db, correlation, web, etc.)
- Reporting-devices
  - Which Devices Are Really Reporting In?
- Number of Incoming Events
  - GREAT Trending Data...When It Is There

Prepare: SIM Configs

- Many initial configurations
- Change Management Policy Model
- SIMS Require a lot of Tweaking

Prepare: Determine Devices

- Determine Type and Quantity of Devices
- Many Devices Log as Syslog
- Some Devices Log Special
  - Determine if this is the Case
  - Find a Solution (e.g. Checkpoint's Opsec Lea Encrypted Log Has Several Free Alternatives)
  - Don't forget to Sync Time!!!!

Case Studies

CS – 6 Months and No SIM

- US Federal Agency Deployed SIEM in 2003
  - Good at the Time, but Became Outdated
- Nov. 2006 Market Survey For new SIM
- Considered 6 Products, Demoed Several
- Ultimately Chose A Recently-Redesigned SIM
- Merrily Deployed New SIM in Jan 2007
- 5 Months of Chaos and Failure
- May 2007 Gave Ultimatum To Vendor

[redacted] 5.6 Installation Timeline

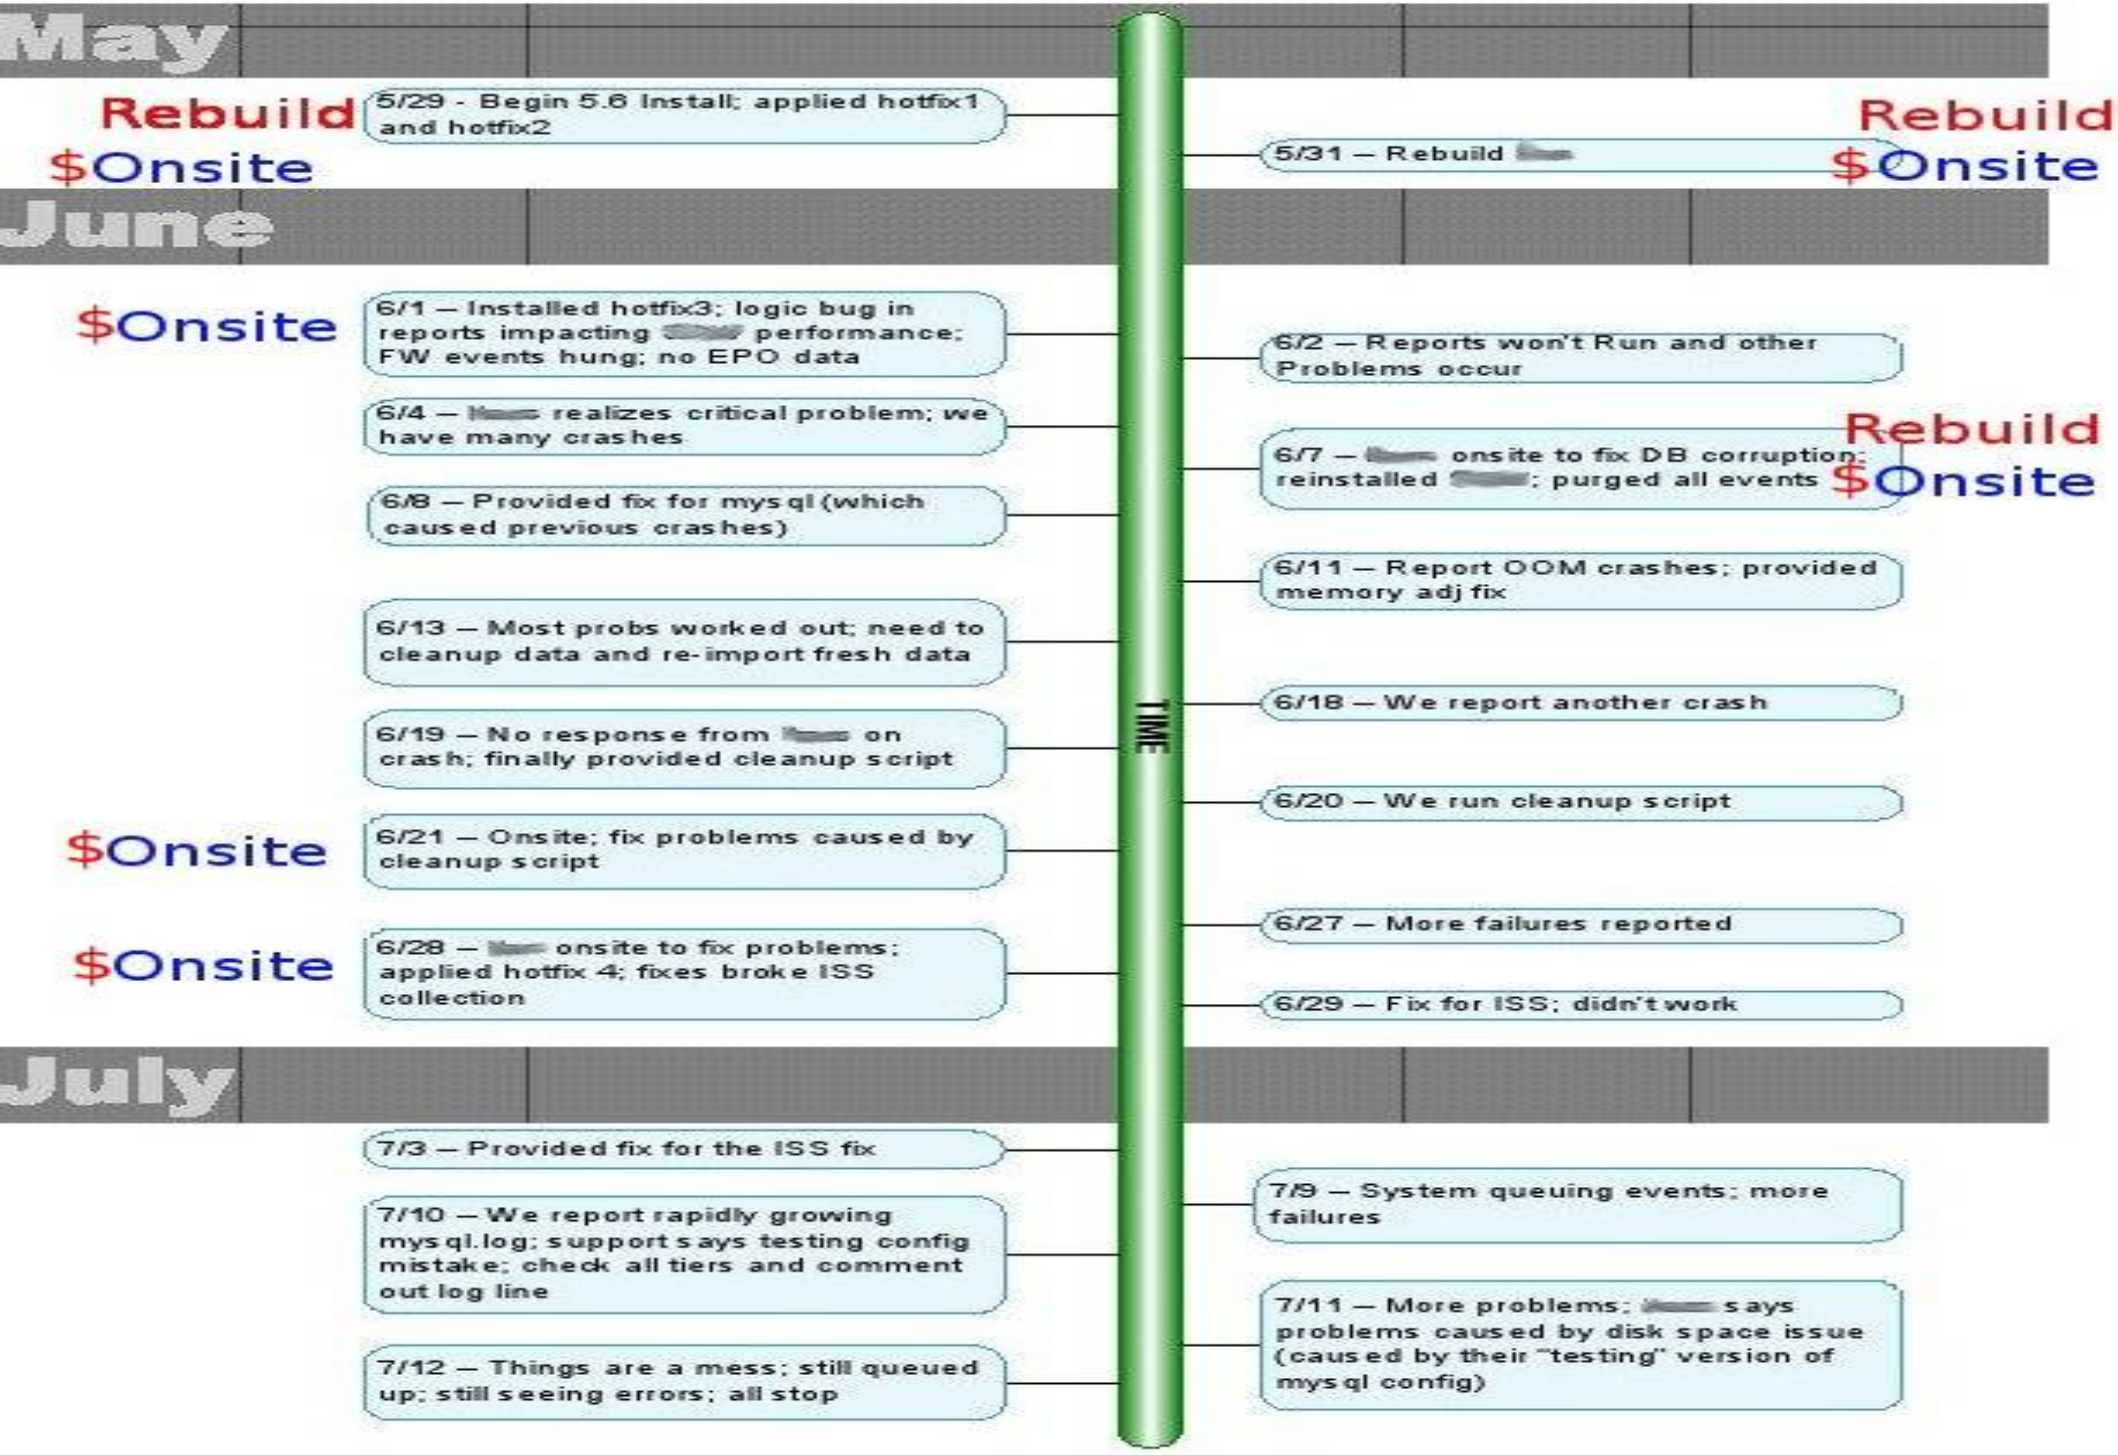

CS: Product TCO

- World-wide US Federal Agency
- Requirements of Central Logging:
  - 140 Enterprise Firewalls
  - 110 NIDS Sensing Interfaces
  - McAfee EPO data
  - ISS Site Secure HIDS data
  - Some syslog data

CS: Product TCO - Product1

- Product Type: Traditional SIM
  - 7 Enterprise Systems (HP DL580)
  - HP SAN - 45 15K SCSI Drives in RAID 10 Config
  - >$2.1M License Investment after 3 Years
  - ~15 hours/week admin cost to maintain
  - ~1 DB rebuild required/year took 120 hours admin time and 40 hours PS
  - Single IP Query across 30 days = 26 - 80min

CS: Product TCO – Product2

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

CS: Product TCO - Product3

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

CS: Product TCO - 1 Year Cost

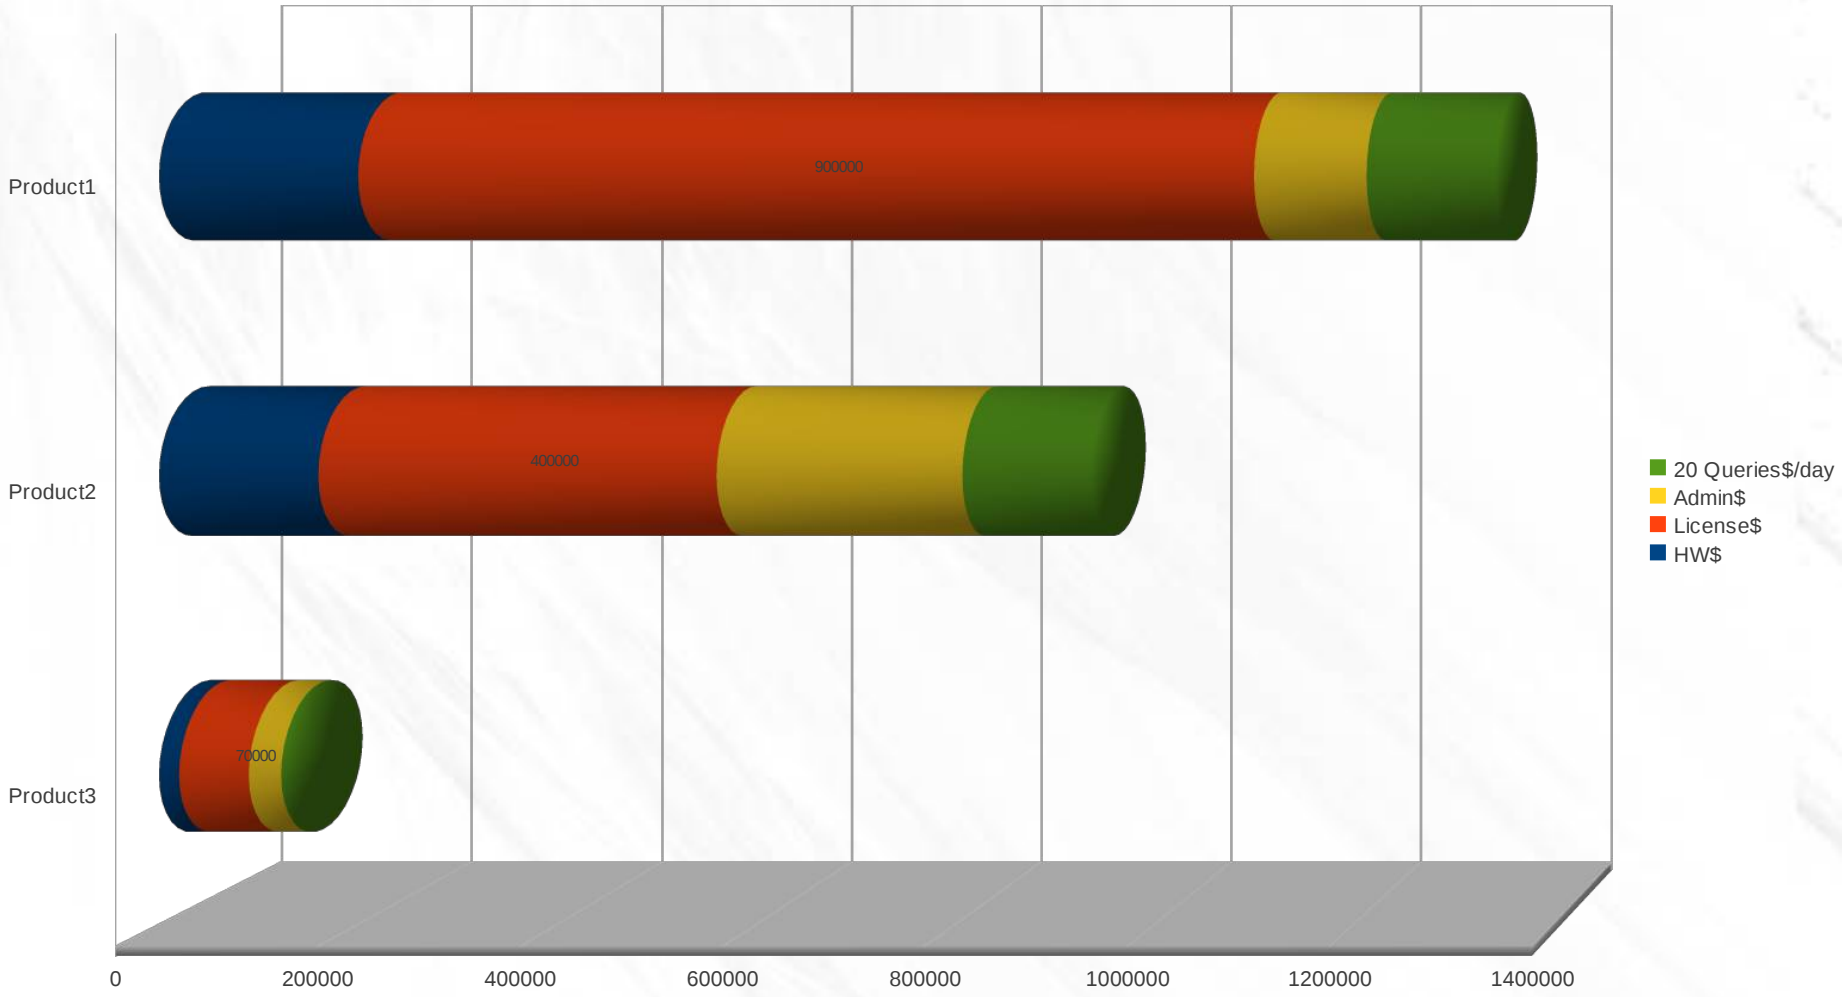

SIM/Log Management Trends

- Flatfile DB Products are Increasing
- RDBS Products are Decreasing
- Cloud Solutions Emerge
- Acquisitions and Rebranding Muddy the Water
- Regulatory/Audits Require Central Log Mgmt
- Threats Prescribe a Nimble Product

Closing

- A Log Mgmt Solution Is Not Just Software, But A Change In Your Information Security Strategy
- Tremendous Drain on Network, System & Time
- Garbage In Is Garbage Out, Not Auto-Secure
- Log Mgmt Can Be A Considerable Asset To A Well-Honed Information Security Program
- Most Organizations Just Need Fast, Reliable Search and Not a Traditional SIM
- A Log Mgmt Solution, at any cost, DOES NOT replace the Skills of a Good Analyst

Presenter:

Sean Wilkerson

Partner - Aplura, LLC

swilkerson@aplura.com

Twitter @sdwilkerson