SIEM 101 Workshop

Optimize IT with Security Information and Event Management

Alex Dow
Chief Research Officer – Mirai Security Inc.
GCIH|SCF|CISSP|OPST

Alex.Dow@miraisecurity.com
https://www.miraisecurity.com

Agenda

- What is SIEM?
- Why buy SIEM?
- Architectures
- Components
- Use Cases
- Townhall Discussion

…A Little Street Cred

- **90's** – Computers & The Internet!, The movie '*Hackers*' was released, NetBus, BackOrifice
- **2001** – School (Boring, but I finally learned TCP/IP)
- **2004** – Bell SOC
- **2008** – Olympic SOC & HoneyNet
- **2010** – Consulting (SIEM, SecOps and ESA)
- **2012** – Co-Founded The Mainland Advanced Research Society (BSides Vancouver)
- **2017** – Co-Founded The Mirai Security Collective (Insert shameless plug here)

Disclaimer

- Generalizations
  - Trying to be as vendor agnostic as possible but there are nuances with each vendor/technology
- Jaded Infosec Warrior
  - The views expressed within this presentation are those of the presenters and do not necessarily reflect the views of their former/current/future employers, clients, partners, friends and/or family members
- Professional Consultation
  - I am a security advisor, but I am not YOUR security advisor (yet)
  - This presentation is for educational purposes only and should not replace independent professional consultation

What is SIEM?

- First: SIEM, SIM, SEM? Huh?
- Logs -> Log management -> SIEM
- Logs vs Events?
- Primary Features
  - Centralized, secure and reliable log collection and retention
  - Fast and easy searching
  - Event correlation and alerting
  - Analytics
  - Dashboarding
  - Reporting
  - Ticketing and automation

The Who's Who of SIEM

- Notable (Unmentioned?) Players
  - Elastic Stack
  - Sumo Logic
  - JASK
- The emergence of Cloud SIEMs
- Death by Acquisition

Drivers for SIEM

- Security
  - Security alert aggregation
  - Anomaly detection via correlations or visualizations
  - Investigation and incident response
  - Situational Awareness
- IT Operations
  - Troubleshooting
  - Alerting on troubles
- Compliance
  - Log retention
  - Audits and real-time risk dashboards

SIEM Component Architecture

1. Event Generation
2. Event Collection
3. Normalization & Enrichment
4. Transport
5. Indexing, Analytics & Correlation

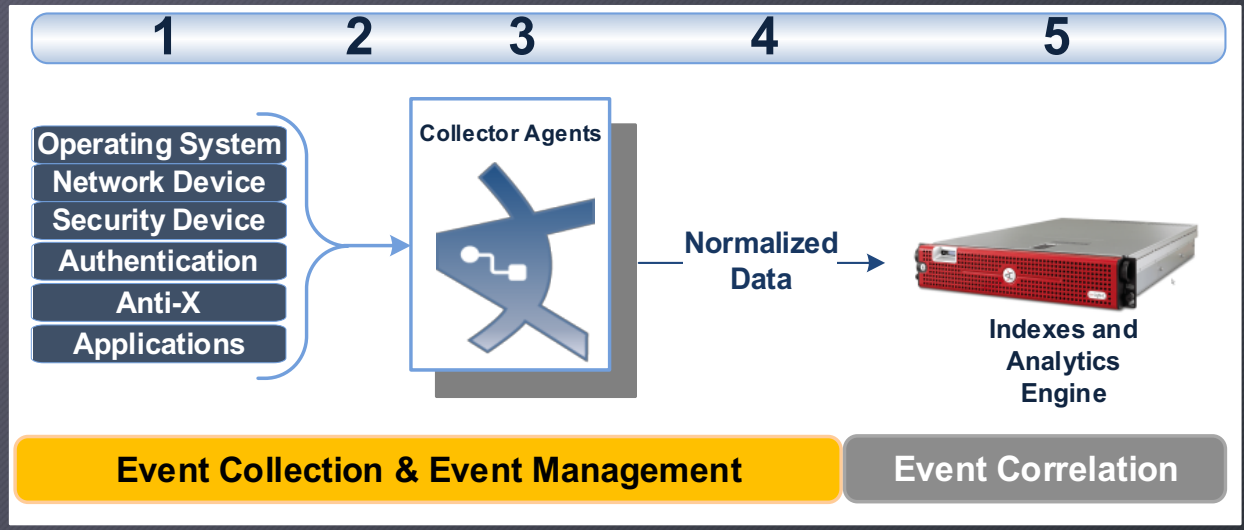

Log Generation and Collection

- Log Sources
  - What: Firewall, OS, DB, application, antivirus, IDS, cloud, packet capture, Nessus Data* and pretty much anything ASCII!
  - How: Configuring logging on your sources
- Collection
  - Agent vs centralized agent
  - Protocols: Syslog, SNMP, HTTPS/API, WMI, SMB/CIFS, FTP, ODBC, etc
  - Real-time vs batching
- To collect or not to collect, that is the question
  - Use case/value
  - Licen$ing
  - Capacity

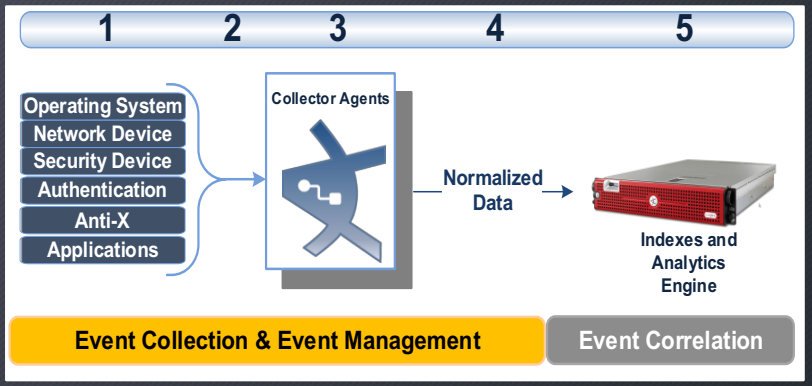

Normalization, Enrichment & Transport

- Parsing and Normalization
  - Structured vs Unstructured Data
  - Disparate logs into one common format
- Filtering and Aggregation
  - Remove noise and save on bandwidth/licen$ing
- Enrichment
  - GeoIP, asset/network models, categorization/tagging, DNS lookups, etc
- Transportation
  - Caching, encryption, compression, bandwidth management
  - Forwards to one or many destinations

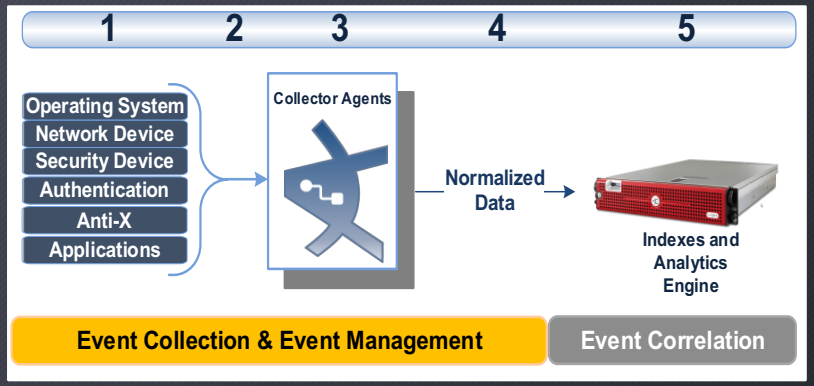

Indexing, Analytics and Correlations

- Indexing
  - Event database management
  - Search management
  - Data retention and archiving
- Analytics and Correlation
  - Asset and network models
  - Dashboards and visualizations
  - Searching
  - (Real-time) alerting and correlation
  - Reporting
  - Ticketing and automation

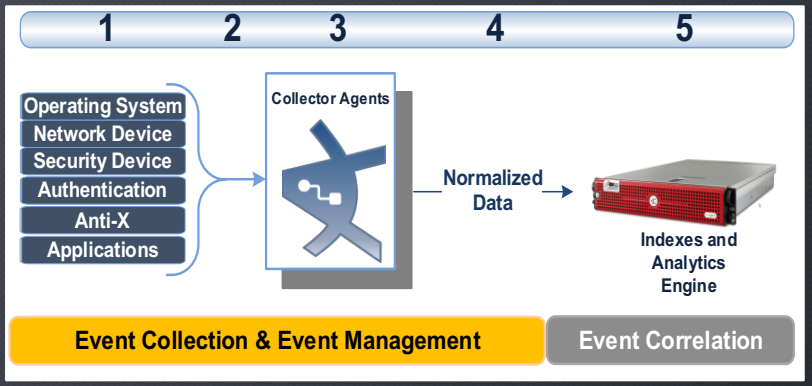

Component Architecture

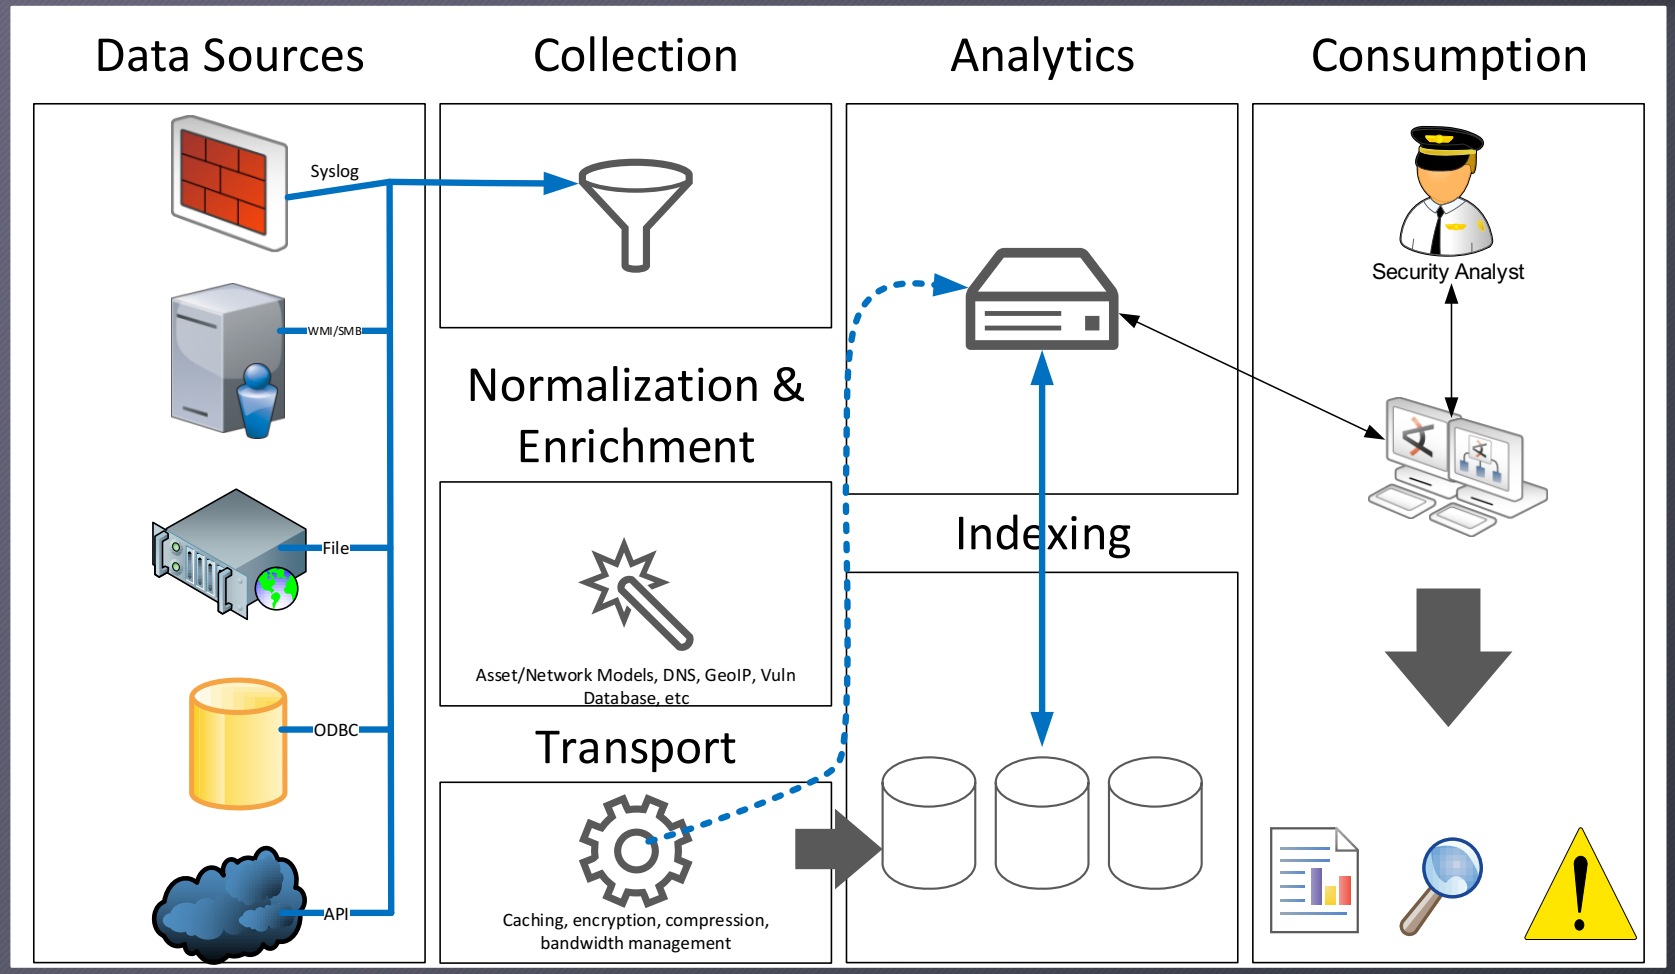

Traditional SIEM Topography

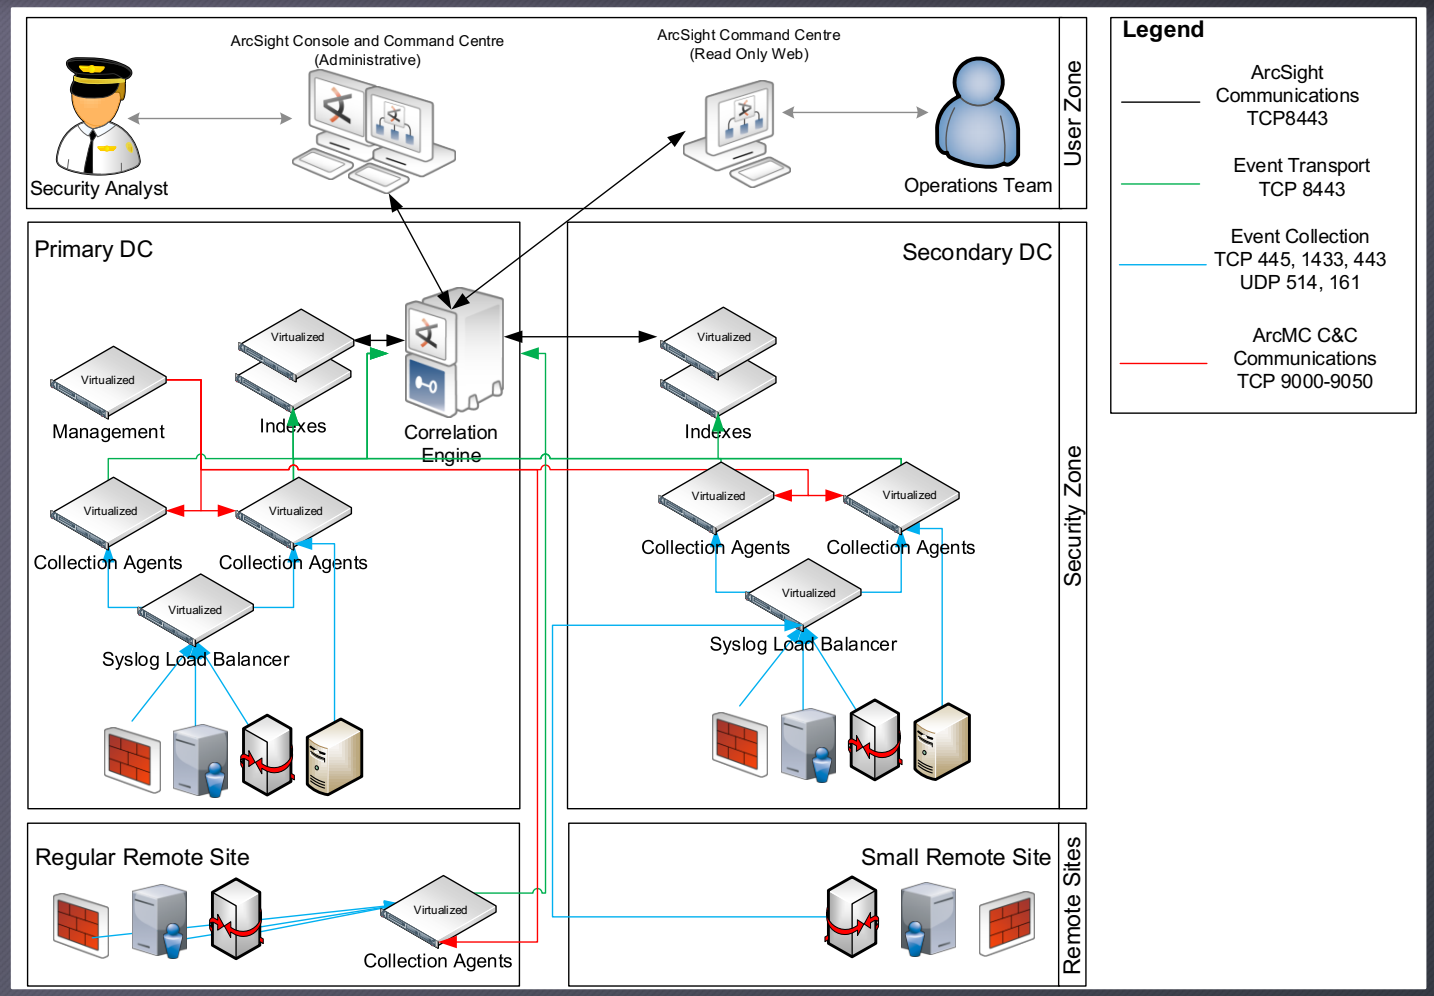

Elastic Stack Topology

- Shippers and Indexers
- Message Bus
- Ingestion Nodes
- Master Nodes
- Data Nodes
- Coordination Nodes
- Tribe Nodes
- Kibana Nodes

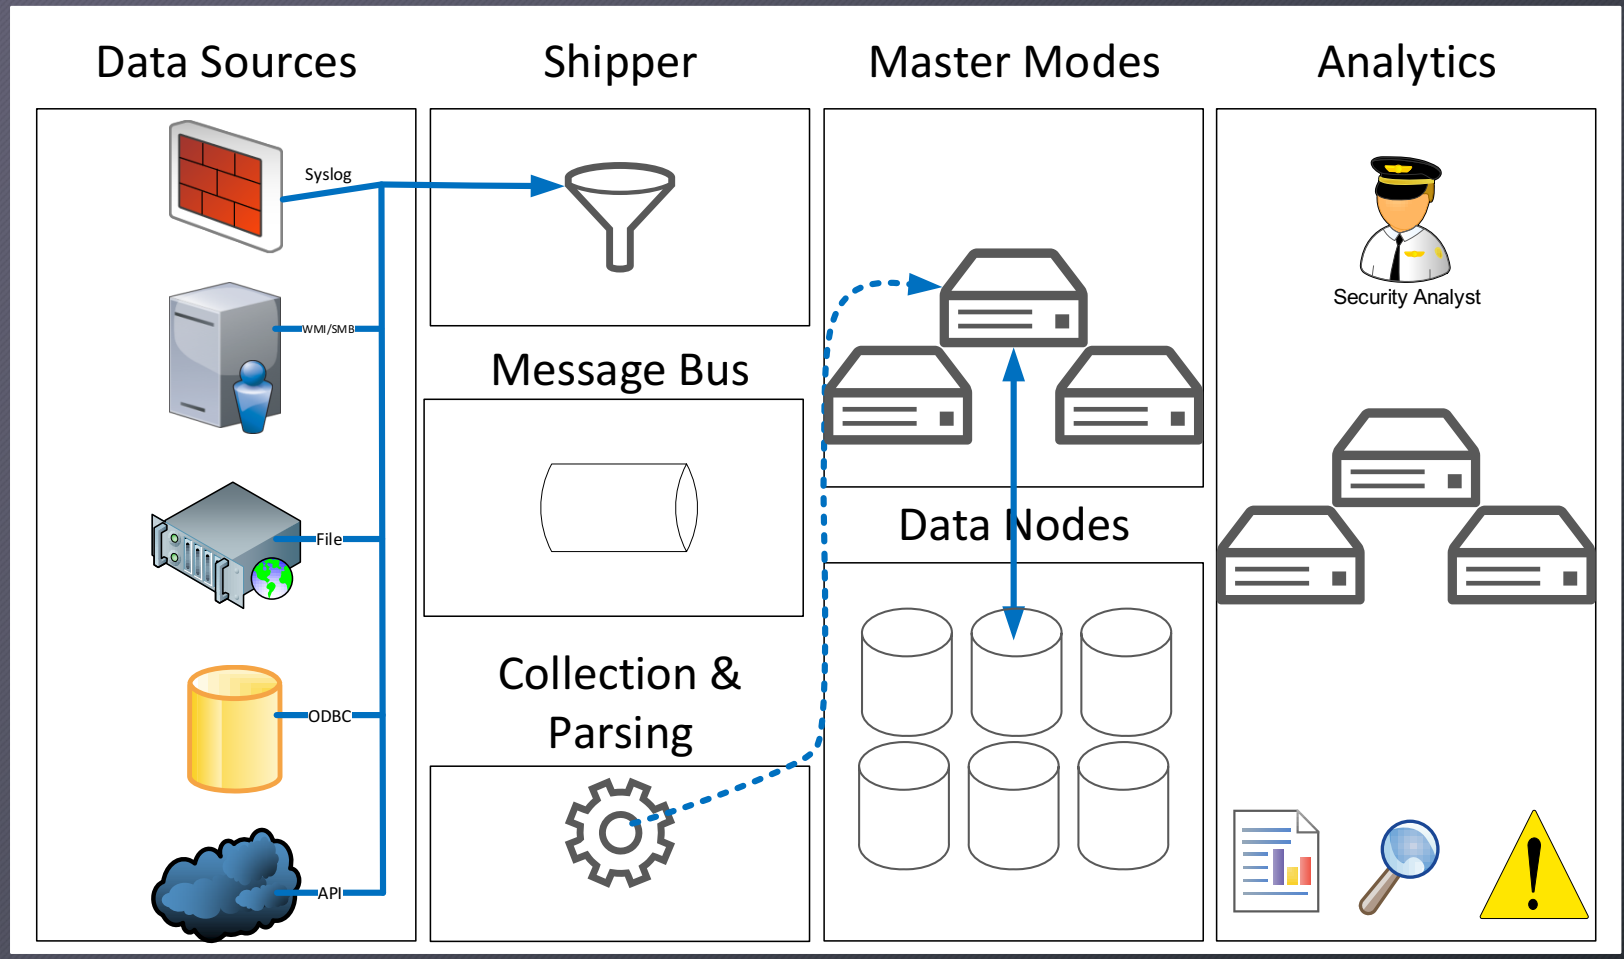

Splunk Topology

- Forwarders
- Indexers
- Search Heads
- ES Search Heads
- Master Cluster Node
- Deployment/License Servers
- Now Cloudy!

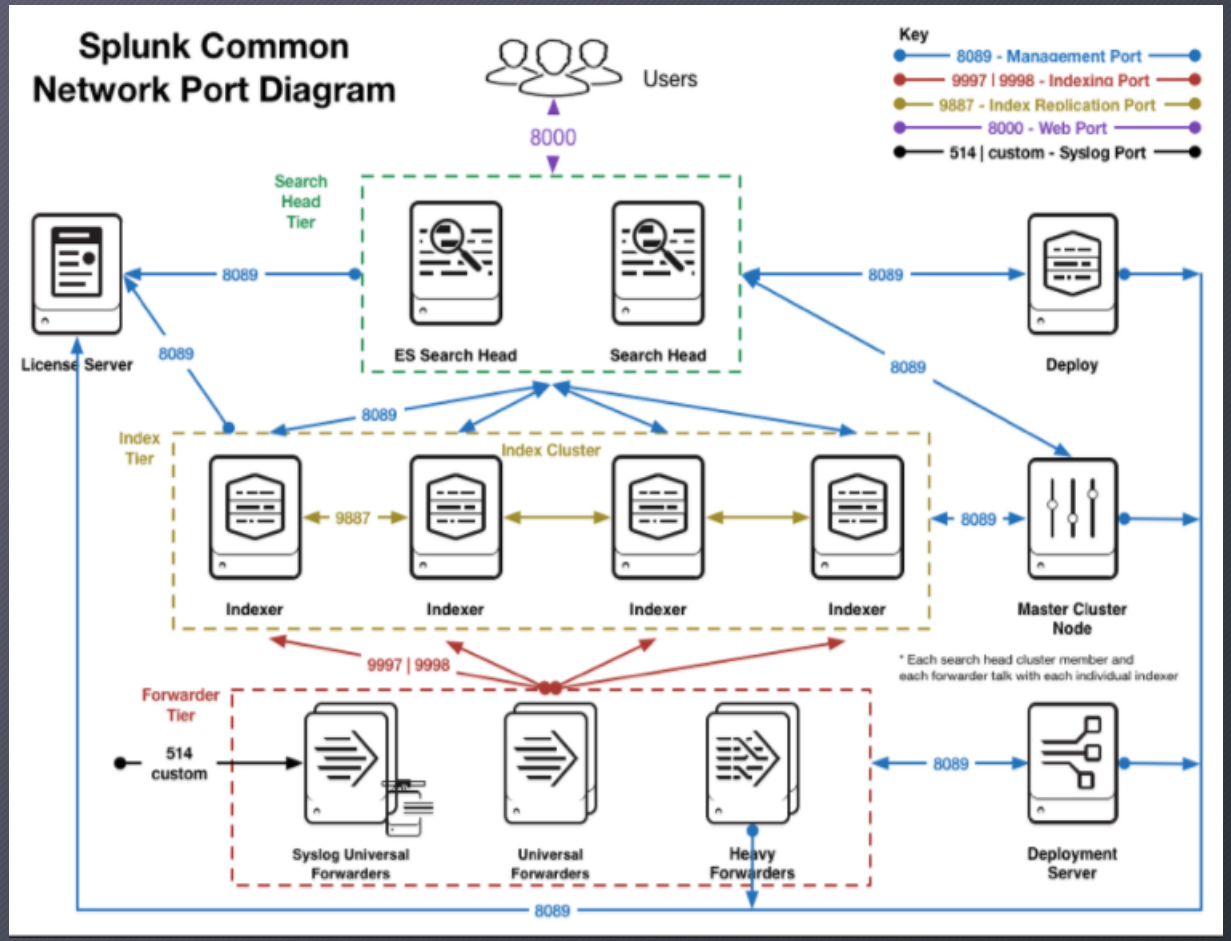

Cloud Topology

- Forwarders
- Indexers
- Search Heads
- ES Search Heads
- Master Cluster Node
- Deployment/License Servers

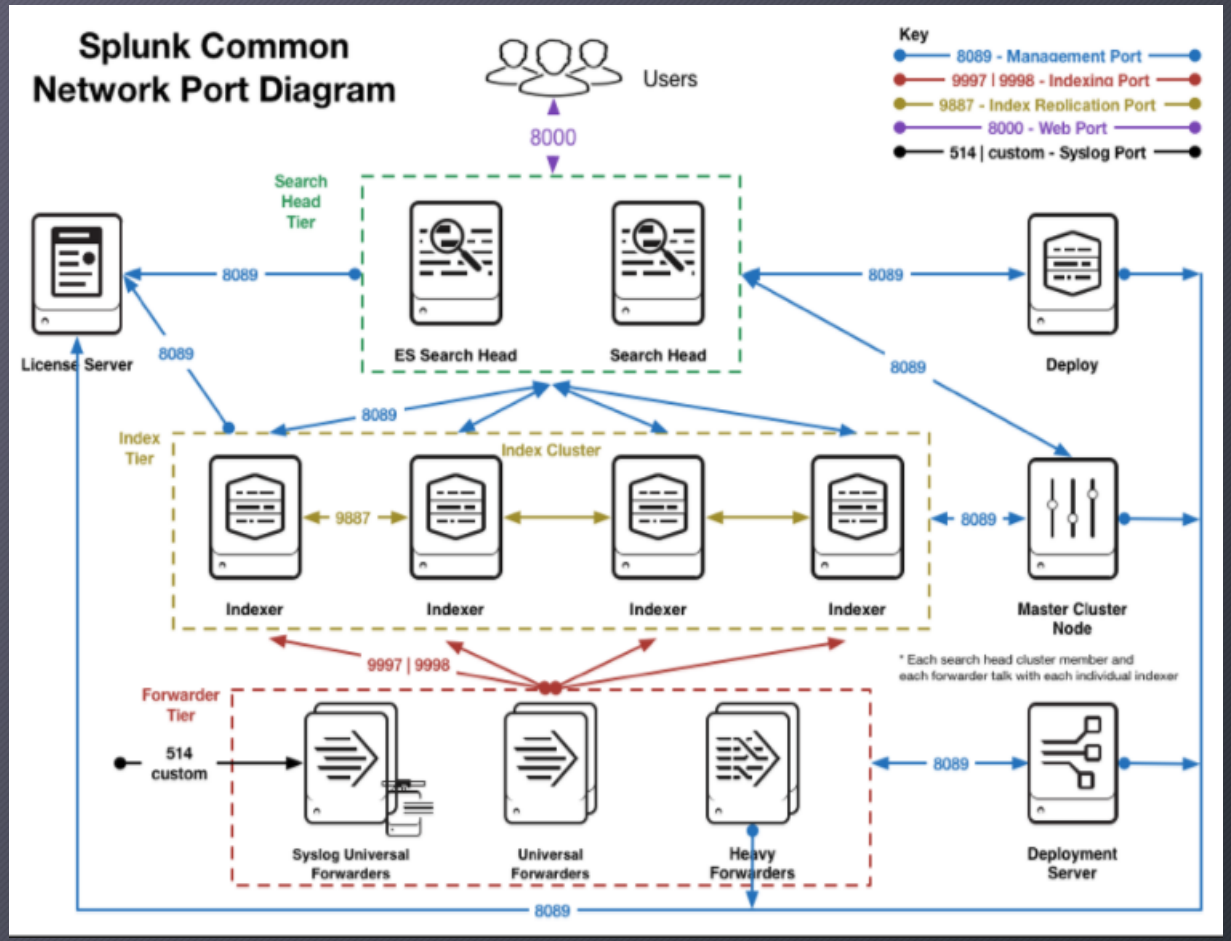

Product Decisions

Traditional

- Pros
  - Security centric
  - Lots of use cases
  - Appliance based
  - Decent documentation
- Cons
  - Appliance based
  - Likely higher costs
  - Scalability concerns
  - Less innovation

Bleeding Edge

- Pros
  - Designed for scale and performance
  - Likely lower costs
  - No appliances
  - Bleeding edge technologies
- Cons
  - Not necessarily focused on security
  - Requires much more knowledgeable staff, less support from vendors
  - Bleeding edge technologies

Advancement and Cool Concepts

- Load Balancing
- Message Bus
- ML, AI
- HDFS and Data Lake
- SOAR

Design Considerations

- Retention
- Performance
- Multitenancy

When implementing a SIEM, goes wrong…

- Sales people suck
- Lack of vision
- Outsourcing 24/7
- Failure to Perform Detailed Planning Before Buying
- Failure to Define Scope
- Overly Optimistic Scoping
- Monitoring Noise
- Lack of Sufficient Context
- Insufficient Resources

Pragmatic Role Out Recommendations

- Day in the life of a SIEM
- Roles and Responsibilities
- Health Monitoring

Use Cases

- Workflow
  - Choosing data sources
- Examples
  - Change management
  - Unauthorized access

Operations

- Roles and Responsibilities
- Health Monitoring and Tuning
- Use case development
  - Atomic, vs correlation, vs advanced correlation
  - Map to other frameworks

Pitfalls

- Parsing
- Stability
- MIA data sources
- Bad forecasting
- Bugs
- What do SIEMs do terribly, stop trying to make it an updown monitor
- Losing data
- WUCS

Town Hall

- What are your drivers?
- Complexity