Event and Information Management and Analysis

Module 14

Significance

- Security Management To-Do List:
  - Security policy – Check
  - Risk management - Check
  - Firewall – Check
  - Hardened platforms – Check
  - Services security – Check
  - Identity and access management – Check
  - User authentication – Check
  - Certificate management scheme – Check
  - Malware protection – Check
  - Patch management – Check
  - Security awareness training - Check

Significance

- How do we answer these questions?
  - How is the IT environment functioning?
  - How effective are our controls?
  - How secure are we today?
  - What activities are we seeing we should we be worried about?

Significance

- Answer: Situation Awareness
- Situation awareness is being aware of what is happening in the IT environment and understanding what available information means to you now and in the future.
- Situation awareness consists of:
  - Identifying goals and objectives
  - Collecting relevant information
  - Interpreting that information
  - Forecasting future status

Significance

- Identifying goals and objectives
  - Cyber defense competition
    - Discussed in future modules
- Collecting relevant information
  - Essentially event and information management is:
    - Identifying useful sources of event data
    - Ensuring that data is being collected
    - Ideally, centralizing and normalizing this data
    - Enabling analysis of this data
    - Managing data retention, security and destruction

Significance

- Event record analysis
  - Perception of the available event data
  - Locating relevant event data
  - Interpreting the meaning of this data
  - Verifying the interpretation
  - Forecasting future status
- Incident communication follows
  - Contacting and informing peers
  - Contacting and informing leaders
  - Contacting and informing customers (subject to approval)
  - Contacting and informing partners (subject to approval)
  - Contacting regulators and/or overarching incident tracking and reporting groups (subject to approval)

Time

- Essential in cyber security operations
  - The only dimension that is shared between physical and cyber space.
  - Yet, it is an artificial dimension
    - Unlike width, depth and height it does not physically exist
    - Time is historically referenced as units of progression of the Earth on its axis as well as orbiting around the sun
      - We actually measure a second as the 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium 133 atom

Time

- Absolute time
  - Time keeping is nearly universal in computing
  - Accurate time keeping is difficult
    - Clock mechanisms have minor defects that produce errors that accumulate over time
  - Constancy is important in time keeping
    - Like a drummer keeping the beat
    - Math can compensate for other imperfections

Time

- Relative time
  - Timeline analysis is very common
  - Evaluating events from multiple sources is as well
  - Forming a single timeline based on multiple sources will hinge on timestamps reported by these sources
  - Time zone differences can be confusing
    - Large organizations tend to pick one as the reference
  - Inconsistent time keeping may foil analysis
    - Did firewall see event A before or after anti-virus reported event B?
      - The answer may influence analysis

Time

- External time servers
  - Operating systems support using networked time servers
    - Time services are provided by Windows domain controllers
    - Time services are provided by Internet time servers
      - With Network Time Protocol (NTP) and related services time keeping sources can be distributed, architecturally similar to DNS name services.

Time

- A common time source improves relative time reliability among event sources
  - If all sources agree when 10:43:23 occurs that is great

Context

- Analysis is highly sensitive to context
  - Where in the environment did the event(s) occur?
  - Relative to the organization's and societal calendar, when did it occur?
    - E.g. Weekend, Labor Day, first day of Q1
  - Are these events particularly unusual?
    - Did we change some technology?
    - Are we running our organization differently?
    - Is a contractor or new employee involved?

Context

- Events with attack signatures relevant to Windows vulnerabilities targeting a Unix system are less troubling.
  - Source may need to be scrutinized, but these events indicate an ineffective attack
- Trust in the event source may also influence analysis
  - Functional reliability
    - Is the absence of relevant events from the source an indication its offline or glitching?
  - Information reliability/accuracy
    - How often does its information accurately reflect “ground truth”
  - Source integrity
    - How likely is the source “telling a lie” at the moment?

Broad Context

- Mission of the organization
- Regulatory constraints on the organization
- Value and sensitivity of information available within the IT environment
- Stakeholders of the organization
  - Customers
  - Personnel
  - Partners
  - Vendors
  - Owners

Absence

- You do not know what you do not know
  - There is no event source in place able to report what is needed
  - Event source may not be capable of reporting relevant events
  - Event source may not be configured to report relevant events
  - Event collection may not know how to process reported events
  - Event content may be missing key attributes
  - Events are being collected but are filtered in the analysis view
- Events not recorded are lost
- Recorded events not collected will be purged
  - Limited device storage is a factor

User Activity

- A sequence of actions taken in order to fulfill user goals and/or objectives
  - Processes may act as a proxy for the user who enabled them
    - Unintended automated actions may be observed possibly resulting from technology design, configuration and possible compromise
- User intentions are not typically discernable in event records.
  - An individual action may be recorded in one or more event records by one or more sources.

User Activity Analysis

- Often there is a lack of metadata or context recorded in events that tie an event to an action and furthermore ties an action to user activity.
- If a user activity is composed of a set of actions, analysts and their tools attempt to:
  - Relate event records from multiple sources to an action that caused them
  - With actions identified and the available context, analysts attempt to associate actions to activities.

Activity Analysis

- Analysts attempt to determine the goals or objectives of these activities
  - The nature of the recorded events assist with determining:
    - whether the actions violate security policy,
    - whether or not the intentions are malicious

Logging

- Event stores:
  - Text files are appended to as new events are recorded
    - Commonly called: log files or logs
    - Common Sources: Operating system services, application processes, firewalls, routers
  - Database tables
    - Common Sources: MS Windows events, database application events, centralized IDS event repositories

Logging

- On Hosts
  - Unix, Linux, Mac OS X
    - Commonly use text logs
      - Some logs may be structured with XML
    - Log locations vary on filesystem by OS
      - Log configuration tends to be centrally managed within OS
      - Common logging service is called syslog, it commonly uses /etc/syslog.conf for configuration
    - Log files can get large
      - File rotation is commonly implemented
      - File compression is also a common practice

Logging

- On Hosts
  - Windows
    - Classes of event types are segregated
      - Application, Security, System and more
    - Access to log entries is via Event Viewer
    - Event Viewer supports viewing remote Windows logs

Logging

- On Infrastructure Devices
  - Log file conventions are less consistent
  - Commonly, device commands and utilities are available to review the log
  - Non-volatile memory is fairly scarce
    - Busy devices will be unable store log entries for very long
  - A “log host” can be established to collect device log entries centrally

Logged-Content Awareness and Interpretation

- Very challenging
  - Many sources of event data
  - Many locations where event data are stored
    - Within one platform
    - Between platforms
  - Each source type uses its own event structure, terminology and operational context
  - Volumes of log entries can be large
    - Needles in a haystack

Security Event and Information Management (SEIM)

- Consolidates events into a single repository
- Normalizes the structure of the events
- Provides dynamic views of the event collection
- Provides cross-source event interpretation
- Labor intensive to install and configure
- Further discussion is outside of scope

Monitoring Sources

- Dynamic sources of operational information
  - Typically these sources complement traditional logging
  - Many sources do not provide a historical record
  - SNMP is a common mechanism for monitoring infrastructure devices
    - Devices with console access often provide tools helpful for monitoring
  - Operating systems and applications provide monitoring tools
    - Unix: top, ps, last, who, iostat, netstat, uptime, du, df
    - Windows: Performance Monitor, Computer Management

Platform Status & Statistics

- Operational health
  - Available services
  - Available resources
  - Patch installation status
  - Security controls operational status
  - Shared resources to network users
  - Operational status of network resources being utilized
- Current & historical performance
  - Load on processing
  - Volume of network I/O
  - Volume of storage I/O
- Remaining capacity
  - Memory
  - Disk space
  - Licenses

Services Statistics

- Apache web server has status reporting
  - mod_status needs to be enabled
- Not all services provide operational statistics
  - Alternate products for the same service (DNS, mail, SSH) may support operational monitoring
    - This feature may differentiate one solution from a popular implementation (BIND, sendmail, etc.)

User Actions

- Operating systems and authenticated services are able to record user actions
  - Monitoring
    - Unix tools: ps, top, lsof
      - Output fields can be intimidating
    - Windows tools: Task Manager
  - Logging
    - Some services will log user actions (e.g. ssh)
    - Windows tools: Event Viewer

Analysis Tools and Techniques

- Log analysis
  - Manual log review
  - Automated via custom or third-party tools
  - Motivations
    - Develop a baseline of what is “normal” user, system and environmental behavior
    - Identify an operational or security problem not previously reported by other means
      - Essentially anomaly detection
    - Verify or locate supporting evidence of an incident

Analysis Tools and Techniques

- Event Correlation
  - Relating and finding greater meaning from events generated by more than one event source.
  - Relating events from multiple sources requires that events share in some fashion a common operational context.
    - E.g. Sources are along a common operational path from threat's point of entry to the target
    - E.g. Sources serve a common function (e.g. anti-virus) and are placed throughout the IT environment. Multiple sources reporting a common issue may indicate the breadth of scope of the incident.

Analysis Tools and Techniques

- Event Correlation
  - Common approach to identifying relationships among multiple sources is to identify common attributes among the events
    - Attributes are ideally produced natively by the event source as an inherent attribute of the threat
    - Supplementing the attribute pool may be necessary in order to improve the operational context of the event information
    - Supplementing the attribute pool may be necessary in order to compensate for discrepancies of terminology used to identify or describe the same semantic (ex. Firewall A issues “permit” and Firewall B issues “allow” for packets that pass policy checking)

Analysis Tools and Techniques

- Event Correlation approaches
  - Manual – labor intensive method of matching, attributes across log entries provided by multiple sources
  - Automation – automated pattern matching of events as they are collected
    - Common feature of Security Information Event Management systems

Analysis Tools and Techniques

- Link analysis
  - Evaluation of relationships among types of things, such as people, hosts, communications, actions
  - Netflow data is a good candidate for link analysis
  - Typically analysis results are presented in an interactive visualization

Analysis Tools and Techniques

- Uses
  - Discovery of unexpected relationships
    - Worm propagation
    - Covert network communications
  - Confirmation of known patterns of behavior

Analysis Strategy

- Determine what decisions you want to enable
- Determine what information you would want to have for those decisions
- Enable and collect that information
  - Prioritizing event sources and types may be necessary
- Differentiate status awareness needs from incident handling
  - Status provides indication of the presence (or lack of) of known potential issues as well as unforeseen issues
  - Incident handling involves, in part, of confirming issues and determining the nature of the issue

Analysis Strategy

- Locate available tools able to enable your analytical objectives
- Determine how to deploy these tools
- Determine who and when and why these tools will be used
- Learn how to use these tools
  - Understand their limitations
  - Understand their use in context of your analytical objectives