Incident Response

Give Me Data or Give Me Death!

Intro

- [~]$ whoami
  - @ab1ff
  - IR / Threat Intelligence / RE / Malware Analysis / Chief Apologizer

- Why does any of this matter?

- What are the key takeaways from this?

Data Source Primer

Know your data, know your threats

Network Data Sources

- Netflow
  - Obtained from networking equipment such as firewalls and routers
  - Lower-level (layer 3 on the OSI model if you care about that sort of thing) network data that typically includes source and destination IPs and ports as well as bytes transferred
  - Not the best quality data, but something is better than nothing
  - Use this to hunt for known malicious IPs, suspicious ports, and anomalous data transfers
- Firewall
  - Firewalls will provide data similar to netflow (layer 3) as well as actions taken for the connection (Allow, Deny) and in more modern firewalls, application identification (HTTP, FTP, etc.) and session metadata (Domain, URI, User Agent, etc.)
  - Data from firewalls that are not application aware can be used much in the same way that netflow is
  - Application aware firewalls, specifically “Next Gen” firewalls can further enable an analyst to hunt for suspicious activities.

Netflow Sample

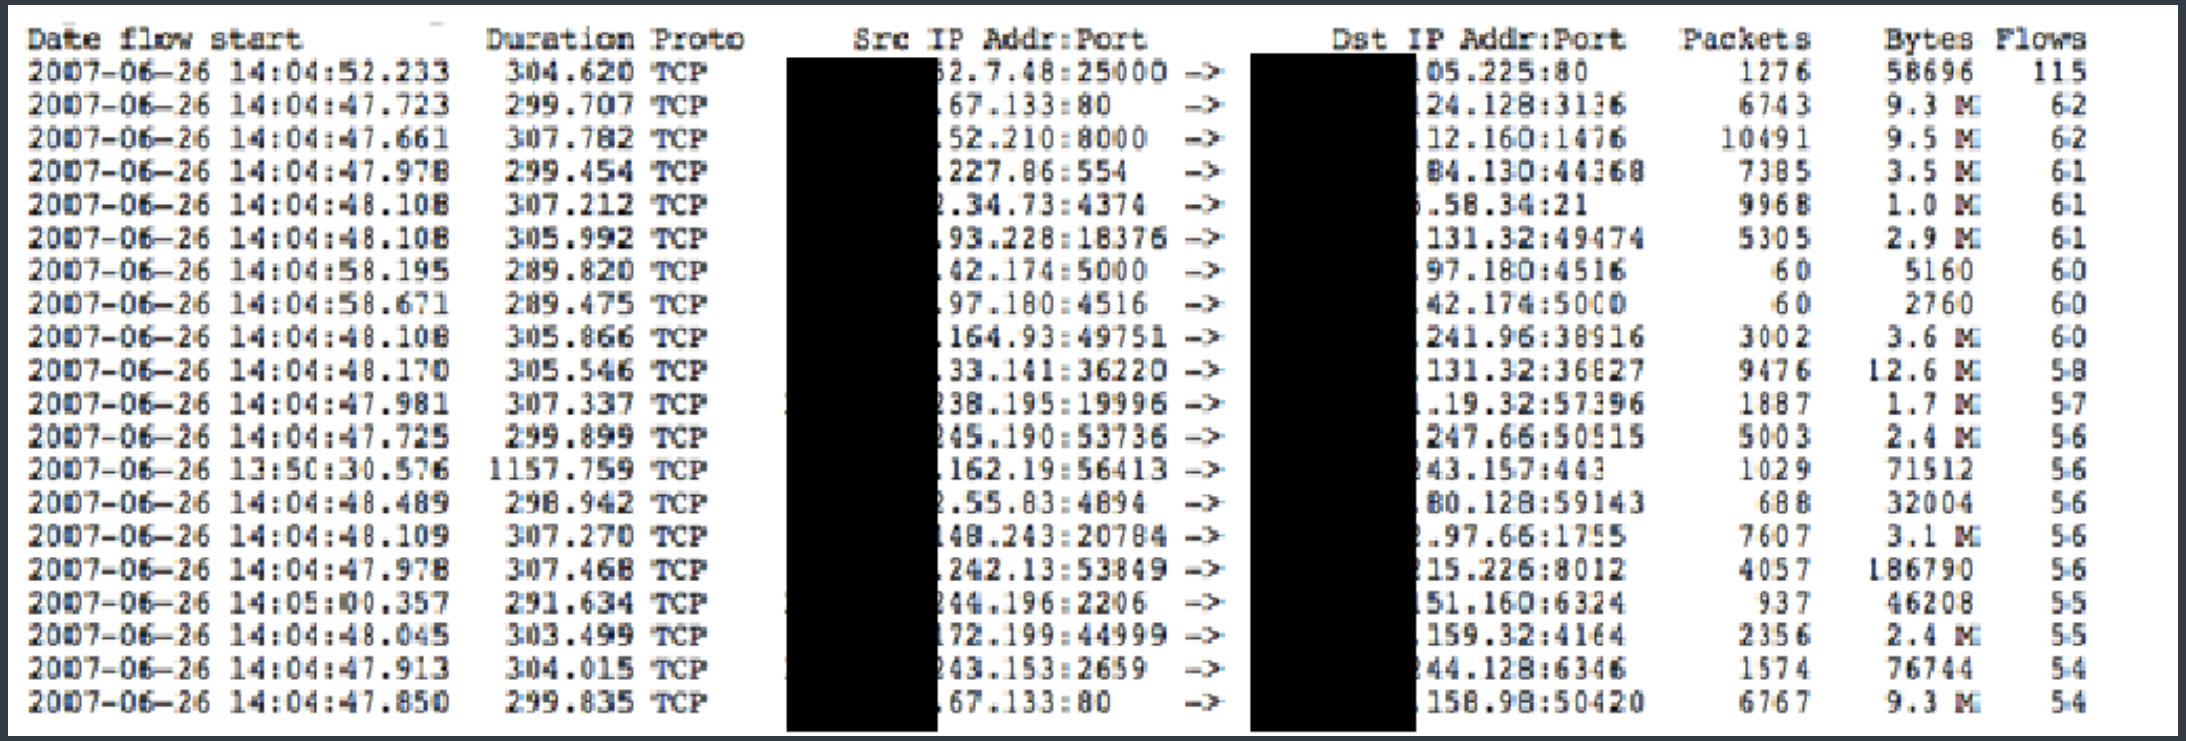

Date flow start	Duration	Proto	Src IP Addr:Port		Dst IP Addr:Port	Packets	Bytes	Flows
2007-06-26 14:04:52.233	304.620	TCP	[redacted]2.7.48:25000	->	[redacted]05.225:80	1276	58696	115
2007-06-26 14:04:47.723	299.707	TCP	[redacted]67.133:80	->	[redacted]24.128:3136	6743	9.3 M	62
2007-06-26 14:04:47.661	307.782	TCP	[redacted]52.210:8000	->	[redacted]12.160:1476	10491	9.5 M	62
2007-06-26 14:04:47.978	299.454	TCP	[redacted].227.86:554	->	[redacted]84.130:44368	7385	3.5 M	61
2007-06-26 14:04:48.108	307.212	TCP	[redacted]2.34.73:4374	->	[redacted].58.34:21	9968	1.0 M	61
2007-06-26 14:04:48.108	305.992	TCP	[redacted].93.228:18376	->	[redacted]131.32:49474	5305	2.9 M	61
2007-06-26 14:04:58.195	289.820	TCP	[redacted].42.174:5000	->	[redacted]97.180:4516	60	5160	60
2007-06-26 14:04:58.671	289.475	TCP	[redacted].97.180:4516	->	[redacted]42.174:5000	60	2760	60
2007-06-26 14:04:48.108	305.866	TCP	[redacted].164.93:49751	->	[redacted]241.96:38916	3002	3.6 M	60
2007-06-26 14:04:48.170	305.546	TCP	[redacted].33.141:36220	->	[redacted]131.32:36827	9476	12.6 M	58
2007-06-26 14:04:47.981	307.337	TCP	[redacted]238.195:19996	->	[redacted].19.32:57396	1887	1.7 M	57
2007-06-26 14:04:47.725	299.899	TCP	[redacted]245.190:53736	->	[redacted]247.66:50515	5003	2.4 M	56
2007-06-26 13:50:30.576	1157.759	TCP	[redacted].162.19:56413	->	[redacted]43.157:443	1029	71512	56
2007-06-26 14:04:48.489	298.942	TCP	[redacted]2.55.83:4894	->	[redacted]80.128:59143	688	32004	56
2007-06-26 14:04:48.109	307.270	TCP	[redacted].48.243:20784	->	[redacted]2.97.66:1755	7607	3.1 M	56
2007-06-26 14:04:47.978	307.468	TCP	[redacted].242.13:53849	->	[redacted]15.226:8012	4057	186790	56
2007-06-26 14:05:00.357	291.634	TCP	[redacted]244.196:2206	->	[redacted]51.160:6324	937	46208	55
2007-06-26 14:04:48.045	303.499	TCP	[redacted].72.199:44999	->	[redacted]159.32:4164	2356	2.4 M	55
2007-06-26 14:04:47.913	304.015	TCP	[redacted]243.153:2659	->	[redacted]44.128:6346	1574	76744	54
2007-06-26 14:04:47.850	299.835	TCP	[redacted].67.133:80	->	[redacted]158.98:50420	6767	9.3 M	54

More Network Data Sources

- IDS/IPS
  - Dedicated software intended to help identify or prevent intrusions
  - Alert-driven: Snort/Suricata
  - Analysis-driven: Bro
  - Snort and Suricata provide notification of activity that meets a rule. These alerts should lead to an investigation and potentially an IR
  - Bro provides a huge amount of metadata from the network to enable threat hunting
- Packet Capture (PCAP)
  - Full copy of network data. This is a lot of data.
  - Great for network forensics, but challenging to hunt through
  - Some IDS solutions such as Snort are able to do selective packet capturing so only the data of interest is collected.
    - The downside to this approach is that you can’t retroactively hunt for previously unknown activities.

The Anatomy of a Snort Rule

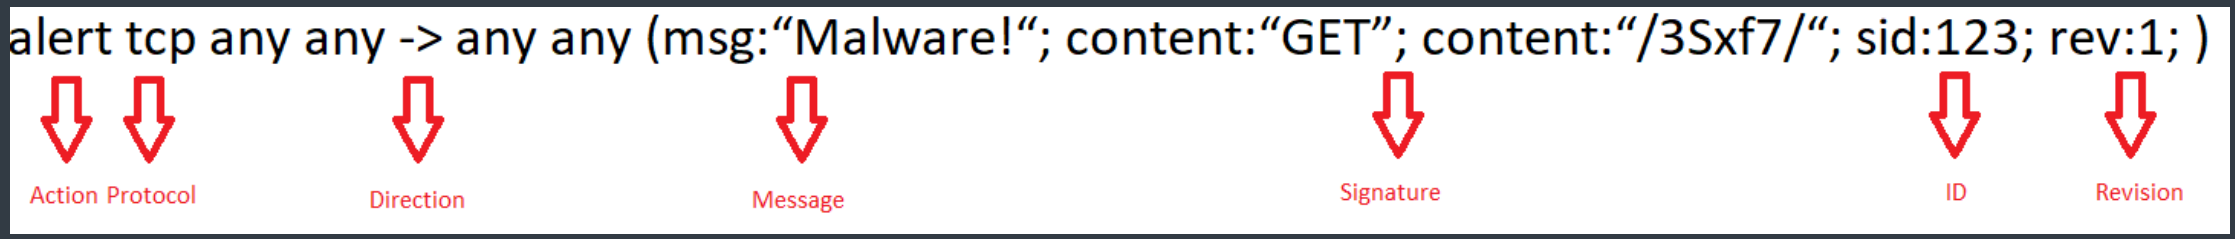

Endpoint Data Sources

- Antivirus
  - Endpoint client that detects and blocks malware.
  - AV Alert can initiate a response, but doesn’t enable hunting
  - Multiple alerts can be an indication of intrusion.
- Realtime
  - Streaming data from the endpoint. This may include information such as file reads/writes, process activity to include command line arguments and spawned children processes, DLL loads, and limited network data such as IP connections and DNS queries
  - This can be achieved through an Endpoint Detection and Response (EDR) product
  - Similar activity can be logged on Windows with an enhanced logging tool such as Sysmon or turning up the logging on default event logs
  - Does not include historical activity

More Endpoint Data Sources

- Logs
  - Windows Event Logs:
    - Security, Application, System, Setup
    - Nonstandard Logs:
      - Microsoft-Windows-TerminalServices-LocalSessionManager/Operational, Microsoft-Windows-PowerShell/*, Microsoft-Windows-WMI-Activity/Operational, Microsoft-Windows-WinRM/Operational, Microsoft-Windows-TaskScheduler/Operational
  - Linux:
    - /var/log/*, /etc/passwd, /home/[username]/.bash_history, /root/.bash_history
- Endpoint Artifacts
  - Windows:
    - Registry, scheduled tasks, services, prefetch, amcache, ntuser.dat, usrclass.dat, MFT
  - Linux:
    - Startup scripts (Init.d or sysctl), cron

Correlation and Analytics

- Correlating and Detecting Across Datasets
  - Security and Incident Event Management (SIEM), User Behavior Analytics (UBA)
  - Solutions in this category don’t create new data, but assist in the analysis and detection of security events across large and disparate datasets
  - Provides event correlation and some include case management features
  - “Single pane of glass” for security operations

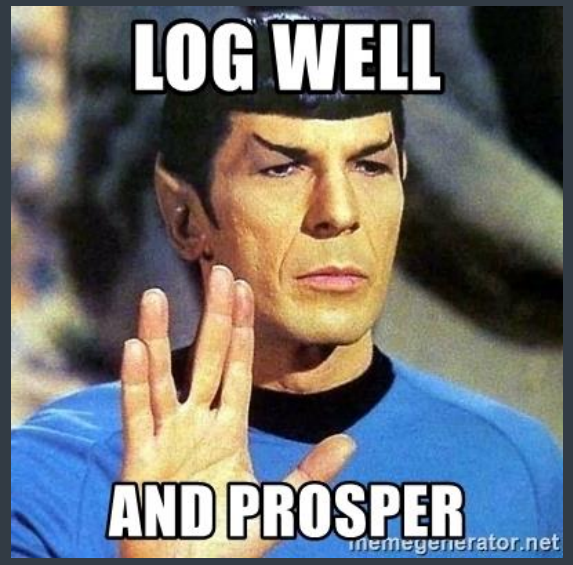

Incident Handling

Putting it all together

IR Challenges

- Things to consider during an incident
  - Scope
    - What has been impacted?
  - Visibility
    - Do we have access to the right data sets?
    - Do we have the ability to effectively search through the data?
    - Is our log retention policy adequate?
  - Timeliness
    - When did the event take place?
    - Are we able to respond or identify impact quickly enough?

Best Practices

- Lights are flashing. We have an incident. Now what?
  - The exact steps to accomplish the above will vary by threat. An incident responder needs to be flexible in their approach
  - There are three primary goals when handling an incident:
    - Determine incident extent
      - What was compromised?
    - Contain the threat
      - Identify all IOCs
        - Network based (IPs/Domains/Protocols)
        - Host based (MD5/Filenames/Staging directories)
    - Conduct a postmortem
      - How did they get in?
      - Where can we improve as a team / organization?
      - What did we do well / What needs improvement?

Determine Incident Extent

- Attackers leave a trail. It is up to the incident responder to follow that trail.
- Start with what you know and work outwards
- Credential dumper ran on a system?
  - What account ran the credential dumper?
  - What else did that account do on the system?
  - Were there any network connections to another internal or external system (Determine C2)?
- Successful brute force?
  - Where was the login from? Internal or external?
  - What did the attacker do after gaining access to the account?
  - Is there any evidence that the attacker pivoted to other systems after the brute force?

Threat Containment

- This can come after or during the previous step
  - If your canoe is sinking, don't just shovel water out. Plug the hole!
- In the event of an attacker on the network, it is important to "stem the bleeding"
  - This is a tricky situation to be in, however, because an attacker who knows they're being hunted may change tactics
- Two approaches to this:
  - Cut access in one fell swoop
    - If you're able to determine command and control methods with high confidence, it may be possible to end connections with a firewall, IPS, or proxy on the network side, or an EDR on the host side
  - Keep the attacker busy
    - Use ACLs on switches and routers to limit what the attacker can access while figuring out how to remove them from the network for good
- Of course, the ultimate goal is to remove all attacker control

Postmortem

- After containing the threat, root cause analysis should be performed
- This will include walking back to the events that transpired before the alert that led to the IR occurred
- Some of this analysis will likely have been completed while determining the extent of the incident
- A security team’s goal should be to continuously improve. Tactics, techniques, and procedures that were observed but not alerted on can be converted into rules to ensure similar activity will be caught

General Dos and Don’ts

- At the start of a response, cast a wide net
  - Start by gathering all possible data that may produce evidence
  - Narrow your focus based on what you know, but don’t be afraid to pivot
- Endpoint and Network data are equally important
  - Don’t neglect either one, they will both be extremely useful when piecing the puzzle together
- Use time as a filter to get a starting point. Follow the trail backwards or forwards as needed
  - Know your timestamp UTC != EST != PST
  - Be wary of timestomping
- Document EVERYTHING
  - Create a timeline, map your evidence to that timeline
- Don’t try to work in a vacuum
  - During an incident inter-team work and communication is vital. Use your team effectively

Helpful Tools

IR Forensic Collection and Analysis: GRR Rapid Response, Autopsy, EnCase, FTKImager, MFT Analyzer, Sysinternals Autoruns, Shimcache parser, Regripper

SIEM: OSSIM, plenty of commercial products

Endpoint Log Collection and Correlation: Wazuh OSSEC, WEF

Windows Event Monitoring: Sysinternals Sysmon

Network Monitoring: Bro, Snort, Suricata, Security Onion

Quick Wins

- Enhance Detection:
  - Enable DNS logging on DNS servers
  - Enable PowerShell logging on Windows systems
    - In Group Policy, enable Module Logging and PowerShell Script Block Logging in Computer Configuration > Policies > Windows Settings > Administrative Templates > Windows Components > Windows PowerShell
  - Enable detailed tracking of process creation events
    - In Group Policy, first enable: Computer Configuration > Policies > Windows Settings > Security Settings > Advanced Audit Configuration > Detailed Tracking > Audit Process Creation
    - Next enable : Computer Configuration > Administrative Templates > System > Audit Process Creation > Include command line in process creation events
- Find Evil:
  - Hunt in key persistence areas (services, runkeys, scheduled tasks). Autoruns can help with this
  - Monitor for network application port mismatch (HTTP over a nonstandard port, SSH over port 80, Tunneling, etc)
  - Monitor for DNS anomalies such as large txt records
  - Look for Windows utility abuse (cmd.exe spawned from anomalous process such as IIS or winword.exe)

Questions?