

Prioritizing Remediation of Enterprise Hosts by Malware Execution Risk

Andrew Chi (Cisco), Blake Anderson (Cisco), Michael K. Reiter (Duke University)

Network Defense: Prioritization is Key


- Each day, a security team must decide:

- Which compromised hosts to **remediate**?
- What vulnerable software to **patch**?
- Which network connections to **monitor**?




Our focus

- Resource constraints: computational, human








Patching vs. Remediation

Agency (and Time)

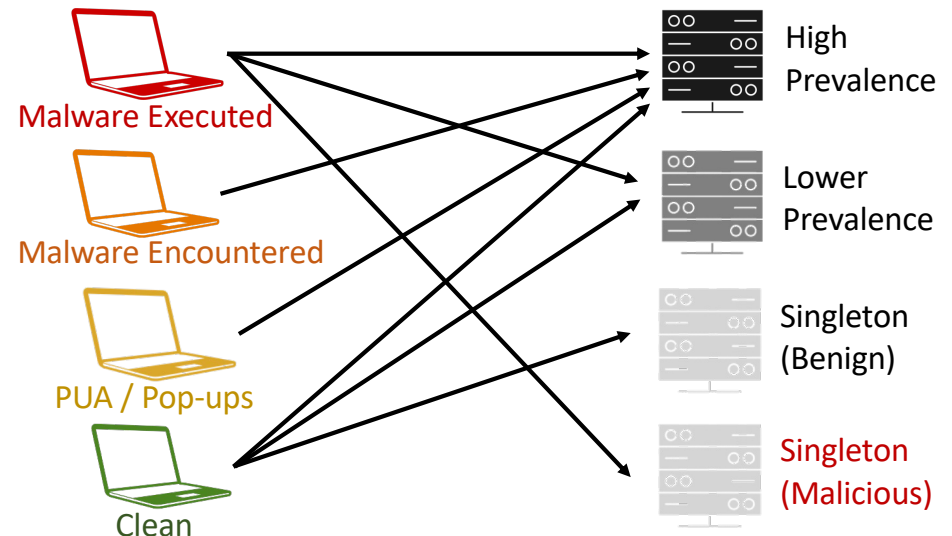


Scope

| | Defender-Driven (relatively static) | <i>Attacker-Driven (dynamic)</i> |
|--------|---|---|
| Global | <p>Universal characteristics</p> <ul style="list-style-type: none"> Vulnerabilities intrinsic to each particular application + version <ul style="list-style-type: none"> Exploitability (attack vector,...) Effect on Confidentiality (C), Integrity (I), or Availability (A) Global patching studies Metrics: CVSS Base  | <p>Public or aggregate attacker activity</p> <ul style="list-style-type: none"> Threat feeds for state of attacker tools and actual attacker activity: <ul style="list-style-type: none"> <i>Exploit code/kit available</i> <i>Actual exploitation in the wild</i> <i>DNS and IP blacklists</i> <i>Reported attack campaigns</i> Metrics: CVSS Temporal  |
| Local | <p>Enterprise-specific environment</p> <ul style="list-style-type: none"> Device & software inventory Host & boundary defenses Security configuration Requirements on (C), (I), (A) Metrics: CVSS Environmental   | <p>Enterprise-specific attacker activity</p> <ul style="list-style-type: none"> <i>Scanning and reconnaissance</i> <i>Malware encounters</i> <i>Compromised hosts</i> <i>Convicted network connections</i> Priority metrics: local, dynamic  |

Goal: Minimize Regret

- Regret is measured by:
 - Failure to remediate (reimage) a host that continues or worsens its malware execution
 - Unjustified business disruption (false alarm)
- Prediction target
 - Hosts that will execute malware (not just encounter it)
 - Actionable time frame: 7 days




NOTE: not all hosts with malware detections require human intervention.

Datasets


Multi-enterprise (2018-2019)

- Millions of hosts
- Thousands of enterprises
- Anonymized
- No interaction with incident response teams
- Endpoint data source:
 - File-based malware detections



Single-enterprise (2022-2023)

- 41,000 hosts
- 1 enterprise (multi-national)
- Worked closely with incident response team
- Endpoint data sources:
 - File-based malware detections
 - Network flows + process ID




Endpoint Data Sources (Agent-based)

- Advanced Malware Protection (AMP)
- File-based malware detections: (malicious events only)
- Network Visibility Module (NVM)
- Netflow[5-tuple]+domain+process (all traffic, but no labels)




```
outlook.exe      report.doc
word.exe (CVE-xyz)
chrome.exe       webpage.html
webex.exe        popunder.js
malware.exe
```



Cisco Secure Endpoint; formerly
Advanced Malware Protection (AMP)

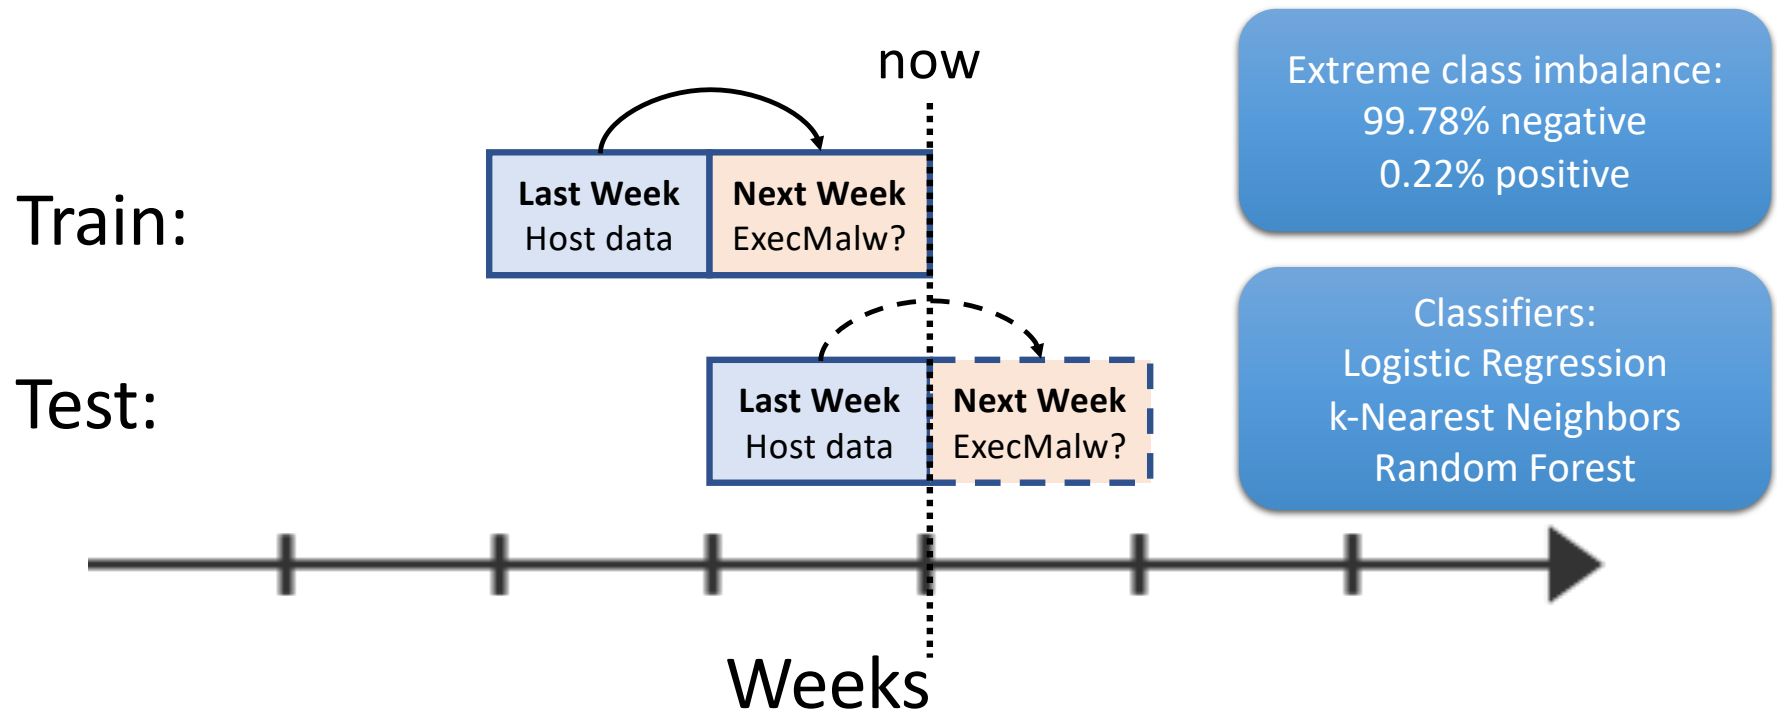


```
outlook.exe -> imap.example.com:993
malware.exe -> maybebad.com:443
chrome.exe -> www.example.com:80
webex.exe -> remote.com:udp/9000
```

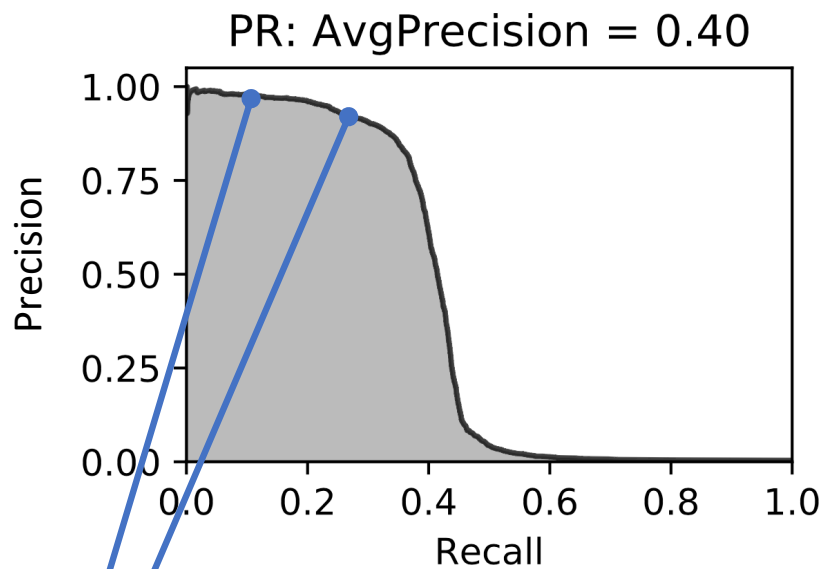


Network Visibility Module (NVM) for
Cisco AnyConnect VPN client

ML Goal: Predict Compromise Next Week



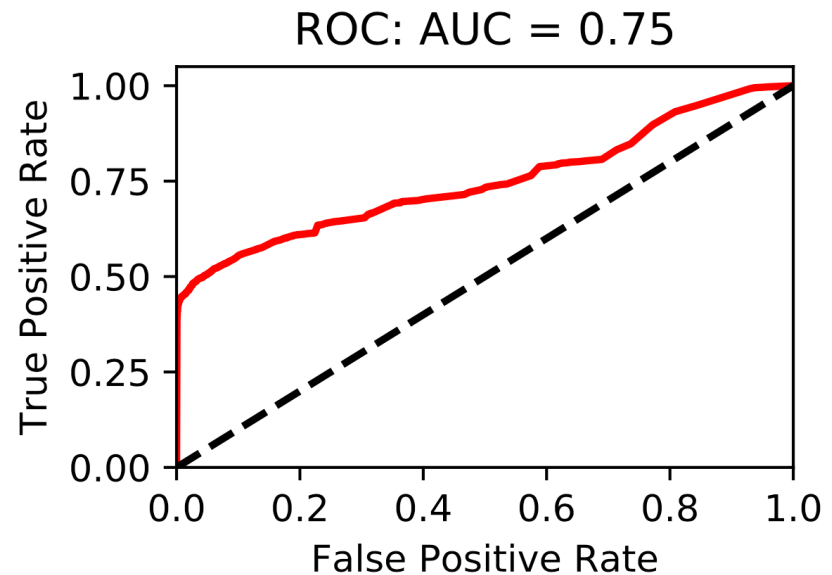
Multi-Enterprise: Typical Classifier Performance



Example points on Precision-Recall curve:

@ Cutoff 1000, Precision=0.98 and Recall=0.10

@ Cutoff 3000, Precision=0.92 and Recall=0.27



WARNING: Receiver Operator Characteristic is misleading in many security contexts, due to base-rate fallacy (Arp, USENIX 2022)


Results: Multi-Enterprise (2018-2019)

| Feature (Event Type) | Importance |
|----------------------------------|------------|
| Executed Malware | 0.2649 |
| Threat Detected in Exclusion | 0.0805 |
| File Detection | 0.0662 |
| Policy Update | 0.0638 |
| Computer Metadata Changed | 0.0543 |
| Failed to Delete from Quarantine | 0.0529 |
| Attempting Quarantine Delete | 0.0524 |
| Low Prevalence Execution | 0.0352 |
| Threat Quarantined | 0.0323 |
| Generic IOC | 0.0301 |

- High precision, moderate recall
- BUT: limited practical use
- Requires too much data
- Data mismatch (sampling bias)
 - Train: multi-enterprise
 - Deploy: single-enterprise

Single-Enterprise Study (2022-2023)

- Malware is a moving target
 - More sandbox-based detections
 - Adapt to represent “typical” case
- **AMP feat.:** malware names
- **NVM feat.:** new singletons (prevalence=1 activity)
- **NVM feat.:** new public suffixes (.com, .co.uk, .k12.nc.us)
- NVM only partially available



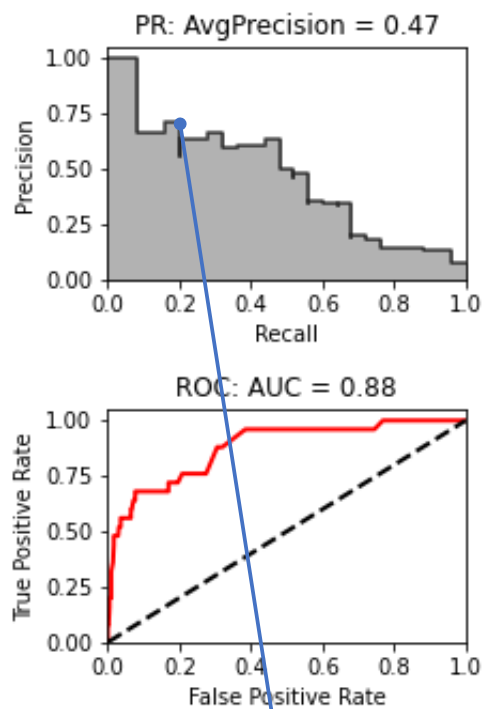
rare ≠ malicious

NVM singleton types

- Process hash: `c1ed4c18...`
- Process name: `rare.exe`
- Dest. domain: `rare.net`
- Dest. IP subnet: `3.1.4.0/24`


Single-Enterprise: Prediction of Malware Execution (1 week ahead)

6-month AMP (22k hosts)




Example Cutoff:
Precision: 0.71 (5 of 7)
Recall: 0.20 (5 of 25)

6-month AMP+NVM (22k hosts)



Example Cutoff:
Precision: 0.83 (10 of 12)
Recall: 0.40 (10 of 25)

12-month AMP (41k hosts)



Example Cutoff:
Precision: 0.92 (54 of 59)
Recall: 0.49 (54 of 111)

Feature Importances

6-month AMP (22k hosts)

| Feature (Past/Current Detection) | Import. |
|----------------------------------|---------|
| current_det_JS:Adware.Popunder.G | 0.059 |
| current_evt_Threat Detected | 0.051 |
| past_compromise | 0.035 |
| past_det_JS:Adware.Popunder.G | 0.034 |
| current_compromise | 0.033 |
| current_evt_Executed Malware | 0.031 |
| current_det_JS:Adware.Popunder.D | 0.030 |
| past_evt_Threat Detected | 0.023 |
| past_Retrospective Detection | 0.022 |
| past_evt_Executed Malware | 0.020 |
| past_det_JS:Adware.Lnkr.L | 0.013 |
| current_det_JS:Adware.Lnkr.L | 0.012 |
| past_det_JS:Adware.Popunder.D | 0.011 |
| current_Retrospective Detection | 0.010 |
| current_det_W32.File.MalParent | 0.008 |

6-month AMP+NVM (22k hosts)

| Feature (Past/Current Detection) | Import. |
|----------------------------------|---------|
| current_det_JS:Adware.Popunder.G | 0.049 |
| past_compromise | 0.045 |
| current_evt_Threat Detected | 0.045 |
| current_compromise | 0.039 |
| new_public_suffixes_count | 0.036 |
| past_det_JS:Adware.Popunder.G | 0.031 |
| current_evt_Executed Malware | 0.030 |
| new_singleton_subnets_count | 0.028 |
| new_singleton_domains_count | 0.028 |
| current_det_JS:Adware.Popunder.D | 0.024 |
| past_Retrospective Detection | 0.024 |
| past_evt_Threat Detected | 0.023 |
| past_evt_Executed Malware | 0.016 |
| new_singleton_hashes_count | 0.013 |
| past_det_JS:Adware.Popunder.D | 0.013 |

12-month AMP (41k hosts)

| Feature (Past/Current Detection) | Import. |
|----------------------------------|---------|
| past_evt_Executed Malware | 0.058 |
| past_compromise | 0.051 |
| current_evt_Executed Malware | 0.038 |
| current_evt_Threat Detected | 0.030 |
| current_compromise | 0.029 |
| past_evt_Threat Detected | 0.022 |
| past_det_JS:Adware.Popunder.G | 0.016 |
| past_det_W32.DFC.MalParent | 0.015 |
| current_det_JS:Adware.Popunder.G | 0.013 |
| current_det_Auto.7DF7E9D.Adware | 0.011 |
| current_det_W32.DFC.MalParent | 0.011 |
| current_det_JS:Adware.Popunder.D | 0.010 |
| past_Retrospective Detection | 0.009 |
| past_det_JS:Adware.Popunder.D | 0.008 |
| past_det_PUA.Win.Dropper.Generic | 0.008 |

Singletons: Benign and Malicious

Benign: normal web browsing


- Author's laptop
- Visited a singleton:
 - farmsidekitchen.com
 - Local salad restaurant
- Singleton (prevalence = 1)
 - No other machines visited the domain that day.
- Web browsing singleton domains are very common
- Conclusion: **No security concern**

Malicious: Sality malware


- duo device health.exe
- Visited many singletons:
 - suewyllie[.]com
 - 724hizmetgrup[.]com
 - pelcpawel.fm.interia[.]pl
 - > 100 IPs in singleton /24 prefixes
- Program hash legitimate, but...
- Many other A/V and network alerts
- Conclusion: **Sality process injection**
- Action: **Reimage machine**

Classifier Degradation if No Retraining

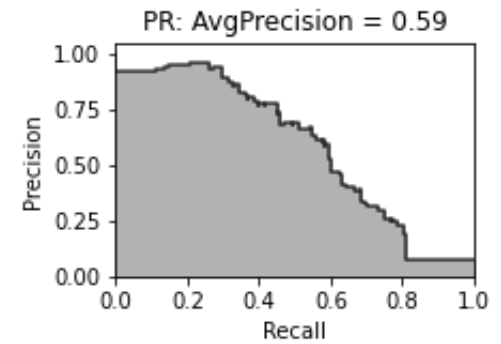
1 week gap



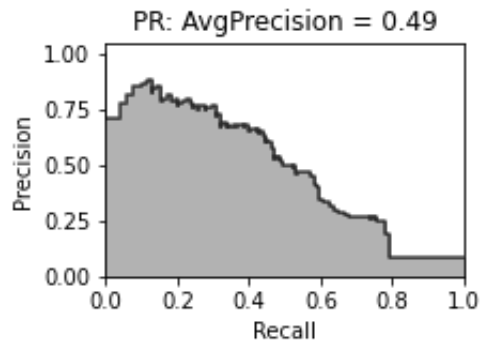
2 week gap



3 week gap



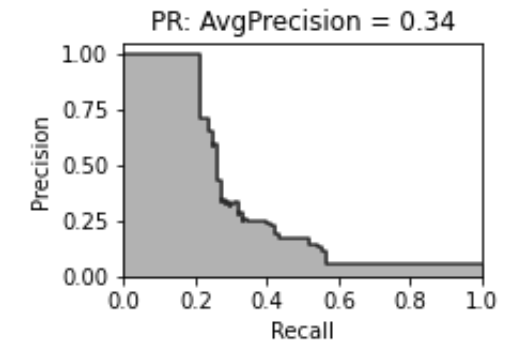
4 week gap



(skip 10 weeks)




14 week gap



Conclusions and Future Work

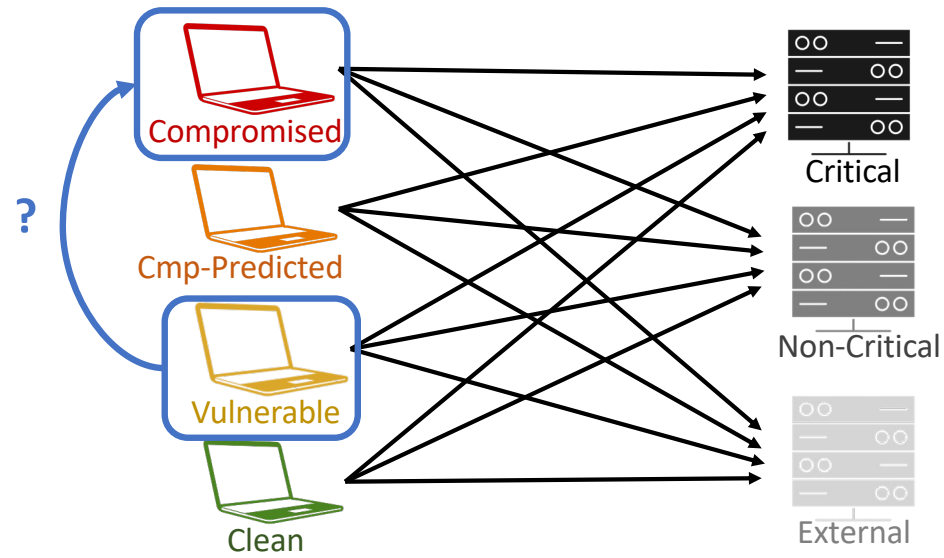
- Single-enterprise malware prediction is possible
 - Not necessarily inferior to multi-enterprise training
 - Local idiosyncrasies
 - Local environment more stable
 - Richer features, combined data sources
- Feature importance: basic explanations
- Future: ongoing collaboration with Cisco CSIRT
 - Threat hunting based on classifier results
 - Gap analysis of current plays vs automated classifier



Acknowledgments

- Cisco CSIRT
 - Brandon Enright
 - Adam Weller
 - Joey Rosen
 - Igor Dobrogorskiy
 - Derek Schmell
- Cisco Secure Firewall
 - Blake Anderson
 - David McGrew

Backup slides





Vulnerabilities → Compromise?

- **CVE-2013-3829**

- Malware execution rate not significantly different between vulnerable hosts and baseline.

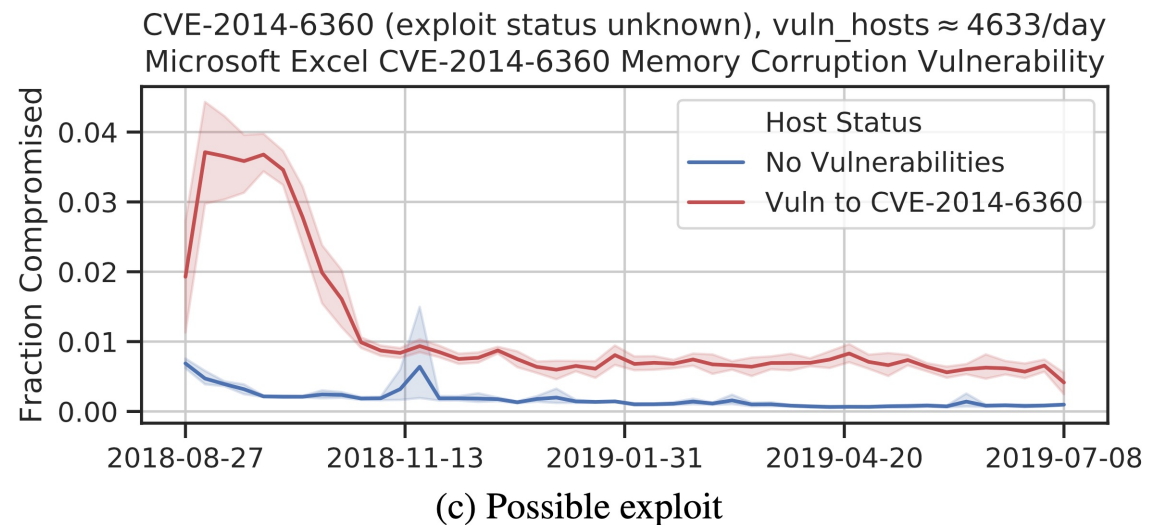
- **CVE-2015-1641**

- Malware execution rate is significantly higher than baseline.
- Consistent with Symantec SecurityFocus, which indicates that CVE-2015-1641 has been exploited in the wild.



CVE-2014-6360

- The malware execution rate is significantly higher than baseline.
- SecurityFocus is not aware of any active exploits of CVE-2014-6360, but AMP data shows above-baseline activity for hosts running a version of MS Excel affected by this vulnerability.
- While correlation is not causation, this may be worth investigating.



These may be hosts whose traffic should be monitored more closely.