
security is not an island
HILTONMALTA
24th Annual
FIRST
Conference
MALTA
17 - 22 June 2012
FIRST
Improving Security Together

Anomaly Detection through DNS Correlations

Michael H. Warfield
Senior Security Researcher and Threat Analyst
IBM Security Services X-Force

24th Annual FIRST Conference
MALTA
17 - 22 June 2012

Why DNS?

- This is still a work in progress but...
- Why look at the DNS? It's just THERE.
  - Aurora showed what can be found.
  - DNSChanger showed what can be done to us.
  - Iodine (a DNS Covert Channel VPN package) should scare the crap out of us.
- Malware is using DNS more and more.
- Maybe it's overdue to take a deep look at what's going on in the Domain Naming Service.

What's on Tap

- Nature of Anomaly Detection
- Nature of Correlations
- Nature and Background of DNS
- State of DNS Deployments and Management
- What Can We Detect Without Correlations
- What Can Correlations Enhance
- Advanced Topics (The Work in Progress)
- Conclusion

Anomaly Detection

- Anomaly Detection is the holy grail of security.
- From a baseline of “normal” behavior, abnormal or anomalous behavior is flagged.
- For select cases of well known baselines, anomaly detection works well.
- Generalized cases are problematical.
- It's primarily statistical in nature.
- Can be prone to false positives and negatives.
- It can catch things nothing else can.

Establishing Baselines

- Baselines are the key to anomaly detection!
- Establishing a baseline is a challenge.
  - A baseline may be “determined.”
  - A baseline may be “managed.”
  - A baseline may be “learned.”
  - Baselines may change.
  - Baselines will have exceptions.
- Baselines for DNS may be determined if DNS is managed properly.

Nature of Correlations

- Correlation is a process of comparing data.
- In math and science there are specific definitions.
- Auto-correlation is comparing the data to itself in some way (time, space, attribute).
- A Fourier Transform is a form of auto-correlation.
  - A Fast Fourier Transform converts time domain data into frequency domain data.
- Other types of correlations are less rigid.
- Correlations provide a method of complex filtering.

Conficker P2P Correlation

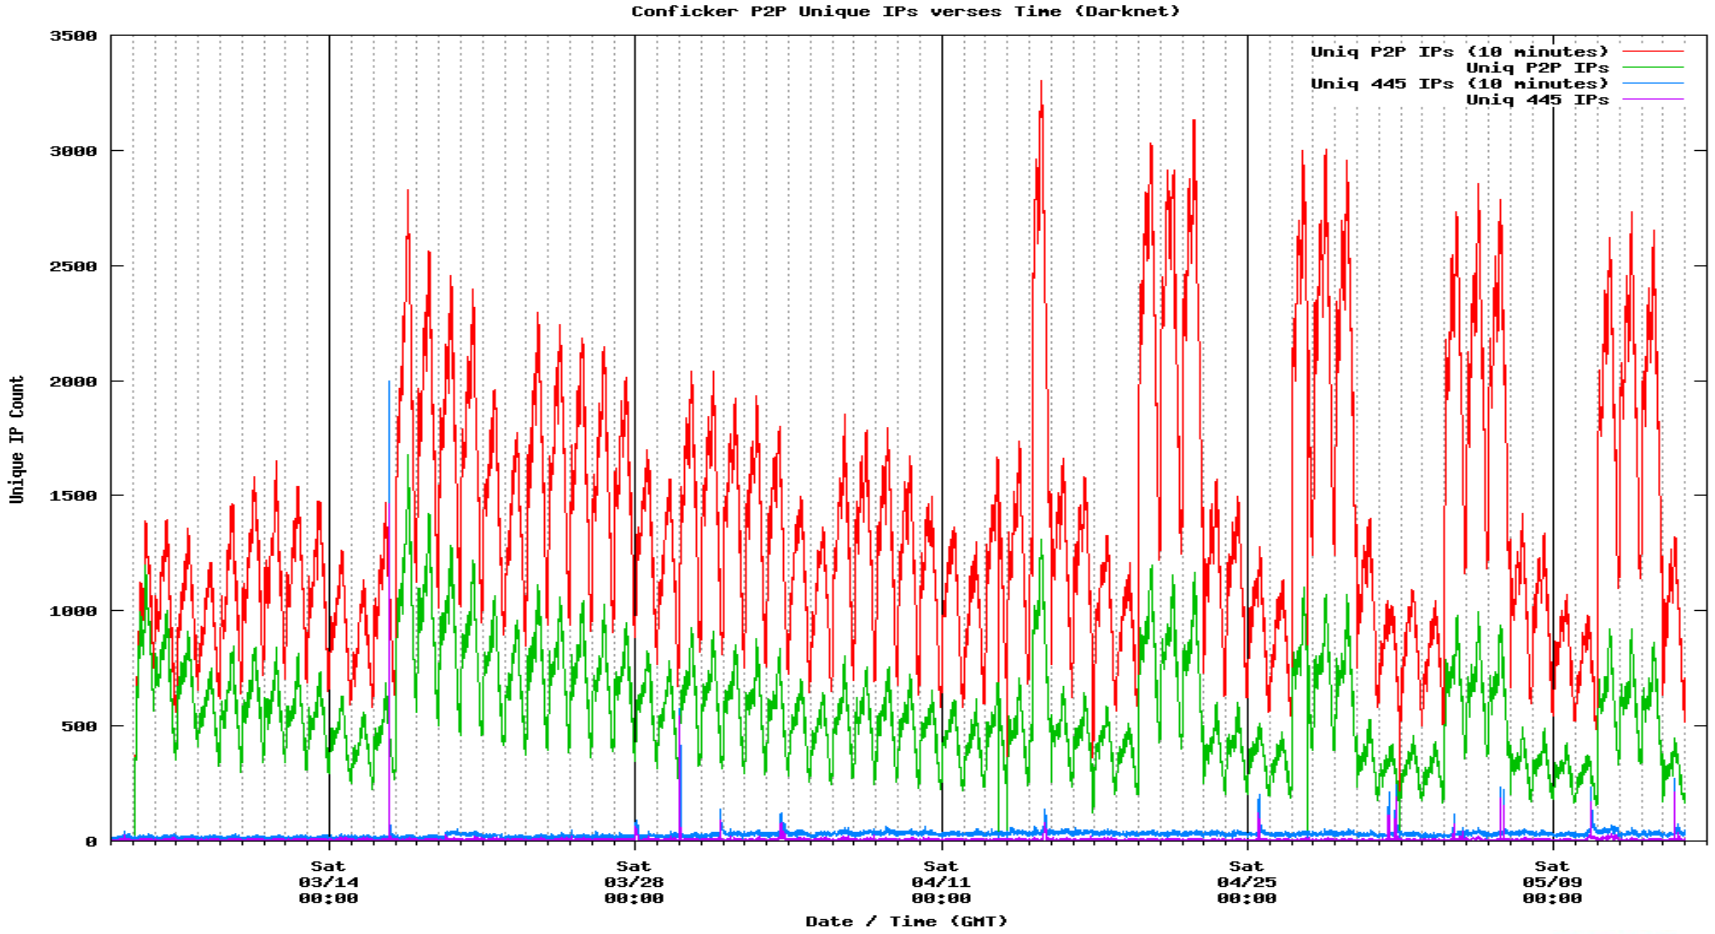

The Domain Name Service

- The domain naming system (DNS) is a fundamental core protocol of the Internet.
- It's mostly UDP based and highly distributed.
- Most of the time it “just works”.
- Organizations rely on it and can be crippled by it.
- IT departments get it working and then are highly reluctant (terrified) of major alterations!
- Many (most?) sites to not adhere to best common practices that have been known for decades!

Managing DNS

- Vast majority of sites do not manage client side DNS at all.
- Unmarshaled, undisciplined outbound DNS is allowed to pass firewalls without monitoring or filtering.
- A very small minority of sites block outbound DNS but they are not any better.
  - They do not alarm on attempts.
  - They cannot evaluate the nature of the activities.

Lurking in the DNS

- Malware beaconing
- Botnet Command and Control
- Data Exfiltration
- DNSChanger style malware
- Covert Channel VPNs
- Advanced Persistent Threats

Malware Beaconing

- Malware Beaconing is just control signaling.
- Malware notifies control sites they are alive.
- Malware receives coded instructions.
- Beacons may be “low and slow”.
- Instructions can be in addresses or text.
- DNS may be the C&C for botnets!
- Malware is increasingly using DNS for control.
- Most beaconing can be detected through simple packet inspection and temporal correlations.

Covert Channel VPNs

- Because DNS is largely unmonitored and unrestricted, it is a prime candidate for covert channel VPN activity.
- OpenVPN works very well over 53/UDP.
- Iodine is a full featured, routed VPN that can even work through DNS caching servers.
- DNSCat works like Netcat only over DNS.
- These can all be readily detected through simple detection yet are not!
- Autocorrelating DNS data can enhance this!

Advanced Persistent Threats

- Advanced Persistent Threats (APT) are not a single type of malware.
- APTs will take advantage of anything available.
- They will use beaconing.
- They will use covert channels.
- They will NOT be spotted by conventional detection.
- They have been spotted through datamining DNS!

Marshaling DNS

- Anomaly detection in DNS depends on managing the baseline.
- Client systems should go through enterprise resolvers and cachers.
- Firewalls should allow established DNS access.
- Firewalls should instrument and monitor all other DNS activities, including packet captures.
- Instrumenting and monitoring unmarshaled DNS does NOT mean merely blocking it!

Filtering vs Instrumenting

- A very small percentage of sites block unmarshaled outbound DNS.
- Sites blocking outbound DNS do little better than unrestricted DNS.
- Most blocking sites ignore blocked traffic.
- Blocking sites cannot evaluate the nature of the traffic.
- Iodine can be detected passing through the firewalls easier than over the DNS servers!
- Covert channels have fallbacks!

False Positives / Negatives

- Some common DNS activities may trigger false positives.
  - Technicians running “host” or “dig”.
  - Engineers with specialized name servers.
  - Individuals needing special forwarders.
- Such activities are valid and should not be prohibited.
- There will always be some false negatives.
  - NOTHING catches EVERYTHING!

Advanced Research

- This remains a work in progress.
- Some areas remain to be explored.
- Correlations against other services and servers.
  - DNS with no correlated other traffic?
  - TCP/UDP/ICMP traffic with no DNS?
- This may qualify other anomalies better.
- Higher false positive rates on their own.
- Has already detected non-security problems.

Conclusion

- The DNS contains a wealth of data to analyze.
  - If managed properly ...
- Correlations on data can improve detection.
  - If we have the data ...
- Anomaly detection is possible and valuable.
- DNS is a vital service for the enterprise.
- IT is highly risk averse for any significant changes.
- These techniques hold much promise.
- How do we get there from here????

Thank you!

Questions?

Feedback?

Answers?

Anomaly Detection through DNS Correlations

Michael H. Warfield
Senior Security Researcher and Threat Analyst
IBM Security Services X-Force

24th Annual FIRST Conference
MALTA
17 - 22 June 2012