Big Data Analytics for Sophisticated Attack Detection

Nikos Virvilis, Oscar Serrano, Luc Dandurand

The number and complexity of cyber-attacks has been increasing steadily in recent years. The major players in today's cyber conflicts are well organized and heavily funded teams with specific goals and objectives, some of which are working under a state umbrella. Adversaries are targeting the communication and information systems of government, military and industrial organizations and are willing to use large amounts of money, time and expertise to reach their goals. It is important to understand the problems and limitations that current technologies face against advanced persistent threats (APT) and the possible benefits that big data analytics could provide.

Since 2006, there have been a large number of advanced, well-orchestrated attacks against industry, military and state infrastructures: In 2006, China was accused of downloading 10 to 20 terabytes of data from the US NIPRNet Military Network. In 2007, Estonia suffered a large-scale cyber-attack that significantly affected the government's online services and financial institutions. In the same year, high-ranking British and US officials publicly accused the Chinese government of penetrating government and business networks in the UK and US. In 2008, at least three US oil companies were targets of cyber-attacks; none of these companies realized the extent of the attacks until they were alerted by the US Federal Bureau of Investigation. Millions of pounds sterling were stolen from British shoppers as multiple chip and pin machines were tampered with via a supply chain attack. In 2010, Google announced that it had suffered a sophisticated attack, named Operation Aurora. This attack affected more than 20 US companies. Google stated that Chinese hackers tried to gain access to specific Gmail accounts and stole intellectual property. In the same year, Stuxnet was detected and classified as the "world's most advanced malware." It was created to target industrial control systems including oil, gas and power industries. In 2011, RSA was attacked and sensitive information for the company's SecurID solution was stolen. This has resulted in further attacks against third-party companies, including Lockheed Martin and other US defence contractors that were also using RSA security solutions. Comodo and DigiNotar certification authorities were also subject to attacks, resulting in the generation of several fraudulent certificates for major companies and organizations. In 2012, another state-of-the-art malware named Flame was discovered, which malware researchers noted as the most complex malware ever created, followed by Red October and, in early 2013, Mini Duke.

While it is believed that these attacks were perpetrated by different threat actors, they share certain common aspects and some of them have been categorized as APTs.

STATE OF THE ART

Almost three decades ago, Butler Lampson described how, in the absence of total isolation, it is impossible to safeguard data from unauthorized access and programs from unauthorized execution. Later, Fred Cohen proved the undecidability of detecting computer viruses, stating that is not possible to build a perfect virus detection tool for a Turing machine. Paul Helman demonstrated that the intrusion detection problem is NP-Hard, which means that it is a decision problem that cannot be resolved in polynomial time in any known way, although it is possible to compute approximations to the solution. Based on these findings, it is assumed that intrusion detection is undecidable and, as such, algorithms are bound to produce a number of decision errors.

The most widely used approach in intrusion detection is signature-based detection, the shortcomings of which have been discussed extensively. A simple testing methodology using known attack patterns revealed substantial limitations in intrusion detection systems while detection rates have been shown to change drastically using small variations of known attack patterns.

From the early days, the research community focused on alternative attack detection methods (behavioral/statistical). Machine learning techniques have been successfully used in certain domains, yet despite the extensive academic research efforts, such systems have had limited success in the field of intrusion detection. Although some of these systems show promising results in controlled environments (detection rates higher than 90 percent and false-positive rates lower than 0.33 percent), they achieve disappointing results in real environments (63 percent detection and 8 percent false positives).

In the last few years, intrusion detection systems (IDSs) have tried to leverage cloud-based approaches to optimize their detection capabilities. Cloud-based approaches allow for the collection of information from a very large pool of users and data analysis at a central point. There has been research on the development of fully distributed IDSs to address the limitations of central and hierarchical systems; however, only small-scale implementations have been realized. Although centralized detection may enable quicker responses to emerging threats (e.g., a new, fast-spreading worm), it offers limited benefits against APT because in targeted attacks the number of infections (usually a handful) is too low to raise an alarm. Similarly, early warning systems (EWSs), which are used extensively for the detection of emerging threats (e.g., worms, botnets), face significant limitations in identifying threats that affect only a very small number of individuals. EWSs also face scalability issues due to the continuous increase of data traffic.

TECHNOLOGY LIMITATIONS

APTs have unique characteristics that significantly differentiate them from traditional attacks:

- APTs make frequent use of zero-day exploits or modify/obfuscate known ones and, thus, are able to evade the majority of signature-based end points and network intrusion detection solutions. In addition, the attacks are generally spread over a wide period of time and, as a result, are often outside the limited detection/correlation window of these systems.
- Attackers focus on a specific target and are willing to spend significant time and explore all possible attack paths until they manage to subvert its defences.

- Based on the analysis of the major APT attacks, it is evident that some perpetrators are supported by nation-states that have significant enabling capabilities (intelligence collection, manufacturing, covert physical access) for cyber-attacks.
- APTs are highly selective. Only a small and carefully selected number of victims are targeted, usually in nontechnical departments of an organization, as they are less likely to identify and report an attack.

Due to these characteristics, current cyber security solutions fail to provide an effective defence against such threats. Following are the main shortcomings of the most common security technologies.

Network/Host-based Intrusion Detection Systems

There are two main detection strategies that are currently used by network or host-based intrusion detection systems (NIDS/HIDS):

1. **Signature-based** is still the most common technique and focuses on the identification of known bad patterns. As with every system that uses a blacklist approach, it is vulnerable to attacks for which the signature is unknown, such as zero-day exploits or use of encoding, packing or obfuscation techniques.
2. **Anomaly-based**, which consists of monitoring system activity to determine whether an observed activity is normal or anomalous, according to a heuristic or statistical analysis, can be used to detect unknown attacks, but despite the significant research effort, such techniques still suffer from a high number of false positives. Furthermore, it is not foolproof, as multiple malware samples use a communication channel that resembles legitimate traffic (e.g. over an SSL/TLS connection) and, thus, can easily evade such systems.

Finally, a major challenge for current IDSs is the limited time window for which the connection state can be maintained. As all modern IDSs are focused on real-time detection, they can only support a short time window (usually a few seconds) in which attacks can be detected for particular TCP sessions. Port scanning is a practical example of this weakness: A quick port scan against a host will trigger an alert from virtually any IDS. However, if this scan is spread over a period of several minutes and, thus, outside of the detection/correlation window of the network intrusion prevention system (NIPS), the attack will pass undetected for the majority of those systems.

Antivirus Products

Antivirus products face the same limitations as NIDS/HIDS, as their detection method is mainly signature-based, supplemented with heuristic analysis. Only a few products offer behavioural analysis. In addition, it is trivial for attackers to test a wide range of antivirus products and modify their malware accordingly to evade detection.

Full Packet Capture

Full packet captures (FPCs) are specialized devices for capturing and archiving network traffic. They are mainly used by analysts to inspect captured traffic after a specific incident. Although they offer the most complete view of the network at any given time, supporting in-depth analysis, FPCs have key shortcomings:

- Limited analysis options are typically provided with the capture system itself, requiring the use of external tools for low-level traffic inspection.
- They offer very limited (if any) integration with other systems (e.g., NIDS/NIPS).

Security Incident and Event Management

Security incident and event management (SIEM) systems collect events from a wide range of sources (e.g., IDS/IPS, antivirus, event logs) and apply statistical correlation to identify potential attacks. The main challenges that such systems face are:

- The limited time window during which these systems will correlate events—usually a few minutes. Events spread over a larger time period will usually not be correlated, and as a result, a carefully orchestrated attack may end up undetected or presented as a series of seemingly unrelated events.
- The correlation is performed centrally and is therefore limited by the available resources.

Overall Assessment

In addition to the weaknesses and shortcomings of the previously mentioned security solutions, perhaps an even more significant factor contributing to the difficulty of detecting APTs is the lack of efficient integration among security solutions. These solutions work as black boxes and tend to only offer (limited) integration if they come from the same vendor. If not, the only possible integration among them is generally through a SIEM system and, thus, suffers from the aforementioned shortcomings. Furthermore, the systems tend to be static, rely on their own (usually proprietary) rules and configuration language, and have their own individual knowledge banks of attack information, with which users must become familiarized.

Also, due to the proprietary nature of these devices and the lack of open standards, an analyst who wishes to write custom rules for detecting specific incidents must do so using a different language for each system (e.g., Snort compatible signature for the NIPS, new correlation rule based on SIEM specific correlation language).

BENEFITS OF BIG DATA ANALYTICS

Taking into account the unique characteristics of APT attacks and the inability of current security solutions to address them effectively, a radical change in the way that such solutions operate is required.

As mentioned, attackers are willing to spread their actions over a wide period of time to evade detection systems. Thus, it is crucial to shift the focus away from real-time detection, which significantly limits the analysis/correlation capabilities. Instead, an approach focused on full-packet capture, deep packet inspection and big data analytics that would enable the use of significantly more advanced algorithms for analysis and correlation, mitigating such evasion attempts, is preferable.

Although offline analysis (analysis of captured traffic) inevitably results in delayed attack detection, it is important to consider that in the majority of APTs, the perpetrators will spend a significant amount of time trying to reach a specific objective (e.g., exfiltrate sensitive data).

There are two main reasons why APT attacks are prolonged:

1. After an initial foothold has been gained, attackers need to explore the network, move across subnets, identify where the information that they are interested in is located and exfiltrate it. As all these steps need to be performed as stealthily as possible to avoid detection, significant time is required.
2. Attackers usually wish to maintain their access and continue to exfiltrate data in the future.

In addition, the correlation of events across large timescales and from multiple sources (e.g., analysis of network traffic, event logs and operating system/application artefacts) is crucial for the detection of sophisticated attacks. Even when attackers manage to successfully evade traditional IDSs, inevitably they generate subtle attack indicators (attack metadata) while exploiting the network. Failed login attempts, increased network traffic from particular host(s), unusual resource utilization and execution of unknown processes can all be correlated and identified as an indication of compromise, even if they are spread over several hours or days. Unfortunately, such indicators are almost always ignored by current IDSs.

Big data analytics focuses on the aforementioned needs and facilitates APT detection by supporting:

- Dynamic and managed collection, consolidation and correlation of data from any number of diverse data sources, such as network traffic, operating system artefacts and event data (e.g., network devices, IDS). This holistic view of the infrastructure enables defenders to correlate sporadic low-severity events as a result of an ongoing attack. In comparison with modern SIEM systems, big data analytics does not have a limited time window based on which correlation can be performed.
- Anomaly detection, based on correlation of recent and historical events. For example, an increased volume of DNS traffic from a particular system for a small time period can be due to legitimate user actions. However, if such a pattern is also identified in historical traffic over a period of days, it is a potential indication of covert data exfiltration. In addition, such correlation can help limit the number of false-positive alerts. Big data analytics solutions increase the quantity and scope of data over which correlation can be performed.

The ability of Big Data Analytics to correlate data from a wide range of data sources across significant time periods will result in a lower false positive rate and allow the APT signal to be detected in the noise of authorized user activities. While processing and correlation does not have to be in real time, it should be completed within an acceptable time window (ideally a few hours) to give the defenders an early warning for potential attacks against their infrastructure.

Further research

Before Big Data Analytics can be used in operational environments for the detection of sophisticated threats, a few obstacles need to be overcome. More specifically, there is need for new detection algorithms, capable of processing significant amounts of data from diverse data sources. As well, there is a need to further progress issues related to the specific problem of malicious activity detection using correlated data sources, such as collecting information from untrustworthy sources, storage and processing performance, time synchronization, meaningful visualization of information, and ensuring the security of sensitive indicators of compromise, amongst others.

Currently, a small number of proof of concept deployments that utilize Big Data Analytics for security event detection exist, and show promising results. We believe that research on this very promising field needs to be intensified, in order to create robust solutions that can address the multidimensional problem of Advanced Persistent Threats.

CONCLUSION

The current industry approach, which is focused on real-time detection with emphasis on signature matching, although effective against traditional attacks, is unable to address the unique characteristics of APTs.

As mentioned, Big Data Analytics currently faces a number of practical limitations and further research is needed for building an operational solution.

Having said that, we believe that Big Data Analytics will significantly enhance the detection capabilities of the defenders, enabling them to detect APT activities that are passing under the radar of traditional security solutions.

AUTHORS

Nikos Virvilis, CISA, CISSP, GPEN, is an information assurance scientist in the Cyber Defence and Assured Information Sharing Division of the NATO Communications and Information Agency in The Netherlands. He focuses his research on advanced persistent threat detection and mitigation. In the past, N. Virvilis has worked as an information assurance consultant/security expert for Encode S.A. and the Hellenic Army.

Oscar Serrano, CISA, CISM, CISSP, has worked for more than 12 years as an IT consultant and researcher for large international companies, including the Austrian Institute of Technology, Siemens and Eurojust. In August 2012, he joined North Atlantic Treaty Organization (NATO) as senior scientist in the field of cyber defence, where he supports NATO efforts to improve cyber defence capabilities of the alliance.

Luc Dandurand works at the NATO Communications and Information Agency where he performs research and development work in cyber defense. He started his career as a signals officer in the Canadian Forces (CF) working as an analyst in the Directorate of Scientific and Technical Intelligence. He then led the CF’s network vulnerability analysis team and founded the team responsible for assessing the security of CF networks by conducting controlled cyber-attacks. Following his departure from the service, he joined the Communication Security Establishment of Canada to lead a team that prototyped novel solutions in cyber defense. He joined the NCI Agency in January 2009.

References

1. Marquand, Robert; Ben Arnoldy; “China Emerges as Leader in Cyberwarfare,” The Christian Science Monitor, 14 September 2007, www.csmonitor.com/2007/0914/p01s01-woap.html
2. Ottis, Rain; “Analysis of the 2007 Cyber Attacks Against Estonia from the Information Warfare Perspective,” Proceedings of the 7th European Conferences on Information Warfare, Plymouth, 2008

3. Brewin, Bob; “U.S., British officials target Chinese as Source of cyberattacks,” Government Executive, 4 December 2007, www.govexec.com/defense/2007/12/us-british-officials-target-chinese-as-source-of-cyberattacks/25874/
4. Clayton, Mark; “US Oil Industry Hit by Cyberattacks: Was China Involved?,” The Christian Science Monitor, 25 January 2010, www.csmonitor.com/USA/2010/0125/US-oil-industry-hit-by-cyberattacks-Was-China-involved
5. Samuel, Henry; “Chip and Pin scam ‘Has Netted Millions From British shoppers’,” The Telegraph, 10 October 2008, www.telegraph.co.uk/news/uknews/law-and-order/3173346/Chip-and-pin-scam-has-netted-millions-from-British-shoppers.html
6. Drummond, David; “A New Approach to China,” Google Blog, 12 January 2010, http://googleblog.blogspot.com/2010/01/new-approach-to-china.html#!/2010/01/new-approach-to-china.html
7. Falliere, Nicolas; Liam O Murchu; Eric Chien; “W32.Stuxnet Dossier,” Symantec Security Response, February 2011, www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf
8. Denning, Dorothy; “Stuxnet: What Has Changed?,” Future Internet, 16 July 2012, p. 72-687
9. RSA FraudAction Research Labs, “Anatomy of an Attack,” RSA Blog, 1 April 2011, http://blogs.rsa.com/rivner/anatomy-of-an-attack/
10. Leavitt, Neal; “Internet Security Under Attack: The Undermining of Digital Certificates,” IEEE Computer, vol. 44, iss. 12, December 2011, p. 17-20
11. Goyal, Ravish; Suren Sharma; Savitri Bevinakoppa; Paul Watters; “Obfuscation of Stuxnet and Flame Malware,” Latest Trends in Applied Informatics and Computing, October 2012
12. Virvilis N., Gritzalis D., “The Big Four - What we did wrong in Advanced Persistent Threat detection?”, in Proc. of the 8th International Conference on Availability, Reliability and Security (ARES-2013), pp. 248-254, IEEE, Germany, September 2013
13. Virvilis N., Gritzalis D., “Trusted Computing vs. Advanced Persistent Threats: Can a defender win this game?”, in Proc. of 10th IEEE International Conference on Autonomic and Trusted Computing (ATC-2013), pp. 396-403, IEEE Press, Italy, December 2013
14. Lampson, Butler W.; “A Note on the Confinement Problem,” Communications of the ACM, vol. 16, 1973
15. Cohen, F. (1987). Computer viruses: theory and experiments. Computers & security, 6(1), 22-35.
16. Helman, Paul; Gunnar Liepins; Wynette Richards; “Foundations of Intrusion Detection [Computer securi-ty],” Proceedings of the Computer Security Foundations Workshop V, 1992
17. Ibid.
18. Denault M.; D. Gritzalis; D. Karagiannis; P. Spirakis; “Intrusion detection: Evaluation and Performance Is-sues of the SECURENET system,” Computers & Security, vol. 13, no. 6, 1994, p. 495-508
19. Spirakis, Paul; Sokratis Katsikas; Dimitris Gritzalis; Francois Allegre; “SECURENET: A Network-oriented Intelligent Intrusion Prevention and Detection System,” Network Security Journal, vol. 1, no. 1, January 1994, p. 22-39
20. Puketza, Nicholas; Kui Zhang; Mandy Chung; Biswanath Mukherjee; Ronald A. Olsson; “A Methodology for Testing Intrusion Detection Systems,” IEEE Transactions on SW Engineering, 1996

21. Chung, Mandy; Nicholas Puketza; Ronald A. Olsson, Biswanath Mukherjee; “Simulating Concurrent Intrusions for Testing Intrusion Detection Systems: Parallelizing Intrusions,” National Information Systems Security Conference, Maryland, 1995
22. Denning, Dorothy; “An Intrusion Detection Model,” IEEE Transactions on Software Engineering, vol. SE-13, iss. 2, 18 September 2006
23. Tsai, Chih-Fong; Yu-Feng Hsu; Chia-Ying Lin; Wei-Yang Lin; “Intrusion Detection by Machine Learning: A Review,” Expert Systems With Applications, vol. 36, iss. 10, December 2009, p. 11,994-12,000
24. Sommer, Robin; Vern Paxson; “Outside the Closed World: On Using Machine Learning,” Proceedings of the 31st IEEE Symposium on Security and Privacy, May 2010
25. Hadziosmanovic, Dina; Lorenzo Simionato; Damiano Bolzoni; Emmanuele Zambon; Sandro Etalle; “N-Gram against the Machine: On the Feasibility of the N-Gram Network Analysis for Binary Protocols,” Pro-ceedings of the 15th International Symposium on Research in Attacks, Intrusions and Defenses (RAID), The Netherland, 2012
26. Snapp, Steven; James Brentano; Gihan V. Dias; Terrance L. Goan; L. Todd Heberlein; Che-Lin Ho; Karl N. Levitt; Biswanath Mukherjee; Stephen E. Smaha; Tim Grance; Daniel M. Teal; Doug Mansur; “DIDS (Dis-tributed Intrusion Detection System)—Motivation, Architecture, and An Early Prototype,” Proceedings of the 14th National Computer Security Conference, 1991
27. Locasto, Michael; Janak J. Parekh; Angelos D. Keromytis; Salvatore J. Stolfo; “Towards Collaborative Se-curity and P2P Intrusion detection,” IEEE Workshop on Information Assurance and Security, June 2005
28. Dash, Denver; Branislav Kveton; John Mark Agosta; Eve Schooler; Jaideep Chandrashekar; Abraham Bachrach; Alex Newman; “When Gossip Is Good: Distributed Probabilistic Inference for detection of Slow Network Intrusions,” Proceedings of the 21st National Conference on Artificial Intelligence, 2006
29. Kijewski, Piotr; “ARAKIS—An Early Warning and Attack Identification System,” 16th Annual FIRST Conference, Budapest, June 2004
30. Op Cit, Falliere
31. Op Cit, Goyal
32. Op Cit, Falliere
33. Op Cit, Goyal
34. Op Cit, Drummond
35. Op Cit, RSA FraudAction Research Labs
36. Big Data Working Group. "Big Data Analytics for Security Intelligence." Cloud Security Alliance (2013).