White paper
Cyber Threat Laboratory

Angelos K. Marnerides[1], Daniel Prince[1], John Couzins[2], Ryan Mills[1], Vasileios Giotsas[1], Paul McEvatt[3], David Markham[3], Catherine Irvine[3]

[1]*Security Lancaster, School of Computing & Communications, Lancaster University, UK*
[2]*Information System Services, Lancaster University, UK*
[3]*FUJITSU Enterprise & Cyber Security, UK*

With the growth of the Internet and billions of connected devices, opportunities are greater than ever for sophisticated and targeted attacks. Gaining an understanding of how attackers operate is crucial in the prevention of further compromise and exploitation. The Cyber Threat Laboratory is an environment built in partnership with Lancaster University which grants unparalleled insight into the emerging threat landscape as it unfolds.

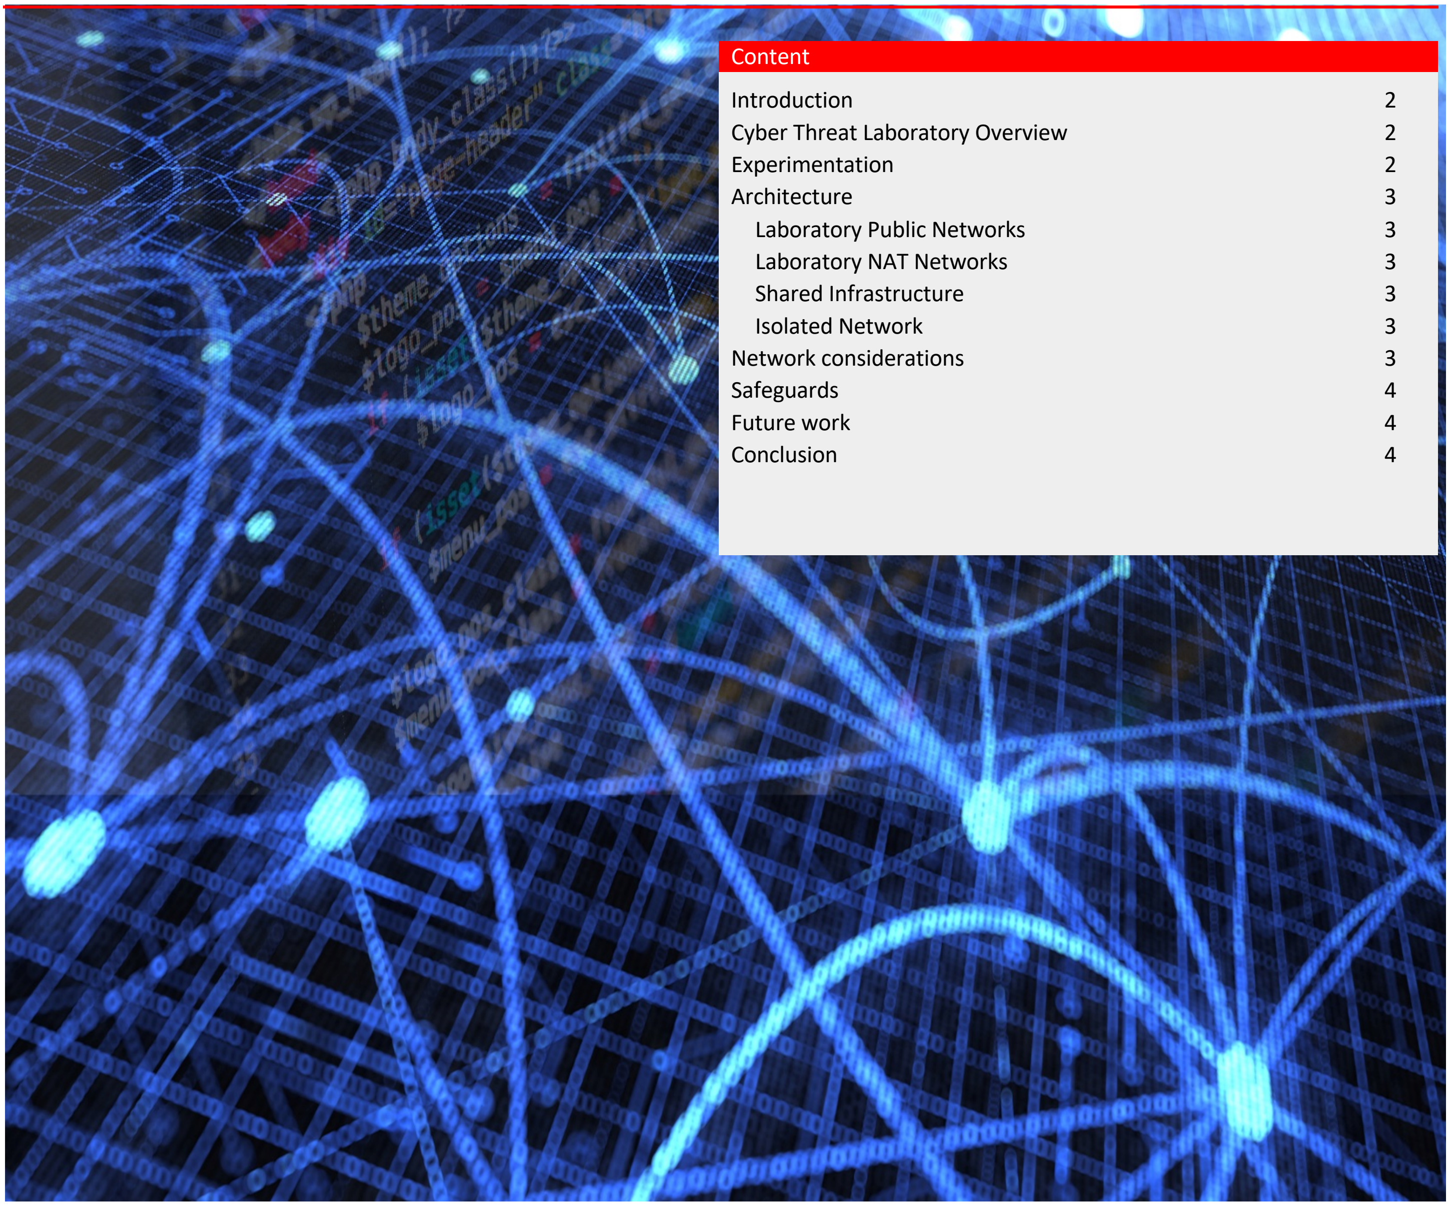

Content

Introduction

Recent trends within the cyber security spectrum indicate an increasing number of threats which employ complex exploits capable of delivering harm to systems on a wide scale. These attacks are non-trivial in nature and aim to siphon proprietary information for monetary or geopolitical gain. It is essential for organisations to have measures in place which protect critical assets or risk financial and reputational damage.

An intimate understanding of how sophisticated attacks operate play an instrumental role in combatting this evolving threat landscape, and by monitoring the communication patterns used by malicious sources this can be achieved. Attracting certain attack campaigns by masquerading as weakly protected services has recently proven fruitful in gaining this data. The Cyber Threat Laboratory enables the collection of these sources of intelligence, allowing rapid attribution of attack campaigns using diverse data sources.

Cyber Threat Laboratory Overview

Fujitsu Enterprise & Cyber Security and Lancaster University have partnered to create a flexible research facility that enables threat analysis and identification of sophisticated attacks to be tested. The Cyber Threat Laboratory hosted at Lancaster University is designed to provide a collaborative platform that allows analysis of threats and behaviour to take place, in a safe and controlled environment. The laboratory provides centralised infrastructure enabling multiple projects and experiments operating simultaneously inside the lab to benefit from mature industry standard tools.

Comparable to any research with unknown volatile outcomes, experiments into cyber threats and malware also needs to be handled in a controlled environment with appropriate safeguards and equipment. Lancaster University with industry input from Fujitsu is now able to offer this to its users though the Cyber Threat Laboratory.

The laboratory consists of multiple inter-connected components which provide a framework for projects to analyse vast amounts of malicious data garnered by myriad sources. This data proves fruitful in understanding the manner in which attackers operate and will ultimately be used in the prevention of such attacks. The method in which the data is retrieved and how this is dissected is dependent upon each experiment's procedure. This allows a wide range of possible techniques which benefit from a shared infrastructure that promotes collaboration.

Various experiments have been conducted in the Cyber Threat Laboratory and many more have been offered as part of academic research such as undergraduate, postgraduate and doctoral thesis proposals. Not only does this foster interest in the security field for young students, but it also benefits the laboratory as a whole by contributing additional data to consider. Active experiments inside the Cyber Threat Laboratory employ various means of capturing raw data relating to campaigns spearheaded by malicious actors. Currently the laboratory hosts honeypots with varying levels of interaction which masquerade as legitimate services exposed to the Internet using Lancaster University's IP space assigned to the lab. These honeypots are provisioned to collect data from all sources attempting communication, therefore, analysis of this data must be undertaken in order to uncover knowledge of any intrusions and reconnaissance attempted.

Experimentation

In an attempt to promote collaboration and further analysis of results, each experiment has access to a shared infrastructure which hosts a range of valuable services. Furthermore, the shared infrastructure is capable of hosting large amounts of data and is primarily used for storing the logs gathered by each experiment, thus reducing the hardware requirements of each. Granting this open access to projects also enables another powerful aspect: the correlation of data from a wide array of sources. This further highlight instances of malicious activity, thus aiding the production of countermeasures using multiple technologies conducted by a diverse set of research members.

The technologies used to assess certain attack scenarios differ depending upon circumstance, however, facilitating the transfer and visualisation of logs via embedded protocols enables each experiment to rapidly analyse rich data sets in a secure environment. The widely acclaimed Elastic Stack was used for this purpose, and currently has a growing number of events shipped from the experiment sensors. Each experiment located within the Cyber Threat Laboratory captures different data as the sensors are located in separate IP space and are configured based upon the researcher's goals. Due to this, the Elastic Stack records data in collections based upon the sensor type, ensuring the encapsulation of data. Laboratory members are also granted access to each collection enabling correlation across data sets.

The sensors, once exposed to the Internet through the networks offered by the Cyber Threat Laboratory, are implemented in different manners and have the ability to masquerade as any legitimate service. Currently active implementations within the lab include sensors acting as popular SSH and Telnet services, tempting malicious actors with the chance to infiltrate University infrastructure. By ensuring these services have weak authentication credentials, the laboratory has captured multiple successful login requests which allow an attacker the ability to further compromise the network. Once authentication is successful an emulated shell session is delivered to the attacker which records all activity within the machine. This is especially useful in the creation of countermeasures for sophisticated attacks as details about the entire attack process are uncovered. While recording reconnaissance activity is useful in attribution of certain conventional (e.g., Mariposa) but also IoT-based large-scale botnets (e.g., Mirai-alike), system level activity administered by malicious actors grants further understanding about how exploits are leveraged to gain increased foothold into infrastructure as well as the type of malware employed by certain campaigns.

In parallel, the virtualization of various experimental testbeds residing within Lancaster's Security Institute is an on-going activity to enable the composition of enriched and dynamic IoT honeypots. Through this activity the development of macroscopic tracking and early detection of large-scale IoT-based botnets is currently in progress. Hence, the Macroscopic Analysis of IoT-based Intrusions (MATI) project has recently started with some interesting initial findings in the context of detecting industrial control systems that are infected with Mirai-alike malware. By correlating data gathered at the Cyber Threat Laboratory with external sources of information (e.g., BGP, Censys scans) we can now identify networks in which devices that manage energy or water utilities are already infected and are members of a greater Mirai-like botnet army.

Architecture

The laboratory is divided into two primary logical zones, a green zone that allows management and entry into the environment and a red zone which hosts the machines and devices where analysis of threats takes place.

The laboratory provides controlled experiments of varying risk levels to take place through a number of segmented networks. These networks employ differing access to local and internet services and are further described below.

Laboratory Public Networks

Public networks are designed to emulate servers and devices that are directly publicly accessible over the internet. This allows analysis of scanning and attacks that are typically seen across webservers and networks.

Laboratory NAT Networks

These networks are designed to emulate personal, enterprise or organisation networks. Internal hosts are all permitted outbound access through a shared public address.

Shared Infrastructure

Hosts on the shared infrastructure network allow other networks located in other red zone networks to utilise central infrastructure such as logging, authentication, name resolution, etc. This reduces the complexity required when setting up experiments and allows collaboration with other users and projects taking place in the environment.

Isolated Network

These networks are used to monitor activity that do not require any external connectivity outside of the lab or have a significant risk of negatively impacting external sources outside of the university. This is especially useful for the analysis of malware garnered from monitoring the experiments in the other networks.

Network considerations

University IP space can often be subject to specialised attacks that provide researchers with a broad range of data. Traffic from most business networks will consist of similar predictable flows, such as connections to online CRM systems, online storage and some personal employee traffic. University networks contain a very broad spread of activity, from large amounts of research traffic from CERN to shops and residential activity.

Academic networks have been at the centre of a number of large attacks over the last five years that have forced Higher Education providers and supporting parties to undertake significant investment into their infrastructure.

Laboratory Benefits

The combination of industry standard technologies and novel architecture design provides users of the lab imperative benefits such as:

- Central placement and overview of all cyber research
- Centralised, mature monitoring
- Considered and tested safeguards
- Real world traffic

Designed with flexibility and security in mind, the Cyber Threat Lab incorporates best practices to encapsulate various experiments within their corresponding network space, whilst allowing access to shared infrastructure components.

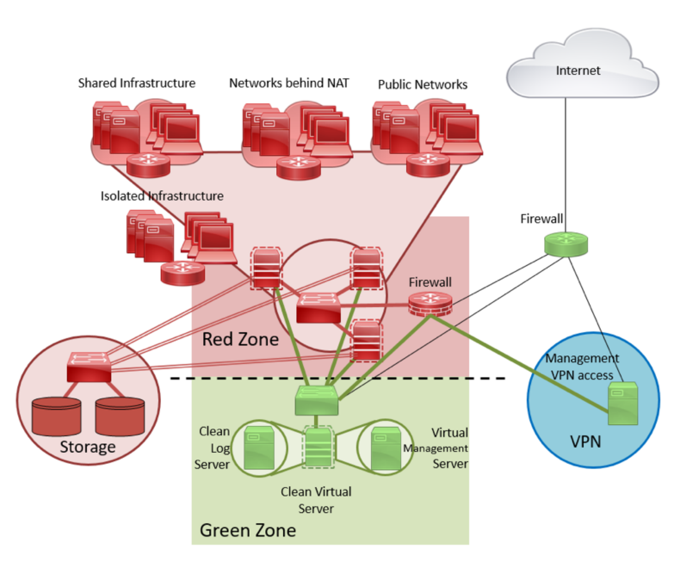

Safeguards

As the network operates inside the university network there is a number of technical and procedural controls that have been implemented to reduce the associated risks. Traffic traversesing the external network from the zones inside the lab is monitored and controlled by the Information Systems Services Team and the Security Lancaster Institute.

Projects are also subject to peer review, risk assessment and require additional sign off to ensure the experiments are appropriate and are placed into the correct locations of the network.

Future work

The Cyber Threat Laboratory is dependent upon research members conducting innovative experiments to gather novel data relating to emerging threat actors. This data is intrinsic in the creation of countermeasures which provide a safeguard against sophisticated attacks that would otherwise wreak havoc among enterprise systems. This perpetual back and forth battle between malware authors and defense architects means that efforts in prevention have to be constant and rigorous, considering all available data. Due to this, it would be greatly beneficial to collaborate with other similar systems to share and correlate between data sets to find previously unseen trends and patterns. Furthermore, the experiments taking place within the Cyber Threat Laboratory must evolve to acknowledge the most recent and impactful threats. Therefore, the Cyber Threat Laboratory will act as the facilitator for a range of cyber threat intelligence frameworks that can adjust over particular stakeholder scenarios. Hence, a diversity of threat scenarios exists within a range of deployments ranging from the IoT up to industrial control systems, critical infrastructures and smart energy systems.

An on-going developing approach in the prevention of crippling attacks on electronic systems is the notion of Active Defence. This technique crucially places sensors within an enterprise network which typically have no purpose other than alerting when they have been interacted with. This interaction suggests that there exists a possible intrusion taking place, and procedures can therefore be put in place, based upon knowledge derived from monitoring, in an attempt to neutralize the threat. The Cyber Threat Laboratory is an ideal candidate for the implementation of such system as it provides a wealth of information garnered from sensors scattered amongst the network.

As the Cyber Threat Laboratory reaches maturity, the technologies embedded within may be deemed ineffective in the monitoring and analysis of the most recent threats encountered. As the threats evolve, so must the way in which they are examined. To combat the most sophisticated attacks which are targeted are highly coordinated in nature, this evolution must also consider the architecture of the network. These threats which incorporate lateral movement between internal services to gain further foothold into infrastructure must have a realistic network topology to encourage malicious actors to utilize exploits and infiltrate multiple times. Due to this, multiple services which are physically located within the same network segment should be deployed and emulate a realistic enterprise network.

Conclusion

This white paper presents the design of the Cyber Threat Laboratory and provides information relating to the experiments which reside therein. The Cyber Threat Laboratory has indexed a vast amount of data emanating from malicious sources and through the efforts of its researchers, there has been substantial progress in the identification of threats which target valuable infrastructure. Through the Cyber Threat Laboratory it was feasible to setup long-term research trajectories in the context of cyber threat intelligence, early detection and attribution of APTs as well as tracking of large-scale IoT-based botnets. Apart from the composition of dynamic honeypots for The Cyber Threat Laboratory monitoring and measurement infrastructure has managed to facilitate the composition of the MATI project, the first project to correlate diverse sets of datasets such as to monitor and identify “on the fly” infected industrial control system devices on a global scale.

In general, the instrumentation and technical objectives within the Cyber Threat Laboratory already indicate pioneering research paths and continued synergy of UK enterprise and academia. As such, it can adequately meet the needs of the UK enterprise sector as well as to establish the UK in general as a world-leader in cyber threat intelligence.

Contact
FUJITSU
Enterprise & Cyber Security
E-mail: askfujitsu@uk.fujitsu.com (Ref:3902)
Website: www.fujitsu.com/emeia/services/security/
2019-07-10