Protecting the Future Enterprise: Active Cyber Defense

by Chris Daly, CISSP

Today's cyber castles are crumbling. Will you champion our defense?

Find out why…

the future of cyber security relies on smarter active cyber defenses with predictive and proactive strategies to stand up against tomorrow's cyber attackers – and what you can do to get started today.

active cyber

www.activecyber.net

PROTECTING THE FUTURE ENTERPRISE: ACTIVE CYBER DEFENSE

The Definitive Guide To Next-Gen Cyber Protections

by Chris Daly, Founder of ActiveCyber.net

Why the Current Generation of Cyber Tools Is Failing and What You Can Do About It

The future of our digital world is in serious peril from an ever-present, ever- evolving cyber enemy. Most of today's static and reactive security mechanisms are crumbling under its rapid, invasive onslaught. But there is hope. Who do ask? You.

In this book, *you* will learn how to use the latest active cyber defenses (ACD) to:

- **Transform your cyber defense strategy to proactively** disrupt and dismantle the attacker's kill chain and more effectively defeat threats both today and tomorrow.
- **Anticipate with predictive analytics** to foresee and forestall your adversaries' next moves.
- **Bolster defenses with context-awareness** so that they can **adapt** dynamically to counter adversarial tactics and deflect new attacks while meeting mission goals.

Ultimately, we must evolve and advance our protections quickly to a new era of **active cyber defense…** or face the consequences of inaction and be left behind in ruins.

TABLE OF CONTENTS

PROTECTING THE FUTURE ENTERPRISE: ACTIVE CYBER DEFENSE

INTRODUCTION

Enterprises face an onslaught of cyber attacks every day. All too often, some of these attacks are successful despite best intentions and even best practices by the enterprises that are impacted by these attacks. Successful attacks are becoming more costly as multiple studies have pointed out. Recent data breaches like Equifax and OPM have impacted hundreds of millions of Americans. There are many reasons for the failures to detect and mitigate these attacks in a timely fashion. In general, **the cyber adversary has "upped their game" while enterprises still scramble to find the right balance of cyber protections** to meet their risk equations. Often, business and government enterprises, large and small, **fail to effectively invest** in the **proactive, predictive, and adaptive cyber protections** that can meet the ever-changing threat landscape of today's cyber adversary. These proactive, predictive, and adaptive protections are collectively known as ***active cyber defenses***.

Current Cyber Defenses Don't Hack It

Generally, cyber experts recommend a multi-layer set of defenses to combat cyber threats. This defense-in-depth best practice approach can be enhanced by continuous monitoring which improves situational awareness of cyber posture and reduces the time to patch systems. However, **defense-in-depth and continuous monitoring** are not adequately addressing vulnerability abatement and threat mitigation for today's IT enterprises. For example, several research reports point out the rising success of phishing and stolen credential attacks as evidence of the shortcomings of the cyber protection regimens in use today. These attacks account for about half of the successful breaches in every year since 2013 as per the *Verizon Data Breach Investigations Report* and are the top threats since 2013 according to the *Microsoft Malware Protection Center*. According to these reports, attackers utilize sophisticated targeting approaches to identify victims. They leverage deceptive tactics and a range of advanced exploits to bypass current defenses and gain access to victims' systems. They use stealthy botnet systems to exfiltrate sensitive data.

The problems are even worse for industrial control systems, Internet-connected consumer devices, medical devices, and other embedded systems being deployed as part of the **Internet of Things**. Many of these systems are unmanaged and unmanageable; however, they are vulnerable to cyber attack. Vendors of these systems claim lack of financial or regulatory incentives to secure them, while, at the same time, system operators choose to connect these systems to the Internet for convenience and business reasons. The result is wide open critical infrastructure systems ripe for cyber attack, as pointed out recently in this *article by SC Magazine UK*. Broadband consumer routers are another prime example of devices targeted by attackers due to their poor, unmanaged security. For example, in March, the *security consultancy Team Cymru* warned that hackers had compromised some 300,000 small- and home-office broadband routers made by firms D-Link, Micronet, Tenda, and TP-Link, among others. As another example, *recent research into medical devices* reveal Bluetooth-enabled defibrillators can be manipulated to deliver random shocks to a patient's heart or prevent a medically needed shock from occurring; and, digital medical records can be altered to cause physicians to misdiagnose, prescribe the wrong drugs or administer unwarranted care. As a *final example*, researchers have shown that it is trivial to hack a traffic light system, thereby causing chaos across an entire municipality at rush hour. These findings

show that cyber attacks have the potential to go beyond financial losses or intellectual property theft to having a much more significant impact on the everyday lives of everyone.

It is apparent from these reports that the current regimens of mitigation and risk management are not as effective as one would hope in stemming the tidal wave of cyber attacks. This situation is particularly disconcerting given the significant investment already made in cyber defenses and the apparent ease by which attackers continue to exploit vulnerabilities in these defenses. In fact, the October 2013 *AFCEA Cyber Committee Report - The Economics of Cybersecurity: A Practical Framework for Cybersecurity Investment* points out that:

> *The fact that the majority of reported damaging attacks have little sophistication leads to the logical question of why the $15 billion annual investment in cyber security protection by United States Organizations is not more effective in successfully blocking these unsophisticated attacks.*

And it's not that cyber defenses are about to turn the corner any time soon. According to the Verizon report: *The trend lines … plainly show that attackers are getting better/faster at what they do at a higher rate than defenders are improving their trade*.

This observation is supported by a recent *FireEye report*, which states that attackers have found a winning formula to evade detection in a number of Advanced Persistent Threat (APT) campaigns in which attack attributes are changed at a faster rate than intrusion detection systems and other defenses can keep up. *"… the threat actors have continuously tweaked the malware by changing its hardcoded strings, remote access commands, and encryption keys*," FireEye said in the report. To put it simply, cyber attackers employ blitzkrieg tactics against a Maginot line of cyber defenses.

The Cyber Adversary Is Incentivized And Organized

Hackers are able to outpace cyber defenses due to a variety of **sophisticated underground ecosystems** as pointed out by the recent Rand report – *Markets for Cybercrime Tools and Stolen Data*, and other sources. The Rand report reveals that cyber criminals trade in many marketplaces that are quite efficient and profitable, using various digital currencies to monetize and trade in exploits, tools, and gains. These markets are also very resilient. Although law enforcement has shown some success in dismantling some cybercrime markets, new ones appear almost instantly, and new advances are being made to create a fully peer-to-peer market system with no central authority for law enforcement to attack.

The Cyber Defender Is Underfunded And Siloed

The cyber defender is under constant downward budgetary pressure. Budget cuts for IT spending have created havoc on the stability of IT service management efforts – efforts which are foundational to security best practices. The most relevant of these IT service management practices: asset and configuration management, identity services, and service operations – are often hit the hardest when budgets start to go south. Data center consolidation and cloud initiatives also create dramatic changes in the landscape of enterprise IT, making it difficult for operators to keep track of assets and secure.

The **responsibilities for cyber defense are often split** among various IT and business units. This division of responsibility, while important from separation of duties and a separation of concerns perspectives also creates silos of activity when trying to maintain security posture and respond to incidents. Like the tale of the blind mice and the elephant, each unit only sees a part of the overall security picture leaving it to hit-or-miss coordination efforts and slow-moving configuration control boards to reconcile differences and organize efforts.

The Cyber Adversary Is Agile And Sophisticated

While the cyber defender is mired in bureaucracy, the cyber attacker is busy crafting **sophisticated APT attacks**. The creators of APTs have the skills and funding to continually evolve their tactics and malware to avoid detection. These technically savvy perpetrators use a variety of sophisticated TTPs such as advanced encryption, anonymizers, steganography, polymorphic and metamorphic exploit code, and zero-day exploits to penetrate the most vigilant defenses and move about undetected.

The Cyber Defender Is Overwhelmed And Reactive

Meanwhile, the tsunami of cyber events that must be adjudicated each day is a key issue facing security operations professionals that are at the forefront in combating cyber attacks. Cyber security analysts must make sense of complex, ambiguous information and must spread their time and attention across multiple sources of information - each source with differing pedigree and distinct format – to identify the nuggets relevant to their particular environment. Some auto-generated alerts do assign an initial priority, often displayed as low/medium/high or a value 1-100. Unfortunately, the incoming priority value is typically fixed per originating sensor and thereby adds little differentiating value. Incoming alerts rarely contain all information necessary for an analyst to judge severity of the threat. Therefore, judging priorities of alerts or events is left to the creativity, knowledge, and experience of the analyst.

According to a *Damballa analysis of Q1 2014 traffic*, the average North American enterprise fields between 10k-20k alerts each day from its security systems, far more than their IT teams can possibly process. Mostly, the Security Operations Centers (SOCs) and Network Operations Centers (NOCs) are barely treading water in dealing with this avalanche of events, resulting in **missing critical indicators** of compromise. For example, *the security firm Mandiant reported* recently that while the average time it took to detect breaches declined slightly from 2012 to 2013, from 243 to 229 days (more than seven months), the number of firms that detected their own breaches actually dropped, from 37% to 33%. A *Ponemon study* from 2015 reviewed the cost of malware containment. Figure 1 is an extract from this report and it depicts the time needed to mitigate advanced malware attacks was a staggering 70 person days. These numbers are comparable to the findings of other research firms for the same period.

Figure 1 –
Estimated Average Hours to Contain Advanced Malware
(Ponemon Institute, 2015)
Extrapolated Average is 587 hours

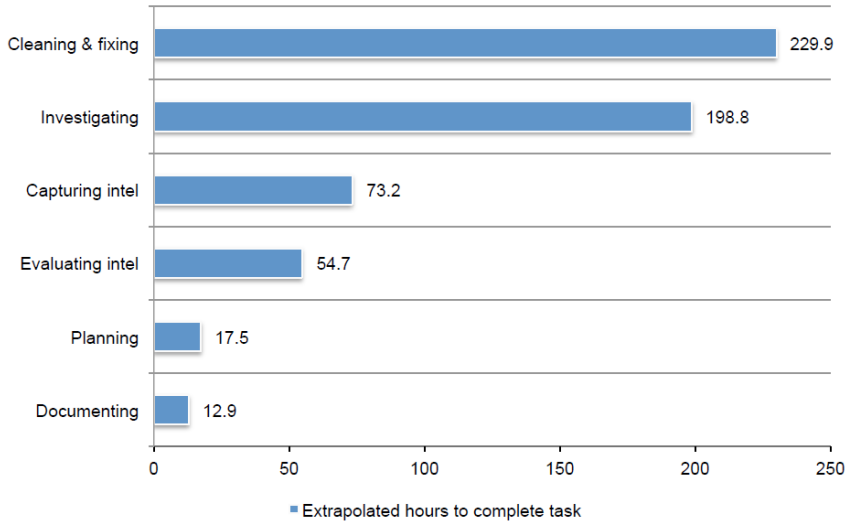

The Adversary Can Choose Targets From A Wide Attack Surface

Adversaries benefit from considerable asymmetric advantages in cyberspace since a single vulnerability may enable widespread compromises. And vulnerabilities are easily identified and obtained through the underground hacker marketplaces or by simply scanning on the Internet. Cyber attackers are known for their ability to quickly reverse engineer patches when they are first released to identify the vulnerabilities for which the patches are intended. Next, they rapidly develop exploits for the vulnerabilities and hunt for targets that are still unpatched. Since many patches may not be installed when they are first released, the cyber attackers' agility enables a fertile hunting ground for exploitation.

Cyber attackers don't need to limit their targeting to just **unpatched systems**. Misconfigured systems, systems with factory default passwords, and unprotected communications are a few examples of low hanging fruit that are easily compromised. For example, Dan Farmer pointed out recently in his report, *Sold Down the River*, that poorly configured and designed interfaces such as IPMI - Intelligent Platform Management Interface – has left significant holes in enterprise infrastructures. Farmer concludes in his report that 90 percent of modern servers could be compromised because of default or weak passwords or weaknesses in the IPMI protocol.

High value, well-protected targets have also shown to be vulnerable through indirect attacks through supply chain partners, or by drive-by attacks and spear phishing attacks. To summarize, cyber adversaries enjoy a never-ending supply of vulnerabilities and targets on which to feast.

The Defender Must Protect A More Complex Baseline Of Assets

Technology churn is an important issue affecting the ability of SOCs and NOCs to turn the corner on the cyber security problem. Cloud computing and mobile computing/Bring Your Own Device (BYOD) are two trends that are at the forefront of this technology churn. These initiatives change the security paradigm, creating a **perimeterless environment** as content and applications move outside of the traditional enterprise boundaries. These changes ripple throughout the entire enterprise, creating significant changes to other impacted systems as well as to traditional risk assessment and mitigation models, tools, and policies.

Security defense strategies are also complicated by the myriad of tools offered by vendors and implemented by users. The tools are used to identify assets, scan for vulnerabilities, patch, remediate, protect the network, manage access control, encrypt data, log events, and much more. Each tool covers some fragment of the overall plan of security protections and must be integrated with other tools, processes, dashboards, and systems to be effective. This abundance of tools increases the complexity of managing the IT enterprise and the corresponding security posture.

As *noted recently by Gene Spafford of CERIAS at Purdue University*, poorly coded software combined with growing network complexity has increased the attack surface at many organizations. This has resulted in "… using all these [security] tools on a regular basis because the underlying software isn't trustworthy." All too often these security tools do not integrate well with each other and only cover a subset of the assets that need to be protected. These coverage gaps result in loss of situational awareness, difficulties in defending against attacks, and slow response processes.

Current Cyber Defense Approach = Game Over

Many enterprises are still mired in a static/reactive protection model whose main focus is on prevention of attacks. As the above evidence suggests, a prevention strategy eventually fails. A new "**proactive**" strategy is needed to detect and contain threats before they can seriously impact enterprise operations. The strategy must

enable early disruption of the attack chain and provide prioritized and automated responses to critical threats to gain positive control over attack consequences. Enterprises need to better predict attacks through reducing the attack surface (which funnels attackers to highly tuned sensors), through deception (to gather information about attackers' TTPs), and by using analytics (to provide early and accurate recognition of anomalous and malicious behaviors). Enterprises need to align defenses to these predicted attack vectors, and remove the cloak of stealth that attackers currently enjoy. The way ahead requires better cognitive systems for threat recognition and attribution, smart systems-of-systems interactions, and autonomous security capabilities. The future of cybersecurity is active cyber defense.

Preemptive Versus Active Cyber

Often there is the misconception that active cyber defense equates to preemptive cyber. The terms "anticipatory self-defense," "preemptive self-defense," and "preemption" traditionally refer to a state's right to strike first in self-defense when faced with imminent attack. According to *Dever and Dever*,

> *The term "anticipatory self-defense," in the context of international law and* jus ad bellum, *is commonly defined as a nation's ability to foresee the consequences of a given threat and to take proactive measures aimed at preventing those consequences. Anticipatory self defense is, accordingly, distinguished from armed reprisal in that the former is protective while the latter is retributive. Some legal scholars, moreover, employ a further temporal analysis to differentiate between anticipatory self-defense and preemptive action.*

In essence, **preemptive cyber self-defense** refers to offensive cyber operations that are initiated based on the criteria outlined in the *Caroline test*. The *Caroline* test has two distinct requirements:

1. The use of force must be necessary because the threat is imminent and thus pursuing peaceful alternatives is not an option (necessity)
2. The response must be proportionate to the threat (proportionality).

The Caroline test originates from a 19th century incident between the United States and Canada known as the *Caroline Affair*; and, to this day, is considered the customary law standard in determining the legitimacy of self-defense action.

Are we nearing the point when a **preemptive cyber strike** is needed? It seems there has been an escalation of state-sponsored cyber attacks against other government entities, as evidenced by the recent OPM clearance database breaches. *Cyber-policies* were put into place in 2013 as a way for the military and U.S. intelligence agencies to deploy cyber weapons against other nations. Some officials have stated that such policies would be put in effect only if an attack might resemble a cyber 9/11. However, given the nature of cyber attacks, it might be difficult to anticipate an attack of this sort. Attribution of an imminent threat could also be an issue.

So is active cyber defense that same as preemptive cyber? The answer is no. Active cyber defenses focus mainly on preventing and containing an on-going attack. Do active cyber defenses help in the preparations leading to a preemptive cyber strike? The answer is decidedly yes. Active cyber defenses leverage cyber intelligence to help predict attacks, capture attackers' Tactics, Techniques and Procedures (TTPs) to support intelligence collection and to help provide attribution, and provide possible vectors to initiate cyber operations through deception and stealthy offensive payloads that could be downloaded by the attacker.

Learn more about active cyber defenses in the following sections.

WHAT IS ACTIVE CYBER DEFENSE?

Active cyber defenses (ACD) go beyond merely best practices. ACD transforms the underlying security infrastructure from a static, fixed, and reactive model to an **agile and proactive capability**. This active defense capability is enabled by several components:

- The rapid fusion of **cyber intelligence** from threat information sources **coupled with vulnerability information, attack status, and asset state** from other sensors.
- This fused threat and sensor data is mashed and mined by **a cyber analytics engine** that outputs **actionable intelligence**.
- **Actionable intelligence** is pumped to the brain of an active cyber defense capability - a cyber command and control **(C3) system** that directs cyber protections and oversees mitigation actions in **real-time**.
- The C3 system orchestrates a **highly automated workflow** and choreographs an intelligent network of security capabilities to constantly sense and respond to security events.
- **Adaptive endpoint defenses** are also enabled through a combination of trust mechanisms and autonomous security capabilities.
- **Deception capabilities** are used in the intelligent network and at endpoints to fine-tune sensors, deflect attacks, and to inform threat information sources about the tactics of cyber attackers.

This ensemble of components, referred to as *active cyber defenses*, forms a **dynamic security umbrella** that can be readily adapted to the cyber adversary's tactics.

Active cyber defenses are organized in a continuous feedback loop called an "***OODA loop,***" shown in Figure 2.

Figure 2 – OODA Loop

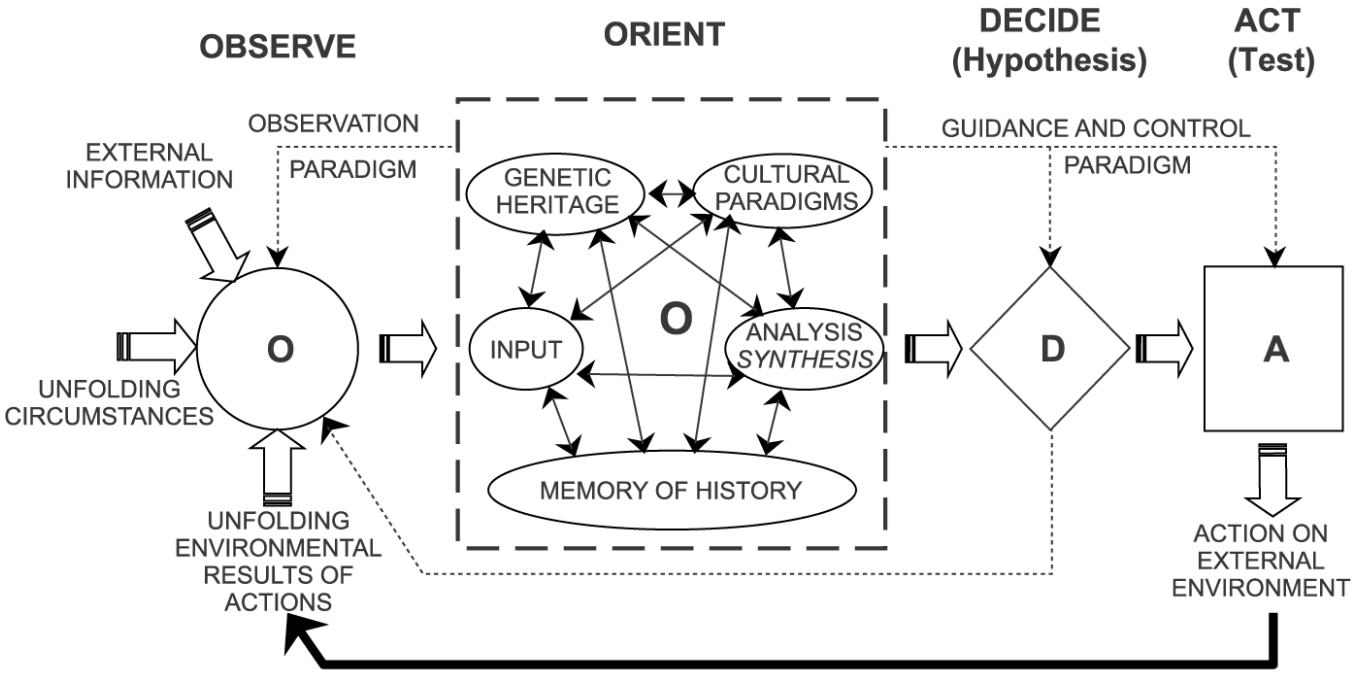

Source: Based on Boyd (1996)

The Active Cyber Defense OODA loop consists of the following capabilities:

1. ***Observe*:** Produces scalable and rapid situational awareness based on coordinated threat intelligence to predict attacks; enables reporting of attack data and information regarding an endpoint's state or status; captures and correlates sensor events along with vulnerability and compliance data to positively identify and quickly pinpoint critical security problems; leverages dynamic methods to deceive attackers while capturing attackers' Tactics, Techniques, and Procedures (TTPs).
2. ***Orient*:** Applies real-time diagnostics to classify, and add context to cyber events; organizes cyber event data and applies analytics to evaluate defensive posture, assesses attack consequences on ability to execute mission, and appraises risk at any point in time.
3. ***Decide*:** Formulates courses of action (COAs) – i.e., adaptive and integrated responses designed to quickly mitigate incidents, to prevent attacks, to patch vulnerabilities, or to contain damage while meeting mission assurance goals; prioritizes and synthesizes COAs using a rule-based decision framework; publishes COAs to control points.
4. ***Act*:** Directs "cyber maneuvers" to reduce the attack surface and respond to attacks; orchestrates and executes COAs across policy decision and enforcement points; monitors failure or success of directed action.

The desired outcome from ACD is a semi-autonomous system that augments the cognitive ability of defenders to detect and stop attacks while minimizing manual intervention in the process. Human cognition of cyber events is facilitated through visual interactive analytics of Big Data.

ACD is built upon the building blocks of "defense-in-depth" and "continuous monitoring." ACD enhances the effectiveness of these building blocks with the following portfolio of capabilities**:**

1. ***Cyber Intelligence and real-time threat awareness*:** ACD entails collection and analysis of indicators of compromise (IOCs) from threat intelligence sources. The IOCs are used to equip active cyber defenses with actionable intelligence needed to disrupt the attack chain.
2. ***Intelligent, software-defined networks*:** Virtualized network security functions can improve cyber defensive posture by fine-tuning protection policies for specific virtual workloads and operational contexts. Or, in other words, virtualized active defenses enable cyber maneuver (i.e., protections move if the workload migrates to another host) which accelerate the ability to respond to cyber threats.
3. ***Orchestration of workflows for incident response and mitigation*:** Contextually-based, security automation of workflows helps to accelerate time-to-respond and rapid execution of appropriate responses. The workflows must have a strong trust foundation for assuring the communications among all valid enterprise and inter-enterprise entities (users, devices, and applications) so that workflows can execute unimpeded from manual intervention.
4. ***Stealthy detective mechanisms and deceptive techniques at endpoints, and delay and counter tactics at the network level:*** Sophisticated cyber adversaries utilize surreptitious tactics and difficult-to-detect malware to penetrate defenses. Active defenses take a page from the adversary's playbook through employing mechanisms designed to deceive and delay attackers into revealing their tactics before damage is incurred, and provide counter-responses designed to disrupt the attack chain.
5. ***Agile cloud computing:*** The software stack flexibility and auto-scaling provided by cloud computing creates an agile platform for deflecting cyber attacks. The on-demand computing resources and big data

analytics offered by the cloud help to provide deep insight into cyber situations and accelerate the OODA loop for mitigation actions.

6. ***Adaptive endpoints*:** Active cyber defenses are all about leveraging context-aware adaptations – that is, the endpoint's protections or detections can morph based on an awareness of the security state of the endpoint and / or an assessment of the threat environment to which the endpoint is exposed. These context-aware adaptations can be used to provide autonomous data protection, or for dismantling malware before it can do harm.

NIST provides a good summary of the overall characteristics of active cyber defense in its description of requirements for an "Adaptive Implementation Tier" in its February 2014-released *Framework for Improving Critical Infrastructure Cybersecurity*, which includes:

> *The organization adapts its cybersecurity practices based on lessons learned and predictive indicators derived from previous and current cybersecurity activities. Through a process of continuous improvement incorporating advanced cybersecurity technologies and practices, the organization actively adapts to a changing cybersecurity landscape and responds to evolving and sophisticated threats in a timely manner.*

The following sections dive deeper into each of these adaptive and predictive capability areas for active cyber defense.

CYBER INTELLIGENCE & REAL-TIME THREAT AWARENESS

Intelligence derives from the Latin verb *intelligere*, to comprehend or perceive. *Merriam-Webster includes a definition for intelligence* of "*secret information that a government collects about an enemy or possible enemy.*" Each of these meanings is relevant to the practice of cyber intelligence. That is, **cyber intelligence** deals with the ability to observe and collect information about a cyber event or cyber adversary that is of interest and previously *unknown*, and then using that information to guide some type of thought or action.

There are three types of "unknowns:"

1. ***Discoverable unknowns*** — unknowns that you don't know about, but somebody else might know about and have disclosed - such as previously detected malware.
2. ***Unexpected unknowns*** — information about something that is unknown to exist, cannot be accurately predicted, and is manifested with minimum warning or no warning at all, examples include an earthquake or a zero day attack.
3. ***Unanticipated consequences*** — this type of unknown stems from a lack of comprehension or understanding of the full effects of an action (or inaction), and enterprises may be the perpetrator, victim, or beneficiary of the unintended results.

Discoverable unknowns are discoverable if you make the effort to discover them. The way to do this is to leverage cyber intelligence services and to scan your network with a variety of active and passive sensors. Unexpected unknowns can be perceived through sensors that detect anomalous events, such as through behavioral-based continuous monitoring approaches. Unanticipated consequences may actually be the hardest to discover since the impact may not be felt by the perpetrator of the action and there may not be any interest in trying to discover the causes and effects. Overall, efforts to discover unanticipated consequences are best directed towards continuous monitoring approaches as well.

Knowing Threats Strategically, Operationally & Tactically

Enterprises should take a **holistic view** of the "discoverable" and "unexpected" cyber threat landscape and how their assets are impacted by threats. Such a holistic view should accommodate 3 levels of cyber intelligence**:** strategic, tactical, and operational. Each level participates in an OODA loop — ACD is about accelerating this loop at each level. Cyber intelligence maps to the observe and orient phases of the OODA loop. Cyber intelligence, therefore, contributes to the ability to make better decisions.

At the *strategic level*, cyber intelligence contributes to a better understanding of the overall threat context, resulting in (hopefully) setting better policies, priorities, and creating better budgets for protection capabilities. At the *tactical level*, the key consideration for cyber intelligence is to inform COAs based on gaining detailed knowledge of adversary TTPs. This detailed knowledge includes indicators of compromise (IOCs) and attack warnings. At the *operational level*, cyber intelligence informs security operators about the current security state and activity of endpoints, resulting in identifying anomalous behavior on the network and at endpoints. By knowing the current and desired security states of endpoints, along with IOCs from the tactical level, security operators are equipped with the actionable intelligence needed to execute appropriate COAs.

Actionable intelligence at the operational level includes indicators and warnings sufficiently detailed to:

- derive IDS signatures,
- support rule changes to firewalls, application level gateways and proxies, and IPSs,
- provide updates to routing tables and VLAN settings,
- add or subtract entries in blacklists/whitelists,
- identify vulnerable endpoint patch levels, configurations and integrity states,
- modify access control and data loss prevention rules,
- enable updates to AV signatures and other filters,
- identify already infected hosts,
- change logging parameters, (e.g. when privileges are escalated, abnormal behavior is detected, or compliance levels are unmet),
- restrict or throttle network access and bandwidth,
- redirect network flow.

The overall purpose of actionable intelligence is to detect (or even predict) an attack attempt and disrupt the kill chain of the attack.

Let's Talk Cyber Football

Football provides a good analogy regarding the relationships between each level of cyber intelligence. The roles of the General Manager, Coach, and Players of the team align well with the three levels of cyber intelligence.

Strategic intelligence encompasses the General Manager (GM) role. The Football GM is responsible for "getting the right players on the field." This responsibility involves **identifying gaps in the capabilities** of the team to execute wins. The gaps may entail the coaching staff, player talent, or knowledge of the capabilities and strategies of adversary teams' and their players. Scouring scouting reports for "intel" on players, coaches, and adversary teams is a key task of the Football GM. Like the Football GM, the Cyber GM uses available cyber intelligence to have an understanding of the general classes of threats and risks that an enterprise faces. The Cyber GM should understand how black markets for cybercrime work - if you know what cyber criminals want, you can start to prevent them from getting it. Knowing what they want can be determined by examining the types and values of cyber commodities that are traded and sold in these markets, and by examining the targets

of emerging malware campaigns. By understanding how their assets are valued in the cybercrime markets, and if the type of asset an enterprise possesses is a target of the cyber adversary, the Cyber GM can begin to classify the threats to their assets and define priorities for protection.

At the tactical or "football coach's" level, there is a concerted focus on "**scouting**" the adversary team's tendencies, specifically what advantages the adversary may have and how they operate — their tactics, techniques, and procedures (TTPs). This tactical intel is essential to preparing the game plan — if you know how the other team calls plays, you can better defend against them. For the Cyber Coach, such a game plan should include:

- threat-based assessments such as red teams,
- subscriptions to multiple threat intelligence information services,
- a tested incident response plan.

In football, the game plan is often **adjusted as the situation changes**. Active cyber defense embraces this same concept as new tactical cyber intelligence is obtained. Indicators of compromise and other warnings are available from a variety of online threat intelligence sources that may tip off the targets and TTPs of the adversary. This intel informs the tactical C3 system which, in turn, may orchestrate workflows to:

- develop new abuse cases,
- reassess the attack resistance of the architecture,
- formulate new observations and event correlations,
- update incident response plans,
- alter mitigation capabilities,
- assess mission options to minimize impacts from the threat.

The operational level is analogous to the play-by-play on the field. As in football, the defensive players (the cybersecurity operations specialists and security automation tools) continuously evaluate their own defensive posture relative to the "offense's" (or adversary) setup to **assess what possible vulnerabilities** the offense may intend to exploit. Tactics also involve disguising their setup to draw out the offense's intent, making defensive adjustments, and making plans to "close the gap" quickly as the play unfolds.

Operational cyber intelligence provides near real-time information through local sensors to determine if an attack is imminent or is already underway. Operational intel informs the active defense C3 system which, in turn, dynamically adjusts defenses based on attack patterns or indicators, the state of assets, and initiates COAs to contain the damage caused by an attack. The COAs are executed by intelligent network capabilities with oversight by security operations specialists.

Standards Are Needed To Accelerate Sharing Cyber Intelligence

Cyber intelligence is most useful if its accuracy and relevancy can be quickly ascertained and it can be translated into action at network speeds. Therefore, it is critical that cyber intelligence gathering, sharing and distribution are accomplished seamlessly, and applied rapidly within the OODA loops at each level. However, this is harder to do than it sounds as pointed out in *a recent report by TM Forum, Managing Data, How to Combat Cyber Threats*:

> *Current threat information exchange systems and methods are anything but high speed in operation.* The report goes on to say that a lack of standards-based tools and support of common data models cause: *a disconnect between security intelligence reporting systems and the activation of security response tools*.

Despite the findings expressed in the TM Forum report, there are initiatives underway by DHS, NIST, MITRE, and industry partners to create data model standards for sharing threat data at all levels, and to provide

specifications to automate the conversion of threat data into actionable intelligence. These standards and specifications include:

- ***STIX™*** — *Structured Threat Information Expression:* an XML standard to automate the sharing of threat intelligence.
- ***TAXII™*** — *Trusted Automated Exchange of Indicator Information*: a set of technical specifications that enable organizations to exchange and securely transport cyber threat information represented as STIX.
- ***CAPEC™*** — *Common Attack Pattern Enumeration and Classification:* a publicly available, community-developed list of common attack patterns along with a comprehensive schema and classification taxonomy.
- ***MAEC™*** — *Malware Attribute Enumeration and Characterization:* a standardized language for encoding and communicating high-fidelity information about malware based upon attributes such as behaviors, artifacts, and attack patterns.
- ***CybOX™*** — *Cyber Observable eXpression:* a standardized language for encoding and communicating information about cyber observables.

STIX, CAPEC, and MAEC each use CybOX to describe *cyber observables*. A **cyber observable** is a measurable event or stateful property in the cyber domain. Examples of measurable events include: a registry key is created, a file is deleted, an HTTP GET is received. Examples of stateful properties include: the value of a registry key, the MD5 hash of a file, the existence of a mutex. The CybOX schema is natively imported and used within STIX, CAPEC, and MAEC to characterize system and network events and behaviors observed within the operational cyber domain, or to use within cyber indicators and patterns, or within **observable malware attributes and patterns**. For example, STIX describes threats using 8 basic constructs that are standardized for threat information sharing, including:

- Threat Actor — Who was doing it?
- Campaign — Why they are doing it?
- TTP — What exactly they were doing?
- Exploit Target — What they were looking for?
- Incident — Where it was seen?
- Courses of Action — What you should do about it?
- Indicator — Why you should care about it?
- Observable — What you are looking for?

The following table lists a sampling of some of the current *use cases targeted by CybOX*.

Table 1 – Use Cases for CybOX

Supported Use Case	Relevant Process	Domain Specific Standard
Analyze event data from diverse set of sensors of different types and different vendors	Event Management	CybOX
Detect malicious activity utilizing attack patterns	Attack Detection	CAPEC
Detect malicious activity utilizing malware behavior characterizations	Attack Detection	MAEC
Enable automated attack detection signature rule generation	Attack Detection	CybOX, MAEC, CAPEC, STIX
Characterize malicious activity utilizing attack patterns	Incident Response/Management	CAPEC, STIX

Supported Use Case	Relevant Process	Domain Specific Standard
Identify new attack patterns	Threat Characterization	CAPEC
Characterize malware behavior	Malware Analysis	MAEC
Empower and guide incident management utilizing attack patterns and malware characterizations	Incident Response/Management	STIX, CAPEC, MAEC, CybOX
Enable consistent, useful and automation-capable incident alerts	Incident Response/Management	STIX, MAEC, CAPEC
Enable automatic application of mitigations specified in attack patterns	Incident Response/Management	STIX
Enable incident information sharing	Incident Response/Management	STIX
Enable explicit and implicit sharing controls for cyber observable information	Information Sharing	STIX, CybOX, TAXII

One example of how STIX/TAXII is implemented is a 2014 DHS project known as "Cyber Flare." This project demonstrated how open standards can facilitate cross-domain threat information sharing. Another implementation example is a product known as Soltra Edge which incorporates native back-end support for all STIX Observable types, including Actors and Campaigns and full TAXII support for upload/ download and discovery. *Microsoft's Interflow Threat Exchange system* is also built using STIX and TAXII standards.

The use of cyber observables extends beyond cyber threat intelligence to cover other cyber measurable events and object states. One example is continuous monitoring tools that focus on detecting and validating observables that reflect the vulnerability posture and patching status of endpoints. Another example is network access control (NAC) tools that monitor the security state of endpoints that are requesting access to enterprise resources.

Sharing Standards Must Be Backed Up By Trust And Authority

Data standards for cyber intelligence increase the accuracy and speed of sharing threat data within an enterprise and between enterprises. Implementing **standards** is facilitated through trusted communities or **circles of trust** such as Information Sharing and Analysis Centers (ISAC) where charters and agreements help to establish the foundational protocols for trusted sharing. For example, FS-ISAC, the ISAC for the Financial Services community, is implementing STIX using Soltra Edge's Avalanche technology as the standard for threat information sharing. In addition to the financial sector, ISACs are in place to support cyber threat exchange for the retail, IT, medical, state and local government, education, utility, and industrial control sectors to name a few.

Facebook announced in February 2015 ThreatExchange, described as an API-based clearinghouse for security threat information. It's really a social platform, something Facebook naturally excels at building, which allows companies to share with each other details about malware and phishing attacks. Pinterest, Tumblr, Twitter, and Yahoo participated in ThreatExchange and gave feedback as Facebook was developing it. Bitly and Dropbox are new contributors, bringing the initial participant list to seven major tech companies (including Facebook). ThreatExchange is built on Facebook's existing platform infrastructure, with layered APIs on top for partner companies to query available threat information and publish to participating organizations. Facebook says early feedback pushed for a platform that lets organizations be more selective about the information they share via a defined set of data types.

The new Cyber Threat Intelligence Integration Center (CTIIC) was also announced in February 2015 and will have the job of integrating intelligence from various federal agencies such as the Central Intelligence Agency and the National Security Agency, and distributing information more broadly to other federal agencies. The creation of the new office suggests a recognition by the administration that the status quo of cyber intelligence sharing isn't working , as multiple federal agencies were involved in monitoring cyber threats, but no one agency played a lead role.

Most perplexing is trying to reconcile what this new center will contribute to the current public and private information-sharing regime. The intelligence agencies have monitored cyber threats for years through the auspices of the NSA/CSS Threat Operations Center (NTOC) which acts as the focal point for mission discovery of cyber threats, characterization and attribution of those threats, creation and sharing of situational awareness, and the development of mitigation strategies.

Perhaps CTIIC is an answer to the increasing array of federal, local and private information sharing centers by providing the ability to fuse and mine threat data across an increasing number of cyber threat sources. Like its cousin, the National Counter Terrorism Center which focuses on mining of terrorist-related information, CTIIC could provide a method to distill classified cyber threat data for distribution at lower classification levels to critical infrastructure providers and other governmental users. CTIIC will draw on the current legislative framework for its authority for sharing. This approach has been questioned by some observers who say that the current authorities fall short in gaining the participation needed by private industry to ensure comprehensive threat information sharing.

Threat information is shared with private industry through different channels such as the Department of Homeland Security's (DHS) National Cybersecurity and Communications Integration Center (NCCIC). NCCIC is a bridge between government, private sector, and international network defense communities. NCCIC houses US-CERT (United States Computer Emergency Readiness Team), ECS (Enhanced Cybersecurity Services) and ICS-CERT (Industrial Control Systems Cyber Emergency Response Team). These teams and programs provide *threat sharing platforms based on STIX/TAXII* and the *Traffic Light Protocol (TLP)* to ensure that sensitive information is shared with the correct audience. TLP employs four colors to indicate different degrees of sensitivity and the corresponding sharing considerations to be applied by the recipient(s). Figure 3 describes the **TLP.**

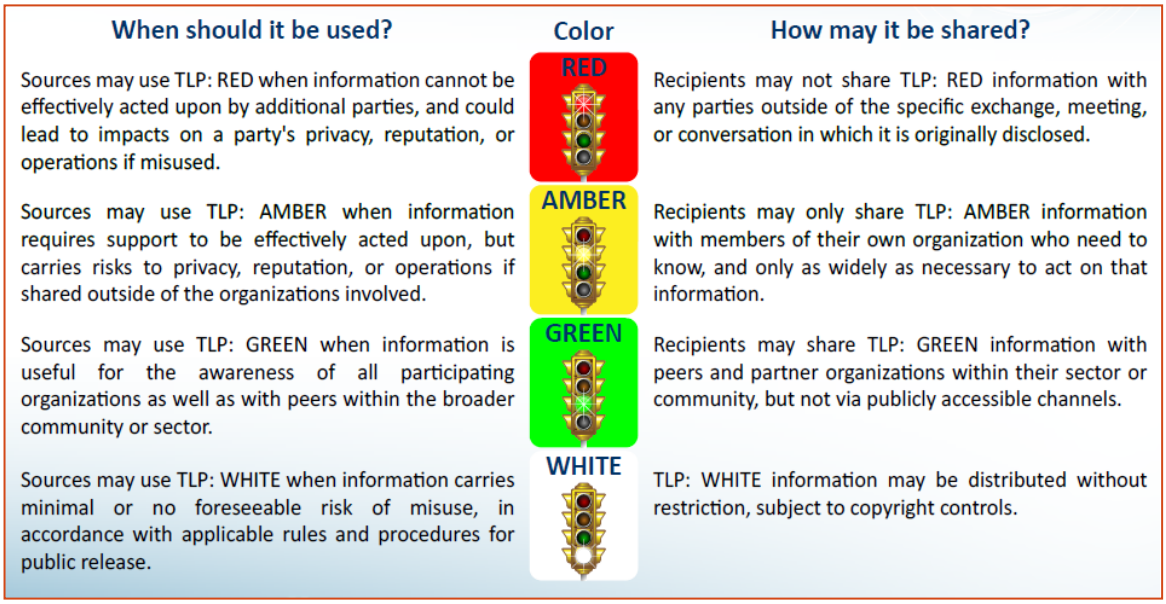

Figure 3 – Traffic Light Protocol (Courtesy of US-CERT)

And there's more—the DHS programs complement DOD programs like the *Defense Industrial Base Collaborative Information Sharing Environment (DCSIE)*,where defense contractors share computer security information between themselves and with the government.

There is also the FBI's National Cyber Investigative Joint Task Force (NCIJTF) which is the focal point for all government agencies to coordinate, integrate, and share information related to all domestic cyber threat investigations. The FBI is responsible for developing and supporting the joint task force, which includes 19 intelligence agencies and law enforcement, working side by side to identify key players and schemes. Its goal is to predict and prevent what's on the horizon and to pursue the enterprises behind cyber attacks. **NCIJTF has also adopted STIX** as the primary standard for sharing cyber observables. Finally, there are the different ISACs mentioned above.

Solving the challenges to rapid, complete, and accurate threat information sharing is only one step towards an effective ACD program. Other challenges remain, such as dealing with cyber information overload and deploying an agile infrastructure that can quickly maneuver to mitigate the threat. These challenges are addressed by other ACD capability areas, such as automated orchestration, agile cloud computing, adaptive endpoints, and intelligent networks.

INTELLIGENT NETWORKS

Today's modern network designs are changing rapidly. Bigger and faster pipes are being deployed at the core transport layer while intelligence is being pushed to the edge network where users connect and content is distributed. The edge network is changing from wired to wireless as mobile devices overtake handsets and desktops as the platform of choice for communication and computing. The network infrastructure is also changing from predominantly **physical appliances to virtual appliances**. Firewalls, wireless controllers, load balancers, switches and routers, are becoming virtualized software network functions. This transformation from physical to virtual allows the infrastructure to easily adjust in both functionality and scale. This new "software-defined network" (SDN) can sense and respond dynamically to changes or events in the network.

Active Cyber Defense And SDN – A Perfect Synergy

The capability of SDN for dynamic network adaptation is a key enabler from an ACD perspective. Virtualized network functions enable a logical network that is independent of the physical network location or state. Virtual workloads can migrate across this logical network without requiring any reworking of security policies, load balancers, etc. New virtual workloads or logical networks do not require (re)-provisioning of the physical network. There is also the benefit of redundancy — nodes in the physical network can fail without any disruption to the virtual workload, but any failures in the virtual layer do not propagate to the physical layer. This type of agile and resilient network allows network engineers and administrators to respond quickly to changing mission requirements and security events.

How SDN Works

In SDN environments, the control plane understands the network schematic and is therefore able to configure the network in response to specified commands. The control plane creates an **overlay** — a temporary logical network put in place to address a demand, a situation, or a security response— without having to touch the physical underlay. The control plane also choreographs the virtual network functions enabling them to be scaled up or down.

The forwarding plane of the SDN environment comprises virtual network elements (vRouters) that carry network user data. vRouters are responsible for forwarding packets from one virtual machine to other virtual machines via a set of server-to-server tunnels. The tunnels form an overlay network sitting on top of a physical IP-over-Ethernet network.

The logical overlay layer defines a virtual network through *service chains* — sets of logical switches, logical routers, logical firewalls, logical load balancers, logical VPNs and more — which are connected to the actual virtual workloads. By creating an abstract layer of virtualized network services, implemented over an SDN control plane, an administrator can easily change a network component's rules when necessary. This dynamic level of control allows an administrator to re-prioritize or even block specific types of packets with a very granular level of precision.

The SDN controller works in conjunction with the Virtual Machine (VM) orchestrator to create logical networks linked to virtual workloads. The logical networks utilize the underlying physical network as a simple packet forwarding backplane. Virtualized network and security services are distributed and attached to workload VMs within the network. As a VM is moved to another host, these services stay attached to the workload VM and move with it. As new VMs are added to a network, policy can be dynamically applied to the new VMs.

Network Virtualization Is On The Rise

Network virtualization is becoming recognized as an important technology for data center and cloud efficiency. Network virtualization is designed to create **virtual networks** within a virtualized infrastructure, which makes the network much more portable and scalable. The physical devices are simply responsible for the forwarding of packets, while the intelligence of the network is delivered by software. The decoupling of the control and forwarding planes delivers superior operational efficiencies and reduces costs, due to hardware independence. In general, a virtualized network can offer all the features and guarantees that a physical network could offer, only with greater agility and flexibility.

These capabilities make network virtualization positioned to integrate and support the increasingly virtual data center environments the network is being asked to connect. Large service providers such as AT&T and Verizon are already embracing this technology, as they try to make their networks more responsive and agile to customer demands. Network virtualization solutions are also being used in the WAN and as part of cloud offerings. Figure 4 highlights the major providers of network virtualization solutions and their adoption levels.

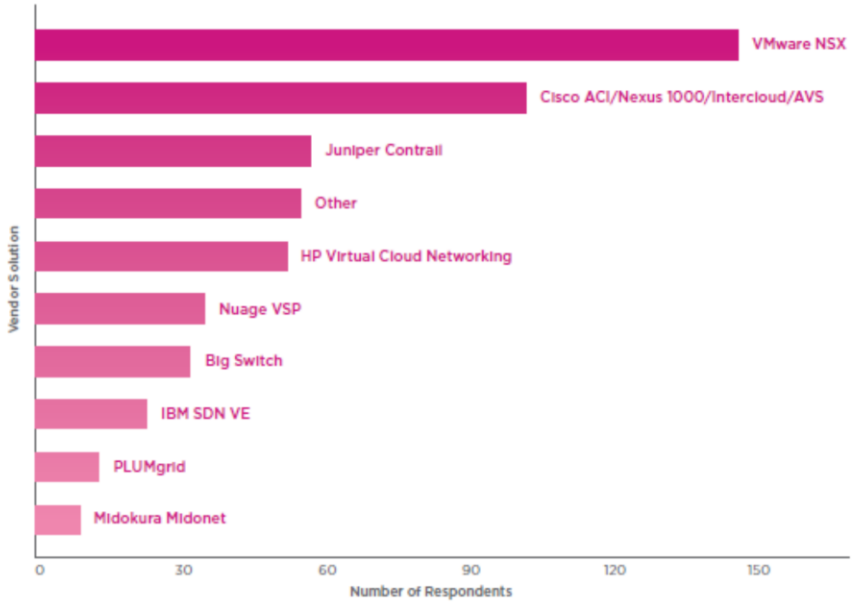

Figure 4 – Network Virtualization Adoption and Providers (*Courtesy of SDNCentral)*

According to SDN Central's 2014 Market Report on Network Virtualization (NV):

> *41% of Service Providers have deployed NV solutions; of those that have not, 85% have plans to deploy NV solutions in the next two years (43% are on an accelerated schedule, looking to deploy within the next 3–12 months). 94% of Cloud Service Providers already have or are planning on deploying NV solutions in their environment within the year. Large enterprises are split – half have already deployed NV solutions, while half have not. 78% of the large enterprises that do not have NV solutions have plans to adopt NV technologies within the next 3–12 months (44%) or 12–24 months (56%). Within the enterprise, size seems to matter, as small to medium businesses are slower to adopt NV technologies.*

As shown in Figure 4, network virtualization has given rise to several different vendors' offerings, each working on a slightly different approach to deliver what they see as the optimal network virtualization offerings. Scalability and ease of implementation lead the key requirements of customers, followed by cloud management platform support and performance. Interestingly, support for L4-7 features and specific hypervisor support trail in requirements. Hyperconverged systems are also leveraging network virtualization support.

Applying SDN To Active Cyber Defense

Dynamic service chaining can be applied in different ways to enable active cyber defenses. For example, service chaining allows for fine-tuning of protection parameters based on specific virtual workloads and traffic flow patterns. Network operators can develop a behavior-based view of network activity, learning and anticipating the "normal" traffic conditions for time-of-day network needs. Security policies can be fine-tuned to these normal behavior profiles and COAs can be developed to possible changes before they occur. Then, when anomalies, security events, or changes in resource utilization actually occur, the SDN controller can sense and analyze these changes and allocate dynamic countermeasures based on service chained security functions. If a failure or a compromise does occur, the network can deploy courses of action consisting of virtualized, service-chained network and security functions to quickly self-heal or contain the threat, determine the root cause, and implement mitigations.

Let's look at a specific threat scenario to see how this dynamic service chaining capability can be applied. Consider that botnets and malware are not only being hosted in the cloud, but also being controlled remotely from cloud servers. The goal of hackers is to disguise their malicious software as regular traffic between corporate end points and cloud-based services. Security operators can pinpoint this type of malicious traffic by funneling traffic into a set of virtual overlay networks consisting of service chains of virtual load balancers, deep packet inspection capabilities, and various traffic filters that are specifically trained to detect select types of abnormal network/application/user behavior.

In effect, the combination of virtualized network security functions and SDN enable *cyber maneuver* where virtual workloads can migrate to different servers or even different data centers, virtual networks can be instantiated with relative ease, and network security protections can be assembled and deployed on the fly. The configurations of networks, hosts and applications can be dynamically modified in a manner that is undetectable and unpredictable by an adversary but still manageable for network administrators.

Cyber maneuver can limit the exposure of vulnerabilities and opportunities for attack by deploying mechanisms that **continually shift and change** over time. This could be accomplished by IP address hopping *w*here a client and server operating in an virtual overlay network share a secret that is synchronized to randomly select (and frequently change) an IP address at the client and the server respectively. (Note: generally, IP address hopping is most effective in an IPv6 environment due to the larger range of addresses available). This dynamic maneuvering capability increases complexity and cost for attackers, while increasing system resiliency to DDoS attacks. Maneuvering should only take place within a contiguous security zone (e.g., the DMZ) and not maneuver from one security zone to another. If maneuvering across security zones is allowed, attacks could be

transferred from one zone to another, which might open up vulnerabilities for the attacker to exploit that were not previously accessible.

New Agile Network Protocols Enable Adaptive Responses

Research is also underway to explore agile, mobile, and secure network protocols. One example is the *telehash protocol*, which enables any app or device to establish private communication channels over a network. The following combination of features in telehash offers distinctive benefits from an ACD perspective:

- all channels are encrypted all the time; there is no unencrypted mode
- because each application instance or device generates its own public/private key pair, they cannot be impersonated and security is not dependent on trust in certificate authorities
- addresses are generated from public key fingerprints, not centrally managed as with IP addresses
- routing is based on a globally distributed hash table (DHT), with no central authority or hierarchy
- channels can be reliable (like TCP) or unreliable (like UDP), and make use of HTTP.

Although telehash apps can run over the current Internet, bindings to Bluetooth, IEEE 802.15.4, and other low-layer transports are also on the way. This ability to run protected without central authority offers agility to applications and devices, especially for users who interact frequently outside the enterprise boundaries with external partners. Telehash could also be well-suited for communications between IoT devices.

AUTOMATED ORCHESTRATION

Many attacks go undetected for months. However, once a compromise is detected, it currently takes most organizations days or even weeks from the detection of a breach to deploy mitigation responses. The overriding causes for delays can be traced to the need by incident managers to coordinate responses across a myriad of roles, locations, and processes. This coordination effort often occurs in different change control boards, which slows the enterprise's ability to respond quickly to attacks.

Automated security orchestration tools provide the ability to automate an organization's **response workflow** (i.e., Course of Action or COA). These tools make possible a controlled and coordinated response in less time and with a higher level of confidence than manual processes. Orchestration tools include manual cutouts along with checkpoints for approvals to coordinate pre-planned and tested COAs. Manual approvals are important since the particular COA taken could have significant impact on critical, ongoing operations. Therefore, these network attack response decisions must at least be authorized by designated personnel knowledgeable about ongoing operations and the subtle relationships between different organizations and mission processes.

Orchestration Tools – The Circulatory System Of ACD

A security orchestrator acts as the "circulatory system" of the ACD C3 system. It takes raw sensor data and pumps it through various analytics to produce context-rich cyber information. This information is used by cyber security operators to decide on various mitigation COAs. Once the COAs are selected by the operator and launched, the security orchestrator choreographs the COAs across the various policy decision points that are involved using the control plane of the orchestration system. Therefore, the richness of intelligent orchestration capabilities is highly dependent on the tool's ability to seamlessly integrate many tools, data repositories, and sensors. Orchestration tools may include:

- connectors to ingest sensor data and threat data from a variety of sources

- data transformation and filtering capabilities to enable format conversion, normalization, and categorization of threat indicators and warnings, security posture data, and affected mission processes
- workflow tools and adapters to move data between various analytic tools and to convey response actions to policy enforcement points, policy decision points, and mitigation tools
- a control plane and rules engines to synthesize and choreograph courses of action
- visualization tools to provide a dashboard of the orchestrated workflow status

The control plane provides the decision framework. It continuously examines relevant data about the changing computing environment, identifies and prioritizes COAs, and intelligently relays COA information in real-time to the devices that need to adapt. The control plane utilizes a knowledge base to understand how COAs are mapped across the network topology, and a rules engine to determine how the different activities of a COA need to be synthesized and what decision points need to be coordinated. The **rules engines** may include different goal-directed behavior algorithms that reflect the risk management approach of the enterprise. These algorithms are used to calculate risk scores caused by a security event and assess how critical mission processes may be affected by a security event. The algorithms will help prioritize COAs that can minimize the risk impact on these mission processes by a security event. These capabilities of an orchestration tool are instrumental in accelerating the orient-decide-act portions of the OODA loop.

Some examples of orchestration tools include FireEye's Security Orchestrator, Phantom, DeMisto, CyberSponse, IBM's Resilient Systems, and ThreatConnect. These orchestration tools combine workflow capabilities, ETL utilities and out-of-the-box connectors, risk scoring, and policy functions to manage and assess the security posture of endpoints. Security information and event management (SIEM) also play a key role in security orchestration. Cisco's Platform Exchange Grid (PXGrid), Trusted Computing Group's Trusted Network Connect (TNC) and Metadata Access Point (MAP) server, Splunk, LogRhythm, ArcSight, QRadar, AlienVault, and ForeScout's CounterACT are all type of SIEM tools that offer repositories that collect, parse, correlate, and store a variety of security event and asset state data. Several of these tools also support capabilities for network access control, threat detection, and other security use cases. These tools come with pre-defined schemas for organizing and presenting the data, and are integrated with security dashboards.

Let's Make A Deal

Partnerships between different vendors are beginning to emerge in the security enforcement and orchestration markets as well. For example, firewall vendors are teaming with threat detection service providers and orchestration tool vendors. In this partnership example, when the threat detection tool identifies new malware, it sends alerts about the malicious code to outside intelligence services on the internet. Subscribers to these services, such as orchestrators, then receive updated indicators of compromise (IOCs) and pass them on to the policy enforcement points (e.g., firewalls) for action. The firewalls translate these IOCs into firewall rules allowing the firewall to automatically disrupt malware communication attempts with Internet-based domains.

These partnerships are being applied across multiple use cases and a variety of response scenarios to tie together related COAs. Some use case examples where security orchestration tools may be combined with other tools include:

- ***IP Address Management (IPAM)*** — to orchestrate DNS, DHCP and directory services to support VM and IP provisioning for ACD techniques such as IP address hopping as well as to flag suspicious IP addresses for blocking.

- ***Identity and Access Management (IdAM)*** — to orchestrate network access control and attribute-based access control approaches where authorization to resources is granted or denied dynamically based on a user's attributes and other contextual factors that could change.
- ***Data-centric access controls*** — to orchestrate changes in data security policy due to changes in context or threat levels at the object level. These controls include policies for databases (row, column or cell levels), Data Loss Prevention (DLP), Digital Rights Management (DRM), need-to-know, originator-controlled, and trusted provenance / reputation types of protection.
- ***Cryptographic and Certificate Management Services*** — to orchestrate and automate the EKCM lifecycle, namely: discovery, enrollment, monitoring, validation, notification, provisioning, remediation, reporting and revocation of certificates and keys.
- ***Traffic engineering*** — to orchestrate topology discovery, path computation, and path installation functions to dynamically enable traffic re-routing "around" security events, such as between an enterprise and a public cloud or between data centers.

As previously mentioned in the section on Intelligent Networks, security orchestration can be applied in virtualized data center environments by combining SDN controllers such as VMWare's NSX-V with VM orchestrators such as vCloud Director to **dynamically provision and direct** virtualized security functions.

Humans Still Need To Conduct the Orchestra

Often it is still advantageous to have a human in the loop when orchestrating a COA. Approval points need to be instrumented in the workflow to ensure that critical processes are not disrupted by active defenses. COAs might cause downtime of systems which must be managed. Collateral effects due to COAs must be assessed. Visual analytics are one method that can provide a rich, interactive experience to sift through large amounts of data, perform what-if analysis, and synthesize actions into coherent and optimized COAs.

Even so, human intervention may be prone to misinterpretation of information or **bias in decision-making**. Training can help ensure that any decision-makers that participate in developing COAs are aware of their cognitive biases, and that measures are in place to account for said biases. A list of cognitive biases includes: anchoring, belief bias, confirmation bias, distinction bias, focusing effect, irrational escalation, and so forth. Decision makers are all predisposed to these biases. Automated orchestration helps to nullify some of the effects of these biases since COAs can be designed and tested ahead of time through scenario analysis, goal-directed approaches, or other planning techniques.

A Common "Sheet of Music" Makes Orchestration Better

Orchestration tools face many challenges in the collection, analysis and dissemination of security event and security posture data of endpoints. Relevant data is scattered across a growing number of sensors and systems, often stored in proprietary formats, and sometimes hidden behind complex APIs, on vendor sites, or even in on-line forums. Event data may be either structured or unstructured and, more often than not, requires transformation to be useful in developing security responses. These data challenges require additional processing by automated orchestration tools or manual processing which slows down the ACD OODA loop. Data standards for security events to **reduce the amount of transformations and manual intervention** required to perform orchestration tasks. Figure 5 shows this funnel of event processing.

Figure 5 – Security Event Funnel

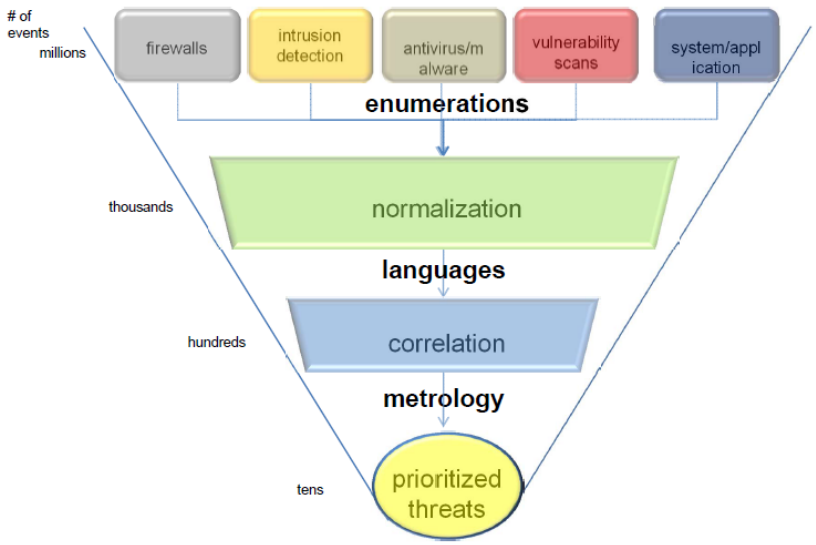

This need hasn't gone unrecognized. A variety of standards are emerging for security orchestration. Several standards efforts center on event management since managing events are an essential task for security orchestration tools. The current lack of event interoperability is becoming an intractable problem as the number of electronic systems and their generated events increase. Some of the standards that have attempted to address this problem include:

- ***CBE –*** *The Common Base Event model,* led by IBM, is a standard that defines an XML event syntax. CBE is described as a common language to detect, log and resolve system problems. CBE has yet to have any noticeable industry impact.
- ***CEF –*** *The Common Event Format* was developed by ArcSight. A CEF message is composed of delimited plaintext strings with optional sets of key-value pairs. It is relatively simple to generate and parse, and is transport independent. CEF is the preferred communication method of ArcSight products, such as the Enterprise Security Manager (ESM), and is supported by several other products.
- ***IDMEF –*** *The Intrusion Detection Message Exchange Format* is an Internet Engineering Task Force (IETF) effort that was designed to enable the communication of intrusion events observed by IDS devices. It consists of two entities: a syntax expressed in XML and the transport protocol (Intrusion Detection Exchange Protocol – IDXP). IDMEF is supported by a very limited number of intrusion detection products.
- ***WELF –*** *The WebTrends Enhanced Log file Format* is similar to CEF in that it is not bound to any specific transport and represents log data using plaintext, key-value pairs. WELF consists of four required and twenty optional syntax fields limited to expressing firewall, virtual private network (VPN), and other simple network-based events.
- ***IODEF –*** *The Incident Object Description Exchange Format* was developed by the IETF to improve computer incident response communications and is often associated with IDMEF. IODEF centers on the human-to-human communication of incident response, not on how the incident was discovered or on the formatting of related log files.

All of these standards have had only limited adoption or only address a portion of the event interoperability problem. Most professionals knowledgeable in this area recognize the need for an accepted, industry-wide event expression standard. Through widespread adoption of a unified language for expressing content of events, such as logs, any orchestration engine can immediately support the logs from any device. This allows timely and automated analysis of the event and choreography of the response without having to perform expensive and time-consuming data transformation functions.

One example of such a standard unified language effort is the **Common Event Expression** (CEE) project undertaken by MITRE. CEE is an event interoperability standard for electronic systems. MITRE recommends a framework to address the various components of an electronic event standard: an open format event expression taxonomy, log syntax, log transport, and log recommendations. Logs contain information about events, which can include device states, monitor readings, and a variety of other information. Logs are often further characterized as data logs, audit logs, alerts, alarms, audit trails, and a variety of similar terms.

In an ideal situation, the logs generated by various devices would reflect a near-real-time infrastructure awareness; they should represent every event that affected a particular device. With all of the necessary event data available, orchestration tools should be able to aggregate, correlate, and prioritize the logs in order to detect any significant events, including anomalous or malicious behaviors, and maintain a constant state of overall system awareness.

NIST has standards efforts underway to extend concepts of **SCAP** to automate the event management space. Where SCAP standardizes the data models of configuration and vulnerability management domains, these efforts focus on standardizing the data models relating to event and audit management, and remediation. To this end, CEE has been integrated with CybOX - The Cyber Observables construct that is intended to capture and characterize events or properties that are observable in the operational domain. CybOX was mentioned earlier in our description of the STIX standards effort for threat information sharing.

NIST is also building on CEE with a proposed specification for an enterprise remediation automation framework in the *Interagency Report 7690 (Draft) — Proposed Open Specifications for an Enterprise Remediation Automation Framework*. The proposed open specifications are intended to enable interoperation among components from different vendors to perform enterprise-wide functions to remediate vulnerabilities. These orchestration actions can be difficult, expensive or even impossible to accomplish due a lack of interoperability among tools and endpoints. As such, NIST is proposing a **Common Remediation Enumeration** (CRE) for standardizing identifiable remediations. The scope of a CRE entry is the set of actions that must be taken to accomplish a distinct remediation objective (e.g., installing a software patch or changing the system configuration). A single CRE could require that multiple atomic actions, such as changing a configuration value and installing a patch, be orchestrated and performed across multiple tools to achieve the desired end state.

NIST isn't the only standards body developing standards and specifications to support automated orchestration. The Internet Engineering Task Force (IETF) is also working on a standards track known as *Secure Automation and Continuous Monitoring (SACM)* . The SACM working group is developing standards focused on asset, change, configuration, and vulnerability management as well as automation interfaces and data formats relating to event management and continuous monitoring. These areas of focus are important use cases for automated orchestration tools.

If security will fail (and it will), orchestration tools must also support forensics, recovery and understanding what happened as constrained by after-the-fact – that is, we must walk backwards in time to replay the unfolding of events that created the security incident. Depending on the scope of the incident, this means that orchestration tools must be capable to extract, exchange, and normalize time-based events across many locations, tools,

perspectives, and time zones. CEE focuses on individual device-generated events, not on whole security incidents.

According to NIST in its *Computer Security Incident Handling Guide* the term incident refers to the collection of information regarding impact, time, cost, or confidence assessments; point-of-contact details; incident mitigation strategies; and any information related to the human factors surrounding incidents and incident response. Incident-related standards efforts such as IODEF are better suited to capturing the holistic information needed for incident response. However, incident reports often include event logs, which may be provided in the CEE format, and a CEE-defined event may be incorporated into IODEF-defined incidents.

DECEPTION, DELAY, DETECTION

Active cyber defense offers deception, delay and persistent detection approaches to disrupt cyber attacks while also increasing the work factor for the attacker. By **camouflaging defenses**, deception can provide real-time insight to a threat actor's TTPs while delaying an attacker's ability to launch a successful attack. Adaptive and persistent detective capabilities can expose stealthy malware operation, while profiling the attacker's device(s) to enable an early indicator and warning capability.

Good Deceptions Rely On Adaptive Defenses

Deception has been a part of the arsenal of cyberwar for the last 20+ years. Over this period, the variety and sophistication of deception tactics for both the attackers and the defenders have significantly increased. However, the three basic requirements for deceptions to be successful have remained constant:

- creating credible deception "stories" that will elicit actions by the attacker
- being able to quickly detect when the bait has been struck
- responding without human intervention.

Good deception stories require good cyber intelligence – cyber intelligence informs deception tactics by identifying what an attacker may find valuable and how attackers operate. Automated responses necessitate a mature security orchestration capability – to manage predefined COAs for maintaining the deception to elicit an attacker's TTPs and targets.

Concealment Versus Simulation – Which Tactic Is Most Suitable For Your Defenses?

In general, there are two main types of defense deception tactics – concealment and simulation. Concealments are masking techniques that inhibit observations by attackers while simulations enhance observations. When used in combination they provide the means for redirecting attackers away from real targets.

Some examples of concealment include:

- altering the environment by creating noise or false traffic in the targeted environment,
- reducing the attacker's opportunity to observe the targeted environment, such as configuring servers to not answer pings, configuring firewalls to prevent traffic flows between certain origins and destinations, and using network address translation (NAT).

Firewalls often use **deceptive replies in response to disallowed packets** to conceal targets and delay attackers. For example, firewalls can simply not reply to disallowed packets. This can delay attackers in their recon phase in three ways:

1. scanners can be slowed down if they wait a long time for a reply, and
2. attackers may interpret the non-response as a dropped packet and retransmit the probe, perhaps multiple times. Another deception used by
3. firewalls send false negative replies to attack probes. For example, the firewall can send an ICMP hostunreachable message in response to a TCP ping.

More sophisticated concealments can be created at the protocol level by using predefined sets of redirection responses to disallowed packets using a rule set similar to router rules. Such a redirection capability could route packets through different interfaces so that the same attacker's IP address goes to different networks depending on measurable parameters in the rule set. Mirroring is also an effective tactic at causing even more highly skilled attackers to become confused. Ultimately, the goal of concealment is to suppress signals emanating from the target environment thus causing the attacker to fail to find a real target.

The goal of simulations is to create the illusion for the attacker that the attack is progressing as expected, using techniques ranging from fake error messages to redirecting the interaction with the attacking computer process to a virtual sandbox.

Some examples of simulations include:

- **honeynets** which impersonate the real environment to lure the attacker to reveal their tactics. Honeynet simulations leverage unused IP addresses to create fake targets. There can be thousands of fake computers simulated (e.g., by using 10.0.0.0), often there can be many more fake computers than real computers.
- **honey tokens** which are files set up to be attractive targets that contain "canaries" which are triggers that generate an alert if an attacker opens or manipulates a file.
- **execution wrappers** which create operating system level deceptions that are invoked whenever a program is executed. The decision on whether or not to employ a deception is based on system state, process lineage, and the respective system call. This type of deception has shown to be capable of successfully deceiving systems administrators who tried to exceed their mandate and access content they were not authorized to see.

A good simulation leads to an attack graph that creates additional alternatives for the attacker. Specifically, simulations induce the attacker to find **false targets**. Successful simulations consume attacker resources, and in some cases cause the erroneous belief that the false targets are real. The result of simulations that are successful is that the attacker goes further through the attack tree in the examination of false targets. This allows the defender to develop greater cyber intelligence on an attacker's TTPs and exploit goals. An additional side effect is that real targets may be misidentified as false targets, thus causing attackers to believe that real systems are in fact simulations or impersonations.

Active Deception Through Cyber Maneuver

The concept of dynamic identity can play a significant role in deception and adaptive defenses. Static mappings of identities create vulnerabilities that undermine protections and facilitate asset theft. For example, static passwords, static credit card numbers, static asset tags, and static IP addresses simplify the tasks of a malicious adversary who wishes to phish individuals and misuse their credentials to gain unauthorized access. By contrast,

dynamic passwords and credit card numbers, dynamic asset tags (e.g., as realized by hash chains), and dynamic Internet Protocol (IP) addresses can substantially complicate the adversary's job and reduce the lifetime of exploits.

From a network security perspective, one of the main assumptions of an attacker is that specific technical details of system operation remain static across all targeted machines and networks. If each machine and network could generate a custom configuration, using a secret not available to the attacker, then a **constantly moving profile** could be achieved to stay ahead of adversaries. Vulnerabilities may be hidden by randomizing the behavior of the software that communicates with potential attackers or associated malware. Such a network maneuver approach is capable of avoiding many attacks, even in the face of zero-day vulnerabilities, and provides a proactive posture for the enterprise while increasing the resiliency of the network. This is the premise of the deception tactic known as cyber maneuver.

Platform Diversity + Maneuver = Cyber Kill Chain Disruption

One method of achieving a constantly moving profile is through the use of artificial diversity. The advantage of artificial diversity is that malware is typically targeted at exploiting vulnerabilities of a particular platform or operating system. The use of **diversity and maneuvering** to different platform alternatives reduces the attack surface available to malware leveraging a particular vulnerability.

Virtual overlay networks provide a method known as service chaining to help implement this type of cyber maneuver. Service chaining allows multiple virtualized network functions, such as virtual firewalls, virtual IPSs, virtual load balancers, etc., to be connected together and to migrate with a specified workload to different virtual machines. Through service chaining, enabled through software defined networking and service orchestration, workloads along with their network security elements can migrate to different virtual platforms with different IP addresses.

Network address space randomization (NASR) using IP address hopping is another example of artificial diversity and has been used as a way to counter malicious attacks. In this case, a large pool of IP addresses is used to make dynamic assignments to hosts. The pattern of changes of the IP address is known to both the client and the server(s), and preferably secret from others. Workloads "hop" to different IP address assignments based on an address destination selection algorithm. Servers could be configured to expect requests to the changed IP address within a certain time threshold. If the subsequent requests do not arrive within a threshold time period, the server system can be configured to terminate further access to the requestor.

Without knowing the pattern of changes of IP addresses, it will be difficult for an eavesdropper to intercept data or to recon the network. To further enhance the security of this approach, the client (user) IP addresses could also be dynamically assigned using DHCP at a gateway and changed on a secret pattern basis as well. By constantly changing the IP address of the client, an additional layer of security is enabled on all transactions. Even if someone were to identify a key to decrypt transactions that use encryption, they would not be able to represent themselves as a client nor would they be able to follow the transactions as they would not know the sequence of IP addresses since the sequence would be a shared secret between the client and server.

The confusion and "noise" generated in the network by the maneuvering activity also **minimize the attackers' ability to observe** the network, thereby increasing their cost and slowing them down. Recognizing that some attacks will succeed, cyber maneuver also provides the ability to disrupt a persistent threat by requiring extra effort by the attacker to continually remap the network and re-establish malware command and control channels. By making the attacker work harder, cyber maneuver can increase the probability of attribution and

detection due to the increased activity required of the attacker. Additionally, cyber maneuver can be combined with cleansing to remove malware that may have obtained a foothold.

Cyber Maneuver + Contextual Awareness = Adaptive Networks

Cyber maneuver decisions can also be influenced by the threat context or the security state of maneuverable assets. A *cyber maneuver decision framework as described here* can leverage this knowledge and create artificial changes to hop intervals, destination targets (including different geographic destinations), and re-direct exploited assets to honeynets for observation. Destination selection algorithms can allow destination LANs to be specified as less or more desirable when network or security conditions change, and support avoidance of particular vendor operating systems, hardware platforms or hypervisors.

Cyber maneuver can also enhance network anonymity and enable new effective responses to many Distributed Denial of Service (DDoS) attacks. For example, a dynamically assigned destination address can be assigned to serve as a dynamic netflow marker making it easier to track suspect traffic. Because an attacker must learn the dynamically assigned destination IP address, a gateway capability knows the IP address of the attacker machine to which it sent the requested destination address. Therefore, the **attacker must expose at least one of its bots** to learn the dynamically issued destination address. While possibly incomplete, this IP address is a useful starting point for a trace-back analysis and filtering for denial of service attacks.

Choose Your Deception

Active Cyber Defense can enable a large variety of specific deception tactics that can delay, confuse, scare away, or tie up an attacker depending on the circumstances and the methods. Table 2 lists several deception techniques based on *studies performed by the U.S. Naval Postgraduate School*.

Table 2 - Evaluation of Deception Methods in Cyberspace

Deception method	Suitability for offense in information systems, with general example	Suitability for defense in information systems, with general example
supertype	pretend attack is something else	0
whole	conceal attack in a common sequence of commands	0
agent	pretend attacker is legitimate user or is standard software	0
object	attack unexpected software or feature of a system	camouflage key targets or make them look unimportant, or disguise software as different software
instrument	attack with a surprising tool	0
location-from	attack from a surprise site	try to frighten attacker with false messages from authorities
location-to	attack an unexpected site or port if there are any	transfer control to a safer machine, as on a honeynet
location-through	attack through another site	0
direction	attack backward to site of a user	transfer Trojan horses back to attacker
frequency	swamp a resource with tasks	swamp attacker with messages or requests
time-at	put false times in event records	associate false times with files
time-through	delay during attack to make it look as if attack was aborted	delay in processing commands
cause	doesn't matter much	lie that you can't do something, or do something not asked for

purpose	lie about reasons for needing information	lie about reasons for asking for authorization data
preconditions	give impossible commands	give false excuses for being unable to do something
ability	pretend to be an inept attacker or have inept attack tools	pretend to be an inept defender or have easy-to-subvert software
accompaniment	a Trojan horse installed on a system	software with a Trojan horse that is sent to attacker
content	redefine executables; give false file-type information	redefine executables; give false file-type information
measure	send data too large to easily handle	send data too large or requests too hard to attacker
value	give arguments to commands that have unexpected consequences	systematically misunderstand attacker commands
effect	lie as to what a command does	lie as to what a command did

The following examples highlight this variety of tactics.

1. Enabling specific abuse detection points in web application code using a library of vulnerability deception modules, and responding to application abuse with session-specific deceptive responses, warnings, and blocks.

Deceptively vulnerable web sites can be employed to identify an attack while minimizing the impact on legitimate users. By embedding seemingly vulnerable abuse points in the web site script and configuration, attackers may be deceived into launching attacks that can be quickly detected. A variety of response actions can then be employed including warnings, throttling the session connection, requiring additional login information, and blocking the attack. Through careful sequencing of the responses, an attacker's TTPs can also be revealed and captured for analysis.

2. Employing persistent tracking mechanisms, such as canvas fingerprinting, evercookie, super cookies, and zombie cookies – to create a "long-term fingerprint" of the attacker and use that fingerprint to recognize an attacker on return visits.

Fingerprinting web site users may cause privacy concerns, however these techniques offer a means to identify and track web site abusers when these tracking mechanisms can be linked to web site abuse. Once identified as an abuser, such a fingerprint could be shared through a threat intelligence service to block, blacklist, or to closely monitor the user's behavior at sites that subscribe to the intelligence service.

3. Employing honeynets and honey files to capture attackers' TTPs and malware profiles, while also providing disinformation to the attacker.

Use cases for **honeynets** assumes an attacker is undertaking an overall attack effort involving intelligence gathering, entries, privilege expansions, and privilege exploitations. Honeynets are used to attract attackers and therefore may be used to detect an attack in progress; an inexpensive canary may be set up such that any traffic to the machine triggers an alarm; false network traffic can lead an attacker into believing that a port or IP address on a honeynet is a valid target, but accessing that IP address and port causes an attack alarm. A variety of attack detection techniques can be implemented with honey files as well, including inclusion of code that will report back to a monitoring server when executed. This can be achieved by using JavaScript for PDF files, the addition of fake entries in robots.txt files for web servers, the use of invisible links, the inclusion of honey-token HTML comments, or remote images that are downloaded when the document is opened; and, inclusion of bait information, such as fake credentials, that attackers may try to use. DNS honey tokens can also complement the use of honeynets. For example, a small number of fake DNS records on the authoritative DNS servers of the organization can be created and configured to initiate an alert when these specific records are requested.

Using Deception With Malware Analysis Tools

Malware sandboxing often includes deception techniques that are designed to **trick malware** into exposing how the malware operates and what it would do on a system if deployed. Attackers recognize that detection methods are becoming more advanced and accordingly have raised their game as well. Since malware often comes with anti-VM or anti-sandbox techniques along with dormant codes, malware detection and analysis tools are becoming more sophisticated by using deception and hiding techniques. For example, a tool from California-based *Bromium leverages micro-virtualization* at the processor level. The unique microvisor architecture creates a sandbox runs on the CPU rather than the kernel level sandboxing that is offered by other vendors. This hardware isolation method emulates a complete system, thereby removing many of the indicators that malware uses to detect if it is running in a virtual sandbox. Bromium is able to stop zero-day attacks caused by routine tasks including browsing the Internet, downloading documents, opening email attachments, and launching files from authorized removable storage devices.

Creating Good Deceptions Requires Investment

A significant effort is necessary to create a good deception. First, **significant planning is needed** to develop objectives and a deception strategy. Second, additional computing resources are needed to create the deception environment. Third, false content can be very difficult to create and expensive to maintain. Finally, a poor deception is worse than no deception at all since it will consume a defender's resources while not aiding in detection or mitigation of an attack. Despite these additional costs, cyber deception provides an adaptive line of defense against today's sophisticated attacks while also helping to capture cyber intelligence about threats.

Deception approaches must be balanced against the additional costs of deployment and maintenance, as deceptions can go stale, and consistency and freshness of deceptions should be maintained. Deceptions such as cyber maneuver must be designed with high precision, with virtualization playing a key role. As such, **cloud computing could be well-suited** to enable this type of active defense. A large address space is also needed for maneuver tactics such as IP address hopping, so IPv6 environments are preferred for this maneuver approach.

AGILE CLOUD

The versatility of the cloud offers some attractive options for delivering adaptive security capabilities, while also providing a platform for new innovations. Here are ways that active cyber defenses can help protect the cloud and, conversely, how active cyber defenses can also, in turn, benefit from cloud computing capabilities.

10 Proactive Cloud Defenses

The very nature of the cloud's **scalability and elasticity makes it adaptive** but how does that translate to security? Cloud computing offers several unique capabilities to enable Active Cyber Defenses. These capabilities include:

1. Big Data analytics for deep insight into cyber events and information flows to detect anomalies, model and assess behavior, and discover threats such as botnets, APTs, and insider threats
2. Remote live forensics and recovery delivered as a service to support incident investigations and recovery operations
3. Cloud-based threat intelligence and application reputation services

4. Cloud-based honeynets / sinkholes that can elastically scale to meet deception needs and to capture / redirect malicious botnet traffic
5. Virtual appliances to automate your incident response COAs
6. Virtual desktops to essentially eliminate patching and provide adaptable virtual personal spaces to securely operate as the risk level changes
7. Trusted Identity Providers and attribute exchanges as ways to offload identity management tasks and make access management processes more agile

At the same time, **adaptive defenses can assist in securing your cloud environment**. For example:

8. Secure enclaves in the cloud, established through adaptive protocols such as single packet authorization and through NFV security elements assembled dynamically to create personal, trustworthy execution environments
9. Secure on-ramps to the cloud that broker security services which dynamically interject security controls for given security contexts between the user and the cloud provider. For example, a cloud security broker could perform data tokenization services for transforming private data into non-sensitive data when stored in the cloud.
10. Self-healing and resilient cloud through proactive secret sharing among a minimal set of critical functions / nodes.

Each of these cloud-based ACD capabilities is highlighted in the following sections.

Big Data, Big Insight

Many large enterprises are flooded with cyber sensor data coming from all directions - log data from servers, IDS and IPS data, netflows and packet captures, endpoint compliance data, vulnerability scan data, firewall alerts, DNS alerts, alerts from outages, and many more. In addition, these enterprises may be subscribed to one or more threat intelligence sources which can also pile in lots of data that needs to be dissected and correlated with other events that are going on in the enterprise. SIEM tools have been used to make sense of all this data, but many of the current tools cannot keep up with the volumes of information that now must be processed and stored by large enterprises. Plus, much of this **data is unstructured** making it difficult to perform the types of correlations that SIEM tools are good at doing.

So enter the cloud and Big Data. The evolution of Big Data tools is enabling **security analytics** to effectively add a level of context and awareness to security incidents that was previously impossible to achieve using traditional SIEM tools. The value of Big Data analytics lies in not just being able to efficiently collect and store large data sets, but being able to make sense of the data over time. Remember that OODA loop - the engine powering active cyber defenses? Well, Big Data lives in that Orient portion of the loop. Big Data analytics can put enterprises in a better position to predict attacks by comparing the current network states to “normal” baselines of network activity. Today's network traffic is of a very different format, of much larger volume and speed, and data moves around and across networks very differently due to virtualized workloads – the ability to detect what is normal versus abnormal is only possible with Big Data analytics.

This need for Big Data tools is amplified by the Cloud Security Alliance (CSA). According to the *Cloud Security Alliance report: Big Data Analytics for Security Intelligence*:

Big Data tools have the potential to provide a significant advance in actionable security intelligence by reducing the time for correlating, consolidating, and contextualizing diverse security event information, and also for correlating long-term historical data for forensic purposes.

Big Data technologies are able to **keep pace with the ever-growing amount of cyber event data** through the cloud's elastic compute and storage capabilities, the application of multi-core and GPU-based processing, and through research and development of more flexible algorithms for analysis. Other unique adaptive features of Big Data technologies are their ability to achieve real-time analysis of streaming data and their ability to munge large amounts of unstructured data.

The CSA report highlights several Big Data use cases for adaptive security:

- Big Data tools were trained to identify **malware-infected hosts in an enterprise network** and the **malicious domains accessed** by the enterprise's hosts. The results showed that high true positive rates and low false positive rates can be achieved with minimal ground truth information (that is, having limited data labeled as normal events or attack events used to train anomaly detectors).
- Big Data tools were used to identify **infected hosts participating in a botnet** through analysis of netflow records.
- Big Data tools were used to sift through massive amounts of data in search of **anomalies** that would be indicative of an **Advanced Persistent Threat (APT)**. By looking for small deviations from the usual patterns of users, and correlating these anomalies to attack patterns, attack indicators, and likely targets, the tools were able to ferret out APTs, exposing these stealthy attacks that could not be identified with more traditional detection methods.

Big Data tools can also be used to partition large datasets for analysis by other tools such as SIEM tools and Splunk, thereby reducing the need to retrain cyber analysts on the Big Data tools. Data provenance tools need to be added to the analytics mix as well. As Big Data expands the sources of data it can use, the **trustworthiness of data sources** needs to be verified. Modeling tools can also take advantage of Big Data analytics. For example, security analysts may wish to model courses of action that a cyber attacker may employ for certain vulnerabilities or attack scenarios to provide predictive indicators or early warnings of attack behaviors. Models can also leverage Big Data analytics to help understand impacts from courses of action that are used by cyber defenders to fend off attacks.

Offloading Security To The Cloud

Security-as-a-service is not a new concept when it comes to the cloud. However, when it comes to cyber forensic services, using cloud-based services is fairly new and unique. The issue with cyber forensics to date is that it is highly manually intensive, requiring specialized tools and skills. Also, there are key issues with respect to preservation of the chain of evidence if prosecution is warranted. Cloud environments can make this difficult due to the transient nature of virtual machines and their multi-tenancy. However, cloud-based cyber forensics provides several advantages since many organizations lack the deep skills needed for this type of work. If prosecution is out of the question, then leveraging these services can be very useful. Also, you don't have to be operating in the cloud to use these services.

One example of this capability is Google's Rapid Response (GRR). GRR is an **open source incident response framework** which is intended to provide a scalable solution for remote live forensics. As everyone is aware, Google publicly disclosed, in 2010, that they were the subject of a targeted attack, commonly referred to as "Operation Aurora." From this experience, Google quickly realized the nascent state of the incident response

industry and tools. Thus, once the smoke cleared, the Google Team began investing a lot of resources into augmenting their own security capabilities and reducing incident response to a search problem.

GRR is now fully-supported and open sourced at *github*. GRR consists of an agent probe that is deployed to a target system, and a server infrastructure that can manage and talk to the agent. One of the key focus areas of Google's work was on identifying artifacts that need to be collected for forensics analysis, and providing the tools to enable the live collection and timeline analysis of these artifacts in a trusted manner (timestamped, hashed, and securely transported). Artifacts are different from indicators of compromise (IOCs) - artifacts are also referred to as cyber observables and reflect the stateful property of an object (e.g., presence of a mutex) or a measurable event (e.g., creation of a registry key on a host), while an IOC is a description of an observation that may be related to an intrusion – the Who, What, Where, When, How, and sometimes why. IOCs tie to observables and artifacts.

Using GRR, a cloud-based security and forensics center can instruct each agent probe to collect events and raw traffic, send them back for deep analysis, and generate new security rules. These new security rules can be enforced by collaborative Unified Threat Management (UTM) tools and the feedback events of such rules can be returned to the security center. By this type of close-loop control, a collaborative, cloud-enterprise network security management system can identify and address new distributed attacks more quickly and effectively.

Another security-as-a-service that seems to be popping up everywhere is **cloud-based cyber threat intelligence** services. What is interesting about cloud-based threat intelligence services is how the cloud helps to efficiently enable the use of honeynets, honeypots, and sinkholing for collecting cyber threats for intelligence processing and sharing, such as: discovering and tracking cyber threat actors, collecting TTPs, and monitoring botnets, as well as proactively denying access to the bots from the botnet herders. The leading organization for honey-based technologies is the *Honeynet Project*. It provides a wide variety of open source capabilities that can be leveraged for honey projects.

The cloud's elasticity and scalability provides many options to exercise honey-based technologies. One example is through the use of low-interaction honeypots such as honeyd. Honeyd is a small daemon that creates virtual hosts on a network. The hosts can be configured to run arbitrary services, and their personality can be adapted so that they appear to be running certain operating systems. For example, Honeyd can appear to the attacker to be a Cisco router, a Windows web server, or Linux DNS server.

There are several advantages to emulating different operating systems. First, the honeypot can better blend in with existing networks if the honeypot has the same appearance and behavior of production systems. Second, you can target specific attackers by providing systems and services they often target, or you can target specific systems and services you want to learn about. Anytime Honeyd sees a connection attempt to an unused IP, it intercepts the connection and then interacts with the attacker, pretending to be the victim. The cloud also allows you to constantly shift the profile of victim systems by provisioning different virtual victim hosts.

Clouds are also useful in enabling honeynets by virtue of their platform-as-a-service or infrastructure-as-a-service offering. Honeynets are architecture, an entire network of computers designed to be attacked. The idea is to have an architecture that creates a highly controlled network, one where all activity is controlled and captured. Within this network you place your intended victims, real computers running real applications. The bad guys find, attack, and break into these systems on their own initiative. When they do, they do not realize they are within a Honeynet. All of their activity is captured without them knowing it. This is done by inserting kernel modules on the victim systems that capture all of the attacker's actions. At the same time, the Honeynet controls the attacker's activity. Honeynets do this using a Honeywall gateway. This gateway allows inbound traffic to the victim systems, but controls the outbound traffic using intrusion prevention technologies. This

gives the attacker the flexibility to interact with the victim systems, but prevents the attacker from harming other non-Honeynet computers.

Using Virtualization And Trusted Computing To Adapt

Cloud service providers offer templates which can provide the foundation to **automate Courses of Actions** (COAs) for defensive responses and mitigation actions. For example, AWS Cloud Formation supports templates for virtual appliances. These templates can be parameterized and come with enterprise features such as governance and automatic tagging of resources. Parameters are a way for users to specify unique or sensitive values in the properties of virtual appliances. This parameterization comes in handy, since security needs can adapt with different threat contexts, as AWS Cloud Formation allows you to change the set of resources that make up a security stack. Therefore, when your threat context changes and you need to spin up a new security policy, you can easily do so through changes in parameters to your virtual security stack.

The network security stacks for cloud environments are composed of different virtual appliances, such as a virtual firewall, a virtual IPS, a virtual application proxy, etc. Cloud users of the security stack need to ensure that the virtual appliances in use are configured according to their needs. One way to make this possible is to use a combination of trusted computing and model checking to provide trustworthy verification. In this approach, a virtual appliance is generated from an abstract model of the functions to be performed by the appliance. Upon bootstrapping the virtual appliance, trusted computing provides guarantees that the virtual appliance in use is bound to a given model through successful attestation. The cloud user can then determine whether the virtual appliance is configured correctly by specifying relevant properties in a usage manifest file, and verifying, using a model checker, that the virtual appliance satisfies such properties. If validation passes, the virtual appliance is deemed trustworthy, and it is ready to be used.

In a similar fashion, **virtual desktop infrastructure** (VDI) can create trusted, tailored, virtual spaces where users can select/create different environments for different activities that satisfy a variety of threat and mission scenarios. Desktop virtualization is leading to the desktop being disassembled. Intrusion detection and prevention, applications, and user personas can be discretely managed and stored, only to be recomposed via a VM orchestrator and the network into a familiar workspace for each user at log-in. Checking and patching vulnerabilities and making updates are easier as well with VDI, since a single gold image only needs to be scanned and patched and then replicated as users need them. This ability to start with a clean computing slate each time a user logs in helps to reduce the attack surface and the exposure to persistent threats in a proactive manner.

Assembling Your Identity In The Clouds

Just as a desktop can be disassembled by virtualizing its components and then reassembling them into personalized virtual execution spaces, a digital identity can also be decomposed into a collection of attributes that can be reassembled into different personas. This Lego block approach to identity enables agility in authentication and authorization processes by:

- allowing the use of different personas (credentials)
- that can be based on different user or machine attributes
- that can be provisioned just-in-time
- for managing dynamic access contexts (privileges, permissions)
- across one or multiple cloud providers.

Attributes can be associated with a person or a resource (such as data or a computer) and have a value. For example, attributes pertaining to a person are traits (the color of eyes, a fingerprint), or obtained from an authoritative source (a skill certification, a role), or based on the environment (location, time of day). Data attributes may reflect the identity and sensitivity of data, while compute resources may have performance attributes and identity attributes.

Attributes can be used in **access management decisions** – for authentication purposes and authorization. This form of access management is known as Attribute-Based Access Control (ABAC). According to the *NIST definition of ABAC*, ABAC relies upon the evaluation of attributes of the subject (persons or non-person entities that are requesting access – role, location, phone number, etc.), attributes of the object (resources being requested - e.g., web URL, data), environment conditions, and a formal relationship or access control rule defining the allowable operations for subject-object attribute and environment condition combinations. As attributes can be engineered to reflect appropriately detailed information about subjects and objects, ABAC ensures great flexibility in expressing fine-grained policies which are increasingly required by applications – especially cloud-based applications.

With an unknown number of potential customers, it is unrealistic for cloud service providers (CSPs) to pre-design and implement all kinds of access control models or design one by one on demand. The cloud platform should provide a flexible and **dynamic access control framework** such that customers can easily configure their own access control policies. Furthermore, in order to distribute their resources (for availability reasons, for example), some organizations may be tenants of multiple CSPs. In this case, such an organization encounters an additional problem of dealing with multiple different access control interfaces and integrating them. Sufficient abstractions can be built on top of ABAC in order to closely mimic the access control abstractions expected by each tenant.

ABAC is well-suited to manage access to multi-tenant cloud resources since the ABAC model can be applied to regulate the fine-grained usage of virtual resources such as disk, RAM, and network. APIs exist that can call the network to execute pre-defined policies using authoritative sources to attest to identities, roles, attributes or other context.

Back-end attribute exchanges and virtual directory systems can also extend the ABAC model across different cloud service providers. Using these capabilities a user can be authorized to access different cloud providers' resources by using different attributes that reflect precise policy conditions required for access. For example, Bob is an analyst working on a joint task force with other government agencies that requires him to access sensitive task force data. Some of the attributes required to access the task force systems are managed by his home enterprise (e.g., identity, organization affiliation) while others are managed by the joint task force (e.g., role), and still other attributes are managed by third party cloud providers (e.g., certifications, clearances). **Back-end attribute exchange systems** and virtual directory systems allow the required attributes to be retrieved from different attribute providers to construct credentials and establish privileges to access the joint task force systems. As this example shows, ABAC is well-suited to federated environments since roles don't have to be remapped between organizations nor does the user information have to be pre-provisioned at the "joint" organization.

As attributes represent information about the users, releasing attributes to a third party ABAC engine is a sensitive activity as the third party may not be trusted. Although trust negotiation and other techniques can be used to regulate the release of information, still, least privilege has to be enforced in sharing attributes. This means the least set of user attributes are released for the purpose of cross-cloud access request evaluations. One method to control the release of attributes would be to allow each user to individually control the release process. *User Managed Access (UMA)* is a profile and extension to OAuth 2.0, which is a lightweight web-based

approach for cross-cloud authorization requests. UMA defines how resource owners can control access by arbitrary requesting parties. Combining *Security Assertion Markup Language 2 (SAML2)* as a trusted container for authenticating attributes from third party sources together with UMA might be one way of allowing individuals to manage their attribute release.

Just as Lego blocks are built to a specific standard so they can be used together, attributes also must be normalized to a common definition and standard canonical form to be useful across different enterprises or access control frameworks. ABAC tools such as Axiomatics come with a dictionary of common attributes to provide baseline definitions. Tools that are built to standards such as the *System for Cross-domain Identity Management (SCIM)* specification also help make managing user identities in cloud-based applications and services easier. The SCIM specification suite seeks to build upon experience with existing schemas and deployments, while applying existing authentication, authorization, and privacy models. It provides a common schema and extension model, as well as binding documents to provide patterns for exchanging this schema using standard protocols. SCIM can provide the Lego block standard for the normalization and exchange of attribute information.

A different approach to managing attributes is provided by Radiant Logic's RadiantOne VDS product. This product utilizes a "smart virtualization layer" to create federated views of attributes across multiple identity silos. This smart layer normalizes the disparate attribute data for "curated" views of attributes that can be customized for each consuming application.

Curation of attributes is a critical function since attributes will be relied on to make access decisions. Since an attribute that is once assigned may no longer apply to a user (for example, a user's clearance level was lowered or suspended) then the ABAC system must periodically check with the authoritative sources to ensure that the appropriate ABAC decision can be rendered. Tools such as Sailpoint's Identity IQ provide capabilities to ensure that attributes are curated and up-to-date.

Overall, ABAC is an important tool for active defenses since it enables secure, fine-grained yet flexible access to resources. ABAC requires an ecosystem of capabilities to work – policy decision points, policy enforcement points, authoritative sources of attributes, etc. ABAC decouples access control from proprietary application silos and works well in federated environments where access requesters may be unknown at runtime. The cloud and ABAC are an excellent combination because of the flexibility of ABAC to support the different access control needs for tenants in the cloud. Enterprises can also offload identity management tasks to trusted third party cloud attribute and identity service providers. By distributing attributes among trusted identity providers in the cloud, enterprises can also increase the accuracy of attributes by checking across multiple authoritative data sources. For example, there are several credit score providers that can be used to provide attributes about a requester's financial standing. Techniques for masking attribute data also exist to render sensitive attribute information opaque while in process by providers and relying parties.

Secure Enclaves

Cloud service providers are attracting more attention by attackers as sensitive enterprise applications and data make their way into the cloud. Cloud users are looking for ways to complement the security protections already offered by cloud service providers to ensure the security of their data. Several active defenses offer unique capabilities to create secure enclaves in the cloud. For example, *Cutting Edge's NetAbstraction* transparently distributes the communications activity within and across multiple clouds and regularly churns the underlying network infrastructure. The **dynamic shifting of communications across multiple commercial providers**, along with the usage of multi-hop transport makes actual user information and origination location and identities a very difficult target for hackers.

Another example of active defenses enabling secure cloud enclaves is the use of Single Packet Authorization (SPA), also known as First Packet Authentication (FPA). This approach is described in detail as part of the *Software-Defined Perimeter initiative by the Cloud Security Alliance (CSA)*. SPA-protected endpoints insert a cryptographic identity token in the first packet of the TCP connection set-up, replacing the initial random sequence number normally generated by the TCP/IP stack services. The server extracts the identity token to authenticate the client before proceeding with the rest of the TCP/IP connection set-up. This "**identity-based networking**" approach provides the following security benefits to the SPA-protected server:

- ***Blackens the server**:* The server will not respond to any connections from any clients until they have provided an authentic SPA token.
- ***Mitigates Denial of Service attacks on TLS**:* Internet-facing servers running the https protocol are highly susceptible to Denial-of-Service (DoS) attacks. SPA mitigates these attacks because it allows the server to discard the TLS DoS attempt before entering the TLS handshake.
- ***Attack detection**:* The first packet from any other host must be an SPA. If a server receives any other packet, it should be viewed as an attack. Therefore, the SPA enables the cloud server to determine an attack based on a single malicious packet.

Servers behind a SPA gateway in the cloud form a protected enclave, invisible to other tenants in the cloud who do not have an appropriate identity token. SPA is based on RFC 4226 (HOTP). One commercial implementation of SPA is provided by BlackRidge Technologies.

Project SECURED is another example of dynamic provisioning and protection of private execution spaces using cloud technology. SECURED is an EU research and development effort based on SDN, NFV, and attestation. SECURED offloads security controls from the user device to a cloud-based network node. This node validates the attestation of an end user device trying to connect to the cloud using Trusted Computing Group's Trusted Network Connect protocol. Once the device is attested, the user identity is validated through the cloud using OpenID Connect. Next the **virtualized personal execution environment is assembled by the cloud**, the user applications and the security policies are brought together through NFV and SDN. SDN moves the secure execution environment to where the user is – thus security protection is independent from the user location. Some types of security applications that can be implemented by SECURED include packet filter, parental control, anti-phishing, content inspection. All of these tasks are executed as NFV tasks at the network node in the cloud. The network node uses hierarchical security policies since the same user connection may be subject to constraints from different tenants in the cloud. In this case, the user is informed of the overall policy applied to her connection (and may refuse connecting to the network).

Cloud Broker

Cloud offerings can benefit in other ways from ACD. For example, a Cloud Service Broker is a structure for organizing an on-ramp to cloud services that are offered by multiple cloud service providers (CSP). The **Cloud Service Broker uses knowledge of the customer's specific needs** with respect to application performance, costs, security, location, ethical standards, and other criteria, and then matches these to the capabilities of cloud service providers. The CSB then creates and deploys the appropriate service within the cloud, subsequently ensuring that the constituent cloud services continually work together to meet the customers required business level objectives.

When *Gartner* first introduced the concept of CSBs in 2011, the research firm laid out the three primary roles of a cloud services broker as aggregation, integration, and customization. With aggregation, a broker packages services from multiple cloud providers to ensure interoperability and security of enterprise data passing

between systems. A CSB focused on integration will help an organization coordinate multiple cloud services, and CSBs focused on customization help IT find and tailor cloud services to meet their unique business and technical requirements. CSBs can also fulfill other important functions, such as helping enterprises define and implement cloud governance policies and analyzing whether to migrate premises-based applications to the cloud.

Cloud service brokers often **offer secure orchestration services** that help customers and CSPs negotiate to an agreed upon security profile for given security contexts and Service Level Agreements (SLAs) along with agreed to costs of entry. An agreement is brokered that allows users and CSPs to dynamically enter shared risk/liability environments and to mutually establish a level of assurance using a risk-based or economic model. Figure 6 illustrates the cloud service broker concept.

Figure 6 – Cloud Service Broker

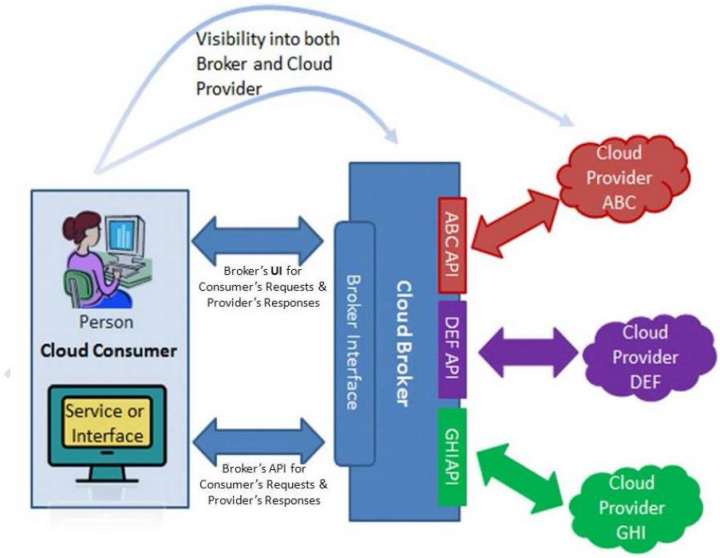

According to Gartner, in addition to brokering access requests, there is an emerging market for brokered services to interject security controls between the users and the cloud services they consume. For example, a Cloud Broker could offer a tokenization gateway that provides a reusable framework for transforming private data into non-sensitive data as it moves into the cloud.

A key technical underpinning of this type of active cyber defense is the Application Programming Interface (API). APIs provide consistent methods for outside entities such as web services, clients and desktop applications to interface with services in the Cloud. More and more, it will be through APIs that cloud data moves; however, the security and scalability of APIs are currently threatened by a problem called the password anti-pattern. This is the need for API clients to collect and replay the password for a user at an API in order to access information on behalf of that user via that API. Active cyber defenses can provide a cloud API broker that can protect the API through proof of possession of a secret or key for each API call. An example of a secret or key could be an OAuth access token or a secret access key. Such an intermediary would provide services such as management of keys and configuration of access to the cloud. ABAC could also be applied to control access to APIs. For example, CA Technologies API Manager can provide ABAC policy enforcement and API services for cloud service brokers.

Secure interoperability standards for cloud service brokers are also important considerations to ensure workloads are moved securely between data centers and the cloud. The *Distributed Management Task Force (DMTF)'s Cloud Infrastructure Management Interface (CIMI) Model* helps define what the scenarios are, and provides interfaces to help with the integration of the different types of service providers. CIMI works with the

DMTF's *Open Virtualization Format (OVF) Specification*, which provides a way to package multiple machines and their requirements, which can be exchanged with the cloud service and cloud service broker. CIMI focuses on Infrastructure as a Service (IaaS), as opposed to Platform as a Service (PaaS) or Software as a Service (SaaS), so users and brokers can move entire workloads.

Overall, the cloud service broker provides an important platform for positioning active cyber defenses for securing access to the cloud. Efficiently brokering secure cloud services is increasingly essential in multi-cloud environments. ACD at the CSB enriches the security capabilities of both the enterprise and the cloud service provider by providing dynamic yet customized security intermediary services.

Self-Healing And Self-Protecting Cloud

Cloud services providers have the ability to elastically scale as workloads grow, while also redirecting network flows as needed. However, to be truly resilient, these activities need to be orchestrated across the data center infrastructure, the wide area network, and the customer campus. What is often the case is that the cloud infrastructure is resilient in silos – the server / storage silo, the data center network silo, the wide area network silo, the security infrastructure silo – however resiliency requires that these silos be coordinated holistically. In this way the cloud providers can provide truly self-healing and self-protecting infrastructures.

The Department of Defense (DoD) is diving into this issue of **cloud resiliency**. DoD funded research beginning in 2012 to create a cloud computing environment that can heal itself after a cyber attack. Researchers at the Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory (CSAIL) are working on a new system that would help a cloud identify an attack and recover from it almost instantaneously, according to MIT. The work is part of the Defense Advanced Research Project Agency (DARPA) *Mission-oriented Resilient Clouds* (MRC) project, which aims to create a cloud network, the resiliency of which is based on its ability to adapt.

MIT researchers at the *Center for Resilient Software* (CSAIL) are trying to develop a system that can tell when something is amiss with a network and defend against it as soon as it happens. For the project to be successful, researchers must have an in-depth understanding of how a cloud-computing environment operates. If researchers can understand how behavior of each silo affects the cloud as a whole, they can prevent future attacks.

Researchers at Orange Labs are also working on cloud resiliency in a project called *VESPA – Virtual Environments Self-Protecting Architecture*. VESPA is an open IaaS **self-protection architecture** and framework. It regulates protection of IaaS resources through several coordinated autonomic security loops which monitor the different infrastructure layers. The result is a very flexible approach for self-protection of IaaS resources. The main features of VESPA are:

- policy-based security adaptation based on a self-protection model capturing both detection and reaction phases
- two level tuning of security policies according to security contexts both inside a software layer and across layers
- flexible orchestration of layer-level self-protection loops using system-wide knowledge to allow a rich spectrum of overall infrastructure self-protection strategies
- a layered, extensible architecture allowing simple integration of commodity detection and reaction components.

The VESPA framework, as shown in Figure 7 below, was implemented over a simple KVM-based IaaS infrastructure and applied to perform dynamic VM confinement. At the bottom, a Resource Plane contains the

IaaS resources to be monitored and protected, i.e., managed elements. Over it, a Security Plane contains commodity detection and reaction components that deliver security services such as resource behavior and/or state monitoring (e.g., an IDS component), or reaction and/or resource state and behavior (e.g., a firewall component). These components are the sensors and actuators of traditional autonomic security architectures. The next plane, the Agent Plane, abstracts away security component heterogeneity by defining a mediation layer between the security services and decision-making elements. This plane is built from two hierarchies of agents, one for detection, and another for reaction. The topmost plane, the Orchestration Plane contains the decision-making logic. It is composed of two types of autonomic managers: Horizontal Orchestrators (HOs) that perform layer-level security adaptation; and Vertical Orchestrators (VOs) in charge of cross-layer security management.

Figure 7 – VESPA Framework

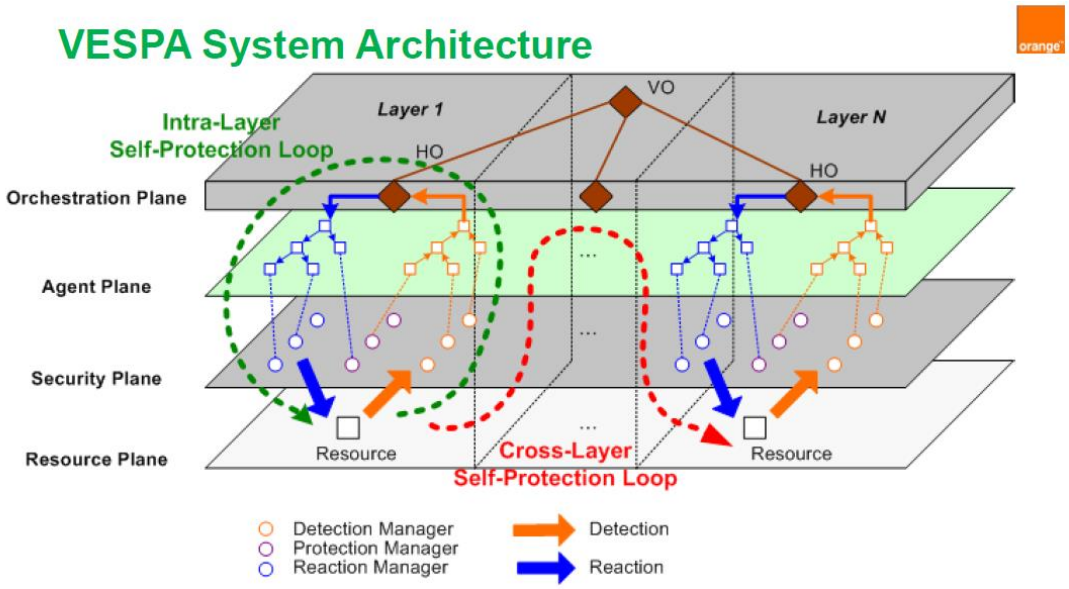

Another research effort focused on resilience uses a cloud-of-clouds paradigm to leverage the availability of multiple or federated cloud environments to create diverse and resilient ecosystems. This architecture aims at providing automated computing resilience against attacks and accidents to complement commodity cloud protection schemes. This enhanced functionality is achieved through specialized "TClouds" middleware standing between low-level, basic multi-cloud untrusted services, and the applications requiring security and dependability.

This middleware is essentially of two classes: **Trusted Infrastructure as a Service** (T-IaaS) and Trusted Platform as a Service (T-PaaS). A user invoking the T-IaaS interface of a TClouds-enabled client-resident software module would transparently obtain resilience of its storage, while invoking a regular storage service interface.

Cloud-COP is still another research effort that builds a trustworthy and resilient, self-healing cloud computing infrastructure out of underlying untrustworthy and faulty hosts. The task supervision and end-user communication are performed by a software mechanism called the Control Operations Plane (COP), The COP leverages provably secure cryptographic protocols that are efficient and robust in the presence of many corrupted participants - such a cloud regularly and unobtrusively refreshes itself by restoring COP nodes from a pristine state at regular intervals. Cloud-COP is distinctive by its judicious use of **proactive secret sharing** (PSS). It implements a cloud control layer that uses the PSS idea to achieve secure self-refreshing resiliency, while

limiting the use of PSS to a minimal set of critical functions to avoid impractical overhead. Included among some of the key features of a COP-protected cloud are:

- Reliable root-of-trust built from unreliable nodes
- Every node is proactively refreshed from a pristine state at regular intervals
- Can tolerate ongoing corruption of cloud nodes, within a limit on corruption rate
- An attacker is unable to target specific jobs since he is prevented from knowing on which node each job is executed.

Cloudlets and IoT

ACD could also be instrumental in securing new "cloud" approaches such as cyber foraging and fog computing or cloudlets.

Cyber foraging allows mobile devices to discover and exploit resources that are around them. One example of cyber foraging is computational offload, which diverts computing from a mobile device to a faster nearby machine, thus saving a mobile device's battery power.

Cloud networks may also improve through *"fog computing" or "cloudlets"*- hubs that can serve as intermediaries between mobile devices and larger cloud services. The hubs could be objects — cars, traffic lights, wireless routers — that are already interacting with the Internet but could give and take data as the need presents itself. This method of data staging could improve data transfers between mobile devices and the cloud by temporarily staging data in transit. Another big advantage of cloudlets would be their ability to **operate despite being disconnected** from the Internet.

Agile cyber defenses would be essential to secure these new cloud approaches. Trusted device identities and on-demand, privacy-preserving access controls would be needed. These approaches could also facilitate a grid of cyber sensors to capture new threat information.

To conclude this section, the cloud offers many opportunities to enable various forms of active cyber defense. Although many of these capabilities are still in research mode, there are still plenty of practical applications that can be employed today. As the tide of cyber attacks rise against cloud service providers, active cyber defenses will need to take a prominent role in proactively mitigating these threats.

ADAPTIVE ENDPOINTS

Endpoints are the new perimeter and special focus needs to be placed on securing them. Poorly secured endpoints can lead to **backdoors for hackers** – especially with the prevalence of phishing attacks where malicious actors use stolen credentials to impersonate legitimate users and bypass enterprise defenses. These hackers leave behind polymorphic and metamorphic malware which pose major challenges in detection and eradication. There are other avenues for attacks that can bypass or overwhelm traditional endpoint defenses, such as difficult-to-detect reflective memory injection and BIOS attacks.

New trends such as BYOD and the Internet of Things (IoT) also raise questions about the ability to secure and manage these new generations of endpoints. As the number of IP endpoints explodes with the **rapid adoption of mobility and IoT devices**, there is an ever-increasing chance of security incidents that exploit them as well. For example, *a recent HP Report* says that 70% of the most popular IoT devices on the market contain major vulnerabilities. There are increasing numbers of wireless exploits that affect mobile users as well. Devices can be fooled into connecting to spoofed networks, authentication to wireless networks can either be cracked or intercepted, and the hackers' ability to capture credentials at a wireless network level has long been established.

Active Cyber Defenses To The Rescue

The cyber defense community has not retreated from this onslaught of threats. These challenges are beginning to be addressed through a variety of multi-layer adaptive security approaches. These approaches include:

- Data-centric protection techniques including self-protecting data
- Behavioral and probability-based methods for detecting polymorphic malware
- Malware dismantling capabilities through kill chain disruption
- Built-in hardware and kernel level protections to combat Advanced Persistent Threats (APTs), such as address space and instruction set randomization techniques, stack guards, and in-line double encryption engines
- Trusted computing techniques, such as measured boot and remote attestation to lock out third-party loaders and bootkits while reliably reporting the security state of the endpoint as part of a network access control capability
- Strongly asserted device identity and device-centric authorization which can make abuse detection easier because of the server's ability to distinguish between your multiple devices and to observe their behavior individually.
- Retroactive security capabilities that can limit exploitation of systems by automatically modularizing software to enable attacker containment.

These proactive defenses can begin to address many of the critical issues at the endpoint if mixed and matched in the right way. Proactive defenses can provide the **runtime security context** which is essential to quickly correlate security issues to defenses. If you can't grasp the context of a security situation quickly, then you are probably toast or on your way to getting burned. Active cyber defenses are all about leveraging context-aware adaptations – that is, the protections or detections can morph based on an awareness of the security state of the endpoint and / or an assessment of the threat environment to which the endpoint is exposed. For example, at a research level, there is *Shield* – an innovative control architecture able to assure E2E security potentially in any application, by dynamically adapting to the underlying systems and using its resources to build the security. The main highlights of this research are:

- The possibility of dynamically discovering and composing the available functionalities offered by the environment to satisfy the security needs
- The possibility of modeling and measuring the security through innovative technology-independent metrics.

Shield leverages a virtual overlay to provide this composability functionality.

There are significant data security and privacy challenges to endpoint security. Active cyber defenses could help in providing **autonomous data protection**, i.e., "self-protecting" data as well as "self-reporting" data. New advancements using Constructive Key Management (CKM), Attribute-based Access Control (ABAC), and secure data compartments (SDCs) that leverage hardware and software controls may provide dynamic but secure approaches to protecting endpoint data.

From a self-protecting data perspective, *researchers at Purdue* and other vendors are developing autonomous security-aware objects (SAOs), which encapsulate sensitive resources and assure their protection. Access to these objects is enforced according to contextual criteria to ensure compliance with location-specific regulations

and service level agreements. SAOs can either use locally pre-loaded policies or securely accept new policies from trusted authorities. Access structures are in accordance with the *Ciphertext-Policy Attribute-Based Encryption (CP-ABE)* schema. CP-ABE supports the notion of attribute-based policies as criteria for encryption. The access structures are embedded as part of the encrypted content. Which user is entitled to decrypt the content, and under which context, are addressed by means of the CP-ABE boolean access structure.

One example of an SAO, being developed by a company called Azos AI, is called CogDat which stands for cognitive data capability. CogDat can sense its situation and autonomously take actions for self-protection — including self-destruction. This method embeds self-protection and intelligence inside the data itself. For example, if CogDat data detects it has been stolen, it can autonomously harvest information about its current environment and send it back to a designated authority and then self-destruct. It also offers protection to data in-use. It dynamically controls computer processes while sensitive data is exposed.

Researchers at Princeton are developing a capability called *DataSafe*. DataSafe provides dynamic instantiation of secure data compartments (SDCs), with hardware monitoring of the information flows from the compartment using hardware policy tags associated with the data at runtime. **Nonbypassable hardware output control** prevents confidential information from being leaked out. DataSafe's software architecture supports flexible, high-level software policies for the data, seamlessly translating these policies to efficient hardware tags at runtime. Applications need not be modified to interface to these software-hardware mechanisms. DataSafe's architecture is designed to prevent illegitimate secondary dissemination of protected data by authorized recipients, to track and protect data derived from sensitive data, and to provide lifetime enforcement of the confidentiality policies associated with the sensitive data.

Encrypting files and the transmission of data help to increase the work factor of attackers. However, it is critical that encryption and key management are implemented correctly and securely for mission assurance. Regular rotation of keys and certificates can reduce exposure to attacks. This rotation can be facilitated by security automation and orchestration. *Constructive Key Management (CKM)* should also be considered as a method of "securing" keys and access to data. CKM is a method where encryption keys are built at the time of encrypting the data and then destroyed. This key is then reconstructed at decryption.

Fighting The Cyber Threat Through Kill Chain Disruption

Endpoint protections against malware are also advancing through the use of **probabilistic / behavioral analysis** and kill chain disruption. For example, *researchers have developed a state machine* that attempts to model a Markov process. The Markov process is hidden, in the sense that it cannot be directly observed. In the context of metamorphic viruses, the hidden process is trained to detect a specific metamorphic family. The training data consists of a sequence of opcodes derived from viruses, all of which were produced by a single metamorphic engine. Once the model is trained, it can be used to score an unknown file, using an extracted opcode sequence, to determine its similarity to the metamorphic family.

Some malware endpoint detection and mitigation vendors are also approaching this challenge through **kill chain disruption.** Figure 8 illustrates the cyber kill chain elements.

Figure 8 – Cyber Kill Chain

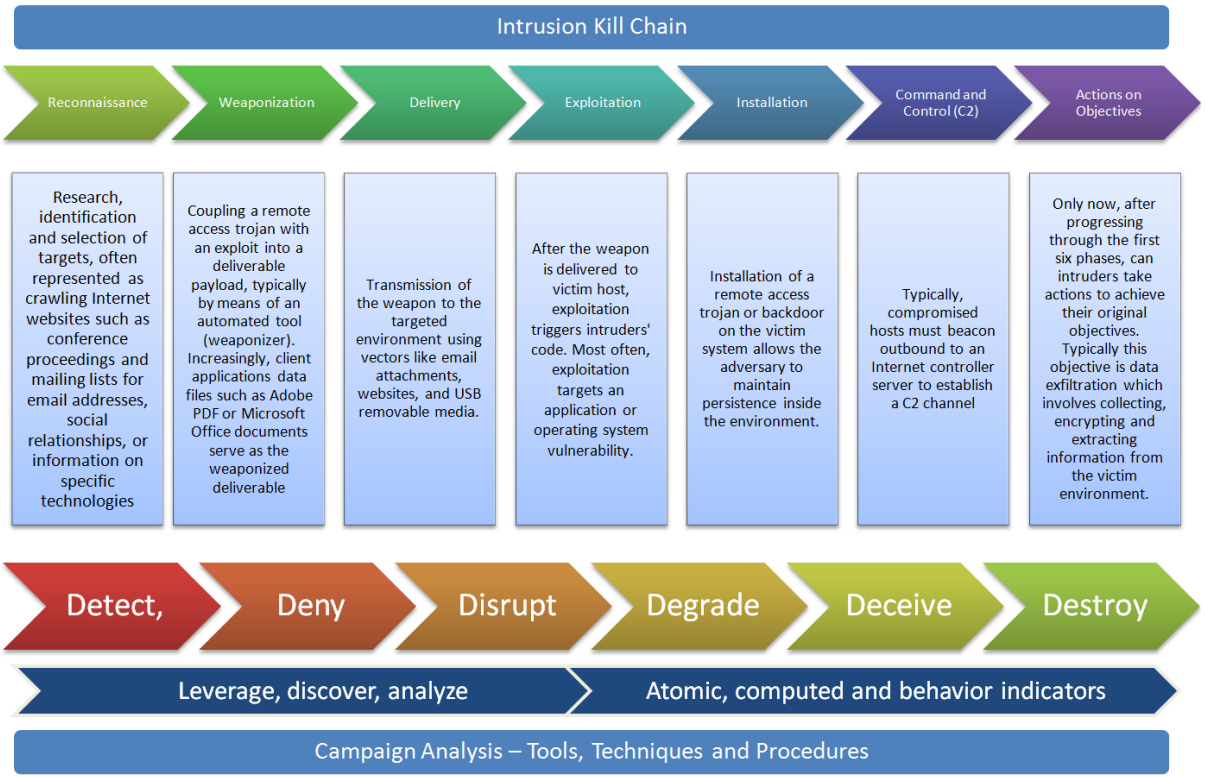

Since an exploit is always based on a chain of techniques, blocking any technique in the chain will block the exploitation attempt entirely. By developing an agent that addresses all the exploit techniques required to execute an attack, the agent can prevent both known and unknown attacks, regardless of security patches or updates on the system. This agent needs to be updated as new exploit techniques are discovered, however, the frequency of updates needed for new exploit techniques is much lower than signature-based methods.

One example of a vendor taking this approach is *Endgame*. Endgame Enterprise uses behavioral analysis, attack chain modeling, and intelligence to detect unknown adversaries and advanced attacks. Sensors inventory thousands of host activities and attributes, including user, system, application, file, and network activity. Endgame's cloud-based analytic engine aggregates application, network, file, and configuration details from the sensors. Using context-based behavioral analysis and attack chain modeling, suspicious behavior is quickly identified and tracked in real time as it evolves. Kill chain visualization identifies malicious behavior and compromised or targeted hosts, enabling an operator to act in time. Full event details are provided to confirm the results of automated analysis. Endgame Enterprise Security Manager guides the security team with advisories that enable fast, effective investigation and response without requiring expert-level skills and knowledge. Endgame Enterprise is powered by a proprietary cloud-based intelligence capability which is fed by a global network of sensors to identify evolving threats providing **early warning** of new techniques and the technology stacks they are being targeted. This intelligence is used by the Endgame Analytic Engine to identify and prioritize known and unknown threats that evade traditional defenses.

Use Hardware Roots Of Trust To Create Adaptable Defenses

By deploying intrinsically secure devices that share provable trust information regarding the security state of the platform, along with infrastructure confirming their trustworthiness, stealthy exploits can be detected. This is particularly important for "unmanaged" platforms such as embedded devices. For example, **hardware-based protections** are useful remedies for memory injection and rollback/replay attacks, as well as BIOS attacks, while supporting secure storage of key material, trusted path, and measured boot. Some examples of these capabilities include Intel's Trusted Synchronization Technology (TST), Intel's Trusted Execution Technology (TXT), Intel's binary instrumentation tool called Pin, Intel TrustLite, *AMD Secure Technology*, *ARM TrustZone*, and *Global Platform*, to name just a few.

At its root, **memory injection is a problem** because processors permit code and data to share the same memory address space. As a result, an attacker can inject his payload as data and later execute it as code using return-oriented programming (ROP) chains. Using a ROP chain to bypass operating system defenses is commonplace and detecting this technique while executing is still difficult. In addition, it has come to light that state actors install implants in the BIOS. However, in practice attackers can install such implants without ever having physical access to the box. Exploits against the BIOS can allow an attacker to inject arbitrary code into the platform firmware. To protect against these sophisticated memory and BIOS exploits, you need to be able to validate all new processes, even those initiated by approved running applications.

Modern CPUs support the detection and resolution of memory conflicts between multiple threads that access the same data: This is called the Transactional Synchronization Extension (TSX) in modern Intel CPUs. Hardware-supported TSX can also be used for security. A special security thread reads protected RAM cells (data or code) in TSX mode; any other (potentially malicious) thread writing to the same cells will cause the CPU to abort the transaction. Changes to memory can also be rolled back. Detecting memory changes with TSX, but without the rollback capability, could also be highly useful for kernel and hypervisor self-protection (such as Microsoft PatchGuard).

The Intel TXT is designed **to combat BIOS threats**. TXT is a set of extensions designed to provide a measured and controlled launch of system software that will then establish a protected environment for itself and any additional software that it may execute. TXT helps to create a Measured Launch Environment (MLE) by taking a cryptographic hash of each process that is launched. One measurement is made when the platform boots, using techniques defined by the *Trusted Computing Group (TCG)*. The TCG defines a Root of Trust for Measurement (RTM) that executes on each platform reset; it creates a chain of trust from reset to the measured environment. Maintaining a chain of trust for a length of time may be challenging for an MLE because an MLE may operate in an environment that is constantly exposed to unknown software entities. To address this issue, the enhanced platform provides another RTM with Intel TXT instructions called a Dynamic Root of Trust for Measurement (DRTM). The advantage of a DRTM is that the launch of the measured environment can occur at any time without resorting to a platform reset.

TST and TXT therefore provide a hardware foundation that assures not only that unauthorized / unwanted applications cannot launch or execute, but also that trusted applications are not modified when launched or while running in memory to compromise the endpoint.

Some trusted hardware capabilities that are based on the TCG **Trusted Platform Module** (TPM) specification provide a cryptographic container to hold attestation keys and other cryptographic keys as well as associated cryptographic algorithms. Over the past several years researchers have attacked implementations of TPMs, such as Microsoft's Bitlocker, through power analysis attacks. These attacks enable the attacker to siphon off the

cryptographic keys from the TPM by examining the power usage while cryptographic operations are being performed. These attacks generally require close proximity to the hardware to execute successfully.

Recently, however, researchers have turned the use of power analysis 180° - from an attack method for extracting secrets to an active cyber defense for detecting malware. The researchers developed the platform - "*WattsUpDoc*" – that can analyze the power consumption caused by the presence of malware in an infected computer. In their experiments, the researchers used WattsUpDoc to detect previously known malware with at least 94% accuracy and previously unknown malware with at least 85% accuracy on several embedded devices—detection rates similar to those of conventional malware-detection systems on PCs. WattsUpDoc detects malware without requiring hardware or software modifications or network communication. This technology may have application in situations where AV software cannot be deployed. For example, there are classes of embedded medical devices that are built with commodity hardware and software and are thus compatible with antivirus or network intrusion detection software (NIDS), but their configurations are commonly off limits to their owners because manufacturers will not support third party software. Some manufacturers explicitly forbid device owners to install OS security updates or antivirus software, under the impression that they cannot certify a device's safety if the software configuration changes. Two hospitals are in beta test with this technology to monitor medical devices for life-threatening malware.

Retro Is Proactive

Another approach that is being developed to better protect endpoints is called **Retroactive Security**. It is apparent from recent breaches that protection-based security mechanisms such as access control and information flow security are not sufficient to deter and prevent attacks. Retroactive Security, which is the enforcement of security or detection of security violations after the execution of a process, is necessary in addition to protection-based mechanisms since not all vulnerabilities can be predicted a priori or managed with prevention alone. One example of Retroactive Security can be found as part of a *research project known as Strata*. Strata aims to automatically modularize software to enable attacker containment. The overall goal is to retrofit a monolithic software system to adhere to two basic security principles:

1. **Privilege Separation**, which posits that resources that require different access rights must execute within different protection domains
2. **Least Privilege**, which posits that each module, running within its own protection domain, must only receive the privileges that it needs to accomplish its task.

Together, these two principles ensure that the attack surface of the modularized software system is minimized, limiting the damage that an adversary can inflict if he were to obtain access to the system. Strata is developing a number of techniques to retrofit software automatically for **attacker containment**.

Active Cyber Defenses For Hard-To-Protect Endpoints

One area that is extremely deficient in security capabilities encompasses embedded systems and the Internet of Things (IoT). As Bruce Schneier pointed out in *The Internet of Things Is Wildly Insecure—And Often Unpatchable*:

> *The problem … is that no one entity has any incentive, expertise, or even ability to patch the software once it's shipped.* Schneier goes on to say: *the result is hundreds of millions of devices that have been sitting on the Internet, unpatched and insecure, for the last five to ten years.*

As the go-to-market race to support the Internet of Things (IoT) begins to heat up, malware producers and hackers are already targeting these pervasive devices. Strong warnings are emerging from governments and threat intelligence sources about vulnerabilities in this target-rich environment. It will be imperative for the IoT

ecosystem of manufacturers, vendors, integrators, enterprises, and consumers to **consider adaptive defenses** to thwart these adversaries in their attempts to undermine and harm this emerging market. As more intelligent devices enter the market, ACD may help improve the state of security for these devices through the employment of lightweight trusted device mechanisms that employ trusted identities and data-centric protections.

As an example of a lightweight trusted mechanism, Intel is developing a capability called *Trustlite*. According to Intel, Trustlite will allow embedded device manufacturers to add features like remote device management, authentication, secure OTA (over the air) updates and **remote attestation** to their embedded devices, regardless of OS and application. The heart of the new platform is something Intel calls an Execution Aware Memory Protection Unit (EA-MPU) — a software-based module that "allows a flexible allocation and combination of memory and peripheral I/O regions without burdening the CPU." Tasks which are designed and believed to implement a particular security mechanism are trusted tasks and are referred to as "trustlets."

Combined with a Root of Trust, such as for Secure or Measured Boot, a range of trusted computing and attestation schemes can be realized depending on the required level of assurance and flexibility. Additional components for Secure Exception Handling, secure Inter-Process Communication (IPC) and secure peripheral I/O can be included to support sophisticated usages like secure user input and secure execution of 3rd party (untrusted) code. TrustLite extends traditional MPU designs by evaluating not only the permissions applied to a particular memory region but also the currently executing instruction address. The resultant **execution-aware MPU** (EA-MPU) allows Trustlite to bind code and data regions into software modules with security guarantees enforced in hardware and independently of the OS. Depending on the desired assurance level, the EA-MPU access rules can be set in hardware, initialized during Secure Boot or by a trusted system service at runtime. In this way, TrustLite enables the secure isolation of software modules independently of the OS or other runtime software. This facilitates the provisioning of various security-sensitive services to platforms with otherwise low security assurance, as it is typically the case for IoT systems.

The concept of having a dedicated microprocessor, with embedded IP to handle the security function at the hardware level is becoming more and more appealing because many, if not most, of future generation IoT devices will function autonomously, thereby being left to its own devices to protect itself. Having one chip that integrates the job of both security guard and controller could be just what the doctor ordered.

One example of such a processor that appeared on the market several years ago from IBM was called *SecureBlue*. SecureBlue was a type of crypto processor referred to as a double encryption device. The IBM rendition offered the ability to protect both the running programs and the data by encrypting both the data and address locations. It put encryptors and decryptors between the processing elements, data storage, and I/O subsystems. All information was decrypted within the secure blocks of the processor and then encrypted before it was stored in memory or sent to an I/O operation. It also had **tamper-resistance key storage** so the keys could zeroize and become virtually invisible to the outside world. It also contained both secure and unsecure I/O channels. The unsecure channels are used for routine I/O operations and maintenance while the secure channels are used for transaction and sensitive data routing. SecureBlue required a few circuits to be added to a microprocessor, taking up a small percentage of the overall silicon real estate, according to IBM. The encryption and decryption happened on-the-fly, without any processor overhead. Although SecureBlue would probably be best suited for use as a cockpit controller for a fighter jet, or at the heart of game console to keep cloners at bay, it generally would be considered overkill for many IoT devices.

Another approach to securing IoT devices was recently announced by IBM researchers and is called *Adept*. Adept combines three components of open source software to produce a self-protecting, self-organizing, distributed infrastructure that may be ideal for IoT. These adaptive components consist of:

- **Bitcoin's block chain** technology – a method to define and secure relationships between entities in a peer-to-peer (P2P) system such as between a device and a user or between different devices,
- **BitTorrent** – a file sharing protocol that supports low bandwidth, distributed P2P environments,
- **Telehash** – a new secure P2P messaging protocol mentioned earlier in the Intelligent Networks section.

Each of these components relies on cryptographic hashes to execute their designed purpose.

The **block chain** is the distributed transaction processing engine that keeps track of **Bitcoin** and other crypto-currencies, however it can be adapted for more than just to support virtual currencies. Basically it's a technology that allows data to be stored in a variety of different places while tracking the relationship between different parties to that data. In practice it uses cryptographic hashes to protect the block chain and to track relationships between devices, between a user and a device and between two devices with the authorization of a user.

BitTorrent is used to help track and distribute the content to be shared between the communicating devices. BitTorrent uses **distributed hash tables** (DHTs) to name and locate objects for sharing in the distributed IoT infrastructure.

Telehash is a messaging protocol built using JSON, UDP and DHTs to send messages between endpoints. It provides an end-to-end encryption library that any application can build on – the end being a device, browser, or mobile app. It works by having every endpoint generate its own unique public key-based address to send and receive small encrypted packets of JSON (with optional binary payloads) to other trusted endpoints. It also provides an automatic routing system based on hash tables to create a full P2P mesh between all endpoints.

IBM is working with Samsung to develop a prototype of this capability. This means someday a smartphone could securely communicate with a door lock without having to rely on a cloud provider to manage these capabilities. Those relationships would be stored on the locks and the phones. The devices would interface through BitTorrent to share files and through Telehash which could send a secure message to unlock or lock the door. By building a platform that keeps the intelligence at the device level, the IoT can operate without a manufacturer's constant attention.

From an architecture perspective, Adept is highly adaptive and secure through its application of cryptographic hashes and distributed hash tables. Distributed hash tables (DHTs) characteristically emphasize the following:

- DHT designs are autonomous, enabling nodes to collectively form the system without any central coordination,
- DHT designs seek to be secure against malicious participants,
- DHT designs are reliable (in some sense) even with nodes continuously joining, leaving, and failing,
- DHT designs are scalable and should function efficiently even with thousands or millions of nodes.

These characteristics seem to be a perfect fit for the Internet where millions if not billions of IoT devices are envisioned as low-cost, low-maintenance devices that should run for years. Having an adaptable but relatively low cost security approach should help to fend off the majority of cyber attacks against the IoT infrastructure. But with this architecture and the use of block chain, one could actually create new business models around sharing more than just data. Devices could share computer power, or bandwidth or even electricity via the block chain's instructions. And while Bitcoins are built to be difficult to mine via computation, there's no need for the Adept platform to rely on scarcity, meaning that the hash tables could track any number of variables.

One catch to Adept is the computational processing power needed to construct block chains. IBM researchers believe that the next generation of IoT devices will be based on more powerful ARM chips rather than 8 bit embedded processors. Or just offload the processing to the Bitcoin network.

To summarize this section: endpoint security can benefit greatly by layering active defenses starting with the use of a trusted hardware layer. Further up the stack, software modularization and data encapsulation techniques can be used to support trustworthy autonomous operations. At the highest layer of the stack, cyber defenses can be self-organized and self-directed by agents that can execute context-aware adaptations – adaptations that are enabled through machine learning and behavioral analysis techniques. These methods provide an **immunity system** that drastically reduces the time needed to sense and respond to today's stealthy APTs.

Autonomous cyber defenses are not just a concept. For example, *MonsterMind* is an alleged US National Security Agency (NSA) program leaked by Edward Snowden in 2014. According to Snowden, it is an autonomous cyber warfare software platform that can watch international connections to identify and "kill" malicious cyber attacks before they hit American infrastructure. This is ACD in a grand scale … or is it preemptive cyber?

SUMMARY

In the past there was a **flawed assumption** that the cyber defenses we built were somehow impenetrable. We, the castle dwellers, sat passively behind our castle walls waiting for an attack to occur. Attackers would find weak spots in our castle walls and exploit these vulnerabilities to attack us. Often, the attack would be long over before any castle dweller would notice the damage inflicted by an attacker.

So the castle dwellers built a moat around the castle with a drawbridge to allow castle dwellers and visitors to come and go. They figured that this perimeter defense would stop the attackers. Unfortunately, the attackers were one step ahead as they disguised themselves to look like castle dwellers to attack the castle. The guards protecting the drawbridge could not distinguish the attackers entering the castle from the normal visitors or castle dwellers and the attackers plundered the castle treasures from the inside.

So the castle dwellers decided to put a **guard** by each room where they kept their most valuable treasures. This worked well until the number of rooms was too many to guard. New guards were hired but they could not be trained fast enough to keep up with the changes in attackers' tactics. And the treasures were often moved to different places inside and outside of the castle. The guards that went with the treasure were **not qualified** to protect it and the attackers took advantage of this shortfall to ambush the guards and steal the treasure.

Finally, a wise wizard visited the castle and told the castle dwellers about Active Cyber Defense. The castle dwellers decided to give it a try and with Active Cyber Defense as their modus operandi, the castle dwellers became vigilant and well-trained. They studied the attackers' methods and built sophisticated sensors and response capabilities that **detect and disrupt the attackers' tactics** before the attacker can establish a foothold inside the castle. The castle dwellers shared information about attacks they saw with the citizens of other castles so new tactics by the roving groups of bandits could be quickly identified and new defenses could be deployed. Some of these new defenses included deception techniques to fool the attackers into revealing their disguises, while other adaptive tactics involved the deployment of ninja guards who were trained to quickly hunt down and subdue an attacker.

With active cyber defense, the castle dwellers established a **highly responsive cyber command and control** (C3) system for all the guards and sentries. This C3 system provides a context-based decision analysis capability that is used by the knights of the castle. The C3 system controls a highly efficient workflow of alerts from the sentries and provides dispatches to the guards to help predict where and when the attacks will come. Now the castle dwellers are not just manning the drawbridge and dumping hot oil after the moat has been breached or the turrets surmounted. With active cyber defenses, the castle guards start the catapults and flaming arrows before the attackers even get out of the woods.

Building active cyber defenses required many meetings of the knights at the Pentagonal Table to develop consensus around defense plans and courses of action against attacks. Sentries had to be trained in new intelligence collection and intrusion detection methods. **Standard methods** of documenting observations were developed. Fast and nimble couriers were recruited to deliver alerts and dispatches from the knights to the guards and sentries. **Specialized squads** of guards were trained to hunt down intruders disguised as citizens. Treasures had to be inventoried and new safeguards put in place when they were moved or accessed. Booby traps, disguises, and camouflage were added to the defenses to **deceive the attackers** and to monitor defenses. New keys were distributed so that only loyal citizens and vetted visitors could access the castle shops.

For the time being, the attackers decided that the castle was too much trouble to bother with, and they decided to go after easier prey. But the knights knew they had to keep looking for new advanced active cyber defenses to keep the attackers at bay.

Do you have an active cyber defense to share with the knights? Will you be the next hero of the cyber war?

Upgrade your cyber security skillset, learn more about active cyber defense, and participate in a growing online community of cybersecurity experts, only at www.activecyber.net.