



Cyber Defense Automation

John Hayes
Founder and CTO of BlackRidge Technology
September 1, 2016

Executive Summary

Cyber defense automation is an imperative. Defensive costs are increasing and automation is necessary to manage the staggering number of devices and endpoints and to better defend our digital systems, resources and assets. Analytics must be integrated with direct policy feedback without requiring human intervention or approval for all responses. Cyber-attacks are automated but still under human direction and control, and therefore cyber defenses must also be automated to have any chance of protecting cyber and digital assets from both the automated attacks and the human element behind those attacks.

In this paper we explore a new approach to cyber defense automation that is realizable and can be achieved in a way that does not burden security operators with false alarms and impede users performing their tasks. By combining statistical process control with network identity and authentication, the number of cyber event errors, aka false positives, can be reduced to an extent that cyber defense can be more automated providing a strong responsive element to your cyber defenses. This also has the effect of increasing the business team's trust in the effectiveness of the security operations to stop attacks and not affect business operations.

Background

Cyber-attacks are a daily occurrence. While the largest attacks make intermittent headlines, the attacks as a whole are unrelenting. Day after day, hour after hour, our computer and network infrastructure, both enterprise and personal, are probed, scanned and attacked in attempts to penetrate and gain a foothold from which subsequent attacks can be staged.

The reasons these attacks are taking place are many, but that is not the focus of this paper. This paper focuses on how automation must be adopted to better defend our digital systems, resources and assets. The attacks of today are largely automated. Computers, as automatic machines, are very good at performing repetitive tasks endlessly. The more computers that can be enlisted and coordinated to work together, the more work can be performed. This is the fundamental principle behind botnets. Botnets are collections of computers that have been co-opted to perform some malicious activity under the control of another computer system. With a herd of computers at their command, botnet masters can instruct their botnets to execute many types of attacks; from simply flooding networks with distributed denial of service (DDoS) attacks, to email spam generation, to exfiltration hosts for storing stolen computer records.

While botnets are not the only type of attack, they do exemplify how computer automation is used for attack purposes. Other attacks may use computer automation as tools in an attack as assists to human operators. These automated tools are used to monitor progress toward an objective, perform evasive maneuvers and detect countermeasures. These tools are often deployed in conjunction with stealthier operations where low observability is the chief priority. Although these attacks may not be as noisy as those performed by botnets, their damage may be more significant. The use of automation versus tools from the attacker's point of view is largely based on the amount of noise generated during the attack.

Detecting and thwarting attacks and cleaning up the aftermath is a difficult task. Most cyber defense systems are neither automated nor integrated. They operate as sets of individual tools which may have aspects of automation incorporated into them. For instance, automating the updating of signature catalogs is necessary but insufficient, because signature based solutions are reactionary

and are unable to detect zero day and polymorphic attacks. An automated cyber defense system must provide better protection than this. The concept of cyber defense automation is not new. The most prominent example of cyber defense automation was the shift from Intrusion Detection Systems (IDS) into Intrusion Protection Systems (IPS). Elements of this can be seen in products by IBM, Cisco, Palo Alto Networks and FireEye.

Automating Cyber Defense

Traditional cyber security solutions are failing. They are unable to scale to the mega- and giga-device populations required now and in the future. They do not provide fine grained attribution capabilities. And most importantly, they are not designed to integrate with security automation systems with direct policy feedback.

A robust automated cyber defense system requires the following characteristics:

- Bounded Operation
- Responsive
- Constraining Behavior
- Topology Independent
- Highly Scalable

Together, these characteristics enable an automated cyber defense system that requires a minimum of operator and administrator actions.

Bounded Operation

Bounded operation is the cyber defense system operating within acceptable error rates. These error rates must account for both false positive (type I) and false negative (type II) errors. Type I errors incorrectly flag an event as being a security event, while type II errors miss classifying an event as a security event. In a statistical analytics system, an event is one that falls outside of the statistical behavioral norm. Taking parameters from statistical process control, an actionable event may be three standard deviations from the statistical centerline, while a warning event is two standard deviations from the centerline. Both actionable and warning events can cause responsive actions. The rate of false positives must be at a level that is acceptable to both security personnel and to end users. Personnel must trust the effectiveness of the security operations to stop attacks and not adversely affect business operations. If the cyber defense system cannot meet these requirements, then no amount of automation will help, because the system will not be used; it will produce too many false alerts and implement inappropriate policies.

There are several approaches to limit errors to acceptable levels; improved sensor data, improved systems design and narrow targeting.

When sensor data is improved, the analytics systems have greater trust in the sensors and are able to make better decisions. The sensor data should be scored, with sensors producing higher quality data leading to higher scored sensor data. One of the easiest methods to increase the quality of cyber sensors and sensor data is to include information that can be efficiently authenticated as part of sensor operation. This is not just authenticating the sensor, but also authenticating the originator of the traffic flows (including the user, device and application generating the flow) indicated by the sensors.

Sensors themselves may have analytical components. Reasoners, machine learning and other analytical components may be directly coupled to raw data sensors. With attached analytical components, sensors can become more accurate in the sensor data they provide and more efficient in the data they communicate to higher level analytical systems.

Systems design can greatly improve the operational bounds of a cyber defense system. The sensors of today are largely existing devices that are reporting their information to a central aggregation system, that were never designed in a control systems sense. These aggregation systems, security information and event management (SIEM) systems especially, have built their business models based on the number of sensors reporting. By providing discovery tools and automatically enabling new network and endpoint devices, these systems increase the number of sensors and the amount of data collected from each sensor. Many times, this additional data is redundant, increasing workload and cost while providing no additional value. At the same time, other areas of the network and endpoints may elude monitoring because they are not easily accessible to the automated discovery tools provided by the SIEM vendors, or the sensors lack the ability or the authority to operate in those environments. Improved system design, where sensors are placed such that they will provide the best quality of information and provide complete coverage is far more effective than randomly adding sensors to produce more redundant data.

The same mechanisms that are used to increase the quality of cyber sensors, such as identity, can also be used to narrow the targeting of actions to specific traffic flows. This allows for more fine grained policy implementation and reporting, increasing the business' confidence in the security operations.

Responsive

Today, most cyber defense responses that result in a change of policy are enacted under the direction of a person. An analytics system issues an alert that is forwarded to an operator who acts on the alert. Often only the highest priority alerts, as determined by the analytics system, are ever seen by the operator. The vast majority of alerts are never seen by an operator. Requiring an operator limits both the amount of cyber security events that can be handled, adds operator latency and adds the potential for operator error. Systems requiring operator interaction have little chance of responding quickly enough to affect the event that triggered the alert.

An automated cyber defense system must be responsive in real time to attacks and anomalous behavior. Policy enforcement must allow implementation of feedback from analytics systems as soon as it is received. An automated cyber defense system cannot require an operator to approve all actions. Operators and supervisory systems should be notified of all actions taken, but must not be integral to most responses.

This is not developing fully autonomous system, as there may be fuzzy factors that would need to be taken in to account for the most extreme events and actions. Certain responsive actions may require operator approval. The goal of an automated cyber defense system is to enable an operation team to respond to automated cyber-attacks quickly and effectively.

Constraining Behavior

One of the challenges of today's security policy technology is that it is generally binary with respect to endpoint traffic; traffic is allowed or blocked, with little middle ground. Constraining behavior is the ability to constrain and limit, on a per flow basis, operations by specific identities without affecting

traffic emanating from the same apparent source address. Good flows and bad flows may emanate or appear to emanate from the same endpoint. The good flows should proceed normally, while bad and suspect flows should be appropriately constrained by redirecting them to alternate resources, applying additional scrutiny to suspect traffic or manipulating traffic flows as appropriate.

This constraining behavior should enable continuous operation, operation where some flows from a device are constrained, while others are allowed to operate normally. Contrast this behavior with the behavior of devices today where an infected or co-opted device is blocked or quarantined until it is remediated. This prevents the device from performing any activity until it has been restored to proper health. In a constrained behavior environment, a compromised device may perform some functions while other functions and their associated traffic flows are constrained.

Sometimes the disengagement of a compromised device would have a more critical impact than allowing a known compromised device to continue to operate. This is especially important when shared resources are used by both trusted and untrusted or compromised systems. Consider a VM environment, where a compromised guest may have also compromised an adjacent guest. Should all of the adjacent guests be terminated, or should those guests be allowed to continue operation, but with their range of actions constrained?

Topology Independent

Network topologies, how the network is operating and routing packets, changes in response to congestion and connectivity failures. Today's cyber defense approaches often couple policy to network topology. Common examples of this are router access control lists (ACLs) and firewall rules. Both of these use policies that are based on network addresses, which may change, requiring coordination with the underlying network to maintain the policy effectiveness. There is also a time element- as the network is changing, topology based policies may be instantiated based on stale and out of date topology information, rendering those security policies, useless or even damaging. When cloud and SDN deployments are included, large portions of the network may be under control of a third party or an orchestration system, making topology coordination even more difficult.

With the advent of cloud computing, the topology of the network is changing constantly and the sensors must change with it. Sensors need to be considered and accommodated by having the sensors incorporated into the underlying cloud infrastructure so that they operate transparently to the cloud orchestration infrastructure.

Automated cyber defense policies must operate independently of the underlying networks, requiring no topological knowledge or coordination for the policies to operate correctly.

Highly Scalable

With the Internet of Things (IoT), network connected devices are estimated to reach populations of 20 to 50 billion by 2020. At these populations, devices outnumber people on the earth by about 2.5 to 5.5:1. The only way to manage everything is with a highly scalable automated system.

Scalability must be enabled in several forms; using computationally efficient algorithms and distributed architectures. Using computationally efficient algorithms may include the increased use of authenticatable whitelisted identities versus catalogs of blacklisted behaviors. Shifting away from a reliance on deep packet inspection and other technologies that have poor computational complexity scaling will assist in this effort. Distributed architectures that increase authentication and verification of correct operation as device populations increase will also assist in the algorithmic

efficiency effort. The architecture of scalability must support not only 50B things in 2020, it must also contemplate 10s of trillions of things.

Another advantage of efficient, scalable algorithms is the ability to monitor and apply policy to cyber traffic and behavior in a variety of environments; enterprise managed data centers, third party cloud providers, remote offices and mobile devices. This enables multiple deployment models; network centric, endpoint centric, cloud and legacy systems which must all be addressed.

Scalability must also be applied to the analytic engines that produce policies. An automated cyber defense system must accommodate multiple analytic systems with varying sensors and decision capabilities. The overall framework that enables this may even have hierarchies of analytical systems with supervisory analytics systems monitoring and adjusting the results from the primary analytics.

Realizing Automated Cyber Defense

There are many challenges to automating cyber defense, but they can be implemented with three types of components; informational, analytical and responsive. Informational components are the various sensors within a cyber system. Analytical components make decisions based upon the aggregated information provided by the various sensors. Responsive components implement the policies specified by the analytical components. This resembles classical control theory in a Multiple Input Multiple Output (MIMO) configuration. A summary of which informational, analytics and responsive component address the characteristics required for automated cyber defense is shown below.

	INFORMATIONAL	ANALYTIC	RESPONSIVE
BOUNDED OPERATION	X	X	X
RESPONSIVE		X	X
CONSTRAINING BEHAVIOR		X	X
TOPOLOGY INDEPENDENT	X	X	X
HIGHLY SCALABLE	X	X	X

Informational Components

Informational components are all about the sensors. Sensors have a calibration that provides information within specific tolerances providing known levels of data fidelity. Sensors may produce different information at different tolerance levels, but the receiving analytic system must know what tolerance to apply to each reported measurement. This enables the sensors to provide the necessary input for bounded operation.

From a network perspective, sensors need to operate independent of the underlying network topology, reporting the network traffic from their deployment viewpoint. At the same time, physical sensors must be physically secured and tamper evident to prevent their removal or relocation. When sensors are deployed on endpoints, they must be cryptographically secured to their host, again to prevent removal or relocation. In order for sensors to provide trusted measurements, each sensor

must have a unique identity that can be easily authenticated. Identity authentication enables analytics systems to discern both missing and inauthentic sensors.

Sensors may have some level of internal processing capability. Intelligent sensors may even have local analytical capabilities to enable them to detect complex signals and filter them before communicating them to a higher level analytical system. A challenge with deploying intelligent sensors is that they must be designed and implemented to insure that the sensor itself is not providing a new attack surface that is vulnerable to compromise. All sensor data, including data internal to the sensor must be protected from attacks.

Sensor authentication and sensor data must be easily authenticated and sensor data should be efficiently collected and securely communicated to the analytics system efficiently and within an expected period of time. To address scalability of large number of sensors, intermediate sensor aggregation points may aggregate and preprocess data on the way to the analytics systems. This is done primarily to limit the volume of sensor data that must be ingested directly by the analytics systems. Most sensors have an upper bound limit for reporting determined by the interface bandwidth supported by the sensor itself.

Analytical Components

Analytical components make decisions based upon the information provided by the various sensors. The analytical components may be a single analytics engine, or may be a collection of analytical engines with different sensors and differing outputs. Analytical components should also take into account the scoring or weighting of each of the sensors. While analytical systems may be topology aware, they must not rely solely on topology information.

Analytical components provide correlation, analytics and statistics over time. Sensors for the most part are stateless, and responsive components are slaves to the analytical components. It is the analytical components that maintain persistence, state and baseline behavioral information over time.

Analytical components must have an inventory of their associated sensors and the tolerance of each sensor to properly analyze their reported information. The tolerance of each informational component should allow the scoring or weighting of the generated information, on the basis of source sensor identity and authenticity, the source sensor location and quality of information produced by the source sensor. This raw data provided by each sensor, along with its scoring information, is used by the analytical components to make decisions.

When anomalous behavior is detected, the analytical components will send policy changes to the responsive components, including the target range of the traffic flow or flows to be affected and the intended constraining behavior. The targeting may be as narrow or as wide as is appropriate. When information is low and the risk to a high value digital asset is high, a broader target may be specified, such as “all identities in the finance group” which may become more narrow and specific as more information is learned.

Analytical components may be arranged hierarchically. This makes the overall system operation easier to scale. For example, a local sensor may have a local analytical component and a local responsive component. The information generated by the analytical component may be communicated to a higher level analytical component and may also receive response control information from the higher level analytical component. This model may be repeated forming a

hierarchical architecture, where each analytical component also functions as an information and responsive component to higher level analytical components.

Responsive Components

Responsive components recognize defined cyber security events and perform an action under the control of the analytical components. Security events may be network traffic emanating from a specific address, physical interface or from a specific identity. It is assumed here that the responsive components have the same ability to narrowly target traffic with the same specificity as the analytical components. By being able to target traffic based on elements that can be authenticated, such as identity, both bounded operation and topology independence is achieved. Constraining behavior can be achieved by enabling traffic to be redirected to alternate network resources, such as resources with limited access or control authorities.

Responsive components, because they directly affect the network traffic, are amongst the most vital components to protect. It is important that responsive components only respond to trusted, authenticated communications from the analytics system. Ideally, the responsive components should be communicated with via an out of band communications from the perspective of the network traffic being responded to. The out of band communications may be implemented physically or cryptographically. In all cases, it is also recommended that the responsive components have the minimum possible attacks surface, including cloaking and the use of non-interactive authentication protocols to establish the responsive communication channels.

In hierarchical environments, responsive components may be coupled with analytical components. When responsive information is sent to the responsive component, it may also effect the tightly coupled local analytical component. Responsive components are also usually sensor components and report when they have performed their responsive task in the presence of matching network traffic, providing additional information to the analytical components.

A Systems Approach to Automation

All of the preceding characteristics must be brought together with a systems approach. A properly functioning cyber defense system cannot be constructed haphazardly with whatever is laying around. It must be thoughtfully conceived, architected and implemented. The analytics systems providing the feedback must be tuned to the quality of the sensor information present. When these characteristics are present and properly implemented, automated cyber defense systems can be realized.

Building an operational automated cyber defense system has a number of system aspects. Unlike traditional SIEM systems that use existing sensors, the sensors in a proper automated cyber defense system must be designed to insure coverage. This design includes the optimum number of sensors and where those sensors should be placed. It includes where redundant sensors should be placed and the informational degradation from losing a sensor versus the cost to have backup sensors. This systems approach also applies to responsive components. The system design should include the optimum number of responsive components and where those components should be placed. It should also address responsive component failures and redundancy.

System Architecture

Networks have natural aggregation points whereas endpoints and applications only see traffic destined for that specific endpoint and application respectively. However, network reliability and redundancy effects may not allow the network to be a reliable choke point – and policy enforcer. Thus these approaches must be designed to support both network centric and endpoint centric deployments, with hybrid deployments having both network and endpoint resident components.

In addition to allowing hybrid deployments with both network and endpoint resident components, there is another hybrid aspect that uses both identity based and behavior based policies. Identity based policies allow for strong authentication and enforcement within a framework of rules and authorities. However, using Identities alone does not address misusing identity within the provided authorities and does not address lost or stolen identities.

By combining authenticated identity with sensors reporting activity to analytical systems, the analytical systems have high trust that the behaviors are being correctly attributed to the correct identity. When deviations are determined by the analytical systems, changes in policy can be sent to responsive components that will affect only that specific identity. This is an example of high quality sensor data (identity) being reported to an analytical system, the analytical system detecting an anomaly and causing a change in policy by sending updated policy information that is narrowly targeted, affecting only the identity in question.

Automation Summary

Cyber defense automation is a realizable goal; it can be achieved in a way that does not burden security operators with false alarms and impede users performing their tasks. In the next section, we show how BlackRidge Technology products can be used with analytics to realize this goal.

Using BlackRidge for Cyber Defense Automation

BlackRidge products provide authentication before allowing the establishment of TCP/IP sessions, stopping the kill chain for both known and unknown attacks. Further, authentication of TCP/IP sessions enables identity attribution before allowing access to network resources.

BlackRidge products are based on Transport Access Control (TAC) technology. TAC uses cryptographic identity tokens to authenticate TCP/IP sessions and apply security policy before each session is established. TAC securely conveys identity at the establishment of every TCP/IP session and authenticates the conveyed identity prior to allowing any response. The TCP protocol that underpins the Internet does not allow identity credentials to be exchanged until after the TCP session is fully established, exposing critical resources. TAC enhances TCP and closes this vulnerability, effectively cloaking any protected network resources from network reconnaissance, port scans and all other forms of unauthorized access.

TAC works with all TCP/IP based applications and is compatible with existing networking and security infrastructure. TAC interoperates with networking and security equipment of all standards compliant TCP/IP equipment vendors, preserving investments that have already been made. TAC can be deployed without requiring all network devices to be TAC enabled, easing migration and deployment. In addition, TAC can be used to communicate Identity across multiple independent administrative domains, without regard to the underlying equipment technology or vendor.

TAC Addresses Requirements for Cyber Defense Automation

There are five requirements for Cyber Defense Automation and we will now review how BlackRidge TAC addresses each of these:

- Bounded Operation
- Responsive
- Constraining Behavior
- Topology Independent
- Highly Scalable

Bounded Operation

Bounded operation is achieved with TAC by using authenticated Identity as the primary object for both sensors and responsive operations. TAC addresses two approaches to bounded operation by limiting errors to acceptable levels, improving sensor data and narrow targeting.

TAC improves sensor data by providing attribution information in the form of authenticated identity information along with the TCP/IP session information to analytics systems. This is the earliest possible time that attribution information can be provided to the analytics system. Additionally, TAC can provide the statistical confidence it has in the attribution information being provided. From a trust perspective, attribution information is higher quality than tradition TCP/IP session information alone, since addresses in TCP/IP sessions can be easily spoofed and should not be used authoritatively for security policy. In a cyber defense automation system using TAC, sensor information provided by TAC that includes attribution information would be scored higher than sensor information that does not include attribution information. A finer grained approach could incorporate statistical confidence into the scoring of the attribution information.

TAC achieves narrow targeting by using Identity instead of using source IP address. Identity is narrower because multiple identities may originate or appear to originate from a single IP address. This occurs when identities are concentrated together in provider networks that use network address translation (NAT) or in systems the use multiple containers that share a single IP stack and address. When targeting is based on source IP address, the affected target is potentially much broader than necessary and often blocks or otherwise disenfranchises good and trusted network traffic along with blocking the intended malicious traffic.

Responsive

TAC has a responsive capability that is accessible to analytics systems called Trust Feedback. Trust Feedback enables an analytical system to change the trust level of an identity or identity group. Trust level is a score of how much authority is allowed to a given identity or group of identities. Trust levels are arranged from highest trust (more authority) to lowest trust (least authority). The policies assigned to each trust level describe the provided authorities and the resource access they allow. Because the policies are predefined in their association with trust levels, dynamically changing trust levels for an identity or a group of identities results in fully deterministic behavior.

An example of trust feedback in operation:

Charlie has an Identity which is inserted into each TCP/IP session using TAC. Charlie is in group Finance and the group Finance, operating at "high trust", has the authority to access the finance

server, the corporate email server, a development server and the server used to initiate corporate banking transactions. All TAC authenticated TCP/IP sessions are sent to an analytics system that also uses sensors from other cyber security systems. Normally, Charlie never accesses the development server, but today the analytics system detects Charlie accessing the development system in an unusual way. The analytics system changes Charlie's trust level from "high trust" to "medium trust" and communicates this change to the TAC system via Trust Feedback. "Medium trust" for group finance has the authority to access the finance server, the corporate email server and a development server. The server used to initiate corporate banking transactions is not accessible at the "medium trust" level. Because the analytics system changed the trust level only for Charlie and not for all of finance, only Charlie is blocked from accessing the corporate banking server. Other authorized persons in the Finance department can continue to access the server.

The policies for "high trust" and "medium trust" are pre-configured for the system. In this example, the behavior change occurs automatically, without any operator intervention or approval. This demonstrates how an automated cyber defense system can be responsive in real time to attacks and anomalous behaviors. Policy enforcement by TAC implements policy change information received via trust feedback from the analytics systems as soon as it is received.

Constraining Behavior

TAC provides constraining behavior by allowing a wider range of policy options that "allow" or "block". A policy can also redirect a TCP/IP session to an alternate resource. The alternate resource may be a server with a reduced authority. An example of reduced authority is a database with read access instead of read/write access. The alternate resource may be a honeypot, designed to gather more information about the behavior and intentions of the given identity, or it may be a remediation service.

In the above example of Charlie:

In an alternate policy for "medium trust", accesses destined for the server used to initiate corporate banking transactions is sent to an alternate server where the transactions are recorded, but never executed. This allows cyber defense operations to monitor the transaction in a secured sandbox, without endangering the real treasury. Operations by other, still trusted, persons are able to be processed by the corporate banking server normally.

Topology Independent

Automated cyber defense policies must operate independently of the underlying networks, requiring no topological knowledge or coordination for the policies to operate correctly.

TAC implements policy based on Identity instead of source IP address like many traditional cyber and network security tools. This enables TAC to operate independently of the network topology and operate where source IP addresses are allocated dynamically and when users operate from multiple locations, such as different offices and campuses.

Not only does TAC operate independently from network topology, TAC does not require any coordination with network topology changes. This is a distinct advantage over systems that use IP addresses to determine security policy. Examples of topology dependent systems include router access control lists (ACLs) and firewall rules. Both of these use policies that are based on network addresses, which may change, requiring coordination with the underlying network to maintain the policy effectiveness. There is also a time element- as the network is changing, topology based

policies may be instantiated based on stale and out of date topology information, rendering those security policies, useless or even damaging. When cloud and SDN deployments are included, large portions of the network may be under control of a third party or an orchestration system, making topology coordination even more difficult.

Highly Scalable

The BlackRidge TAC products have been analyzed with populations exceeding 100,000,000 tokens. BlackRidge TAC products handle millions of concurrent TCP/IP sessions. These numbers are for a single BlackRidge TAC instance. Multiple TAC instances, can support arbitrarily large Identity populations.

TAC products are available in physical network appliances, software instances for cloud deployment and as integrated software modules for OEMs, enabling monitoring and cyber policy enforcement in a variety of environments; enterprise managed data centers, third party cloud providers, remote offices and mobile devices. This enables multiple deployment models; network centric, endpoint centric, cloud and legacy systems.

Analytical Systems Integration

BlackRidge TAC products provide high quality informational sensors and constraining responsive behavior. BlackRidge does not provide big data and behavioral analytical solutions. Instead, BlackRidge TAC products are designed to integrate with a variety of analytical systems, including traditional Security Information and Event Management (SIEM) systems, behavioral systems and knowledge based systems. By providing trusted attribution information in standards compliant formats, it is easy for analytical systems to intake the BlackRidge sensor information. And by using BlackRidge Trust Feedback, which uses RESTful interfaces, constraining, responsive behavior can be easily implemented. This approach allows sophisticated customers the flexibility to choose the most appropriate analytical systems for their automated cyber defense deployment.

Conclusion

BlackRidge TAC makes Cyber defense automation a realizable goal without burdening security operators with false alarms and impede users performing their tasks. BlackRidge TAC provides trusted attribution information to analytical systems and provides a deterministic mechanism for responsive, constraining behavior. BlackRidge TAC also protects against a number of threats and attacks, is interoperable with network and security equipment from multiple vendors, provides centralized or distributed policy, supports and spans multiple simultaneous administrative domains and provides strong, authenticatable identity in all deployments. TAC is software designed to be ported to various hardware, virtual and embedded platforms.