# Active Cyber Defense: A Vision for Real-Time Cyber Defense

MJ Herring, KD Willett

*Information Assurance Directorate*
*National Security Agency, Fort Meade, Maryland, United States*
*E-mail: JIWfeedback@nsa.gov*

**Abstract:** *Cyber operations consist of many functions spanning cyber management, cyber attack, cyber exploitation, and cyber defense, all including activities that are proactive, defensive, and regenerative in nature. A subset of cyber defense, Active Cyber Defense (ACD) focuses on the integration and automation of many services and mechanisms to execute response actions in cyber-relevant time. ACD is comprised of a set of logical functions to capture details from enterprise-level architecture to operational realization with the primary objective to become a living part of DoD cyber operations to help defend the nation from cyber-based adversaries.*

**Keywords***:* *Active Cyber Defense (ACD), Cyber Defense (CD), Sensing, Sense-Making, Decision-Making, Message Fabric, Shared Situational Awareness, Automated Response Action, Coordinated Response Action*

## Introduction

In their *2013 Data Breach Investigations Report*, Verizon notes that while 24% of the initial compromise stage of data intrusions takes minutes or seconds, the predominant number of initial compromises take hours. These breaches consist of a series of actions performed in real-time that lead to a persistent malicious presence in the targeted network. Per the Verizon report, discovery of malicious activity by network owners is currently on the order of months, meaning that malicious actors have time to exfiltrate terabytes of data and perform other malicious acts that are unlikely to draw attention in a timely manner.

Recognizing the need to accelerate detection and response to malicious network actors, the United States (US) Department of Defense (DoD) has defined a new concept, Active Cyber Defense (ACD) as “DoD’s synchronized, real-time capability to discover, detect, analyze, and mitigate threats and vulnerabilities” (DoD 2011). Among the many needs of war-fighter operations, there is the need to be secure, which includes the concepts of hardening, protecting, attacking, and defending among the war-fighter domains of land, sea, air, space, and cyber. Cyber is an integrating capability for the other domains, as well as a standalone domain that has its own unique needs for cyber defense.

Cyber defense includes three complementary categories: ‘proactive’, ‘active’, and ‘regenerative’. ‘Proactive’ activities harden the cyber environment and maintain peak efficiency for cyber infrastructure and mission functions. ‘Active’ activities stop or limit the damage from adversary cyber activity in cyber-relevant time. ‘Regenerative’ activities restore mission effectiveness or efficiency after a successful cyber attack. These categories form a continuum of cybersecurity activities occurring continuously and simultaneously on networks, integrated by a common framework of automation that includes ACD as a subset of integrated cyber defense. The focus herein is on ACD.

ACD is purposely designed to be applicable across the U.S. Government (USG) as well as Critical Infrastructure and Key Resources (CIKR). This facilitates the reuse of ACD-related solutions across the USG and CIKR. Such leveraging is fiscally responsible to the U.S. taxpayer and ultimately minimizes the total cost of ownership for ACD across the USG. The Information Assurance Directorate's (IAD) role in ACD is the design of a deployable infrastructure that will allow a properly authorized defender to set up and initiate a defensive response to the threat. Implementation of the platform will be dependent upon collaboration and agreement with network owners. Our specific focus is defending DoD internal networks (for example, from the network boundary into and including the host) through the integration and automation of existing cyber-security solutions. The motivation behind the IAD's focus is derived from

- National Security Directive (NSD) 42 (The White House 1990) that "establishes initial objectives, policies, and an organizational structure to guide the conduct of activities to secure national security systems from exploitation; establishes a mechanism for policy development and dissemination; and assigns responsibilities for implementation",
- Executive Order (EO) 13587 Independent Assessment (The White House 2011) states "the Secretary of Defense and the Director, National Security Agency, shall jointly act as the Executive Agent for Safeguarding Classified Information on Computer Networks" and,
- the IAD's decades of experience defending DoD networks.

Attacks in the non-cyber domains require physical proximity and time to execute (for example, a bomb must be close to a target; a bullet must physically hit its target). Cyber is unique in the lack of need for physical proximity to execute an attack (that is, anyone with an Internet connection is a potential participant in this worldwide battle space) and in the vastly reduced time required to perpetrate an attack (for example, bits on a wire travel much more quickly than traditional troops or munitions). ACD addresses this vastly reduced time necessary for a successful attack by integrating many solutions to provide response actions in cyber-relevant time.

Cyber-relevant time is a purposely vague term that accommodates the needs of the battle space. If the battle space is a Central Processing Unit (CPU) and Random Access Memory (RAM), and the combatants are software applications vying for control, the cyber-relevant time is nanoseconds to microseconds. If the battle space is between two computers of close physical proximity, cyber-relevant time is milliseconds to seconds. For a battle space between two computers on opposite sides of the world communicating via satellite links, cyber-relevant time is seconds. With live operators and delays inherent in cognitive processing, key strokes, and mouse clicks, cyber-relevant time is seconds to minutes. The requirements for ACD increase as the adversary becomes smarter and quicker.

Cyber defense includes employing non-real-time big-data analytics to find trends in historical data repositories; likewise, cyber defense includes actuarial-like predictions of future events. The ACD monitoring activity may provide data feeds to these analytics, and the ACD sense-making activity may take influence from these analytics in the form of decision support algorithms;

however, these historical and future analytics are outside the scope of real-time processing and, therefore, outside the scope of ACD.

## ACD as a Capability

A comprehensive ACD solution requires the integration of many tools. The complexities of ACD can never be entirely captured in a single tool. ACD functionality may occur within a single platform, but this is one example or one thread through ACD and not the entirety of ACD. Moreover, ACD functions may be geographically dispersed: sensing may occur in Hawaii; sense-making may occur in Washington, D.C.; decision-making may occur in U.S. Cyber Command (USCYBERCOM), and acting may occur in the European Command (EUCOM). The ACD design must accommodate a wide spectrum of such scenarios with performance occurring in cyber-relevant time. Therefore, the approach is to design ACD as a capability expressing desired results that may consist of an indeterminate number of tools that provide those results.

The primary beneficiary of ACD is the decision-maker. Decision-making is the act of selecting the best choice(s) among available options. Each decision-maker receives guidance from decision-drivers in the form of externally imposed authoritative mandates (legislation, regulation, directive, instruction, Executive Order), negotiated mandates (contracts, service-level agreements), or self-imposed mandates (internal policy, standards, procedures). Deriving ACD requirements includes decomposing each decision-driver into data elements necessary to make the decision. For example, operations standards for a particular mission may include requirements for specific values in a series of Windows registry entries. The ACD administrator decomposes this operations standard to identify all the parameters that represent compliance with that standard. The parameters represent requirements for decision-support, which is the information necessary for the decision-maker to decide.

The next step is to translate the decision-support parameters into data sources in the asset space and source data on each respective asset. Once the data needed and their location in the cyber-asset space are identified, appropriate sensors can be deployed to retrieve the data. Decision-makers provide requirements for the content necessary to collect (as just described) as well as the necessary frequency of data collection. The guiding principle to determine frequency is the question, ‘How old is too old to make an effective decision?’ For a compliance decision that occurs annually, collecting once a year is adequate. For critical operational decisions under emerging threat conditions, collecting once every X seconds or minutes may be more appropriate. This variety implies the need to dynamically configure sensors to accommodate changing mission needs. An increase in collection data volume and frequency will affect cyber environment performance. Because the cost of increased security can be decreased mission performance, there is a tradeoff between security and ease of use/performance.

ACD follows the principle of ‘collect once and reuse many’. This means that data elements collected for ACD will be used for multiple purposes, including decision-support to areas outside of ACD. For example, the same data elements used to ensure operational security may also help to determine Federal Information Security Management Act of 2002 (FISMA) compliance. In some cases, the primary motivation to collect data may be for ACD, and these same data may be useful in these other areas, such as certification and accreditation decisions, Command Cyber Readiness Inspection (CCRI) decisions, or FISMA compliance decisions. In other cases, the primary motivation to collect data may be for some other reason (for example, to validate operational service-level agreements), and these same data elements may be useful to ACD. Whatever the initial motivation to collect data from cyber assets, there is an even greater motivation to reuse that data in as many decision paths as it is applicable. This goal promotes smart and efficient workflow. Achieving this efficiency requires aware and intelligent management of integrated cyber defense to understand the needs and function of the individual parts in the context of the entire operation.

## ACD Constituent Parts

ACD consists of six functional areas shown in **Figure 1:** sensing, sense-making, decision-making, acting, messaging and control, and ACD mission management. These logical functional areas are used to delineate role, fit, function, and dependencies. Sensing is ongoing observation with intent to provide awareness. Sensors are devices or people who make these observations and obtain a snapshot of current operational states. Sense-making uses analytics to provide understanding in a particular context (for example, mission, operational need, local security configuration). Each decision-maker will have a unique context within which to make decisions. ACD accommodates the automation of decision-making as well as the cognitive supplement of human decision-makers (that is, ACD may provide decision support).

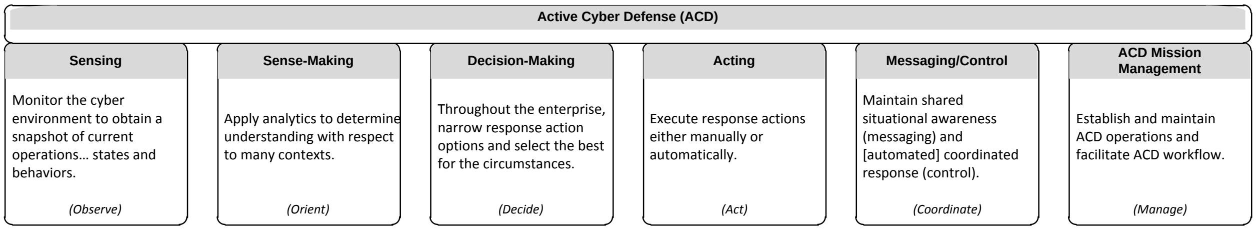

Sensing
Monitor the cyber environment to obtain a snapshot of current operations… states and behaviors.
*(Observe)*

Sense-Making
Apply analytics to determine understanding with respect to many contexts.
*(Orient)*

Decision-Making
Throughout the enterprise, narrow response action options and select the best for the circumstances.
*(Decide)*

Acting
Execute response actions either manually or automatically.
*(Act)*

Messaging/Control
Maintain shared situational awareness (messaging) and [automated] coordinated response (control).
*(Coordinate)*

ACD Mission Management
Establish and maintain ACD operations and facilitate ACD workflow.
*(Manage)*

**Figure 1:** ACD Functional Areas

The objective of decision-making is to select among available response actions. ACD accommodates the ability to execute ACD-internal-automated-response actions as well as the ability to prompt external actors with action recommendations. The action decision ultimately resides with the cyber-asset owner. ACD does not impose *de facto* actions out of the control of the asset owner. While ACD calls out a decision-making function, decisions are made throughout the ACD workflow and outside of ACD using input from ACD. Decision-maker roles include, but are not limited to ACD administrators, cyber-asset owners, mission commanders, and security-operations personnel. Each has his/her own context and decision-drivers. There remains the challenge to establish real-time precedence and adjudication to resolve inevitable conflicts (for example, national policy requires X, but local operating needs require Y) to determine who wins and provide defensible justification.

Messaging and control is the heart of situational awareness and coordinated response actions. Key operational gaps in ACD are the lack of a common communication medium (for example, message fabric) to interconnect all ACD-related tools at speed and scale, the lack of a standard interface for tool connection to the common communications medium, and the lack of a standard message set understandable and actionable by all connected tools. Upon successful realization of messaging and control, all ACD-related tools will have the ability to make each other aware of current activity (that is, to achieve shared situational awareness). Similarly, messaging and control will enable the tools to coordinate response actions that may include disseminating the same response action among similar assets (that is, a vulnerability mitigated in one is a vulnerability mitigated in all) or a more sophisticated combination of defense actions that hit multiple layers of network defense to preempt adversary attack and/or prepare the enterprise to weather an active attack.

ACD mission management provides ACD internal control of workflow where the scope of control is limited to the operating environment of that particular instance of ACD. There is no universal management of ACD. The ACD Reference Architecture (currently in draft) is to guide many instantiations of ACD—some of which will be standalone operations, and some of which will connect in varying degrees of coordinated operations. Participation in any semblance of a federated ACD operation is purely voluntary, and participants choose their level of participation. Participation is not imposed and certainly does not take place without the knowledge of the cyber asset owner(s).

To reiterate, ACD is not a single solution; it is a capability to provide context and interoperability among many solutions under the six functional areas. An integrated, cohesive ACD solution implies the use of many sensors, analytics, and displays to support many decision-makers. For example, ACD may accommodate any number of analytics from any number of perspectives. The type and focus of the analytic is dependent upon the needs of the decision-maker who will use the results of the analysis. ACD intends to accommodate what is available today (current tools) and what will be available tomorrow (future tools yet unknown). This leaves room for new, better, faster, and cheaper solutions across all functional areas.

## ACD Operational Concepts

No single government entity will own ACD. ACD is a capability within cyber defense with the unique differentiator of providing situational awareness and response actions within cyber-relevant time. The only way to achieve this is to integrate many dozens of tools across the ACD functional areas. Sensing will include sensors and sensor subsystems (that is, sensor-management systems). ACD may have some native sensors (that is, sensors controlled by an ACD instance), but ACD will more likely interact with sensor-management systems that, in turn, directly control the sensors. One example of this is the Host-Based Security System (HBSS). ACD does not intend to directly touch any HBSS sensor; rather, ACD communicates with the ePolicy Orchestrator (ePO) server that, in turn, controls many HBSS sensors. A variety of sensors and sensor subsystems are necessary to monitor cyber assets (for example, Windows desktops and servers; UNIX; mainframes; network infrastructure, including routers and switches from many vendors; phone equipment; mobile equipment; industrial equipment, including Supervisory Control And Data Acquisition (SCADA)). No one sensor or sensor subsystem can watch them all.

Sense-making includes a wide spectrum of analytics that convert raw data into information for decision-support. Decision-makers are at every organizational level and include computer operators, system administrators, operations managers, leveraged security services, program managers, investment managers, policy makers, and governance. NASCAR provides a useful analogy. In NASCAR, the car driver is the operator and is looking for details on speed, proximity of other cars, fuel level, tire pressure, laps-to-go, and position relative to the leader. The pit crew provides technical support in refueling, changing tires, and repairing the engine. The race track owner is concerned with track facilities, racetrack schedules, parking, and attendee safety. The racing commission is concerned with marketing, maintaining policy, coordinating all race schedules, ranking drivers, and providing overall governance. Each of these decision-makers is related under the sport of NASCAR. However, each has very distinct roles and distinct decision-making needs. Moreover, as NASCAR is not the only type of car racing, car racing is not the only type of sport, and cyber is not the only war-fighter activity. ACD will support a subset of these decision-makers. In fact, data collected by ACD may support more decision-makers outside of the scope of ACD.

An important distinction between contexts in the above decision-making analogy is that the racing commission really does not have a need to see all the speedometer readouts of every car actively racing on every track. Likewise, it makes no sense to display an average speed of every car actively racing on every track. There is no operational decision to be made from this information, and the racing commission should not reach out and step on the gas pedal or stomp on the brake of any particular car as that is the job of the driver. However, the racing commission is very much interested in collecting information on car performance, pit crew performance, and race-track results to make policy decisions that further the overall interest, performance, and safety of stakeholders. Of particular interest to the racing commission is profitability, which implies overall cost management and operational efficiency. While USG use of ACD is not concerned with profitability, ACD does contribute to cost management and operational efficiency via automation.

Upon receipt of decision support from sense-making, decision-making selects the best choice(s) among available options that ultimately leads to some action. Acting is the performance of a sequence of steps resulting from choices made in decision-making process. Actions may be manual or automated. A key principle in ACD design is for actions to be automatable and not inherently automatic. Automatable may be considered similar to a plane’s automatic pilot. Under certain circumstances, automatic pilot is useful; however, the live pilot makes the choice to use it or not. ACD operates under the same principle. If an operator is comfortable with certain actions being automated, then that operator can flip the switch on and let ACD do its thing from beginning to end. If operating conditions are such that any automatic change to the operating environment is undesirable, then the operator will flip the switch off and queue ACD-recommended actions for review, approval, and execution by the live operator.

Messaging and control is the ability for ACD to provide shared situational awareness via standard communication methods and to provide coordinated response actions via standard control signaling with a standard message set. Shared situational awareness is informational only; that is, ‘Here is a heads up on what we see’ or ‘Let us talk if this is of interest’. Coordinated response actions are technical devices informing other technical devices of recently performed or imminent actions and requesting/directing these other devices to take action as part of an overall coordinated response.

ACD mission management covers the establishment and maintenance of overall ACD, and it facilitates workflow through the functional areas. For example, if a sensor needs updating, ACD mission management handles the update; if a new analytic becomes available, ACD mission management inserts it into the appropriate area. Moreover, ACD facilitates workflow through sensing, sense-making, decision-making, acting, and messaging and control. By separating out the management functions, each functional area may focus on its particular role. This separation increases design and operating efficiency by isolating common functions under mission management versus duplicating them in each functional area.

## ACD's Role in Broader Cyber Operations

ACD is a part of overall cyber defense that itself is but part of the broader cyber operations in support of mission execution. From a DoD perspective, the overarching concept is Computer Network Operations (CNO) that consist of Computer Network Attack (CNA), Computer Network Exploitation (CNE), and Computer Network Defense (CND). The current DoD Instruction 8530 (2001) addresses CND, and the latest draft revision renames CND to cyberspace defense. Any addition of a qualifier by nature reduces the scope (for example, the set of red cars is a subset of cars: all red cars are cars, but not all cars are red). Similarly, ACD is a subset of cyber defense: all ACD is cyber defense, but not all cyber defense is active.

## The Operating Environment

From a DoD perspective, ACD exists within a federated operating environment in which most decisions with regard to the cyber assets are made by the respective asset owners. ACD may recommend and may facilitate automated responses that the asset owners agree to; however, ACD will not override asset-owner decisions. An important point is that ACD does not come prepackaged with foregone conclusions that will take over local network operations.

## Operational Goals and Objectives

The cyber attacker's goal is to generate a desired effect brought about via three objectives of get in, stay in, and act. The cyber defender's goal is to minimize threat efficacy brought about via three objectives of keep out, throw out, and restrain. The desired effect may be tactical (for example, destroy a database server), or it may be strategic (for example, disrupt troop deployments that result from the database server output). ACD provides support to defend against both by first identifying and defending tactical targets and providing situational awareness to tip and cue mission-assurance-related activities.

## General Cyber Attack Sequence

A generic attack sequence to get in is to enumerate the cyber environment, find vulnerabilities, gain access, escalate privileges, insert malware, and operate malware to the desired effect.

A generic set of stay-in activities are proliferate, avoid detection, and persist. Proliferate implies malware duplication with the intent of hedging attacker odds against detection of any single malware instance. Avoid detection implies hiding (for example, rootkits that insert hooks and modify operating system commands or common processes such as dynamic link libraries). Persistence is surviving through various conditions of rebooting, software patches, and other system modifications.

A generic set of act activities are designed to perform function and produce results. Perform function is the running of the exploit. The nature of exploits varies widely and includes every kind of malware that may attempt unauthorized disclosure of data, denial of service, or unauthorized modification of data. The production of results includes both a tactical result (for example, destruction of data) and a strategic result (that is, mission implication). ACD predominantly addresses the tactical result; however, parts of ACD work in complement with mission assurance to detect and respond to the adversary's intended strategic result.

ACD intends to monitor for the presence, state, and behavior of attacker attempts to get in, stay in, and act. For example, ACD sensing looks for enumeration behavior on the network (that is, activity that is mapping the network). Sensing looks for behavior that is attempting to identify vulnerabilities and monitor behavior as well as states that indicate unauthorized access, unauthorized privileged-user presence and activity, the presence of malware, and the activity of malware. Upon detection, the ACD workflow continues through sense-making, decision-making, and acting.

## ACD Operational Example

An operational ACD capability might be a system of individual cyber-security solutions already deployed on a network (for example, Intrusion Detection System (IDS), Intrusion Prevention System (IPS), virus scanner, adaptive firewall) integrated together to produce an enterprise-spanning, holistic ACD capability. ACD may also be a set of unique tools integrated in a single platform to provide active defense against a specific threat vector in a specific portion of a network (for example, network or enclave boundary, host). Additionally, while mature solutions will ideally be vendor agnostic, based on common open standards and common command-and-control messaging, early instantiations of ACD will by necessity have proprietary elements.

As an early example of a platform-based ACD capability, the SHARKSEER solution is a collection of best-in-breed commercial products integrated into a single suite that provides active defense against zero-day attacks. This innovative capability includes a non-signature-based network sensor that identifies known malware in near real-time using government-enhanced, commercially developed signature and heuristic cloud technologies, and two behavior-based sensors where one sensor focuses on identifying real-time malicious human driven behaviors and the other on malware behavior. SHARKSEER is integrated on a state-of-the-art platform that provides high-speed, low-latency communication among the component parts and adds the ability to block malicious connections. This government enhanced integration of commercial products is an early instance of an operational ACD system designed by the IAD for defense of DoD networks and is capable of identifying and defeating rapidly evolving, previously unknown attacks that the individual products working alone cannot defeat.

To measure the fit of the SHARKSEER solution, sensing is first examined using the ACD functional framework. The SHARKSEER solution uses network flow sensors to route live network traffic through three streaming analytic capabilities. In sense-making, malicious-human-behavior analytics identify antecedent behaviors on incoming connections that are related to attempts to compromise the network. Real-time, signature-based analytics examine incoming traffic for indicators of known malware or files with bad reputation scores. Malware-behavior-based analytics examine incoming files and Uniform Resource Locators (URLs). In decision-making, alerts from the sense-making analytics stimulate simple courses of action passed to acting. In acting, SHARKSEER really performs two basic functions: blocking and passing packets. The messaging and control function is provided by the hardware platform hosting the SHARKSEER solution.

SHARKSEER is intended for deployment at the network boundary (between internal and external networks) or at enclave boundaries (subdivisions of the internal network) and is designed to integrate with enterprise-wide holistic ACD solutions as other capabilities emerge. SHARKSEER is already integrated with enterprise email sensors and host-based systems via two deployed private clouds. Also part of the DoD Joint Information Environment Single Security Architecture and the Joint Regional Security Stack, SHARKSEER is slated for near-term deployment to provide defense for critical DoD networks.

## Conclusion

The IAD's work on ACD is complete when ACD becomes a living part of DoD cyber operations. At that point, ACD will be in a maturity model that upon reaching the highest level of maturity will need to evolve the functional areas to sustain itself at that level. Part of this maturity considers that future requirements for ACD will evolve that cannot all be anticipated today. For this reason, the ACD architecture and systems engineering are capability-based and not tool-based. Capabilities are expressions of desired results, agnostic of the solutions that produce those results. Tools come and go with changing technology; capability needs are more enduring. The tools of ACD today will very likely be different ten years from now; however, the desired results from ACD should be close to the same. The IAD's ultimate vision for ACD is a capability that becomes a living part of how DoD operates and is conducive to emergent behavior so that the results of ACD as a whole are more than the behavior of the sum of its parts.

## Way Forward

The way forward for ACD from a functional perspective involves dozens of parallel activities spanning all ACD functional areas. These activities occur predominantly in commercial development, some with and some without government sponsorship or explicit government requirements. A key activity to produce a cohesive ACD capability is integration. This implies the need for a common communication medium (for example, message fabric), standard interface, and standard message set. The goal is to adopt, adapt, or develop this common communications medium, standard interface, and standard message set for vendor use to create products that may become part of standardized ACD. Each individual product brings its unique value, and the whole of ACD becomes more than the sum of the parts as tool interoperability provides for the realization of cohesive and adaptive ACD operations.

Achieving the vision of a standardized ACD includes engaging the National Institute of Standards and Technology (NIST), commercial vendors, industry leaders in security concepts and technology research, and appropriate USG governance bodies across DoD, civil agencies, and the intelligence community.

ACD is a first glimpse at the broader capability of security automation, which is to maintain a state of being free from danger or threat within acceptable risk-tolerance boundaries with little or no human intervention. The benefits include workflow efficiencies, process coordination, priority task execution, and intelligent resource allocations. The potential pitfalls include race conditions, gridlock, thrashing, and subverting parts of security automation for means other than their designed intentions. A clever adversary may turn poorly designed security automation into an attack tool that works against itself. These examples do not discourage security automation; rather, they raise awareness for careful security automation design, including the design of ACD.

## References

Department of Defense 2001, *Department of Defense Directive 8530.01, Computer Network Defense,* Department of Defense, Washington, D.C., United States.

—2011, *United States Department of Defense Strategy for Operating in Cyberspace*, Department of Defense, Washington, D.C., United States.

The White House 2011, *Executive Order (EO) 13587, Structural Reforms to Improve the Security of Classified Networks and the Responsible Sharing of Classified Information*, The White House, Washington, D.C., United States.

—1990, *National Security Directive (NSD) 42, National Policy for the Security of National Security Telecommunications and Information Systems*, The White House, Washington, D.C., United States.

Verizon 2013, *2013 Data Breach Investigations Report*, viewed 5 February 2014, <www.verizonenterprise.com/DBIR/2013>.