From: 800-171comments@list.nist.gov on behalf of [redacted]
To: 800-171comments@list.nist.gov
Subject: [800-171 Comments] Comment for Initial Public Draft of NIST SP 800-171, Revision 3
Date: Thursday, July 13, 2023 1:28:29 PM
Attachments: 800-171_R3_DKI_Comments.pdf

To whom it may concern,

Thank you for the opportunity to respond to the NIST SP 800-171, Revision 3. Recently, I have started my own cybersecurity company and this feedback is based on working with a number of industry participants to include service providers, infrastructure operators, cybersecurity companies, and companies across critical infrastructure sectors and Defense Industrial Base (DIB).

Best regards,

John Lewington
President and Founder
DARKNIGHT INDUSTRIES

PROTECTING CONTROLLED UNCLASSIFIED INFORMATION IN NONFEDERAL SYSTEMS AND ORGANIZATIONS

SP 800-171 Rev. 3 Comments

John Lewington
President and Founder
DARKNIGHT INDUSTRIES, LLC
[redacted]

DARKNIGHT INDUSTRIES
[redacted]

National Institute of Standards and Technology (NIST)
Attn: Computer Security Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8930)
Gaithersburg, MD 20899-8930

Dear Ron Ross and Victoria Pillitteri,

Thank you for the opportunity to provide feedback on Protecting Controlled Unclassified Information (CUI) in Non-federal Systems and Organizations - SP 800-171 Rev. 3 (Draft). This feedback is based on working with industry participants to include service providers, infrastructure operators, cybersecurity companies, and companies across critical infrastructure sectors and Defense Industrial Base (DIB). Comments, feedback, and recommendations are focused on increasing the specificity of security requirements to remove ambiguity, improve the effectiveness of implementation, clarify the scope of assessments, and help organizations better manage CUI risk across a variety of different factors and environments. As the threat environment and cloud environment continues to evolve, the government may want to further address how controls:

- Need to go beyond practicing good cybersecurity hygiene by continuously and formalized CUI content in the form of data within Microsoft Word, Excel, PowerPoint, and PDF files from unauthorized access and release.
- Associated with continuous monitoring, compliance and security for servers, containers, and cloud assets across hybrid and multi-cloud environments can be implemented and audited while balancing other system needs.
- Make incident response management and data loss prevention for CUI feasible at scale across hybrid and multi-cloud environments, and supply chain systems and organizations.

A whitepaper has been provided upfront to better explain the 800-171 security comments and recommendations. Detailed comments in the NIST provided template are at the end of the document.

Thank you,

John Lewington
President and Founder
DARKNIGHT INDUSTRIES, LLC

Table of Contents

DARKNIGHT INDUSTRIES
5304 Carlton St.
Bethesda, MD 20816

Background

A critical factor for achieving success is the ability to share CUI and collaborate effectively and efficiently while satisfying the security and privacy requirements for protecting that information. Federal agencies and contractors routinely generate, use, store, and share content and information applications, databases, and in the form of data within Microsoft Word, Excel, PowerPoint, and PDF files that, while not classified, still requires protection from unauthorized access and release.

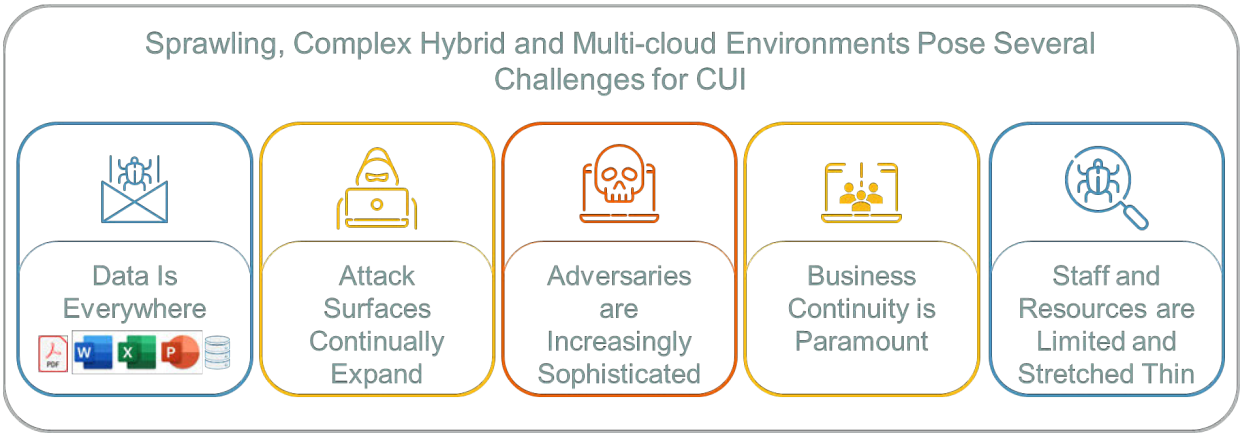

Today, information systems and content are becoming more dispersed, cloud-centric and containerized, mobile, and shared across different IT infrastructures and multi-cloud environments. These environments are subject to different types of stewardship to include protecting CUI. How is security and compliance for all servers, containers, and hybrid and multi-cloud assets unified?

For example, the 800-171 states the data encryption should happen at rest and in motion if an organization is to have good hygiene, but all encryption methods aren't equal. Technology decision-makers must evaluate each approach against the threat models for the environments they manage. For instance, whole-disk encryption defends against physical theft of the drive. Moving up the stack to network protection measures like Secure Sockets Layer (SSL), Transport Layer Security (TLS), or virtual private networking (VPN) also can pose potential issues. Data is encrypted at one end, only to be decrypted at the other end and may still expose CUI to unauthorized entities. What happens when one of those files is removed, whether maliciously or unintentionally, from those existing access controls?

If CUI is exposed to unauthorized entities, what are incident response tactics, techniques, and procedures (TTPs)? How can the government or contractors immediately remediate the situation and prevent further data loss? How can the protections implemented by one contractor be shared with other organizations?

Thoughts on Protecting CUI and Other High-Value Assets

More must be done to protect CUI content than practicing good cybersecurity hygiene and we recognize there is no "Silver Bullet" for cybersecurity. Data is everywhere: on devices (e.g., laptops, desktops, mobile devices), in applications running in both on-premises and outsourced

environments, and in the cloud. This distributed nature of data complicates the process of establishing and maintaining CUI. In responding to the draft, we kept in mind that security requirements need to work across:

- A variety of devices, on different operating systems, and different computing infrastructures.
- Legacy systems that don't always work well with newer ones.
- Different file formats to enable flexibility and ease of use for those with authorized access.
- Systems and organizations to make security management for monitoring, preventing data loss, and tracking electronic documents that contain CUI feasible at scale.

Our recommendations for the 800-171 take an evolutionary path and cut across the listed security requirements below and are detailed in the NIST provided the template to follow.

- Audit and Accountability (3.3)
- Incident Response (3.6)
- Maintenance (3.7)
- Security Assessment and Monitoring (3.12)
- System and Communications Protection (3.13)

As the government and contractors modernize their cybersecurity architecture and look towards adopting a zero trust architecture, identity and access management and data-centric security become front and center. The challenge with data-centric security technologies such as Digital Rights Management is they can be complex and time-consuming to implement and maintain as it depends on organizations knowing what data they have, what its metadata, and what security and privacy requirements it needs to meet so the necessary protections can be achieved.

Therefore, organizations may want to move up the technology stack to the network layer and consider an Open Extended Detection and Response (XDR) security solution to provide a holistic data loss prevention (DLP), visibility, detection and response across hybrid and multi-cloud environments. As organizations adopt zero trust architectures, an Open XDR solution can provide important visibility and analytics in assessing the cybersecurity posture. Key benefits may include:

- Using a network TAP so that the Open XDR system itself cannot be detected or attacked during network monitoring.
- Integrating with 3rd party endpoint agents, which can provide automation around responses.
- Providing a deceptive tactics capability that shifts the advantage from the attacker to the defender by using decoys and lures to trap attackers in the deception layer, giving analysts the time to detect attackers earlier, study their moves, and defeat them before damage can be done.
- Ensuring protection of CUI using session reassembly services. DPL processing may be not sufficient, as malicious insiders can easily circumvent these systems by obfuscating or embedding sensitive data in a large benign payload.

The mapping out of data security features across network, endpoint, cloud platforms, SaaS apps, operating systems, plus web, email and cloud gateways becomes the modern-day challenge. No one vendor is the best-of-breed answer for protecting CUI and other data within hybrid environments. The modern-day security stack consolidates legacy security silos into features to

improve content, context and enable automation for improved detection and response. Data loss prevention (DLP) solutions need to be content and context aware for effectiveness, not a large source of alerts and noise impeding security analysts. For data theft and loss, regulatory compliance, intellectual property protection, sensitive data use monitoring, advanced threat detection, sandboxing, investigations, retrospective analysis and hunting, network traffic analysis at high speeds for all ports and protocols may need to be considered as an elevated as a critical security requirement.

Furthermore, the adoption of cloud infrastructure has left many organizations managing security requirements across hybrid deployments—environments in which some applications and infrastructure have been moved to or built in the cloud, while others remain in a data center. Research shows that 87% of organizations that are in the cloud use two or more cloud service providers requiring organizations to start aligning their business systems and processes to hybrid and multi-cloud environments while ensuring seamless security and visibility. However, running numerous applications can create large, fragmented environments where IT teams struggle to keep track of the CUI data. Most of these incidents stemmed from misconfiguration or a lack of visibility or access management. Meanwhile, adversaries and malicious insiders strike as often as every 2 seconds. Therefore, we recommend NIST consider incorporating hybrid and multi-cloud security requirements to help prevent CUI data loss.

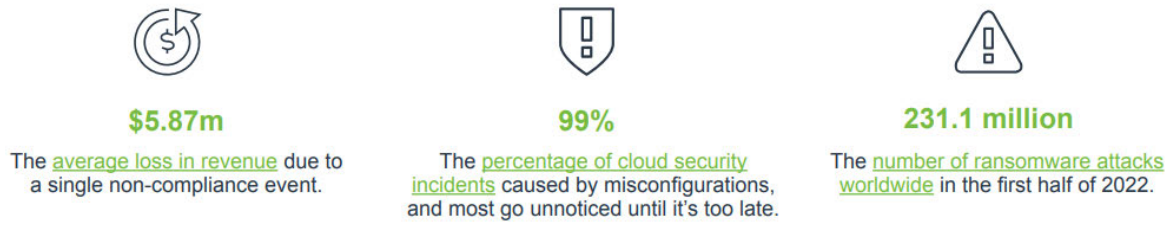

Cloud-native application protection platforms (CNAP or CNAPP), a Gartner analyst defined term, combines two groups of cloud-native security solutions. Typically, cloud security posture management (CSPM) and cloud workload protection platforms (CWPPs) are separate tools. CNAPPs offer the capabilities of both in a single, unifies security and compliance across the infrastructure asset lifecycle to help protect, detect, remediate, and continually improve security for on-premise servers, containers, hybrid and multi-cloud assets. CSPM platforms provide automated monitoring for potential cloud vulnerabilities like misconfigurations or compliance

violations. CWPPs also scan for vulnerabilities, but they focus hardening and configuring applications before implementation, then actively monitor for threats when these workloads are running. Whereas CSPM platforms protect cloud environments from the outside, CWPPs manage internal protections.

CNAPPs combine these functions to provide visibility and security control management across all cloud functions, inside and out. They use extensive automation to monitor and respond to threats from development to end-use, including extra protections like identity management. Combining the capabilities of CSPM and CWPP solutions allow CNAPPs to address both their limitations. While traditional security can be a tremendous help, their limitations can create holes in businesses' cloud security, especially when they rely on multiple devices from various services. Hybrid infrastructure requires that both existing and new security compliance expectations and operational requirements be addressed. Cloud and container requirements add new security and compliance challenges, as achieving security compliance may focus on new components and the exposure of data center assets to cloud-hosted resources.

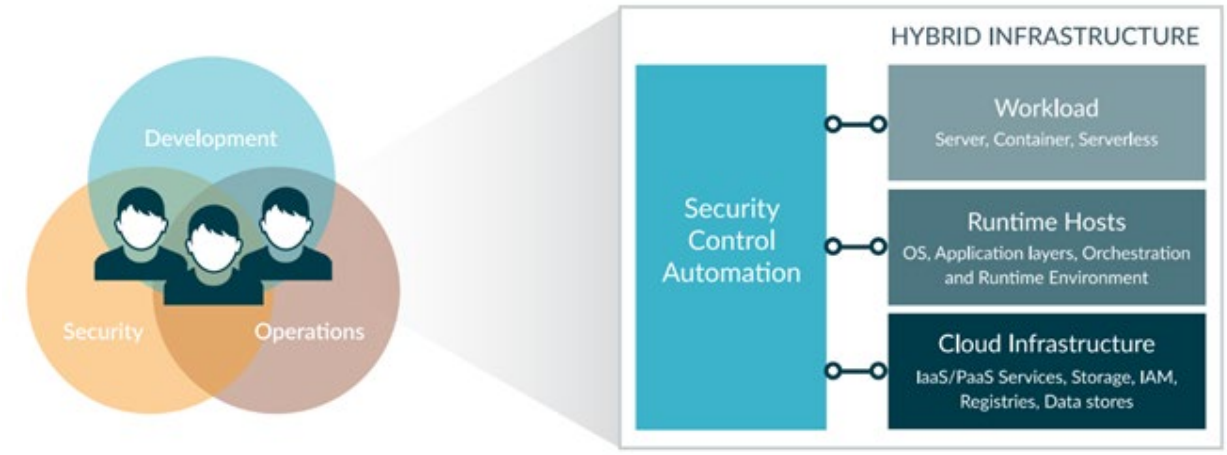

Traditional on-premise enterprise applications adhere to comprehensive security and compliance control standards listed in the 800-171. Attack surfaces for hybrid cloud deployments span multiple layers of abstraction and control—including the Cloud Service Provider (CSP) control plane, container orchestration and management, server and container workloads, and DevOps pipelines.

The number of asset types available has expanded with IaaS and PaaS resources offered by CSPs, and the diverse technical components used in cloud applications introduce their own unique security concerns. It is imperative that these concerns be addressed in alignment with existing security requirements for data center assets. Policies which consider and cover all resource asset types are a necessity, and security controls must be embedded at each layer of control. Unification of these security functions within one platform can provide complete visibility across all layers of a hybrid deployment, protecting against oversight which could leave CUI exposed.

Data Loss Prevention for CUI

In today's dynamic infrastructure environments, networks provide countless potential entry points for cyber adversaries. These breaches can cause a network to shut down as soon as the adversaries

gain access, most notably the theft of CUI and other intellectual assets, and the threat of ransomware. To avoid falling victim to an expensive data breach, security requirements may need to proactively protect systems, provide deep visibility, and actively calculate cyber risk of all assets and communications in, out, and through the network, to stay ahead of the next attack.

Network DLP – An Integrated Feature of the Network

Network DLP is a key security requirement within a larger security stack for detection and response for CUI. The leading industry analysts recommended defense for detection and response of advanced threats and data theft is the combination of endpoint and network visibility augmented with sandboxing, threat intelligence, anomaly detection and full metadata capture for real-time and retrospective analysis. DLP alerts are also an important data source for machine learning based anomaly detection of insider threats and anomalous user behaviors. Thus, security requirements of the past are being addressed by integrating features of more comprehensive security solutions. Cloud, SaaS and mobility have driven DLP towards an integrated feature where data is always in motion.

For regulation compliance, network DLP is a core functionality to validate secure transmission of data on approved channels, plus an important visibility perspective for monitoring sensitive information use. Combined with endpoint visibility, intellectual property and confidential information protection can be addressed on and off networks. Network DLP is also the fallback and most stringent set of controls after first pass analysis of integrated DLP features within email, web and cloud gateways. While monitoring and detection alerts are low risk and popular, preventing data loss and theft as it happens has more value given prevention controls do not impede business operations.

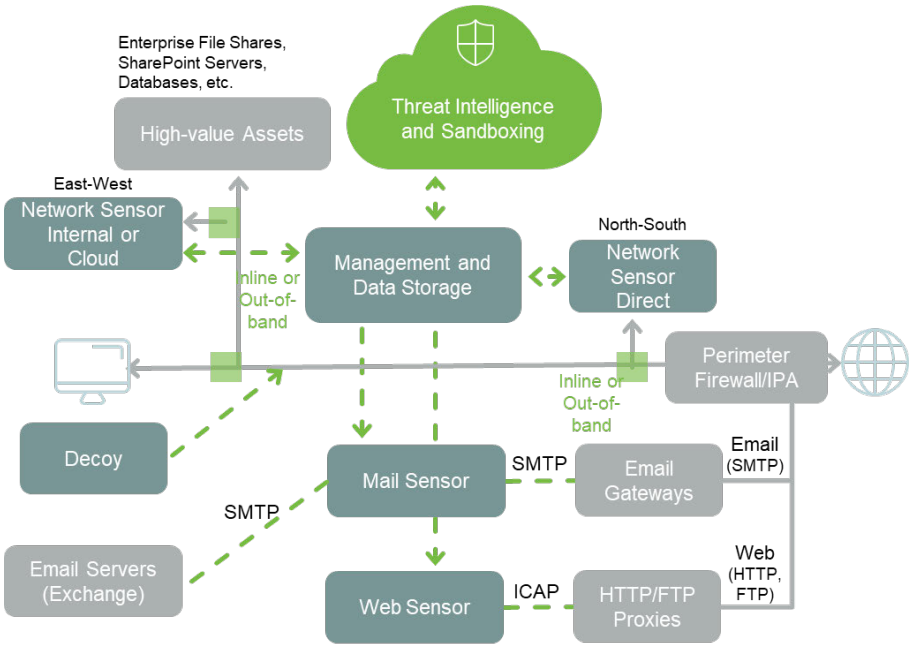

Security requirements should address data loss and theft across all ports and protocols, which is important as most proxy server deployments do not always have this breadth of visibility and

control. Also, content inspection engines should go beyond exact pattern matching and use data profiling, plus avoid sampling or dropping traffic to achieve performance levels. Enterprise scalability for multi-gigabit networks is required with minimal to no impact on network performance as well. While DLP solutions can prevent CUI leaks without requiring a web proxy or ICAP integration, the requirement to have a web sensor with ICAP integration does improve encrypted traffic visibility. Security requirements may also include an email sensor and metadata storage and analysis for up to 360 days.

Ability to Stop CUI Loss

The "P" in network DLP is for prevention. Unfortunately, many solutions claim to be DLP solutions are data loss detection or alerting solutions with no prevention capabilities, and many others have very limited prevention capabilities (often dependent on 3rd party integrations). Detection of data loss or theft is necessary but doesn't protect an enterprise against the harmful consequences of data leakage. Many organizations and technology solutions lose sight of the security goal: detection enables a report on the state of the organization's compliance position and describes when policies were breached — it doesn't stop the breach from occurring, only prevention can.

Organizations must demonstrate they are exercising adequate care to avoid data leaks of sensitive or protected information while stopping data loss and ensuring that digital assets are secured. While one can be compliant with certain regulations by solely reporting a breach after it occurs, that does nothing to protect the organization from the ensuing loss of CUI, customer mistrust, potential civil and regulatory penalties, and the massive distraction of the investigation and reporting of a breach.

When looking at DLP security requirements, it is important to understand what type of traffic can be prevented by the solution. Internet Protocol (IP) networking defines 65,535 ports, which are sub- addresses or logical locations allowing two computers to connect simultaneously over a variety of protocols. While some protocols still observe their official defined port, the port/protocol paradigm is unfortunately no longer dependable for security controls. Some network services have been deployed on non-standard ports to enable multiple services on the same machine. However, a majority of the traffic running on non-standard ports is likely to be traffic attempting to evade controls. As applications have evolved, they have implemented port hopping to find ways to work through enterprise firewalls. This is particularly true of social networking, online conferencing and text messaging and this presents a significant risk for data leakage. Recent evidence shows that usage of online gaming and streaming platforms can be used to exfiltrate sensitive data as these services were not given the scrutiny common web and email services are.

Network DLP solutions can address the requirement for all 65,535 ports unlike proxy-based solutions. Covering all ports is required to prevent data leakage for all outbound network traffic. It is also important to evaluate the security architecture used for prevention. There are two main categories of prevention requirements:

1. Network sniffer, data link layer of the OSI model, capable of preventing traffic on all 65,535 ports:

- Out-of-band prevention via session poisoning with TCP RST (reset) packets.
- Inline dropping of packets.

2. Proxy/Gateway, application layer of the OSI model that views traffic on a specific protocol on a specific port:
   - Proprietary gateway with dedicated proxy/gateway.
   - Integration into existing proxy/gateway solution.

Each of these architectures addresses different prevention requirements, impact to the network infrastructure, and user experiences. Any solution should address both requirements to provide the benefits of each approach.

Accurate Inspection of CUI

While a DLP solution must be capable of preventing data leakage across all ports on the network, the solution must also be able to accurately detect sensitive or protected information to prevent leaks of digital assets. If the system cannot detect the information, it cannot alert on or prevent the unauthorized disclosure from occurring.

When performing content analysis, the possibility of false positives and false negatives appearing must not be ignored. Significant false positives can make a system difficult to manage. Hence, a lot of discussion is focused on or around false positives. But false negatives present a much greater risk for data leakage and compliance than false positives. False negatives mean that protected information is disclosed without the system generating an alert, thus bypassing the prevention and remediation processes required under internal policy and external regulation. A false negative could expose an organization to finding out about a leakage of digital assets from a third party, law enforcement or public exposure — what the DLP solution was deployed to avoid in the first place.

Many will attempt to discuss false positives, or the incorrect generation of an alert, without discussing false negatives and the total cost of ownership. All three are significant when addressing network DLP security requirements.

	Definition	Error Type	Impact
False Positive	"When a test incorrectly reports that it has found a positive result where none really exists." (Wikipedia)	Type 1 Error	An alert is generated for valid network traffic that does NOT violate policy.
False Negative	"When a test incorrectly reports that a result was not detected, when it was really present." (Wikipedia)	Type 2 Error, Miss	An alert is NOT generated for unauthorized network traffic that does violate policy.

Visibility Is a Key Component of Accuracy. Before discussing content analysis methods, it is important to understand the level of visibility required for a network DLP solution for network traffic. If the solution cannot analyze or understand the network session because it doesn't

understand the protocol/application or isn't looking for it outside of standard network ports — it will trigger a false-negative condition anytime sensitive information is sent in this fashion.

First, it is important to understand what traffic the network DLP solution understands. If a DLP solution advertises "all channels", "all protocols", or "all applications", it is important to understand what that means. Most network DLP solutions have an unknown protocol decoder that attempts to find information. However, these unknown decoders are highly inaccurate, often suffering from both significant false positives from noise on the network and false negatives from the inability to recognize the information in the session. The list of actual decoded protocols provides a much more realistic view into the information understood by the network DLP solution.

Second, it is important to understand where the network DLP solution looks for the protocols it understands. To avoid the risks of false negatives, a network DLP solution must scan all the channels on all 65,535 ports. Unfortunately, many network DLP solutions only look for the protocols on well-known ports (e.g., SMTP on port 25) or require the customer to tell the solution where to look without providing the ability to cover all ports. A comprehensive network DLP solution needs to be able to analyze all communications, both official and unofficial, in the organization to avoid false negatives.

Content Analysis Methods. A variety of content recognition methods exist, and they usually always fall into two general categories: registration and profiling.

- Registration requires that protected content be enrolled in the system. This system then generates algorithms to detect an exact match or fingerprint of the actual content that has been registered with the system.
- Profiling uses rules that describe information, typically statistical, pattern, and/or key attributes that the system uses to evaluate information. It does not require that the actual protected information be provided.

Registration. The first generation of network DLP solutions invested heavily in registration technologies, specifically exact matching technologies. Exact matching has an attractive message of incredibly low false positives. However, for exact matching to be successful, all protected information must be enrolled. Any information that would violate policy, but was not enrolled, would leave without being detected — leading to a false negative and a resulting data leak.

DLP Enterprise Scalability

A successful network DLP solution must be able to scale to keep up with the constantly updating network and computing infrastructure of the organization. When evaluating scale, there are two critical areas to consider. First, network performance — or the ability for the network DLP solution to analyze the traffic at network wire-speed. Secondly, it must support zones of control — defined as the locations in the network where the network DLP solution can be deployed.

Network Performance. To prevent data leakage on the network, the network DLP solution must be able to capture and analyze information in the network session in real time at wirespeed, which is not a trivial feat to accomplish. It must be efficient enough to analyze all network traffic, all without introducing any form of network delay.

As organizations further leverage networking and the internet, the speed of network connections is increasing. Additionally, when moving beyond the network perimeter to positions inside the enterprise network, performance becomes even more critical. It is crucial to select a network DLP solution that can meet an organization's connection requirements today — one that is built upon an architecture that will scale to meet future bandwidth requirements.

Unfortunately, most first-generation network DLP solutions struggle above 100 Mbps of bandwidth. It is important to note that solution providers may attempt to disguise their poor performance by providing the performance of the fastest supported network interface card, versus the supported throughput their analysis engine can deliver. Many also claim to support a gigabit Ethernet interface; however, their analysis engine may support 100Mbps or less of bandwidth.

It is critical to know the performance capabilities of the analysis engine, as once that capability is met the solution will either begin to fail, drop traffic, or sample. Failure tends to be easy to detect, as the system no longer functions. However, dropped traffic and sampling are often not advertised, but always lead to critical issues. Dropping traffic is when a system discards, and therefore does not analyze, any traffic over what the threshold the analysis engine can handle. Sampling is similar, except the system attempts to be systematic about selecting which sessions it analyzes and which sessions it does not. In any case, if the network DLP solution fails, drops traffic, or samples, the organization is exposed to the risk of data leaving the network undetected (a false negative) and the risk of finding out about an unauthorized disclosure of sensitive information from a third party or law enforcement.

Supported Zones of Control for CUI. Organizations have deployed network security controls beyond the internet perimeter across the enterprise network, creating internal zones of control within the network. These internal perimeters, where technologies like firewalls and intrusion prevention have been deployed, are also likely locations where network DLP also needs to be deployed. Examples of internal zones of control typically include:

- Edge of the data center controlling information leaving to endpoints and departmental servers
- Inside extranet connections controlling access from business partners
- Inside network connection to outsourcing providers controlling information extracted (versus systems access)
- Inside the VPN concentrator controlling information leaving to remote endpoints and employee home computers
- Between divisions of companies (e.g., manufacturing division to financing division) or between division and enterprise backbone
- Protecting secure networks (e.g., HR or engineering network).

Support for internal zones of control requires additional capabilities beyond the requirements at the internet gateway. First, performance is critical. Internal networks are typically significantly higher performance, further stressing the importance of wire-speed performance (above). Additionally, different protocols are typically seen inside the network than at the internet gateway.

File sharing between servers and database traffic are examples of network protocols seen inside the enterprise and typically not at the internet gateway. To control this important internal traffic,

the network DLP solution must understand these internal protocols. The ideal DLP solution should be designed from the beginning to support multi-gigabit-speed networks without sampling or dropping traffic. Unlike the first generation of DLP solutions that only support 100Mbps of analysis, solutions today need to support up to 10Gbps of analysis in a single appliance without sampling or dropping traffic. Load balancing with multiple appliances can support higher network traffic analysis. DLP sensors should also provide support for internal protocols like database traffic and file sharing. The ability to provide these critical internal zones of control enables organizations to get more value out of their DLP.

Open Architecture & 3rd Party Ecosystem

Practically every IT purchase has impact on other assets in the enterprise. Minimizing or eliminating negative impact, while creating synergies with existing deployed assets, enables any solution to add more value than it does as a standalone system — so the network DLP solution should support that same goal.

Minimizing Operational Impact. The network DLP solution must be able to be deployed in a manner consistent with enterprise architecture standards and be minimally invasive to the enterprise architecture to lower the solution's operational impact and therefore maintain the lowest possible total cost of-ownership. The implementation should not negatively impact network performance, add additional points of failure, and/or require desktop or server reconfiguration.

Adding Value to Existing IT Assets. While minimizing impact to existing assets is often enough for deployment approval, adding value to existing IT assets — particularly other security assets — should be a goal of the network DLP solution. The most impactful way to add value to other systems is integration and information sharing to simplify management of risk. Network DLP, while critical to protecting an organization's CUI, is only one component of the security infrastructure. It is important to integrate the knowledge gained from network DLP solutions into other security and risk management solutions. Unfortunately, many DLP providers aren't interested in sharing information with third-party products, as many of the DLP "suite" providers are motivated by the sale of their own security solutions, regardless of the quality of the solution. The adoption of cloud, SaaS apps and mobility have driven DLP into an integrated feature status that requires information sharing.

DLP should not be a silo. Network DLP provides the content-aware mechanism to understand what information is leaving the enterprise or crossing internal zones of control. However, to maximize risk reduction this information should be correlated with other information present in the security and risk management infrastructure including network security/intrusion prevention and identity access management systems. Security Information and Event Management (SIEM) and Security Analytics (SA) tools are popular places for aggregating this information.

Additionally, network DLP provides information on how the organization communicates and where sensitive information such as CUI is flowing. This is an important intersection with other security technologies, especially encryption and systems management. DLP helps an organization understand compliance with their encryption policies and encryption can provide the capability to send protected information securely to approved parties. While an encrypted tunnel may not be

inspected, validating the secure transmission of sensitive data on approved channels from known sources to known destinations is important for compliance with regulations. DLP solutions do provide high-performance network solutions designed to deploy within an organizations infrastructure with no impact.

Evolution of Hybrid and Multi-cloud Security

Adoption of cloud infrastructure has left many organizations to include the Federal government, contractors, and regulated industry managing hybrid deployments— environments in which some applications and infrastructure have been moved to or built in the cloud, while others remain in a data center. Security for hybrid infrastructure must account for both cloud and data center resources. Existing security and compliance requirements for workloads must be met wherever they reside, without relying on network boundaries—and additional consideration must be given to the exposure of data center assets to cloud resources. At the same time, the cloud introduces significant new requirements for securing the service provider control plane, new administrative credentials, a broad range of infrastructure-as-a-service (IaaS) and platform-as-a-service (PaaS) resources, and a more dynamic operational model.

The challenge for security requirements that both existing and new audit expectations and operational requirements need to be addressed. Past audits may dictate expectations of server controls, and tooling may need to integrate with existing security management workflows. Cloud and container requirements add new security and compliance challenges, as auditors will focus on new components and the exposure of data center assets to cloud-hosted resources. Thus, the integration of security with new processes and a broader set of stakeholders is critical, as cloud and associated trends drive responsibility for procurement, configuration, and management of resources to individual DevOps teams.

Hybrid and Multi-cloud Security Considerations

When looking at CUI environments, infrastructure choices are driven by the requirements of individual applications and may involve resources across multiple clouds and data centers. The security team will require portable tooling, effective in both the data center and cloud. Different applications will have varying degrees of interaction and integration—which means security implementation needs to account for network connectivity between data center and cloud assets, and for workloads that may move between environments. Therefore, evolving security requirements for CUI may ask organizations to provide security awareness across diverse locations, architectures, and asset types. Location-agnostic assessment and monitoring help avoid assumptions about network boundaries, and potential growing into zero trust for the surrounding environment.

As cloud architectures are specifically designed for rapid incremental changes and on-demand scaling, the security requirements and the tools used need to be equally fast and elastic. Software-defined resources enable routine change to cloud-based assets, so security assessments must occur in tandem. For many applications this is underscored by ephemeral workloads, which can be spun up for as little as a single task before being decommissioned. These workloads remain in scope for compliance, and security controls must be applied and tracked, with associated data preserved for audit.

Therefore, security automation becomes critical to the DevOps process in a hybrid environment. Development, operations, and security are best implemented as tightly integrated functions for each application. The CI/CD pipeline manages change and deployment of the entire application stack, and it can also integrate automation of security deployment, testing, scaling, and monitoring. To achieve this, security controls must evolve to be embedded as code, enabled by API access to security data and programmatic configuration and initiation of assessments.

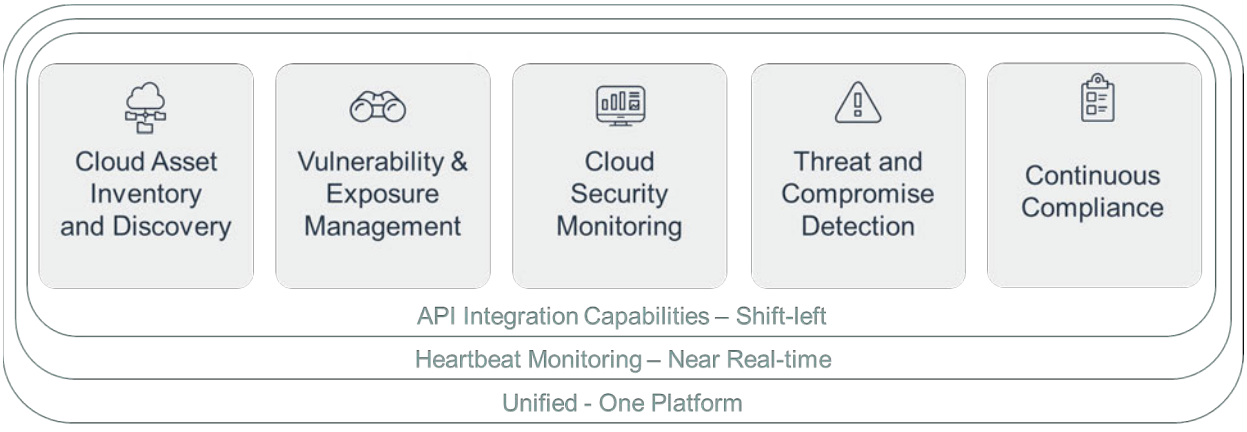

Unify Security Controls Across Diverse Assets and Locations

In reviewing the 800-171, we recommend considering the use of point solutions for different security controls and asset types be reconsidered, as it complicates security management and integration, and may inhibit an organization's ability to gain or maintain a complete picture of their security posture as it relates to CUI. Instead, a successful hybrid cloud deployment will consolidate security functions across all locations and resource types.

When reviewing the security requirements in the context of hybrid environments, new solutions such as CNAPP platforms of hybrid environments, new solutions such as CNAP platforms offer automation of discovery, inventory, vulnerability management, configuration hardening and survivability, drift detection, monitoring, and compliance auditing—all within one unified platform. These functions should be provided at the workload or resource level, without making assumptions about protections offered by network boundaries.

Hybrid cloud deployments have multiple layers of control and configuration—including the CSP control plane, hosting and orchestration for containers, and application workloads. These layers interact and work together to support the application, and each requires different approaches for security assessment and monitoring. Security requirements may want to address unifying visibility from all layers to achieve comprehensive coverage and the approach to securing hybrid deployments should mirror these layers. The CSP control plane, and the IaaS/PaaS resources it manages, require agentless API-based connectors. This is critical for serverless infrastructure, whose components are fully managed through the cloud account.

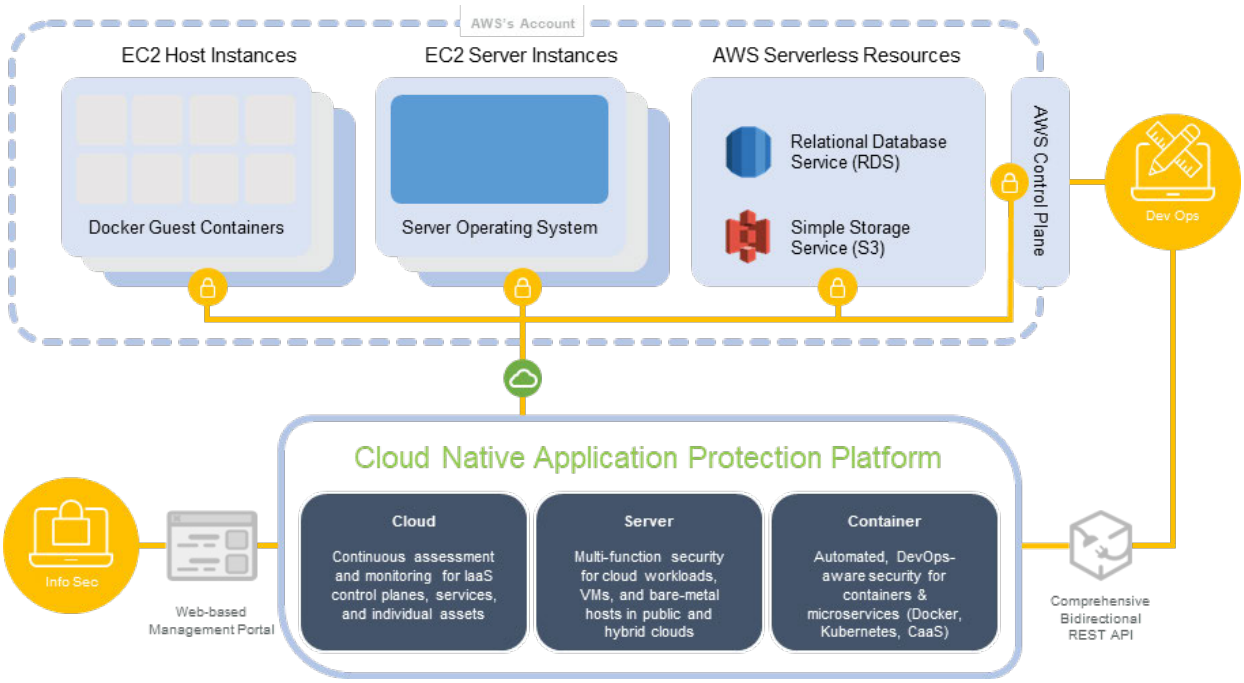

Cloud-hosted servers also depend on this assessment of the computing services on which they run (e.g., AWS EC2). However, servers in all environments require deep introspection, using cloud-friendly agents lightweight enough to be embedded with a negligible impact on performance. Container security requires analysis of both the container image, and its supporting infrastructure. This can include hosting environments in the data center, CSP resources (e.g., ECR, ECS), and cloud instances running Docker, Kubernetes, or other technologies. These diverse asset types, which can all be integral and interdependent parts of the application, have different security concerns and requirements but need to be covered in a unified fashion.

Automating Security Functions for Speed and Scale

The ability to effectively automate security configuration, assessment, monitoring, and response is becoming a foundational security requirement for DevOps and security teams, especially when handling CUI. Programmatic, API-driven access to security platform data and configuration is essential—as is the ability to trigger assessments, manage response, or export and manage security policies using version control systems, such as Git. These functions are even more powerful when individual API access is provided to each DevOps team. This aligns security implementation with the mandate for incremental improvement of automatic testing, quality control, monitoring, and responses to operational issues.

Integrating Security with Development and Operational Workflows

Cloud applications are generally supported by a CI/CD process managing frequent automatic deployments. Shifting security assessment "left" and into that pipeline is becoming more critical when handling CUI. However, to "shift left" the security requirements may need to have organizations have their security platforms integrate with development automation servers such as Jenkins, Atlassian Bamboo, or GitLab CI/CD to build security evaluation into continuous integration tests and fail builds that introduce vulnerabilities or misconfigurations. Security validation should also be added to staging environments, especially when using infrastructure as code, to catch any issues introduced when the release-ready application is instantiated.

Furthermore, the DevOps pipeline itself presents an additional attack surface, as it serves as the supply chain for an organization's application. Code repositories, image repositories, and test automation servers—whether instantiated as data center servers, cloud instances, or CSP managed resources—are effectively part of the production environment. This makes them attractive targets, emphasizing the need for appropriate security controls.

By building security into DevOps workflows when handling CUI, agencies and contractors begin the process of maturing their security implementation over time. CI/CD automation applied to security fixes allows them to be tested and deployed quickly, reducing response time, while security testing built into the DevOps pipeline prevents introduction of new issues. Built in security functionality streamlines collaboration between security teams and DevOps, supporting implementation of guardrails around critical security concerns, triggering immediate feedback, and even enabling automated remediation. CNAP platforms also allow organizations to start from the tens of thousands of out-of-the-box policies and rules that encompass Center for Internet Security (CIS) benchmarks, regulatory standards, and security best practices. This can help simplify operations for small and medium size businesses by grouping assets and applying policies and rules in a way that makes sense.

Summary

The protection of CUI resident in nonfederal systems and organizations is critical to federal agencies and can directly impact the ability of the Federal Government to successfully conduct its essential missions and functions. While multi-cloud infrastructure improves reliability and scalability while reducing technical overhead and risk, it creates complexity for managing CUI and critical controls across a broad range of components. For many organizations looking to

leverage the benefits of the cloud, data center assets need to be retained for the short or long term, leaving them to manage the complexity of hybrid cloud deployments. Security control requirements remain critical and maintaining them requires a solution that works seamlessly in hybrid and cloud environments, covers instances and cloud services in a unified fashion, and integrates with automation, new technologies, and process improvements that enable long term incremental benefits from digital transformation. Therefore, the Government may want to consider Open XDR and CNAPP security requirements to make makes securing CUI public, hybrid, and multi-cloud environments faster and easier.

DARKNIGHT INDUSTRIES

#	Submitted By (Name/Org):*	Type (General / Editorial / Technical)	Source (publication, analysis, overlay)	Starting Page #	Line #	Comment (include rationale)*	Suggested Change*
1	Fidelis	General	Publication	17		**Audit and Accountability (3.3)** Hybrid deployments utilize a broad range of components, requiring varying approaches for discovery, security, and monitoring. There are minimal security requirements for server and container infrastructure hosted in both the cloud and data center, securing both cloud and hybrid deployments, and cloud environments, including serverless applications and resources, as well as the IaaS and PaaS services that support cloud-based server instances and containers.	Audits are an important driver for security—and an opportunity to verify that all necessary controls are in place and operating correctly. Especially in the cloud, organizations need to establish and maintain compliance through continuous assessment and feedback, identifying compliance issues as they happen. Audits really need to automate the collection, storage, and monitoring of the exhaustive data points needed during an audit. Because of its ongoing, rapid discovery and assessment, they will track and maintain data on even highly ephemeral assets.
2	Fidelis	General	Publication	42		**Incident Response (3.6)** Organizations need actionable security findings and remediation advice, enabling system owners to eliminate vulnerabilities and mitigate exposures—hardening servers, containers, and cloud resources against attack.	Organizations need to continuously assess and monitor the cloud infrastructure system integrity and detect drift. Exposures can be created by newly discovered vulnerabilities, updated policy rules, or configurations that have been innocently or maliciously altered. Hybrid cloud environments requires customizable, automated alerts that notify team of any activity or changes that may represent threat or compromise.
3				44		**Maintenance (3.7)** Security automation is critical to the DevOps process in a hybrid environment. Development, operations, and security are best implemented as tightly integrated functions for each application. The CI/CD pipeline manages change and	Security controls may need to be embedded as code, enabled by API access to security data and programmatic configuration and initiation of assessments. The DevOps pipeline itself presents an additional attack surface, as it serves as the supply chain for the application. Code repositories, image

						deployment of the entire application stack, and it can also integrate automation of security deployment, testing, scaling, and monitoring.	repositories, and test automation servers—whether instantiated as data center servers, cloud instances, or CSP managed resources—are effectively part of the production environment. This makes them attractive targets, emphasizing the need for appropriate security controls.
4	Fidelis	General	Publication	55		**Security Assessment and Monitoring (3.12)** Container security requires analysis of both the container image, and its supporting infrastructure. This can include hosting environments in the data center, CSP resources (e.g., ECR, ECS), and cloud instances running Docker, Kubernetes, or other technologies.	Especially in hybrid environments, organizations need to analyze and monitors access controls within individual servers, across cloud accounts. Administrative accounts must be properly configured, as should the use of roles and cross-account trust, to mediate access. Security groups and VPCs may also be used to manage network segmentation. Organizations need to assess security groups and VPCs, as well as access control settings on cloud assets such as S3 buckets, to ensure they are configured securely. They also need deep insight and control for servers and container hosts, monitoring their local accounts, open ports, and network connections.
5	Fidelis	General	Publication	58		**System and Communications Protection (3.13)** The simultaneous use of cloud and data center infrastructure introduces unique challenges for security. This is especially true for cloud and data center resources that integrate directly, requiring strategies for securing communication between assets hosted across the two environments. Information security teams must maintain support for existing data center security and compliance requirements, while restructuring tooling to account for cloud	Different applications will have varying degrees of interaction and integration—which means security implementation needs to account for network connectivity between data center and cloud assets, and for workloads that may move between environments. Security for hybrid cloud must provide security awareness across diverse locations, architectures, and asset types. Location-agnostic assessment and monitoring help avoid assumptions about network boundaries, maintaining zero trust for the surrounding environment. Cloud

						environments and DevOps processes. Hybrid infrastructure requires a re-evaluation of security implementation—both to ensure coverage and to unify visibility across two very different environments.	architectures are specifically designed for rapid incremental changes and on-demand scaling. Security for hybrid cloud needs to be equally fast and elastic. Software-defined resources enable routine change to cloud-based assets, so security assessments can occur in tandem. For many applications, this is underscored by ephemeral workloads, which can be spun up for as little as a single task before being decommissioned. These workloads remain in scope for compliance, and security controls must be applied and tracked, with associated data preserved for audit.
6	Fidelis	General	Publication	17		**Audit and Accountability (3.3)** Cloud, SaaS and mobility have driven DLP towards an integrated feature where data is always in motion. Security stack architects need to consider use cases and the broader reality that DLP by itself is not the end all answer to prevent data theft and loss. For regulatory compliance, network DLP is critical to verify data in motion is transmitted securely and on approved channels. For intellectual property protection, endpoint DLP is critical for data in use on and off networks. Monitoring how users handle sensitive information requires DLP, however, broader context beyond DLP is required to determine high risk users. DLP alerts are also an important data source for machine learning based anomaly detection of insider threats and anomalous user behaviors.	The Who, What, When, Where, and How of CUI - Provides multiple sophisticated content analysis technologies to detect sensitive and/or protected information. - Provides ability to combine multiple content recognition methods in a single rule. - Able to go beyond exact matching for ease-of-deployment and solution scalability and provides analyzers to profile or describe digital assets without the need for intensive registration and maintenance processes to identify sensitive information. - Provides cross session analysis and content tagging of metadata. Capability to provide cross-session traffic analysis using scheduled or automated rules. - Provide a method to analyze behaviors that span multiple network sessions.

							- Ability to apply new threat intelligence automatically against historical network session metadata and generate new alerts for past events based on application of new threat intelligence. - Analytic rules that support event rate, event set, sequence, and frequency
7	Fidelis	General	Publication	58 & 44		**System and Communications Protection (3.13) & Maintenance (3.7)** When protecting CUI, it is important to understand what type of traffic can be prevented by the solution. Internet Protocol (IP) networking defines 65,535 ports, which are sub- addresses or logical locations allowing two computers to connect simultaneously over a variety of protocols. While some protocols still observe their official defined port, the port/protocol paradigm is unfortunately no longer dependable for security controls. Some network services have been deployed on non-standard ports to enable multiple services on the same machine. However, much of the traffic running on non-standard ports is likely to be traffic attempting to evade controls. As applications have evolved, they have implemented port hopping to find ways to work through enterprise firewalls. This is particularly true of social networking, online conferencing and text messaging and this presents a significant risk for data leakage.	When protecting CUI in a hybrid and multi-cloud environment, the government may want to consider the following security requirements: - Conducts session-level (not packet-level) inspection of network traffic across all 65,535 network ports in a single layer 2 network appliance. - Provides visibility into the protocols, channels and applications in use on the network including, at a minimum, the following channels: SMTP, POP, IMAP, FTP, HTTP and HTTP/2, Jabber, CIFS /SMB, DB2, Oracle, LDAP, social media services, and webmail platforms. - Performs port-independent protocol inspection that inspects all 65,535 ports for all supported protocols. - Extracts enterprise human-readable content and related meta-data contained in the session and any attachments and compressed files for analysis. - Inspects SSL/TLS encrypted sessions via integration to an ICAP-enabled proxy server. - Ability to understand database and file sharing applications to establish internal zones of control.

							- Collect metadata and analytics on-premises or cloud for real-time and retrospective analysis for up to 360 days.
8	Fidelis	General	Publication	58		**System and Communications Protection (3.13)** The 800-171 states the data encryption should happen at rest and in motion if an organizations is to have good hygiene. Yet all encryption methods aren't equal. For instance, whole-disk encryption defends against physical theft of the drive. Moving up the stack to network protection measures like Secure Sockets Layer (SSL), Transport Layer Security (TLS), or virtual private networking (VPN) also can pose potential issues. Data is encrypted at one end, only to be decrypted at the other end and could be expose CUI to unauthorized activities. What happens when one of those files is removed, whether maliciously or unintentionally, from those existing access controls?	Recommend expanding security management requirements for monitoring, preventing data loss, and tracking electronic documents that contain CUI feasible at scale. Data-centric security management necessarily depends on organizations knowing what data they have, what its characteristics are, and what security and privacy requirements it needs to meet so the necessary protections can be achieved. And while data-security management must be rigorous, it can't be complicated or time-consuming. Mapping out data security features across network, endpoint, cloud platforms, SaaS apps, operating systems, plus web, email and cloud gateways becomes the modern-day challenge. No one vendor is the best-of-breed answer for protecting CUI and other data within hybrid environments. Data loss prevention (DLP) solutions need to be content and context aware for effectiveness, not a large source of alerts and noise impeding security analysts. For data theft and loss, regulatory compliance, intellectual property protection, sensitive data use monitoring, advanced threat detection, sandboxing, investigations, retrospective analysis and hunting, network traffic analysis at high speeds for all ports and protocols should be an elevated as a critical security requirements for CUI.