InformationWeek
:: analytics
Analytics.InformationWeek.com

November 2009
$199

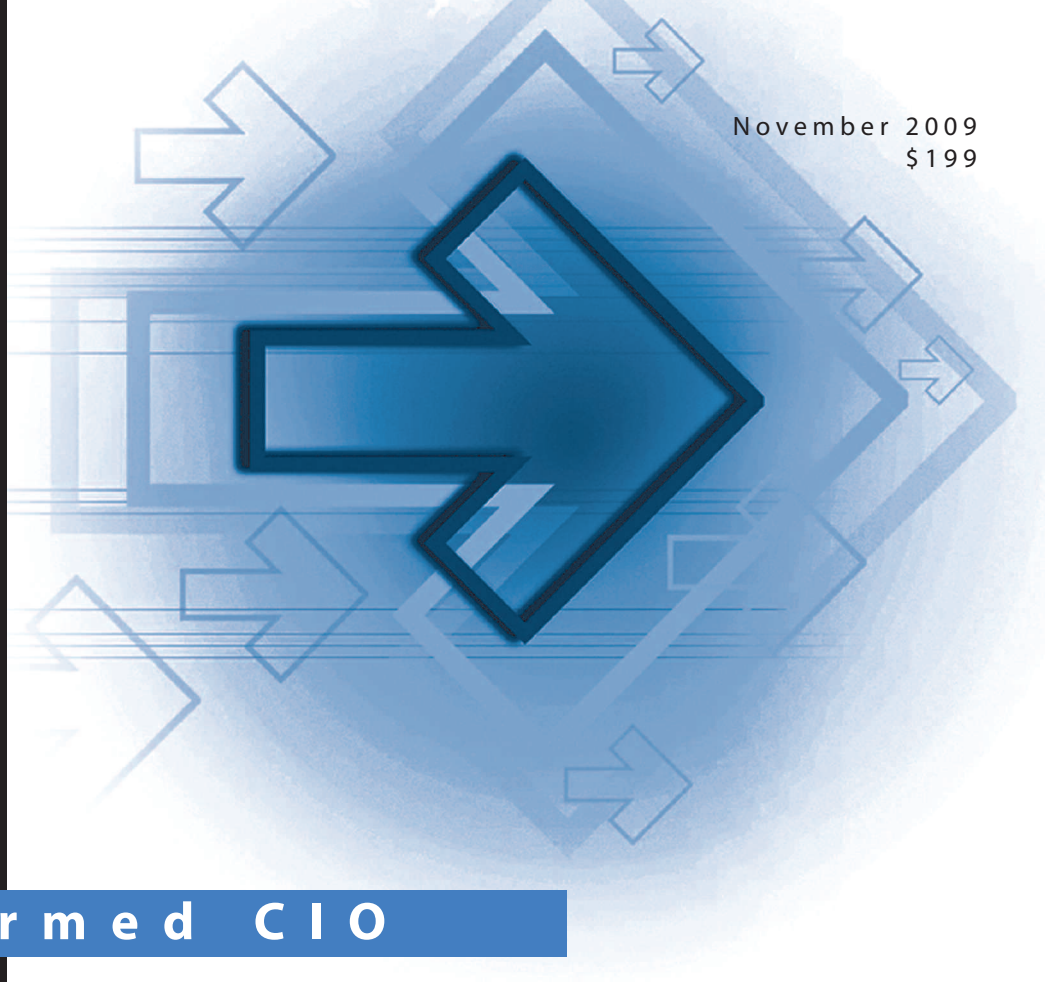

Informed CIO

In Association With
InformationWeek Government

# Integrity Check: 5 Steps To Data-Centric Cyber Security

It's becoming progressively more challenging to protect sensitive data and systems, and agencies that outsource major applications to the private sector are further abstracting system boundary and perimeter concepts. The answer is a shift in thinking away from yesterday's security approaches and toward data-centric protection via technologies like encryption, data loss prevention and strong access controls.

By Richard Dreger

# TABLE OF CONTENTS

**Richard Dreger**
CISSP, CISA, CWNE
*InformationWeek Analytics* Contributor

**Richard Dreger** is president of WaveGard, a vendor-neutral security consulting firm, and an *InformationWeek Analytics* contributor. Rick has significant, broad-based technology experience with extensive skills in the information assurance, security and wireless networking fields. He has consulted for a wide breadth of clients in both the public and private sectors, and his professional background includes over 15 years of experience in *Fortune* 100 companies as well as smaller technology consulting firms.

Rick has complemented his hands-on consulting experience by leading courses such as the CWNP wireless curriculum and the (ISC)2 CISSP review. In addition to being one of the 11 founding members of the Certified Wireless Network Experts (CWNE) roundtable, he is also co-author of the *Certified Wireless Security Professional (CWSP) v2* study guide and numerous *InformationWeek* articles. Rick obtained his BSE from Duke University and his Masters from Villanova University.

# Executive Summary

**Change is the only constant.** This is particularly true in the cat-and-mouse world of information security. With a constant flow of zero-day attacks and malevolent—albeit not always innovative—thinking on how to best exploit hardened systems, data defenders need to be ever vigilant. Certainly, public- and private-sector CIOs are constantly bombarded with new silver-bullet applications, appliances and techniques aimed at providing enhanced protective controls. But even the most sophisticated tool is of limited value if we don't understand a key tenet: Sensitive data can still be vulnerable even when placed within a well-protected infrastructure.

It's the age-old problem of having strong, solid exterior walls and limited additional inside defenses. The analog in the information technology realm is that of very strong perimeter defenses (firewalls, IPS, hardened border routers) at interconnection points, but only limited supplemental controls at the "trusted" core of the enterprise.

Although it's an archaic assumption that firewalls alone constitute an adequate defense, in our practice, we still see the occasional IT group that subscribes to this approach's effectiveness. More progressive organizations, often with significant investments in information assurance technologies, may be better protected, but even they can be lulled into a false sense of security when their systems are surrounded by sophisticated network appliances, intimidating physical security controls and exhaustively documented security policies.

In our recent *InformationWeek Analytics* Government IT Priorities survey of federal technology decision-makers, cyber security was the No. 1 IT initiative within respondents' organizations in terms of importance and current leadership focus (ahead of data records management and DR planning). For most of these shops, cyber security means dealing with the Federal Certification and Accreditation (C&A) process required by FISMA. This

# Executive Summary

mandated approach is highly proscriptive: There are 17 separate control families with which to comply, each bringing its own specific directives. Although some of these can be deferred by using common controls for the organization (for example, information security policy or incident handling) others cannot be—and rightly so.

The upside to FISMA and the ensuing NIST documentation is that agencies have a consistent and broadly applicable standard for how information security should be applied to systems that are deemed to warrant a given classification level. The downside is that the true goal of adequately securing sensitive information and preserving core mission processing sometimes gets lost in a maze of requirements. By proposing a highly data-centric approach, we'll help agency CIOs and CISOs refocus their security programs back to the essential precept of protecting information.

## Render Unto the Common Controls...

To focus on data security, we need to start by understanding the control environment. Each Federal Major Application (MA) and General Support System (GSS) must be certified and accredited prior to being put into production. This process involves defining information system impact levels by formally determining security categories for confidentiality, integrity and availability—C, I and A—per FIPS 199 guidelines. Based on these results, a high-water mark of low, moderate or high is selected for the information system as a whole, and a specific set of security controls outlined by FIPS 200 and NIST SP 800-53 are applied.

Figure 1

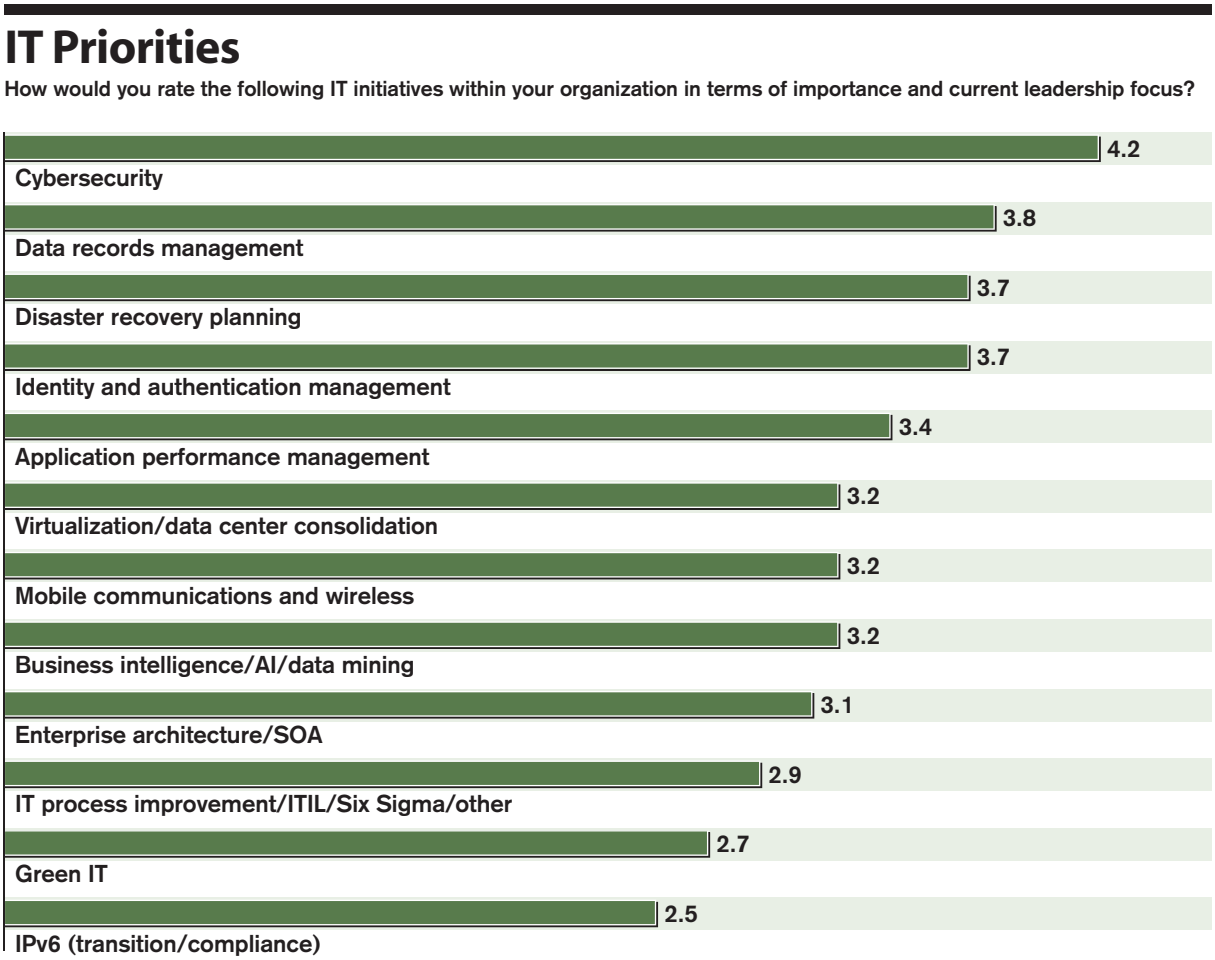


Note: Mean average ratings based on a five-point scale, where 1 is "not at all important," and 5 is "extremely important"
Data: *InformationWeek Analytics* Government IT Priorities Survey of 309 government technology professionals

So far so good. However, anyone who has had to define a system security plan for a moderate- or high-category MA or GSS knows the complexity of this process. First, comprehensive security policies, standards and common controls must be referenced, or created where they do not yet exist. Then, there are detailed steps that must be followed during the creation of the system and the implementation of supplemental controls. Then, technical controls must be selected, configured and tested for efficacy … and so it goes. Given that security programs based on FISMA/NIST special publications tend to be fairly rigid and monolithic, imposing a fair degree of overhead on the system owner, the appeal of deferring controls to supporting systems or leveraging existing security program documentation is understandable. And, in fact, selectively integrating system-specific protections and shared infrastructure controls can provide smart economies of scale while still providing proper security.

Problems arise, however, when agencies defer too much to the common-control pool and don't focus enough on the unique needs of their individual information systems and the data contained therein. Specifically, CIOs and CISOs need to consider the following:

Figure 2

## Decline in Government Reliance on Contractors?

**Do you believe that the government may rely less on contractor personnel for IT projects in the future?**

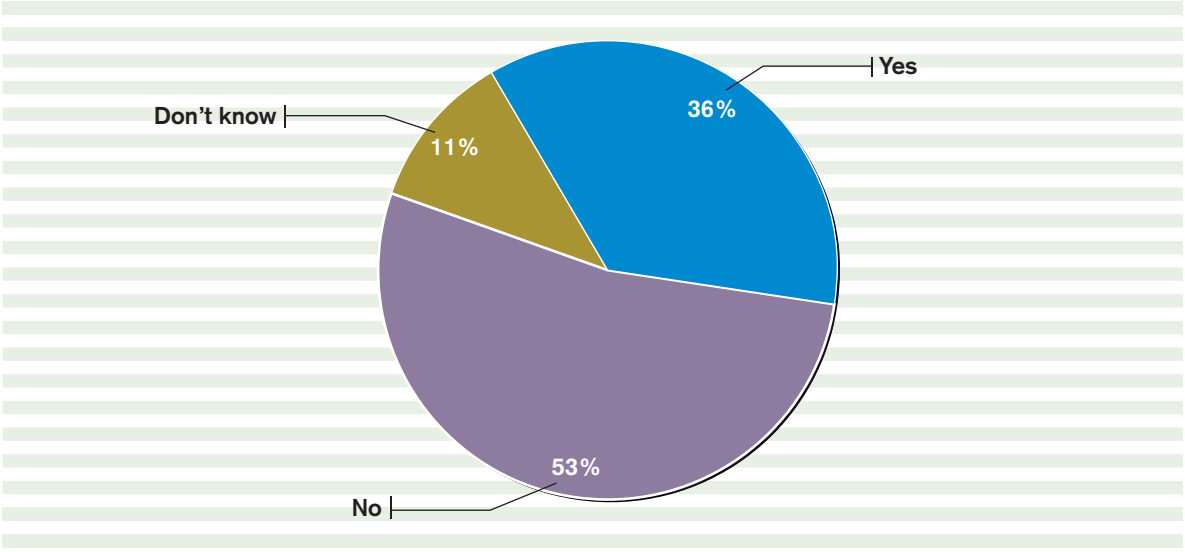


Data: *InformationWeek Analytics/InformationWeek* Government Survey of 177 federal government technology professionals

**Utilize common controls** where applicable for policies, SOPs and similar documentation. Unique, hybrid documentation will still likely need to be created to properly document the system and its control suite.

**Leverage infrastructure controls** and standards for robust networks, well-designed directories and intelligent perimeter controls where feasible. Supplement when necessary.

**Apply data-centric system controls** to cover the unique requirements of your applications and the information they contain.

Now, these three bullets comprise an outwardly straightforward, layered approach to providing smart, tailored protection for information systems. Problems arise, however, when government CIOs choose to extend the system boundary into the cloud or outsource systems to private-sector service providers—a path that is not set to decline in popularity, according to our recent *InformationWeek Analytics* survey of federal government technology professionals (see Figure 2, previous page).

By adding mobility to where systems reside, the list of available supporting controls is blurred and becomes much more subjective, as different groups start interpreting what really constitutes "adequate" security. Certainly, reasonable people can disagree as to what "adequate" means for the various NIST SP 800-53 controls, and this calculation will become even more subjective as private-sector perspectives get more strongly mixed into the federal methodology.

Figure 3

## Federal Initiative Success

**How would you rate the success of these recent federal initiatives?**

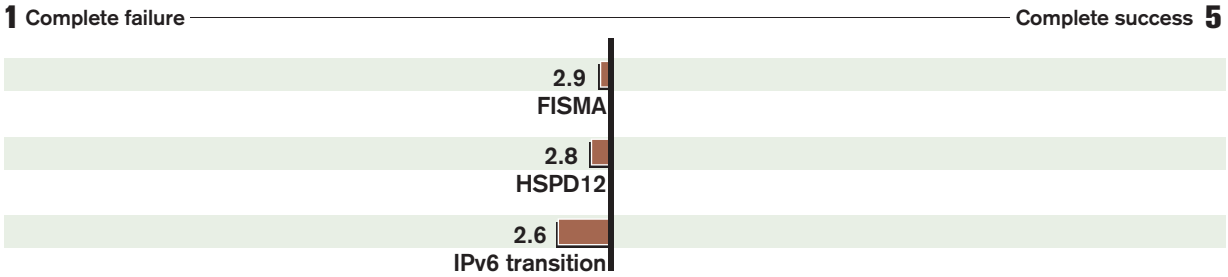


Note: Mean average ratings

Data: *InformationWeek Analytics*/*InformationWeek* Government Survey of 177 federal government technology professionals

Even with the great equalizer of requiring C&A accreditation of systems prior to going live, it is absolutely essential that strong data-centric controls exist so that we can rely less on compensating environmental or infrastructure controls as the environments in which devices and data reside become more dispersed.

In this report, we'll focus on ensuring that data-centric controls are in place and providing strong assurance levels to our confidentiality, integrity and availability requirements—regardless of the supporting controls implemented at any given time or where data ultimately resides.

### 1 | Master controls are out. Think data-centric instead.

A defense-in-depth architecture relies on pulling together a series of integrated, overlapping controls that work together seamlessly to form a strong, homogeneous whole. This approach moves away from using a single master control or appliance that can "do everything" and promotes a more distributed and tailored security posture.

Any multifaceted defense must include documented policies, robust infrastructure controls and protections focused on the crown jewels—the data contained within the system. Just because an MA is secured within a strong general support system (GSS) environment does not necessarily mean that the MA's data is protected. Rely on the supporting infrastructure and security program controls for some protection, then tailor system- and data-specific controls to address the remaining risks.

Information needs to be secured regardless of whether it is at rest or in transit. Data at rest represents information that is sitting on some type of media, such as a computer hard drive or backup tape. Data in transit represents, quite literally, information that is moving from one place to another, such as in an e-mailed attachment or transmission over a WAN. To properly protect data, we must understand how it needs to be used and then ensure that appropriate controls have been put in place to protect it for those functions.

When we think about protecting information and preserving the C, I and A characteristics, a few essential issues come to mind:

We need to **encrypt mobile data,** particularly on devices that may be anywhere on the globe at any time.

We need to **uniformly implement strong authentication controls** to prevent unauthorized access to information.

Data in the hands of users can become slippery and capricious, so we must **inventory everywhere sensitive information resides** if we're to have any hope of preventing its accidental (or intentional) loss.

**Protecting data integrity** to help ensure proper system operation is paramount to minimize the effects of sophisticated malware.

Next, we'll delve into technology approaches that can serve as supplemental controls for protecting data.

### 2 | Embrace data encryption.

Because large databases or sensitive file stores can be quickly copied onto thumb drives or portable devices, the protections afforded by strong physical controls are muted. Think of the situation this way: If a system resides in a secure data center protected by very strong network controls, the data itself is likely fairly safe. Only those persons with access to the physical system or logical access to remotely connect can access the information. Once data becomes portable, and then resides on a physically insecure system, such as a laptop, we have a whole new game.

Physical access to a device trumps a huge range of painstakingly applied logical controls. For example, if a non-encrypted laptop containing sensitive information is stolen, attacks can be made to quickly extract the information off the system.

Common controls used to protect systems can be worked around—for example:

**Alternate boot disks** can be used to boot up the laptop and then mount and extract stored data.

Tools can be run to **subvert administrative accounts** on the system to take full control and obtain complete file access.

**Removing the hard drive** and mounting it on another system enables attackers to get past BIOS or OS controls.

What can we do to mitigate these risks? First, we need to understand that the physical security portion of our defense-in-depth approach really no longer applies once data has been moved to a mobile device—such as a laptop, thumb drive or smartphone—that is easily lost or stolen.

Second, we want to ensure that any information leaving our systems has been removed by an authorized individual and placed onto an approved portable device. Now, we admit that this

Figure 4

## Current and Planned Encryption Use

**Which of the following uses of encryption are currently implemented in your organization? Which do you plan to implement within the next 12 to 24 months?**

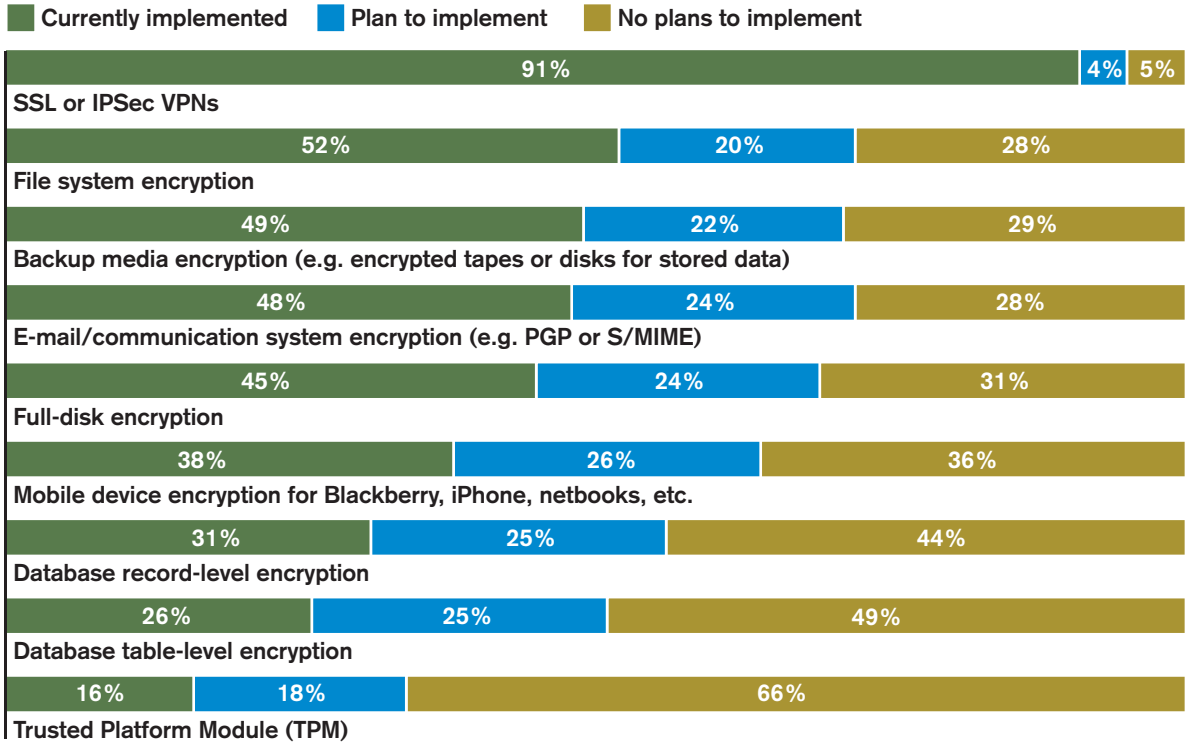


Base: 430 respondents at organizations using encryption

Data: *InformationWeek Analytics* Data Encryption Survey of 499 business technology professionals

second point may be a big assumption in current environments; however, in a data-centric security model, controlling where your data resides, who can access it and how the information can be moved is essential. More information on these points will be provided later in the article.

So, back to our situation where information has been moved to an approved mobile device; we still need to understand how to compensate for the loss of physical security controls. One of the first defenses is encryption using FIPS 140-2-compliant modules and a strong algorithm, such as AES. If a laptop is secured using an approved whole-disk encryption system, or even if the data resides in separate encrypted "canisters" on the drive, additional authentication credentials are *required* prior to accessing the data. If the stolen laptop drive is taken and mounted, the contents are unusable unless the proper decryption credentials are provided.

Think of it this way: Numerous controls are implemented on servers and systems to protect the data they store. For some operations, this sensitive data must become mobile and reside on devices such as laptops. When this happens, we must ensure that **comparable or better** security controls are being deployed to preserve required protection levels. Thus, disk encryption can bolster our confidentiality and integrity controls to compensate for a loss of physical security and to help support those NIST 800-53 controls requiring data encryption.

To make disk encryption more usable and manageable, look for these types of features within candidate products:

**FIPS 140-2 compliance** assures that the cryptographic modules have been tested to work properly. Controls such as NIST SP 800-52 AC-3 Access Enforcement require NIST 140-2 compliance if encryption of stored information is going to be used for access enforcement.

> Centralized management and directory-integrated authentication.

> Help-desk procedures for lost credentials.

> Encryption-key recovery tools.

> Secure, remote disk wiping.

These approaches can be further extended to highly portable and easy-to-lose thumb drives, too. Depending on the product's capabilities, workstation USB ports can be controlled by agent software that mandates actions, such as:

> Encryption of all files on the drive.

> Multifactor authentication to access stored files.

> Centralized USB drive management.

> Limiting USB drive mountings to authorized systems.

There are logistical issues that must be addressed throughout the organization to mesh data-centric controls with broader group policy settings and security program elements. For example, settings that prevent unauthorized USB drives or systems from receiving data will need to be applied to prevent copying information to improperly secured devices.

Further, user education, system operating procedures, and policies covering data sensitivity and handling will need to be layered in to support the more technically minded encryption controls. But as we said before, this is a defense-in-depth approach with a particular focus on data-centric security controls, and data encryption plays a valuable role in safeguarding our information.

## Look for the Seal

Federal Information Processing Standards (FIPS) 140-2, "Security Requirements for Cryptographic Controls," provides a level of assurance for using algorithms and cryptographic modules that are compliant with this standard.

There are different levels of approved security, but the main idea is that if a module or algorithm is considered FIPS approved, the user knows that it has passed a rigorous set of tests and will deliver a given level of protection. This is essential, since agencies do not have the time or ability to adequately vet various products or algorithms themselves. NIST maintains lists of FIPS approved modules and vendors at *http://csrc.nist.gov/groups/STM/cmvp/validation.html*.

**3 | Implement strong authentication controls.**

Authentication involves that most subjective of concepts: attempting to prove that you are indeed who you are asserting yourself to be. Once users have established their identities, role-based access controls can be applied to limit their actions to only those authorized for a given job role. Authentication controls typically deal with three families: that which you know, that which you have and that which you are. This typically translates into passwords as something you know, smartcards or key fobs as something tangible you have, and biometrics as something you are. By combining two or more of these, we have a valid multifactor authentication control.

When we ask users to authenticate, most organizations use username and password controls, which while ubiquitous, can also be quite weak. Some access schemes, such as the control enhancements listed under NIST SP 800-53 IA-2 Identification & Authentication, require the use of multifactor authentication to help minimize abuse and raise the bar against those trying to impersonate an authorized user.

Figure 5

## Defining User Roles and Responsibilities

**How well defined are user roles and responsibilities, including the required access to different systems and data within your organization?**

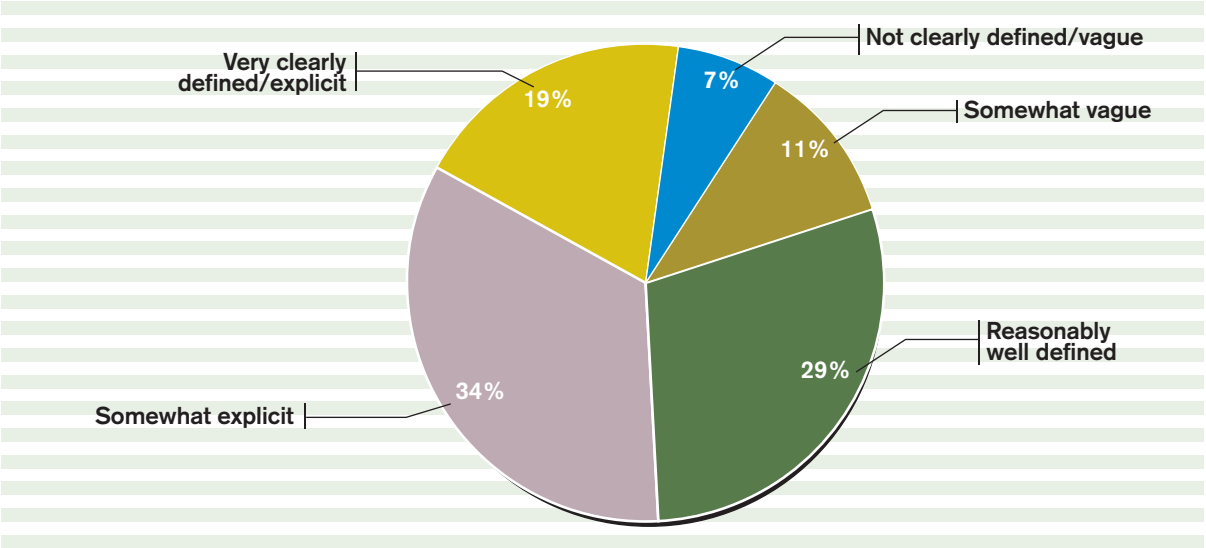


Data: *InformationWeek Analytics* Regulatory Compliance Survey of 379 business technology professionals

Let's go back to the idea of data protection. Important information likely resides on databases, servers, major applications, workstations, laptops and other such devices. Interacting with this data in an appropriate way is essential for core mission processing, and thus we need to regulate exactly how data is being accessed and used. If unauthorized individuals can view data, we have had a major breach of our confidentiality requirement. If this data can then be further changed and/or removed, we've just affected the integrity and availability requirements.

Clearly, limiting access and properly authorizing users are key to data handling. One major way of restricting access to data is by combining strong authentication controls with tailored role-based access controls. If we consistently apply these robust identity management requirements, then we can take advantage of some significant benefits:

> Users can be uniquely identified.

> Multifactor systems are harder to subvert than single-factor systems.

> Roles can be defined by job function and permissions, custom-tailored to the actions required to properly execute these functions.

> Audit logs can become far more usable by tying actions to a user's identity. This is in contrast to tying actions to more abstract identifiers, such as system names or IP addresses.

Unfortunately, merely providing strong authentication and targeted encryption is not enough to fully safeguard our data and ensure its proper use. Sometimes, it is those users whom we trust the most who cause the biggest problems. To this end, user background checks are useful and in some cases mandated as per NIST SP 800-53 PS-6 Access Agreements, where various checks must be made prior to granting users access to certain sensitive information.

However, even with some assurance that our administrators and power-users are straight shooters, problems will still occur, and when they do, it's essential that we have visibility into the various application flows and data manipulations that are affecting our sensitive information. Data loss prevention is an emerging technology that provides very strong data-centric controls, though many systems are still expensive.

**4 | Use data loss prevention to "watch the watchers."**

Since data loss prevention (DLP) has many different potential definitions and focus areas, let's take a minute to define this technology space. At a gut level, think of DLP technology as an information-vetting system that reviews data content with an eye toward possible threats or policy violations. If a potential problem is found, appropriate actions can be taken to stop the data flow before it exits the system or trusted perimeter.

A main impetus for DLP is the growing threat posed by end-system exploits. User systems can be exploited by simply opening an infected attachment not flagged by antivirus. Think it won't happen? We constantly see zero-day attacks, where malware scanners have no definition and therefore can't stop the exploit. Or, there's the threat of a user visiting a malicious Web site that exploits problems in the browser's code. Once the system is compromised, a back channel can be established to one or more rogue Internet hosts allowing for the exportation of files, data and sensitive content. A DLP system not only aims to stop these flows, but also can prevent the accidental transmission of information via e-mail, instant messenger or other protocols.

Figure 6

## Identity Authentication

**What is the primary measure you take to authenticate identity?**

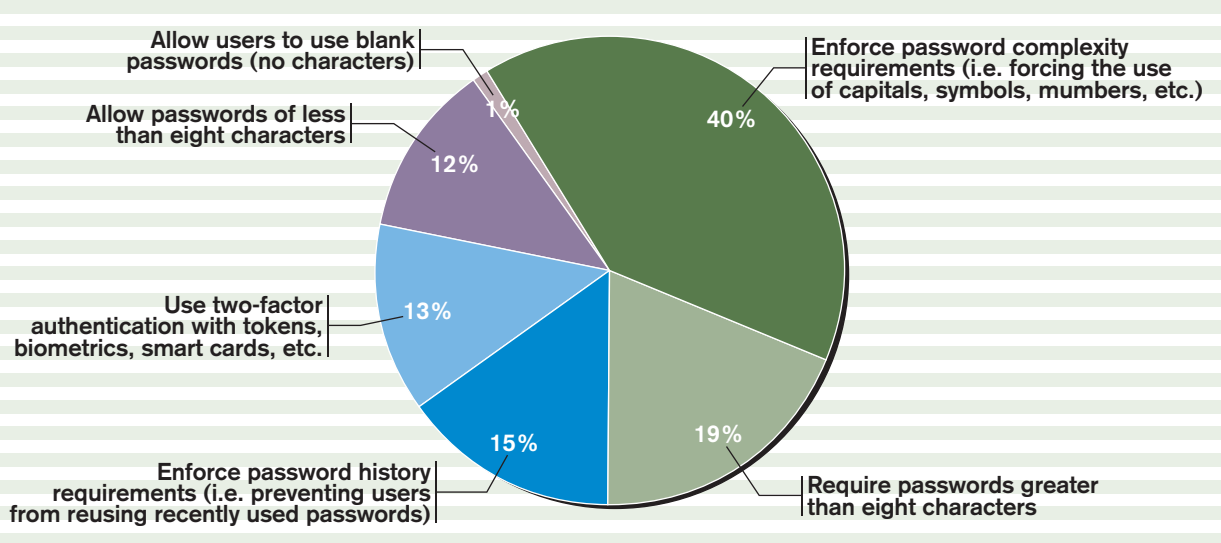


Data: *InformationWeek Analytics* Data Loss Prevention Survey of 218 business technology professionals

For example, say someone is sending a report that contains classified information or data from a moderate/high security system to other managers and accidentally "fat-fingers" an e-mail address. The DLP engine could intercept the message, see that classified data should not be transmitted outside the domain and block the transmission. Alternately, someone may wish to copy files from a server onto a USB device to facilitate working at home. This action can be blocked by policy, prevented and then alerted on to help identify a potential problem.

DLP systems are an emerging market, and each product has its own relative strengths and weaknesses. Generally speaking, these systems use a flexible set of triggers to identify and classify information. Keywords, dictionaries and dynamic rules help identify and mitigate threats. Both host- and network-based implementations are available to address different facets of the data loss threat. DLP deals with data in motion, data at rest and data at the endpoint (for example, on portable devices such as USB drives). A huge range of network protocols can be supported for analysis, including IM, HTTP/S, SMTP, FTP and others, and flexible policies may be

Figure 7

## DLP Product Capabilities

**Whether or not you are currently using or planning to use DLP, what capabilities would you consider most important to include in a DLP product? Please rank the list of capabilities from 1 to 7, where 1 is the most important capability to have and 7 is the least important.**

| | Rank |
|---|---|
| Content security: The ability to scan e-mail and attachments for content that violates policy, and take action as necessary | 1 |
| Malware protection: The ability to prevent malware, bots and viruses from stealing critical data over open communication channels | 2 |
| Endpoint protection: The ability to report and control data leakage on PCs, laptops and smartphones | 3 |
| Enforcement: The ability to block or quarantine actions that would violate policy; for example, stop an e-mail from being sent, or stop data from being copied to removable media | 4 |
| Archival: The ability to archive conversations and prevent leakage over non-standard communication channels, such as instant messaging clients or cellular text messages | 5 |
| Enterprise data discovery: The ability to crawl all databases, data sources, file shares, e-mail databases and endpoint hard disks for information deemed vital for corporate and customer security | 6 |
| Reporting: The ability to report and alert on breaches centrally, with the ability to map certain breaches to regulatory requirements or custom business rules broken | 7 |

Data: *InformationWeek Analytics* Data Loss Prevention Survey of 218 business technology professionals

applied to data to help ensure compliance with documented security program requirements.

For DLP suites to work properly, the organization must have a strong, well-defined security program—in particular, a detailed data classification and handling policy along with supporting procedures. Given the many options that are available for classifying data and then defining handling requirements, having a well-thought-out framework into which the technical controls can be integrated is likewise essential.

Note that DLP technology maps to numerous facets of the NIST 800-53 controls and spans multiple control families. Some specific items include Information Flow Enforcement (AC-4), Media Protection (MP) for ensuring data encryption is performed and boundary protection (SC-7). The full list of controls covered will vary depending on how comprehensive the DLP system is and whether both network- and host-based controls are being employed.

**5 | Layer on data integrity controls.**

When systems and applications start breaking or acting in an unusual way, right away we ask, "What changed?" A seemingly simple question, but one that can be very difficult to answer conclusively. The operational environment of large agencies and organizations is dynamic and sophisticated, requiring ongoing changes to be validated and implemented to ensure that our systems are providing us with the tools we need to effectively conduct business. With all of these approved changes—and an occasionally frenetic deployment tempo—mistakes can be made, and system security may be unwittingly compromised.

Think of data integrity controls as helping to ensure that information, system settings and file configurations are as you expect them to be. Essentially, integrity validation helps guarantee that files do not change from approved baselines and norms without making it clear to operators that something has been altered. Thus, if a data file, driver or system file is now different from the trusted version, it's a clear indication that something is going on that needs to be investigated. Without integrity checking, these modifications could easily fly under the radar and open a serious gap in our security stance.

For example, let's assume that we have a series of servers that all run on a preapproved, hardened OS. The application for our MA is then placed on top of this platform, tested, approved and rolled into production. A snapshot is taken of the approved and secured system and tagged as our trusted baseline.

Now let's further assume that as this system is operating, it somehow becomes infected with malware that bypasses other controls and subtly changes key files to force the system to operate in a flawed and less secure way. Even though some vulnerability was exploited and the system was changed, our integrity validation suite will find the unauthorized change, alert the appropriate parties, and allow us to contain and mitigate the problem—ultimately returning the system to a trusted state.

## Think Outside the FISMA Box

Even when we've thought through and provided a strong suite of security controls, inappropriate access may still occur. To help compensate for this, deploy strong audit trails on systems that log, track and monitor data interaction events to help identify problems and provide solid forensic evidence if needed. This evidence can be used to identify unforeseen problems, potentially identify abusers and close holes that might have gone undiscovered.

The idea of defense in depth is a solid one that can promote a robust yet operationally flexible security program. When integrated, strong management, operational and technical controls form a mesh of protection. In the federal space, FISMA-mandated control frameworks as outlined by NIST provide a series of clearly defined protections, but as we've discussed, the expansive nature of the guidance can also have a negative impact and sometimes lead us away from adequately protecting system data.

To this end, our goal is to help refocus CIOs' efforts back on a data-centric approach to security and away from simply following a checklist. The techniques we've discussed, including encryption, strong authentication, data loss prevention and data integrity controls, make an impressive package when coupled with a well-documented security program and robust infrastructure controls, but they by no means describe the full universe of possible control options. Regardless of the control suites that CIOs ultimately deploy, remember that the end goal is always the protection of our systems and the information contained therein. There are no hard and fast rules as to what specifically must be selected, but hopefully after reviewing this article, smart data-centric security controls will play a key role.