



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Serving Sensitive Unclassified Data and Unclassified Data: Comparing NIST and DoD Guidelines

By

Joseph McGill

For

GIAC Security Essentials Certification, version 1.4

Abstract

A 1994 report to the United States Secretary of Defense and the Director of Central Intelligence states "It has been estimated that as much as seventy-five percent of all government-held information may be sensitive." (Joint Security Commission, Chapter 2, "Dealing with Sensitive but Unclassified Information") Sensitive data, while not classified, does have a requirement for confidentiality. Unclassified data has no requirement for confidentiality. In private companies, there is also the need to protect confidential information. Recommendations and requirements for managing systems containing sensitive and unclassified data have been defined by various components of the federal government. These guidelines can be used in private industry as well.

This paper will present a model for serving sensitive data and unclassified data to a LAN, focusing on the design elements that pertain most closely on the classification of the data. Requirements of two United States federal government organizations, the Department of Defense (DoD), as stated in DoD Directive 5200.28 (DoD 5200.28) and the more recent Memorandum of Guidance and Policy for Department of Defense Global Information Grid Information Assurance (GIGIA), and the National Institute of Standards and Technology (NIST), in the Special Publication 800-series documents, will be referenced. Many of the requirements are similar between the two organizations and are based on the same principles and practices. But some of the key differences are in the area of having a single machine serve both classifications of data, and these differences have a significant impact on configuration and cost of servers.

Definitions

Sensitive Information

Sensitive is a classification of data that usually defined to be not classified but requiring confidentiality. The Computer Security Act of 1987 defines "sensitive" information this way:

any information, the loss, misuse, or unauthorized access to or modification of which could adversely affect the national interest or the conduct of Federal programs, or the privacy to which individuals are entitled under section 552a of title 5, United States

Code (the Privacy Act), but which has not been specifically authorized under criteria established by an Executive order or an Act of Congress to be kept secret in the interest of national defense or foreign policy; (Computer Security Act of 1987, Sec. 3.d.4)

This definition is commonly referenced by United States federal government documents addressing sensitive data but unfortunately it is not universal (see below). For perspective, GIGIA adds, "This includes information in routine DoD payroll, finance, logistics, and personnel management systems"(GIGIA, page 18) though the term is not limited to these types of information.

Of course, an organization need not be part of the federal government to have sensitive data. Private companies also have data that should be kept confidential. The principles defined in the NIST and DoD guidelines can be applied to private corporations too.

Relating these definitions to the fundamentals of computer security, it is clear that confidentiality is important with sensitive data, as are integrity and availability. There are a variety of terms used to refer to sensitive data. These include "sensitive", "sensitive unclassified", "unclassified sensitive", and "sensitive but unclassified".

Unclassified Information

DoD Directive 5200.28 supplies this definition of "unclassified information": "Any information that need not be safeguarded against disclosure, but must be safeguarded against tampering, destruction, or loss due to record value, utility, replacement cost or susceptibility to fraud, waste, or abuse." (DoD 5200.28, page 20)

Note that this definition has no requirement for confidentiality, but integrity and availability are required. And again, this definition can easily be applied outside the federal government.

Historically there has been some confusion about just what is considered sensitive and how it should be handled. Some of this confusion probably results from the fact that "sensitive" has a generic English language meaning but is also a specific, government-defined label for a particular class of information. One person may say "sensitive" (i.e., "My personal financial records are sensitive.") and mean information that should be kept private, and another person may "sensitive" and mean a U.S. government class of data that has a list of defined requirements mandated to preserve confidentiality. It is also possible some people hear the term "sensitive" used to describe information that intuitively seems like it should be kept private and assume that is all that is being said, when in fact, the term is being used in the specifically defined way shown above.

It also doesn't help that different parts of the government use the terms in different ways. For example, the United States Department of Energy (DoE) defines "Unclassified Controlled Nuclear Information" (UCNI) as "certain unclassified government information prohibited from unauthorized dissemination under section 148 of the Atomic Energy Act-As Amended."(DoE Order 5635.4, section 5e) This confusing definition indicates that this kind of unclassified information does require confidentiality. It does not mention the word "sensitive" and makes it sound like some UCNI data might be made public (if authorized) but some might not be authorized for publication. Andrew Helyer wrote a good exploration of the federal government's view of sensitive data in a paper titled "*Sensitive But Unclassified*"(Helyer).

While the terms "sensitive" and "unclassified" usually refer to government data, and the

government has specific requirements relating to these types of data, this paper will use them more universally so as to apply the principles in a way that makes them useful to non-government organizations. When using these terms outside of the government, "sensitive" can refer to any data that requires confidentiality and "unclassified" data can mean any data that doesn't require confidentiality.

A Sample configuration

In this example, an organization needs to serve sensitive data to one set of users and unclassified data to a different community of users. There are any number of examples that this problem describes. For example, a company may want to have a web server to publish their personnel policies, job openings, and so on, while at the same time maintaining the private information about employees, including social security numbers, performance reviews, payroll information, and so on. Another example would be when a machine serving the sensitive community has a particular feature, say, a large robotic tape library and hierarchical storage management software, that would be a great benefit to the unclassified community, but which is too expensive to purchase and maintain two copies of that resource. Is there a way to securely serve both communities on one machine?

Assumptions about the environment

In order to focus on the particular aspects of design that pertain to the classification of the data, some assumptions will be made in the other areas considered part of a "Defense-in-depth" strategy. For example, whatever the configuration is, and whatever number of machines may be a part of it, these assumptions apply:

- Physical controls exist to limit physical access to the system (i.e., the system console, hardware) to only those authorized;
- There is a suitable access control policy in place to confirm the identity of the user prior to accessing the system;
- Systems are configured to guarantee accountability with proper auditing functions enabled;
- Systems are configured to ensure integrity of data. This includes proper backups, permissions, contingency planning and so on;
- Systems are maintained with the latest appropriate patches to prevent compromises and maintain availability;
- Users are trained to be aware of their responsibilities regarding system security;
- Procedures exist for handling security incidents;
- Risk management analyses have been performed to assess the value of additional security measures vs. the increased cost of those measures;
- Security planning and implementation are practices performed throughout the life-cycle of the systems;
- The security posture is periodically reviewed to assure consistent application of good practices;

These assumptions are based on good practices commonly cited in government guidelines. To explore these in more depth, see the NIST Special Publication SP800-14, "Generally Accepted Principles and Practices for Securing Information Technology Systems." (Swanson and Guttman, NIST) Similar principles can be found in (GIGIA). The NIST Computer Security Resource Center (NIST CSRC) Special Publications include a number of documents that give guidance on particular topics.

In this example, we will assume that the community of users whose data is sensitive operate on a need-to-know basis. Authority to access the system containing the sensitive data does not imply the authority to view all the sensitive information on the system. In other words, there may be groups of people working together on projects, and they quite properly need to share data. But at the same time on the same machine, there are other groups working on other projects, and neither group has need to see another groups data. Examples might include an accounting group, a personnel group, and so on in a corporate environment, or groups working on an aerodynamic simulation of an aircraft, and electromagnetic wave propagation in the upper atmosphere, and so on in a defense contractor environment. None of these groups needs to know what others are doing, and thus it is important that access, whether unintentional (accident) or intentional (espionage), be prevented.

In the DoD regulations, this is referred to as the "Least Privilege" principle. DoD 5200.28 states:

Least Privilege. The AIS [Automated Information System] shall function so that each user has access to all of the information to which the user is entitled (by virtue of clearance, formal access approval), but to no more. (DoD 5200.28, page 24)

Finally, if it isn't obvious already, the sample configuration proposed will be described in terms of the Unix operating system (or some variant thereof). But the principles should still apply to other multi-user operating systems.

Implementation Details

The key question here is how much separation is needed to keep the sensitive data secure, while at the same time serving unclassified data to a more open community.

Defining access policy

The system with sensitive data should, as a policy, prohibit world-read access to any user data that is considered sensitive. Procedures should already exist to verify that world-read is prohibited on certain files (such as the shadow password file, ftpusers, and so on) and those procedures should either be extended to cover sensitive data owned by users or similar procedures should be written specifically for this task.

On a permission-related note, world-write access to user data files should be prohibited whether user data is sensitive or not, as world-writable files are vulnerable to intentional or accidental corruption by any process on the system. To maintain data integrity, world-write can only be allowed when no other option exists. This is important to the sample configuration too. If there are world-writable files, it is possible that a user who owns sensitive data could accidentally write a copy of a sensitive file over the top of a world-writable file in the unclassified domain. Opportunity for this type of action should be minimized.

Because there may be groups of people legitimately working together on sensitive data, sharing may still be allowed by assigning those people to groups defined on the system and permitting group-read and group-write access. Group permissions would be left up to the groups themselves to manage as they see fit. But at no time would they be allowed to make their files world-read or world-write.

Systems containing sensitive data typically do not require mandatory access controls. Therefore,

an obvious concern in having a system with sensitive data is that a user may mistakenly change permissions to allow world-read or world-write. To guard against this, automatic procedures should be in place to frequently scan the filesystems and report any violations of this policy. Any violations found should be corrected as soon as possible, and the owner of the data should be contacted to make sure (s)he understands the policy, and confirm that it was a mistake that the violation occurred.

Defining security domains

NIST defines a "security domain" this way: "A set of subjects, their information objects, and a common security policy."(Stoneburner, et al, page B-3) NIST defines a "subject" to be: "An active entity, generally in the form of a person, process or device, that causes information to flow among objects or changes the system state."(Stoneburner, et al, page B-3)

A convenient way to define a security domain is to have a particular filesystem, or set of filesystems, in which only sensitive data is allowed. Any new user or group of users who need to work with sensitive data would be given directories in these filesystems. The filesystems serve to define a boundary as described in NIST SP800-27: "Clearly delineate the physical and logical security boundaries governed by associated security policies." (Stoneburner,et al, page 7)

The obvious feature to go along with a filesystem dedicated to sensitive data is a separate filesystem or filesystems to be a security domain for unclassified data. User data that has no confidentiality requirement can be placed in this filesystem and the world-read permission may be allowed.

By dividing the filesystems into sensitive and unclassified security domains, it is easier to build tools that check for proper file permissions as the tools don't need updated every time a user is added to the system. The process that checks for improper permissions will always check that all files in a given security domain meet the permission requirements for that security domain.

Accommodating users who need to work in more than one security domain

Should a person have a need to work on separate projects, some of which are sensitive and others which are unclassified, that person should be given separate accounts for each project, and the projects should be located in the filesystem appropriate to the classification of the data, or in other words, an account in each security domain.

It could also happen that within a single project, there might be need to have sensitive data and unclassified data. The same model would be used here: separate accounts for the sensitive and unclassified sides. Those accounts would reside in the security domain appropriate to the classification of the data.

Procedures for declassifying sensitive data

If a user wishes to unclassify sensitive data, or in other words, make the data publicly accessible and therefore no longer sensitive, as the owners of sensitive data commonly are authorized to do, they can use a file transfer utility, such as running ftp from the sensitive account to the unclassified one. This is a more explicit way of unclassifying data and it makes it more clear that there was an intent to move the data from the sensitive domain to the unclassified domain, as opposed to setting world-read, which could happen by mistake.

Comparing the configuration to the government guidelines

At this point, some readers might be concerned that a system with sensitive data is open to users whose data requires no confidentiality. This concern may stem from the fact that in some environments, users who are authorized to handle the sensitive data may need some sort of administrative authorization (i.e., have passed a background check, or be above a certain rank or level of management). Having this authorization might imply approval to access a system containing sensitive data (though, as stated above, it doesn't necessarily authorize access to all the sensitive data). There are multiple ways to address this concern.

The DoD view

The DoD answer to this is to treat all data on the system as if it was at the highest classification of any data on the system. From DoD 5200.28: "The level of control and protection shall be commensurate with the maximum sensitivity of the information and shall provide the most restrictive control measures required by the data to be handled..."(DoD 5200.28, page 23)

This means that, in the DoD view, any system with sensitive and unclassified data shall be treated as if all the data was sensitive. This view is also expressed in GIGIA: "All GIG [Global Information Grid] information systems shall employ protection mechanisms in accordance with the level of concern (i.e., high, medium or basic) that satisfy corresponding criteria for high, medium, or basic levels of robustness." (GIGIA, page 2, section 4.3)

And further explains:

GIG information systems processing sensitive information... are assigned a basic level of concern and shall employ IA [Information Assurance] products that satisfy the requirements for at least basic robustness when the information transits public networks or the system or network handling the information is accessible by individuals who are not authorized to access the information on the system.(GIGIA, page 3, section 4.3.3)

"Basic" in GIGIA terms is a defined level of concern that mandates certain security measures be in place. If the unclassified data was, say, an open-to-the-internet web server, the lack of confidentiality makes the system not meet the "basic" requirements. Another feature of the basic level of concern is that users must authenticate themselves for all accesses, even through a web page. (GIGIA, Implementation Guidance, page 21) The GIGIA would require this:

GIG information systems that allow open, uncontrolled access to information through publicly accessible web servers or unregulated access to and from the Internet shall employ mechanisms to ensure availability and protect the information from malicious tampering or destruction. Such systems shall also be isolated from all other GIG systems. (GIGIA, page 3, section 4.3.4)

In essence, the DoD is defining the security domain to be the computer system. Any overlapping security domains will be treated as if all security domains are at the highest security level. These DoD regulations leave only one possibility: Split out the unclassified data that is for public consumption onto machines physically separated from the sensitive machines. This serves the mission of keeping the sensitive data confidential, and the public data public. But it fails to use minimal resources as there is the requirement that multiple systems be used.

One possibility is to further subdivide the people needing access to the unclassified data into the group who maintain the data, who clearly need write-access to the data, and the group who only needs read access to the data. The set of users needing write access to the data may be allowed on the system (we'll assume they have met the requirements for authorized access to the system), but the filesystem may be exported read-only and root-as-nobody, using the Network File System (NFS) or a similar product, to another system that serves the data to the community only needing read access to the data. This is a possibility more open to interpretation: the users who only have read access to the exported filesystem 1. do not have access to the server (at least they can't log on), and 2. do not have any access to the sensitive data on the server. And exporting a read-only filesystem that contains only unclassified data from a server that contains sensitive data on other, physically separate filesystems, does not present any greater threat to the confidentiality of the sensitive data. (Unless there are vulnerabilities in NFS, which there could be. But these vulnerabilities might be present on the server anyway if, for example, the server was serving sensitive data over NFS to another machine accessible only to users authorized access to the sensitive data. If the NFS daemons are running at all, they may be a vulnerability no matter where they are exporting data.) A strict interpretation of the DoD regulations would still forbid this: the exported, unclassified filesystem is still part of the sensitive system (that is, the sensitive security domain) and therefore cannot, under any conditions, be allowed public access.

The NIST view

NIST recognizes that there are times when a single system may contain both sensitive and unclassified data. In SP800-27, the third principle presented in the paper states "Clearly delineate the physical and logical security boundaries governed by associated security policies." (Stoneburner, et al, page 7)

The discussion of this principle has valuable insight into this problem:

Information technology exists in physical and logical locations, and boundaries exist between these locations. An understanding of what is to be protected from external factors can help ensure adequate protective measures are applied where they will be most effective. (Stoneburner, et al, page 7)

Relating this to the example, it is important to keep in mind that what is being protected is the sensitive information. The key isn't who has access to the system, but rather who has access to the security domain. So the boundary should be defined where it best protects the information. The discussion of this principle goes on to say:

Sometimes boundary is defined by people, information, and information technology associated with one physical location. But this ignores the reality that, within a single location, many different security policies may be in place, some covering publicly accessible information and some covering sensitive unclassified information. Other times a boundary is defined by a security policy that governs a specific set of information and information technology that can cross physical boundaries. Further complicating the matter is that, many times, a single machine or server may house both public-access and sensitive unclassified information. As a result, multiple security policies may apply to a single machine or within a single system. Therefore, when developing an information system, security boundaries must be considered and communicated in relevant system documentation and security policies. (Stoneburner, et al, page 7)

Accepting the fact that both sensitive and unclassified data, with the unclassified being accessible to the public, is a clear difference of philosophy with the DoD. But NIST believes with proper application, such an arrangement may be acceptable.

Some of the other principles serve to define how NIST recommends such decisions be made. Principle number four states "Reduce risk to an acceptable level." (Stoneburner, et al, page 8) This seems obvious, but the point here is that beyond a certain point, the expense of reducing risk exceeds the cost of losing the data or the confidentiality of the data. Cost-benefit analysis should be applied.

Principle 9 states "Strive for simplicity." (Stoneburner, et al, page 10) and principle 11 states "Minimize the system elements to be trusted." (Stoneburner, et al, page 11) Which is simpler, a single system in a more complex configuration for serving unclassified and sensitive data, or two separate systems in simpler configurations, one serving sensitive data, the other unclassified? And are there situations where a single system need not even be much more complex to serve both classes of data?

Principle 15 speaks directly to these questions: "Formulate security measures to address multiple overlapping information domains." (Stoneburner, et al, page 12) The discussion of this principle states:

An efficient and cost effective security capability should be able to enforce multiple security policies to protect multiple information domains without the need to separate physically the information and respective information systems processing the data. This principle argues for moving away from the traditional practice of creating separate LANs and infrastructures for various sensitivity levels and moving toward solutions that enable the use of common, shared, public infrastructures with appropriate protections at the operating system, application, and workstation level. (Stoneburner, et al, page 12)

In this statement, NIST is moving beyond the DoD model of strict physical separations of sensitivity levels, to a model that assesses the individual organizational security needs and missions, tailoring the solution in an efficient and cost-effective manner. Clearly, this does demand careful design and implementation.

Relating the NIST guidance to our example, it appears that it is an acceptable one. The cost-benefit is likely better than the DoD model because a single system is being used with multiple security domains defined. It is arguably more simple, as there is one system instead of two. The same goes for the number of elements to trust - one system vs. two. The security policies regarding the data is well defined as are the boundaries between the classifications of data.

Conclusions

In this paper, a sample configuration has been put forward for serving both sensitive and unclassified data from a single server. The benefits of this are greater economy, simpler environment, and fewer items to trust and maintain. The requirements to do this are clearly defined policies that are effectively communicated and enforced, clear boundaries between the security domains, and effective defense-in-depth management of the system.

While this model appears to comply with the NIST guidelines, it appears to conflict with DoD guidelines. In the DoD view, the model fails to sufficiently separate the publicly accessible data

from the sensitive data. Doing so would require separate systems for the types of data and would be more expensive to do under DoD guidelines. But the DoD may have good reasons for such a mandate. One possible reason is that these requirements mandate a more uniform environment that requires less assessment when a new mission is proposed for the system.

The goal of this paper is not to condemn one set of guidelines or the other, but rather to present the differing philosophies. The organization facing a task similar to that presented in the paper must come to their own conclusions about which model better serves their needs (or is mandated) for their environment.

List of References

1. Joint Security Commission. "Redefining Security, A Report to the Secretary of Defense and the Director of Central Intelligence." February 28, 1994. URL: <http://www.dss.mil/search-dir/seclib/jcs.htm>.
2. United States Department of Defense. "Directive Number 5200.28. Security Requirements for Automated Information Systems (AISs)". March 21, 1988. URL: http://www.dtic.mil/whs/directives/corres/pdf/d520028_032188/d520028p.pdf
3. United States Department of Defense. "Guidance and Policy for Department of Defense Global Information Grid Information Assurance" June 16, 2001, URL: <http://www.c3i.osd.mil/org/cio/doc/gigia061600.pdf>
4. United States Congress. "Computer Security Act of 1987." January 8, 1988. URL: http://cio.gov/Documents/computer_security_act_Jan_1998.html
5. United States Department of Energy. "Order 5635.4, PROTECTION OF UNCLASSIFIED CONTROLLED NUCLEAR INFORMATION", February 3, 1988. URL: http://www.fas.org/irp/doddir/doe/o5635_4.htm
6. Helyer, Andrew. "Sensitive But Unclassified." April 3, 2002. URL: <http://rr.sans.org/policy/sensitive.php>
7. Swanson, Marianne and Guttman, Barbara, National Institute of Standards and Technology, "NIST Special Publication 800-14, Generally Accepted Principles and Practices for Securing Information Technology Systems", September, 1996. URL: <http://csrc.nist.gov/publications/nistpubs/800-14/800-14.pdf>
1. United States National Institute of Standards and Technology, Computer Security Resource Center Special Publications, URL: <http://csrc.ncsl.nist.gov/publications/nistpubs/index.html>
8. Stoneburner, Gary and Hayden, Clark and Feringa, Alexis, National Institute of Standards and Technology, "NIST Special Publication 800-27, Engineering Principles for Information Security (A Baseline for Achieving Security)", June, 2001. URL: <http://csrc.nist.gov/publications/nistpubs/800-27/sp800-27.pdf>

