Server Security Standard

Information Technology Security Standard #5, v1.1

Standards Name: Server Security Standard

Original Release Date: 3/29/2006
Latest Revision Date: 7/03/2007

Authority: This standard was approved by the Information Security Council, acting as representatives of the CIO and the Technology Leadership Committee (TLC). It derives its authority from the CIO and the campus Information Technology (IT) community.

Summary: This standard provides the minimum security requirements for all University owned and operated Servers, where the term Server has a specific meaning, as provided in the Definitions.

Applicability/Scope: Required for all who operate a University-owned or operated Server. It should also be followed by all campus organizations that offer IT services to the campus or Internet.

Responsible Office: ITS/CIO.

Definitions

Server: any network-connected device that offers resources to some or all of the campus or to the Internet. This includes traditional Windows, Netware, MacIntosh and UNIX servers, along with network appliances, printer/copiers with server-like capabilities, and a wide range of other, similar devices.

Unit Head: While Syracuse University is the legal owner or operator of all IT Systems, it delegates oversight of particular systems to the head of a specific unit, department, or office of the University ("Unit Head"), or to an individual faculty member, in the case of IT systems purchased with research or other funds for which he or she is personally responsible. A Unit Head will be a dean or vice president or someone operating at that level.

Unit IT Leader: The Unit Head may designate an individual to manage the IT Systems within a unit. The Unit IT Leader oversees the technical staff, including Systems Administrators and Application Developers, who design, build, operate and maintain the IT Systems.

Reason and Background for this Standard

Servers usually provide access to each unit’s data and IT services and are a critical component of the academic, business and research missions of the University. Poorly secured servers will lead to exposure of University Data, disruption to the business, and a general uncertainty about the integrity of the data and systems. It therefore is vitally important to apply well-accepted security standards and practices to those Servers. This standard provides both generic information that applies to all Servers and additional, platform or service-specific information in the Appendices. Certain kinds of network services need to be carefully regulated on the campus network or else serious disruption to campus services will occur. Such services are described in Appendix A.

Standard for All Servers

This portion of the standard applies to all Servers, regardless of their operating system. Servers that fit one of the types described in the Appendices must also apply those platform-specific standards. For example, a Windows 2003 Server must also meet compliance with the requirements in the appendix for Windows Servers.

Those Servers that house Confidential University Data, or that provide access to it, may be required to meet additional requirements, as defined by the appropriate Data Custodian or by the ISC or CIO. Standards for these situations are shown below as (HS), for “High Security.” Operators of such Servers need to be aware of any relevant laws and policies, which are outlined in the Security Practices for Technical Staff.

As with all IT Security Standards, exceptions may be requested using the IT Security Exception Process.

All Servers that offer resources and services to the campus or the Internet must serve the mission of the University and operate in accordance with all University policies. When they are operated in a fashion contrary to this mission, they are subject to disconnection and the owner may be subject to disciplinary action as defined in other University policies and procedures.

The following standard describes a generic server that is built to a “bastion host” configuration: that is, it begins life in the most secure state possible. From there, needed features and services are enabled.

1. Change Control

a. Server Deployment**:** All Servers must be registered and audited before offering services to the campus and/or the Internet. Registration and auditing procedures change over time and are covered in the on-line pages covering security procedures.

b. System integrity and security is always impacted by changes. Changes need to be made with proper thought to the possibility of new risks being created by the change.

c. A major part of security management for Servers is the application of patches, updates and other fixes. Such updates need to be part of an overall maintenance process and need to be applied in a timely fashion, weighing the risks from not applying the update against possible disruptions due to the effects of the update.

i. Security updates that mitigate malware able to propagate automatically (e.g., worms) must be applied within 1-2 days of their release and sooner if there are exploits in the wild. Servers lacking such updates that pose a risk to the campus may be disconnected.

ii. If an update would impact normal services provided by the Server, other methods of mitigating the risk may be used, but must be cleared with the Information Security Officer.

iii. Changes should always be tested to the degree that is practical. Units should work cooperatively to determine if a change has been adequately tested within a similar environment elsewhere on campus.

iv. Updates must be pulled from sources known to be secure. Any internal update servers should be designed for this purpose and carefully controlled, so that updates can be trusted. An update server that receives feedback from the Servers on successful installation of updates can be a useful tool for verifying the update process.

2. **Access Control**

a. Physical Access

i. All Servers must be in physically secure locations.

ii. Those who require access must be authorized by the Unit IT Leader. All staff within the unit must be trained on the procedures by which people can be admitted to the Server room.

iii. (HS) Access must be controlled by a mechanism which records an identity and all entries and exits must be recorded.

b. Media Access.

i. All removable media containing non-public University Data (which includes any credentials) must be stored in a secure location, following the standards listed above for Physical Access (1a).

ii. (HS) All removable media must be inventoried and verified at regular intervals.

iii. (HS) Media containing Confidential Data must be encrypted or housed in a site with the access controls described in Physical Access (2a).

c. Authentication

i. User and Administrative Authentication processes must follow the Authentication Standard. That standard is considered a sub-component of this standard.

d. File system and Data Access (authorization).

        i. Any operating system that stores non-public University Data must store that data on a filesystem that provides user and/or group-based access controls. Those controls must be configured to prevent access to the non-public data by unauthorized users, including any anonymous or “guest” users.
        ii. Similarly, any applications provided by the Server must ensure that only authorized users may access files on the Server. This usually means that the Server must either authenticate the user, or extend trust to another Server that has performed the authentication. Trust relationships need to be understood, documented, and managed, since they can easily turn into avenues of attack.
        iii. File shares that provide access to large user populations across unit boundaries, or to “any” account, must be regularly monitored for abuse and unwanted changes, particularly if users have write access.

3. **Logging and auditing**
    a. Operating System Logging.
        i. Operating systems events related to security must be logged, preferably to an external, secured log server that meets all the requirements of this standard. Locally stored security logs may provide some evidence of an attack but are usually cleared after a successful attack.
        ii. Any authentication or privileged access events should be logged. If such logging produces unmanageable quantities of data, reducing the logging to authentication events is acceptable.
        iii. Logs should be configured to be large enough to hold at least a week’s worth of data.
        iv. Accurate timestamps are essential to good logging; a properly managed time service must be used.
    b. Log Review.
        i. Security logs must be reviewed regularly, where regularly means at least weekly. Use of automated tools that can examine logs and provide alerts when unusual events are seen can greatly reduce the burden.
    c. System Baseline and Regular Auditing.
        i. After the initial build of the Server is complete, a security audit must be performed that both verifies the basic security of the Server and provides a baseline against which the Server can be checked for unwanted changes.
        ii. A similar audit must occur at regular intervals, with the intent of finding unauthorized changes to system security.
        iii. ITS Security will perform scans of registered servers from time to time. Any serious vulnerabilities that are found will need to be rectified within a reasonable timeframe; the responsible administrator can work with the ISO to determine a suitable plan for remediation. Vulnerabilities that pose an immediate risk to

confidential data may require taking a service off-line until the risks can be addressed.

4. **Reducing the Attack Surface**
    a. Unnecessary services and programs should be disabled or removed.
    b. As a general rule, test and sample scripts should be removed. Many of these offer weak points or opportunities for reconnaissance to an attacker.
    c. Additional protections should be used to limit unwanted traffic sources from reaching the services. For example, a firewall can protect the system management interfaces (e.g., SSH, RDP) from access by sources other than the system administrators.
5. **Detecting Malware**
    a. Servers that transfer or store User files and data *must* run malware detection software. See the Compliance section for more specific information.
    b. Servers should also run regular scans of the files and system settings to look for unauthorized changes.
6. **Recovery**
    a. Administrators must assume that a Server will eventually fail or be compromised. Proper backups of the system and its data must be made regularly and well documented. Since prior backups may be the only reference for determining the sets of data involved in a compromise, all Servers housing non-public University Data must be backed up at least weekly.
7. **Trust Relationships and Dependencies**
    a. A trust relationship is one in which one machine trusts another implicitly: certain actions are allowed by virtue of the trust relationship and the trusted machine does not have to authenticate (or has lesser authentication requirements) for certain actions. Too much reliance on trust relationships can allow an attacker to move easily from machine to machine. For this reason, trust should never be granted to administrative accounts from a less secure machine (e.g., a desktop) to a more secure machine (e.g., server). Similarly, a less sensitive system can depend on a more sensitive system, but not the reverse.
8. **Administrator Practices**
    a. The daily practices of Server Administrators may be the single most important factor in the security of the system. Important security practices discussed more fully in the compliance section.

Part 2 -- Complying with the Server Standard

The following checklist provides detailed information on why certain elements of the standard exist and a section by section list of steps that must be performed. Some steps are treated as optional, but encouraged, and are noted as “Optional” or “Guideline.”

The ability to provide server registration and auditing is still being developed. The following is an expected timeline for these processes, and shows which groups are impacted:

1. July 1, 2007: All ITS Servers offering Internet services must be manually registered with IT Security and must pass a network vulnerability scan before offering Internet or campus service. Units may also request vulnerability scans by contacting itsecurity@listserv.syr.edu.
2. July 1, 2007. By this date, auditing tools will have been developed that can snapshot the initial build of any AD-based Windows server and can also be used for subsequent audits.
3. July 1, 2007. Servers offering services that are blocked at the Internet edge may provide direct Internet service if there is a justified need and the risks are limited. Server administrators can request an exception through the usual exception process, which is documented on the ITS security pages.
4. January 2, 2008. The default Internet filtering policy for ITS server subnets will change from “default allow” to “default deny.” This marks the completion of the ITS server security zone known as “the bubble.”

Section 1 – Change Management

Registration

- As capabilities are developed, Servers will have to be registered with ITS and must indicate the network services they are offering (web via HTTPS, SMTP, file sharing via SMB, etc.). This will ensure that network defenses can be properly aligned with the needs of the Server and that no unnecessary services are exposed to the Internet. See the ITS web pages related to registration for details about exact registration processes.
- Once a more automated audit process is available, Servers must receive a full audit, including a vulnerability scan, before providing service to the entire campus or the Internet. Any serious vulnerabilities must be corrected before the service is allowed to reach a production state.

Changes

- Units must have a process for testing and applying patches and other updates related to security flaws. Any security vulnerability that can be remotely exploited must either patched or otherwise mitigated within a few days of its announcement, or the server is likely to be compromised. ITS may be able to assist units in providing special network defenses for servers that cannot be fully protected by the Server’s own defenses.
- In cases where a Server cannot apply a critical security fix because of its impact on a vital application, units should discuss other options with IT Security. If a particular vendor develops a pattern of not repairing security flaws in a timely

manner, the unit or campus should reassess the overall costs of continuing to use that product.

- Note that changing the name or IP address of a Server, or removing a Server from service is a change similar to registering the Server. Such changes will eventually need to be part of the registration process, so that unnecessary holes are closed and unnecessary monitoring removed.

Section 2 – Access Controls

Physical Server and Server Room Security

- Servers must be housed in a dedicated, locked room. Placing servers in open areas or in shared access rooms creates uncontrollable risks.
- All potential entry points, including windows, must be adequately resistant to physical forcing that special tools and extended time would be required to defeat them.
- (HS) Lock cores must be separate from the building master.
- (HS) Entry points must be in visible locations that can be monitored (server rooms should not be in remote corners of a building where one could spend hours attempting to gain access without being noticed).
- A list of people who require access, including any needing emergency access, must developed, approved by the unit head, and communicated to the unit staff.
- Controls must be implemented that only allow those staff to have access.
- All staff must be trained on how social engineering attacks work and procedures must be developed that indicate how staff may allow other (e.g., maintenance) personnel into the room.
- Combination locks or keypads must have their combination changed whenever someone who knows it leaves an administrator role or at least once per year.
- (High Security) Each access must be logged and the logs reviewed at least weekly.
- (High Security - option) Accesses must be recorded to a secure video system and reviewed weekly.

Media and Console Access

- Initiating a reboot of the server from alternative media (CD, floppy, USB drive) should not be possible without providing a password.
- (Optional) CD/DVD drives and other attachable drives should have auto-run disabled.
- The server’s GUI/text system console must be left in a logged-off or password-protected screen-locked state.
- The last logged-on user’s name must not be shown when a login sequence is initiated.

- The server must be attached to a network switch/hub that is located in the server room or a secured room.
- Media used for booting and repairing the servers must be labeled and stored within the server room or other secure room.
- Backup and removable media must be similarly secured.
- (High Security) Inventories should be performed at regular intervals to verify that no thefts have occurred.
- (High Security) Backup media must be encrypted or stored in a room with the High Security controls described under Physical Security.
- Any media being reused for other purposes or being disposed of must be sanitized according to the Data Security Standard. Remember that even boot media may contain valuable passwords or configuration information.

Authentication

- Use the Authentication Standard when implementing authentication mechanisms (login controls).

Filesystem and Data Access Control

- Most operating systems offer a filesystem that has user and group-based access controls. Whenever non-public University Data is stored on the Server, it must be housed on such a filesystem (NTFS in Windows) and the permissions for these data must not include “guest” or “anonymous” or “everyone” or whatever the equivalent may be for the operating system. This means that personal folders, mail stores, departmental shares/files holding this kind of data must be properly protected. Sloppy management of “groups” and group permissions often allows people to gain access to files and folders they are not authorized to see.
- (Guideline) Any filesystem that used in conjunction with web or other public services must be carefully protected so that the front-end service doesn’t have unnecessary write access. Services like IIS and Apache should never be run with the administrator/root rights or rights that would allow them to modify system files and settings.

Section 3 – Logging and Auditing

- Logging of security events needs to consider the value of the information. All servers need to record any login/logout activity on the administrative interfaces, such as SSH, web management consoles, Windows network logins, etc. Where possible, the logs should be sent directly to a common server, where they can be analyzed and reported on using automated tools. Manual inspection of log data is tedious and unlikely to be done regularly, but having the raw logs available can be useful for forensics.
- For Windows, the Audit Policy under the Local Security Policy should enable these settings:
    - Account logon events: success, failure

  - Account management: success, failure
  - Logon Events: success, failure
  - Policy change: success, failure
  - Privilege Use: failure
- For UNIX, all SSH and other administrative logins should be sent to an external syslog server, as controlled by syslog.conf. Any use of inetd/xinetd should also be logged, along with console login activity.
- Actually analyzing the UNIX and Windows logs is a challenge. While many syslog tools are available, each has its own strengths and weaknesses. In the Windows world, use of commercial tools is almost mandatory. A good resource for finding logging tools is http://www.loganalysis.org. Make sure the tools you select have good ongoing support and are not a personal project that someone has abandoned.
- Accurate timestamps are essential for obtaining accurate logs. NTP-based time syncing must be run on every server. Until SU has a more reliable NTP source, we recommend the use of one of the NIST NTP servers. For less critical systems and for UNIX servers, clock.syr.edu will usually work. We have seen problems using this server with certain appliances and Windows systems.
- Logs that can grow without limit should not be placed in partitions where filling the partition could cause system or service failure. Denial-of-service attacks can rapidly increase the size of logs. Setting up a separate partition for log data and other “variable” information can reduce outages.
- Periodic audits of a server provide a chance to detect unexpected changes. An external vulnerability scan can reveal backdoors and other open ports that may indicate a compromise has occurred. Locally-run audit tools can verify the state of the machine, provided that stealthing tools (e.g., rootkits) have not been installed by an attacker.

Section 4 – Reducing the Attack Surface

- Service Minimization. A key principle in building any Server is that it should not run programs and services that are not essential to its purpose. This notion of “service minimization” is one of the most important components in creating a secure system. Each additional service running on a Server gives an attacker another point of entry to leverage against some other part of the system. Services not running are not vulnerable and may not even need to be patched; overall maintenance effort decreases. Administrators need to be particularly careful of web services, which are generally wide open to the Internet.
- Wherever possible, client machines and even servers should use smtp-server.syr.edu for sending e-mail messages. This machine includes anti-virus and anti-spam filtering.
- The campus provides Internet-facing MX hosts that provide anti-virus/spam filtering for SMTP traffic. Most servers do not need to accept inbound SMTP traffic and should avoid doing so. Units should use the MX hosts for reception of SMTP traffic from the Internet, even if they provide the ultimate destination for

their staff's e-mail. Continuing to use Internet-facing SMTP servers that lack malware filtering leaves small holes through which viruses and spam can leak.

- Services that are not needed must be disabled, particularly the more commonly attacked services such as ftp, http, snmp, and file sharing. By reducing the number of entry points, there is less chance that a configuration mistake or security vulnerability will be exploited.
    - In the Windows environment, be careful of Windows 2000 Server, which installs far more by default than Windows Server 2003.
    - In the UNIX world, try to reduce or eliminate inetd/xinetd services, many of which are obsolete.
    - (Optional) Any use of SNMP should be limited to a version that encrypts the strings (v2 or v3).
    - Outdated and dangerous programs/protocols such as tftp, rcp, rsh, and the "Simple TCP/IP services" should be removed or disabled.
- Administrators need to be aware the both UNIX and Windows machines expose a number of RPC-based services and these often cannot be fully disabled, particularly on Windows.
- Many systems now include IPv6 functionality in the TCP stack. This can provide avenues of attack that are not easily detected; disabling unnecessary IPv6 support is a good idea, but test to verify that doing so doesn't have side-effects on other network components.

Section 5 – Detecting Malware

- A well-built server that has tightly controlled user access may not need any special detection tools, but since many servers allow users to modify locally stored files, it has to be assumed that malicious code will be stored in some of those files. In those cases, running an anti-virus package is required. The software must pull any detection signature updates on at least a daily basis and must scan the user files at least daily. Certain kinds of files, such as databases, may be exempt from scanning, if the files are immune from infection or the scanning would impair system operation. An alternative to server scans of user data is to perform the scans on the end user's device; this is acceptable, provided that a high percentage of the devices perform the scans.
- As discussed under logging and auditing, routine checking for changes and rootkits is strongly recommended. Administrators must assume that defenses will fail from time to time. Without detection, the problem won't be found until serious damage has occurred. Note that modern stealthing programs can actually run a backdoor on top of any open port and the backdoor will be invisible to external scans, unless they know how to open the door. The newer rootkit detection tools are improving in detecting malware from a running system, but the best way to verify the security of a machine is from a bootable CD (e.g., ERD Commander, Knoppix, etc.)

Section 6 – Recovery

Recovery from a compromise must include an assessment of which University Data, credentials, or other valuable information may have been stolen, or whether unauthorized changes to the system were made. This usually cannot be determined without some forensic analysis and will often require that prior backups be used for comparison. The amount of analysis must be in proportion to the potential loss; a server housing nothing but public information does not warrant any serious analysis, beyond discovering how the attack occurred. For systems housing sensitive University Data or holding credentials that could be used against other sensitive systems, a complete analysis is necessary and in some cases law enforcement may even become involved. In all cases, the unit or system administrator should report the incident to the CIRT (abuse@syr.edu or itsecurity@listserv.syr.edu) immediately; the CIRT can help make these determinations and advise on the best course of action. Units need to understand that incidents in their area may have implications for the rest of the campus, so sharing of information is in everyone's best interest.

Section 7 – Trusts and Dependencies

Excessive use of trusts can create a pathway from less secure systems to more secure systems that requires no additional authentication. Cases of using the old "rsh" services on UNIX systems have been known to allow a single machine to access all other UNIX servers without any passwords, simply by chaining requests. Similar goofs are possible with SSH trusts and other tools for remote management.

Section 8 -- Practices

None of the above security measures will help at all if the administrator practices sloppy security. A smart administrator is the best defense a server can have. Here are some simple practices that can help.

1. Keep your environment used for e-mail and other Internet activity separated from the environment used for accessing servers. In other words, use a second machine for systems administration. In a sense, this creates a firewall between your more exposed activities and your sensitive, internal work. An alternative approach is to lock your desktop down to the point that exploitation is nearly impossible: work with an unprivileged account, load only trusted software, install a personal firewall and anti-virus and anti-spyware software, etc.
2. Guard your administrative-level passwords at all costs. Never enter these passwords from a machine you don't absolutely trust.
3. Do not share administrative passwords from a highly secure system with those for less sensitive systems or desktops. Consider a two or three tier approach to the sensitivity of your resources and use separate accounts and passwords for each tier.
4. Passwords for databases and applications are just as important as those used for operating system access. Be careful about how these are stored and protected.

5. Passwords must be changed whenever someone knowing them leaves the organization or changes roles. If you cannot achieve this change in less than a day, you have a structural or procedural problem that must be resolved.
6. Stay sensitive to behavior that seems incorrect for a system. Unusual behavior can be a telltale sign of an attack or a change to the system. Know your systems and what is normal. Take snapshots of settings, perform scans, check your logs, etc. Monitoring is every bit as important as protective defenses. Build some burglar alarms: for example, a script that watches for changes to executable files in the system32 folder can be tip-off that something bad is happening (or that you installed a patch!); the possibilities are endless.
7. Understand what traffic is normal for your servers. Using a sniffer to examine your server traffic can be enlightening and even help find problems. IT Security will be available to help with traffic collection and in interpreting what the traffic means.
8. Understand what your truly valuable resources are and protect them accordingly.
9. Make an effort to understand the business processes in your area and how people handle the University Data. While you may not be able to know everything that happens, understanding what staff are doing with the data will provide opportunities to point people to more secure methods for moving and storing the data.

Appendix A – Windows Server Standard

These requirements are in addition to the basic Server Security Standard.

1. **Platform**. Windows Server 2003 is the standard platform and must be used wherever possible. Use of Windows 2000 Server may be used, provided the following occurs:

a. All anonymous enumeration of accounts and shares is disabled.
b. There is a plan for migrating to Windows Server 2003 (or its replacement) that will be implemented before any end-of-support by Microsoft occurs.

Use of Windows NT is prohibited, due to the lack of Microsoft support (unpatchable). Anyone needing this platform must file an exception request and will be expected to provide an external firewall for protecting the server**.**

Appendix B – Restricted Network Services

The following network services must not be run on the campus network without prior approval from ITS, since they can easily cause major disruptions on the campus network.

1. DHCP.
2. Routing protocols, such as RIP and OSPF.

3. Multicast IP.