Linux Server Configuration Guidelines

This document is meant to be a living document and intended to accompany more detailed, step-by-step resources. Suggestions in this document are taken from administrators around campus and are provided to serve either as a checklist and reference for those who may work with servers on a regular basis, or a starting point to those who are new to the task. Content in this document is provided as a suggestion or reminder with emphasis on best practice.

ITS has a license for Qualys Vulnerability Management scanning. Contact us if you would like to learn more and potentially configure scans of your servers to help identify potential vulnerabilities.

Table of Contents

Do You Really Need a Server?

ITS offers many hosted solutions to the rest of campus. Listed below are some data technologies that ITS offers as a service.

Web Hosting – Ektron CMS, Apache, Tomcat, IIS
Application Server Technologies – ASP, JSP, PHP, CGI
Database Technologies – MySQL, Oracle SQL
File Servers – Windows Share Drives (Samba), SFTP

When possible, it is best to leverage these services instead of standing up a new server.

A Note on PCI Compliance

PCI Compliance dictates certain standards for any machine that comes in contact with payment cards of any type. These same standards apply to any machine that shares a network with any other machine that comes in contact with payment cards. ITS must be contacted before any servers that may come into contact with payment cards are stood up.

Planning

It is a good idea to formally identify the goals, scope, and needs of a new server before beginning. This will form an idea of what the end product will look like and help identify pit falls before they come up.

Server Role

After defining goals, determine what roles the server will perform. Servers should be purpose-built devices. This means that a server should run the services that are necessary to perform its role and no more. For example, a machine running as a webserver may not need to run FTP, cups, or postfix. Some common server roles include webhosting, file hosting, database, directory, and print management.

Further down is the “Role-Specific” section containing guidelines for some of the more common server roles.

Vulnerable Services

Some services have inherent security vulnerabilities, often because they do not encrypt sensitive traffic. Use more secure alternatives when possible. For example, SSH can (and should) replace Telnet for administrative tasks. SCP is more secure than FTP and can fulfill most of the same needs. Most modern FTP clients also support SFTP, which relies on the server's sshd service.

Some services may not have a secure alternative. If legacy or insecure services must be run, use compensating controls to manage them. For example, VNC can be run through a secure, encrypted SSH tunnel to keep VNC passwords from floating down the wire in plain text.

Picking a Platform

At this point, it is time to pick a distribution of Linux. Use a mature server-oriented system when at all possible. ITS tends to use Red Hat Linux and may be able to provision licenses upon request. Other examples of server-oriented Linux distributions include CentOS and Ubuntu Server. BSD (Berkeley Software Distribution) is also a Linux-like option for some, but is generally considered a higher learning curve than common Linux distributions.

Network and Access Control

It is important to plan network architecture and configuration ahead of time. Using a "least access necessary" model, plan what network segments need access to this machine. For example, should the entire world be able to SSH in to the server? Or should just IPs on campus. Maybe only IPs from the VPN with the System Administrator's on-campus workstation as a backup would work. Machines with port 22 open to the world are generally attacked with weak and default password checks all of the time. Limiting access to ports and services can lower or eliminate these sorts of attacks from the outside world.

If only a limited number of users need to access the server, put it on the 10-net. If the server needs to host web pages to the rest of the world, consider putting it on the 10-net anyway and requesting a NATed IP address.

Once the location on the network has been decided, determine what the firewall rulesets should look like. For example, on a webserver, iptables can be configured to allow all traffic to port 80 (http) on the server, but only allow connections from a few ip addresses to port 22 (SSH). If the server is going to be a virtual machine, be sure that the host secured and that there is no network connectivity between the host hypervisor and the virtual server.

Plan the use of a kernel-level firewall, such as iptables or ipfw. Use these to dictate what ports are available to what IPs. For example, on a webserver, iptables may allow all traffic to port 80 (https) on the server, but only allow connections from a few ip addresses to port 22 (SSH).

Lastly, make sure any DRAC, hardware controller card, hypervisor console, or other management interface is on a strictly firewalled, private subnet. This subnet should not be routable to non-NAU IPs. Use an NAU VPN if there is a need to manage remotely.

Users and Authentication

Identify potential users and access levels ahead of time. Determine what users need what roles and what groups will be needed to manage those roles. For example, perhaps student workers should be able to power down or reboot machines, but should not be granted any further administrative powers. A “student_worker” group or a “power managers” group may be necessary so that student workers are not granted full root access just for this purpose.

Setting Things Up

Setting up Accounts

The first task that should be completed is changing the root password to something long and preferably complex. Either make your team remember it, or store it in a secure location, such as an encrypted password safe.

Next, set up administrator accounts. Create accounts for administrators, set initial passwords, and then set up sudo. From here on, use `sudo su -` to become root from a administrator user instead of logging in directly as root. Remember to use of the root account to should be limited to a need-basis; don’t use admin privileges where they are not absolutely necessary. At this point, test your new administrator accounts and consider setting up SSH Shared Keys.

Make sure that there are no service accounts with passwords. This will prevent service accounts from being brute-forced via ssh or other login mechanisms. To do this, check that in the "/etc/shadow" file, there is a "!!" in the field where a password hash would belong and that the user’s shell is set to either "/sbin/nologin.”

Addressing Defaults

Once accounts are configured and administrator accounts have been tested, disable root login via ssh. In "/etc/ssh/sshd_config" make sure the line "PermitRootLogin no" exists and is uncommented. Disable ssh protocol 1. In "/etc/ssh/sshd_config" make sure the line "Protocol 2" exists so that protocol 1 will not be enabled.

Next, set the root password on mysql. This can be achieved with the following command: `mysqladmin -u root password NEWPASSWORD`.

Disable IPv6 unless there is a need for it.

Remove unwanted/unused packages that were not specified at install time. Some packages cannot be uninstalled, in which case disable corresponding services that are not needed. Servers typically do not run X by default. If use regular use of the server's GUI is not anticipated, change the default run level to init 3 in "/etc/inittab". `startx` can be used to launch the desktop environment after logging in on the console if really necessary, but some may find Xserver unnecessary altogether. Below is a list of commonly removed packages:

"smartd bluetooth libvirt-devel libvirt-java-devel libvirt-client libvirt-java xinetd ypbind ipa-client oddjob-mkhomedir sssd"

Once all unnecessary packages have been removed, begin identifying and disabling unnecessary service. Use "/sbin/chkconfig --list" and look for anything that is "on." There are some services that need to run for OS/hardware to work. Services like "haldaemon" and "messagebus" are examples of these. Check before disabling a service when unsure.

Following that, use "netstat –ln" to make sure all unnecessary services are disabled. Any remaining unnecessary services should be hunted down and disabled or uninstalled.

This is optional but a good security choice – set the sticky bit on the /tmp directory so that non-root users cannot delete or move other users' files.

Access Control

Begin configuring TCP Wrappers. TCP Wrappers are blunt-force instruments to hard-limit incoming and outgoing connections based on service. Use a strict "/etc/hosts.deny". Often, hosts.deny may contain only one line: "ALL:ALL except localhost". Be as specific as possible with "/etc/hosts.allow" entries by specifying exact IPs and by avoiding wildcards.

Afterwards, further configure iptables to limit incoming and outgoing traffic. Red Hat provides some documentation on configuring iptables.

Make sure any open ports are protected by both tcp_wrappers when possible and iptables.

Automatic Updates

If using Red Hat Linux, for example RHEL 6, run “/usr/sbin/rhn_register" and configure RHN registration to use the NAU proxy server located at royal.ucc.nau.edu to enable updates. This was a change from RHEL 5 which used “/usr/sbin/rhnreg_ks” instead. Some changes in December 2014 may result in further changes with the licensing as Red Hat moves to the Satellite subscription service. Turn on auto-erratta in RHN unless there is a compelling reason not to do so.

For other systems, configure the package manager to automatically keep your system up to date. If some packages need to be excluded from auto-update for business reasons, exceptions can be made. However, the default should be to auto-update packages.

Logs

Use an external syslog server on the 10-net to maintain log information for multiple servers whenever possible. This way, if a host does get compromised, logs that the intruder cannot modify will exist on a different machine.

The syslog external logging may require opening a port on the client sending out logs or conversely, on the syslog server end. Ports will depend on the specific communication mechanism as well as whether using syslogd or (starting in RHEL 6) rsyslogd.

Log iptables DROPs and REJECTs. Use the "iptables -m limit" flag to prevent DoS issues. Run logwatch reports from here (and we uninstall this elsewhere unless it is needed for things like Apache and Tomcat logs). Make sure that "/etc/syslog.conf" (or for RHEL 6, "/etc/rsyslog.conf") has an entry pointing to the correct syslog server. Make sure that the hostname of the correct syslog server is hardcoded and defined in "/etc/hosts".

Various Local Services

Set up ntpd to use "ntp1.nau.edu”, “ntp2.nau.edu”, and “ntp3.nau.edu” as local time servers.

Back up the system on a regular basis. ITS offers secure NetBackup services.

If the machine is a virtual machine, be sure to install and maintain the Latest VM Agent , such as VMWare Tools, XenTools, etc.

Server Hardening

SE Linux

"Security-Enhanced Linux (SELinux) is an implementation of a mandatory access control mechanism in the Linux kernel, checking for allowed operations after standard discretionary access controls are checked. It was created by the National Security Agency and can enforce rules on files and processes in a Linux system, and on their actions, based on defined policies." – Red Hat

Information Technology Services tends to not use SELinux for business-use reasons. It is however always an option to further harden a server and can only enhance the server's security.

SSL/TLS

Test or learn more about your Web Server's default SSL/TLS protocol configuration, cipher suite configuration, and certificate status using the following free, online, resources. In particular, 2014 saw several vulnerabilities to SSL and current best practices include disabling SSL, enabling TLS, where impacts to services will not be experienced.

- https://www.ssllabs.com/ssltest/
    - Be sure to check the "Do not show the results on the boards" box
- https://sslanalyzer.comodoca.com/

General server side SSL/TLS best configuration practices can be found here:

- https://wiki.mozilla.org/Security/Server_Side_TLS

Two-Factor Authentication

When planning what IP addresses to allow for administration, consider restricting access to a two-factor VPN. ITS hosts a two-factor VPN service using secure tokens. If this is something of interest, contact ITS for more information.

Antivirus

ITS offers antivirus for all platforms – including Linux. Sophos Antivirus for Linux is available upon request. Contact ITS for more information.

Linux Malware Detect

"Linux Malware Detect (LMD) is a malware scanner for Linux released under the GNU GPLv2 license, that is designed around the threats faced in shared hosted environments." - R-fx Networks

Malware Detect has proven to identify web-based malware (such as php scripts that users are redirected to from phishing emails) that traditional anti-malware will usually miss. ITS has recently found use for this among some of its webservers.

tripwire

Consider tripwire - this creates a hash of a configurable list of important files in a database. Periodically, the hashes are re-generated and compared to what is saved. This will detect changes to those files that could be a result of the machine being compromised and can be set to notify the administrator.

fail2ban

Consider implementing fail2ban. fail2ban monitors log files for failed login attempts and after a configurable number of tries, will block the offender's IP address permanently or for some configurable amount of time.

Disposal

Media destruction takes several forms, including physical destruction or electronic destruction. Physical destruction of media can take several forms, such as drilling holes in the physical drive. Electronic destruction needs to follow certain guidelines, and ITS suggests the standard DoD and NIST guidelines for 3-pass wiping of a drive. More details about the NIST Special Publication 800-88 can be seen here: http://www.nist.org/nist_plugins/content/content.php?content.52.

One such tool for performing the multi-pass drive wipe is found here: http://www.dban.org/