
Systems Configuration and Administration Best Practices

Abstract

An article covering some methodologies and best practices in architecting and administering IT infrastructure.

© 2006 Lesley Mitchell

The information contained in this document represents the current view of the aforementioned group on the issues discussed as of the date of publication. Because these are a variable within the discussion, it should not be interpreted to be a commitment on the part of the aforementioned group and that group cannot guarantee the accuracy of any information presented after the date of publication.

This paper is for informational purposes only. NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT



Attribution 2.0 UK: England & Wales

You are free:

- to copy, distribute, display, and perform the work
 - to make derivative works
 - to make commercial use of the work

Under the following conditions:



Attribution. You must give the original author credit.

- For any reuse or distribution, you must make clear to others the licence terms of this work.
- Any of these conditions can be waived if you get permission from the copyright holder.

Your fair use and other rights are in no way affected by the above.

This is a human-readable summary of the [Legal Code \(the full licence\)](#).

[Disclaimer](#) 

Introduction

Administering large IT systems is frequently a labour intensive task. This can often be traced to insufficient planning of the system as a whole, rather than as a group of individual machines.

Areas that should be addressed, include, but are not limited to the following:

- Consistency
- Naming
- Authentication
- Access Control
- Logging
- Email
- Security
- Remote access

Consistency

Large IT systems often grow organically, leading to a collection of machines each with its own OS configuration and application configuration. In a system containing multiple machines with multiple configurations, administration time is exponentially increased.

An infrastructure based on a single configuration is more easily managed. Each machine becomes an interchangeable building block, which can be replaced at any time by an identically configured system, in the event of planned or unplanned downtime. Patches and upgrades need only be verified and tested once for the entire system. Disaster recovery time is improved, and capacity expansion simplified.

Provision of system-wide identity and home directories offers the users a consistent view of the infrastructure.

Naming

Centralised name services, such as DNS and LDAP, are important keystones in producing coherent infrastructure.

DNS should be used to remove the tie between services and physical servers, by using service oriented CNAMEs rather than IP addresses as the primary identifier of the service location.

Services accessed by a DNS name are easily moved between servers. By simply changing the reference in DNS, the change is propagated throughout the entire infrastructure. This allows for improved response time to increased capacity demands, and unplanned downtime, as well as reducing actual outage time during planned downtime.

LDAP should be configured to provide consistent UIDs and GIDs across the entire

infrastructure, thus allowing for coherent access control and improving the ease of auditing of system usage.

Both DNS and LDAP services should be configured across multiple physical servers replicating data between these points, thus reducing risk of single points of failure.

Authentication

It is vital that every action can be traced back to the originator. Authentication provides system users with an identity in the system. Without unique identities it is impossible to implement any form of access control.

A networked authentication system is provided by LDAP, often in association with Kerberos, which offers strong encryption of authentication tokens, thus allowing for each user to hold a single identity across the system.

The superuser should be localised to each machine, using a separate authentication token per system, which should not be distributed by the networked authentication system. A separate network superuser account should be created to provide root access.

Direct root logins are less easily traced to a specific user and thus break the usual auditing processes, however, as they are occasionally necessary, shell command histories should be logged to `/var/tmp/root.date-PID`

Access Control

Authentication and access control are tightly linked. It is impossible to maintain coherent access control without unique identities for each system user, which are provided by the authentication system.

Access should be controlled at both the system and application level. This can be achieved using netgroups which may be distributed through the system using LDAP.

For tasks which require superuser permissions, access should be granted through tools such as `sudo` and `SMITTY`. Access to these tools should be limited using netgroups and the AIX Role system, which provide logs of accesses and commands run and thus an audit trail for the use of these powers.

All applications should be owned by a group specific to the application. Administrators of the application will need to be members of this group.

Further control can be implemented using TCP Wrappers, which, together with netgroups can be used to restrict access not only to specific users, but also such that those users may make connections from specific hosts.

Logging

All systems logs should be collated to a central log host through the SYSLOG facility. Where possible, application logging also should be configured to use SYSLOG, either directly configured in syslog or by using named pipes.

Rolling local logs may be kept for 7 days, with automated housekeeping.

The log host provides a central location for the whole system audit trail, which can be easily accessed, archived and backed up. System and service logs remain accessible regardless of whether the host machine is running, and can be maintained across service relocations and host re-installations, etc.

Without accurate time, logs become confused and less useful. As such, all systems should be configured to use NTP, a standard system to provide synchronised time across networks. The time source may either be provided by a local radio clock, or using the various atomic clock backed time sources provided on the Internet.

Email

SMTP should be configured to allow system monitoring logs to be emailed to administrators, however, delivery should not be permitted locally. All email should be relayed to the central corporate mail host and delivered to mail boxes named for the application.

As with centralised logging, this allows for monitoring and auditing to be removed from individual machines and collated in a central location for archiving.

Security

No system shall act as a router. Systems which act as routers are not under the control of the network and can provide diverse routes.

No system shall expose ICMP redirection. While ICMP is an important protocol for proper communication between servers and/or networks, ICMP packets, including ICMP redirect, are extremely easy to fake. ICMP redirects are used by routers to specify better routing paths out of one network. With spoofed ICMP redirect packets, an attacker can then alter a host's routing tables and divert traffic towards external hosts on a path of their choice; the new path is kept active by the router for 10 minutes.

XDMCP access should be disabled and access to X Windows restricted to SSH TCP/IP port forwarded connections. XDMCP is an inherently insecure protocol, which can be used to get around TCP wrapper and netgroup restrictions for console logins.

Remote Access

Remote access to systems should be offered using both SSH and telnet. This allows remote access to all systems regardless of the client system.

Remote access to the superuser should only ever be allowed through SSH, since telnet allows clear text passwords to be sent across the network.

Public key authentication is an alternative method of authentication to a login server, which does not require a password. A pair of keys are generated for each host. One is Public and allowed to be known by anyone, the other is Private and is only known by the host. The Private key is able to generate Signatures. A Signature can only be created by

the holder of the Private key; however, these Signatures can be verified by any Public key holder.

This mechanism should be used to allow a management system to execute commands across the entire infrastructure from a single point.

To configure this, a key pair should be generated for the management user on the management console. The Public portion of this key pair should then be distributed to all systems. This should be placed in the home directory of the user to be accessed in the `.ssh/authorized_keys` file, for example, this could be placed in the automounted home directory of the network superuser account.

Control of which processes may run on which systems should be added using netgroups.

Operating Systems

All Operating System installs should be identical. This can best be achieved using an automated installation system to provision systems, as machine built systems can be guaranteed to be identical unlike hand built systems.

The OS build should include the Trusted Computing Base (TCB). This is a set of files that must be trusted if the rest of the system is to have security and integrity. This includes items such as the OS kernel, all login handling programs, and the authentication programs.

Several commands are also provided to help ensure files continue to remain trusted after installation and to monitor changes made to the system.

Every system that requires a non-standard build, for whatever reason, must be individually accepted via the change control process.

Applications

All effort should be made to ensure that applications can be run without change from any node, to ensure that systems remain interchangeable. Which applications run on which systems can be controlled through the use of SRV records in DNS and by restricting access to applications by netgroup and TCP wrappers.

Applications should be installed against the standard base image and the resulting system compared to the base to determine what changes are made during the installation. These changes should then be isolated to a piece of storage, and the application run from there.

A specific user and group should be created for every application. Applications should, as far as is possible, be run as the application user and not the root user.

To administer an application a user must be a member of the application group.

Storage

It matters that, on login, a user's home directory is available, however, the physical location should not be not important. This can be achieved using NFS and automounts.

Automount maps should be distributed through the entire network using LDAP, thus allowing a user's home directory to follow them regardless of the system they are logged into.

By classifying data as volatile and non-volatile it is possible to make decisions about the nature of the storage it will require. For example, a trusted gold image of an Operating System is non-volatile and should be located on a piece of storage that is generally read-only, to reduce the risk of changes, however, a user's home directory is volatile and requires both read and write access.

Applications will have non-volatile components, e.g. the application binary, and volatile components, e.g. the application data, and storage choices should be made to reflect this. Again, NFS and automounting, combined with netgroups, should be used to allow any host to run the application, thus furthering the de-coupling of services and applications from physical hosts.

Access to NFS exports should be restricted by netgroups and TCP wrappers, thus ensuring not only that exports are only mounted by authorised users, but also only on authorised hosts.

Maintenance

The system should be maintained and patched in line with vendor recommendations.

Patches, whether they are for the Operating System or applications, are system changes and as such must be dealt with through the change control process, however, it is important that maintenance and patching occur in a timely manner.

Maintenance and patches must be applied to all relevant systems within the infrastructure, in order to maintain the interchangeability of hosts.