

Access Control Lists in Linux & Windows

Vasudevan Nagendra & Yaohui Chen

Categorization: Access Control Mechanisms

- **Discretionary Access Control (DAC):** Owner of object specifies who can access object (files/directories)
 - Control access on discretion of owner
 - Access privileges decided when file created
 - Ex: Windows, Linux, Mac, Unix
- **Mandatory Access Control (MAC):** system specifies which subjects(users/processes) can access which objects.
 - Based on security labels mechanism
 - Subjects are given clearance
 - Objects are given security classification
 - Matches clearance of subject with classification of object.
 - Examples: secret, top secret, confidential

Access Control List (ACLs)

- **Filesystem Access Control mechanisms:**
 - ACLs
 - Role Based Access (RBAC) - Can be Implemented as either DAC/MAC
- **ACL:** Fine-grained discretionary access rights given to files & directories.
 - Specifies, which users/processes are granted access to objects.
 - Access rights tied with objects.
- **RBACs:** System access on basis of authorization
 - specific roles are permitted to perform certain operations
 - Access rights not tied to objects
 - Example: Roles created for various job functions.
 - Consider multiuser systems with users of different roles are accessing.

ACLs Continued..

- **Network Access Control Mechanism:**
 - Netfilter
- **Netfilter (NACL):** network traffic filtering framework for Linux
 - Set of hooks in kernel to handle packets.
 - Intercept calls, events or messages
 - Between s/w components of OS or Applications.
 - Registers callbacks with n/w stack, called for every packet.
 - Access Controls / Filtering rules applied here.

Background: 9 bit permission Model

- Every file system is associated with:
 - 3 set of user groups(classes),
 - 3 set of permissions
 - 9 bits are used to determine the characteristics
 - Also called as base/*minimal* ACLs.
- Example: `ls -la file.txt`
`-rwxrw-r-- 1 root cse506 2 Nov 19 05:55 file.txt`
 - **Owner class** with read, write & execution access
 - **Group class** with read & write access
 - **Others class** with read only access.
 - For changing the file permissions we use the `chmod`.

Background: Other Access Control Options

- **Setuid:** Allows subjects to run executable with permission of file owner.
 - When subject doesn't have adequate permission
 - Examples: passwd/gpasswd/sudo/chsh/mount/ping/su/umount
- **Setgid:** Equivalent (as setuid) property for groups.
 - No matter which user starts it, program runs under group ID
 - All files & directories created in the setgid directory, will belong to the group owning the setgid directory.
- **Sticky bit:** Assigned to directories, prevents users from deleting each other's files.
 - Example: /tmp where any user can store files, but only owner of file has rights to modify or delete the file.

UMASK

- Consider default behavior of file and directory creation
 - 666 & 777 respectively.
 - To change this default behavior – use *umask*
- Defines the permissions to be masked while object is created.
- Examples: `umask 002`
 - File creation: $(666 - 002 = 664) = rw- rw- r--$
 - Directory creation: $(777 - 002 = 775) = rwx rwx r-x$

Drawbacks & Limitations of 9 bit permission model

The price of playing tricks with this permission model:

- Setuid-root - Allows even ordinary users to perform administrative tasks.
 - Buggy application easily compromises system
 - Increase complexity of system configurations.
- Limitations of the base/9 bit permission model:
 - No fine grained control access to non-class users

Extended ACLs for finer-grain control

- **Extended ACLs provides:**
 - beyond simple user/group/other ownership.
 - more than 3 base classes
 - contains any number of named user & groups
 - contains mask entry.

Utilities/Library functions:

- **getfacl:** Check the current state of ACL on file/directory.
getfacl test-dir
- **setfacl:** Modify/add ACL to additional user or group.
setfacl -m user:student1:rwx,group:osclass:rwx test-file
- **chacl:** changes the ACL of file or directory
chacl u::rwx,g::r-x,o::r- test-file

Access Control Entries (ACE)

- Set of entries that defines permissions for user or groups

Example of an ACL Entry in Linux system:

Type | TextForm

owner user::rw-

owning group group:rw-

/*Base Class*/

other other::--

named user user::vasu:rwx

named group group:vasu_grp:rwx

/*Extended Class*/

mask mask::rw-

default:user::rwx

default:group::r-x

default:group:vasu_grp2:r-x

/*Default class*/

default:mask::r-x

default:other::---

More details of Extended ACLs

- **Default ACL:**
 - defined for a directory
 - the objects in directory inherits it.
- Extended ACLs contains entries for additional users or groups.
 - What If permissions are not contained with in owning group?
 - Solution: Solved by virtue of Mask entry.
- **Mask Entry:** maximum access rights that can be granted for users and groups.
 - Mask applicable on:
 - Named user,
 - Named group &
 - Owning Group

Extended Attributes (EAs)

- Typically stored in separate datablock, referenced from inodes.
 - Attributes: Defines Properties of files

Examples:

1. For ext4 fs in linux,
 - inode has a field *i_file_acl* (type `ext4_fsblk_t`),
 - *i_file_acl* -> references to filesystem block with EAs stored
2. For Solaris with UFS file system,
 - inode has a field *i_shadow*
 - References to file system block with EAs stored
 - files with same ACL points to same shadow inode.
 - Implementation dependent optimization.

ACL Implementations

- How ACLs passed between user and kernel space?
 - FreeBSD, Solaris, Irix & HP-UX have separate ACL system calls.
 - Linux: Uses Extended Attributes.
 - Huge Performance degrade for file access at first.
 - ACL Caching is provided by some file system.
 - some filesystems limits # of ACEs. (Implementation Dependent)

Access Check Algorithm

- Subject's access request to object –

Step 1: Select ACL entry that closely matches requesting process

- ACL Entries Looked up in following order:
 - owner
 - named users
 - (owning or named) groups
 - Only single entry determines the access.

Step 2: checks if matching entry contains sufficient permissions.

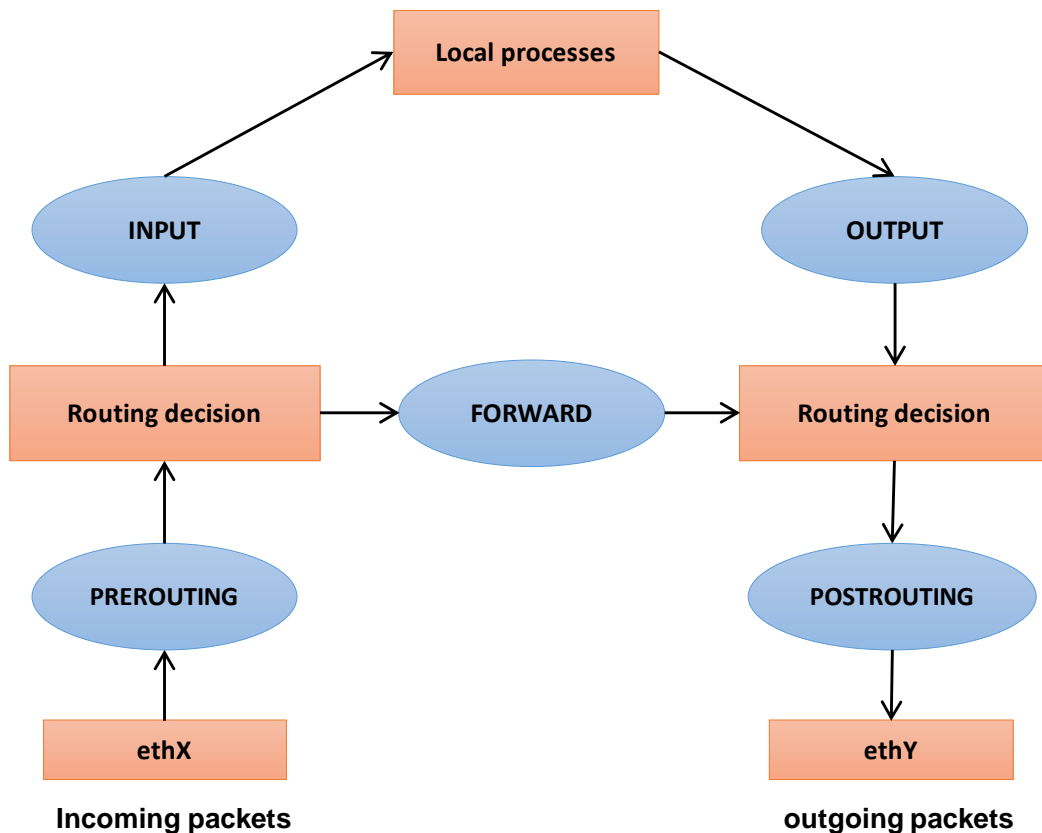
Netfilter - Network ACLs for Linux.

- Packet filtering framework inside Linux kernel.
- Enables following main functions:
 - Packet filtering: ACCEPT/ Drop / Log & other actions
 - NAT: Changing IP/Port (Source & Destination)
 - Mangling: Changing packet contents, ToS, Labeling, etc.,
- Support: Both stateless & stateful packet filtering
 - stateless: No track of the state of packets
 - stateful: Keeps track of packets
- Supports both IPv4 & IPv6

Netfilter Architecture for Network ACLs.

- Hooks & Custom Functions:
 - provided at several points of kernel network stack
 - Hooks: exploited to define custom functions
 - Manipulating Packets headers & data.
 - Actions on packets itself.
- Purpose of Hooks:
 - Debugging
 - Extending functionality
 - Intercepting keyboard/mouse events
 - Monitor system calls to analyze system behavior

Architecture: Netfilter Architecture



Kernel path for Incoming packets

Figure: Netfilter Architecture


- **PREROUTING:** Functions triggered before routing decision
- **POSTROUTING:** triggered after routing decision.
- **FORWARD:** Action on forwarded packets - “ACLs”.
- **INPUT:** Action on Incoming packets
- **OUTPUT:** Actions on Outgoing packets.

Improving the Granularity of Access Control in Windows NT

Yaohui Chen

Access Control In Windows NT

- Access Control Model
 - SubjectObject



*Storage Resource Management (SRM)


Access Control In Windows NT -- Explained

Hello mate, I want to read the password file, here's my **access token**

User SID: Chen
Group SID: Black hats



NO!!! One of the **access control entry** in the **Security Descriptor** says you as user Chen should be **denied to read** this file.




Hold on, let me check..

Password file

Security Descriptor

ACL

Entry1 :
SID: Chen
Type: Access deny
Access Mask: Read



Access Control Entry (ACE)

Type	Inherit Flag	Access Mask	SID
Allow	Inherit_only	Read	Users (Chen)
Deny	No_Propagate	Write	Groups (admin)
Audit	Object_inherit	Execute	
	Directory_inherit	Create.....	

Types of ACEs

- Access-denied
 - Used in an ACL to deny access
- Access-allowed
 - Used in an ACL to allow access
- System-audit
 - Used in an ACL to log attempts to access.

Access Control Entry (ACE)

Type	Inherit Flag	Access Mask	SID
Allow	Inherit_only	Read	Users (Chen)
Deny	No_Propagate	Write	Groups (admin)
Audit	Object_inherit	Execute	
	Directory_inherit	Create.....	

Inherit Flags of ACEs

- **Inherit_Only** (For containers)
 - Only used for inheritance, not apply to this object
- **No_Propagate** (For containers)
 - Only Inherited onto sub-objects, but no further
- **Object_Inherit** (For objects)
 - Inherited onto sub-objects
- **Container_Inherit** (For containers)
 - Inherited onto sub-containers.

Access Control Entry (ACE)

Type	Inherit Flag	Access Mask	SID
Allow	Inherit_only	Read	Users (Chen)
Deny	No_Propagate	Write	Groups (admin)
Audit	Object_inherit	Execute	
	Directory_inherit	Create.....	

Access Rights of ACEs

- Access Mask
 - Jointly-Used with the field **ACE types** and field **SID** when checking
 - **16-bit** long, can be turn on and off
 - Each bit corresponds to a specific access right.
- Example

/Chen's phone book

ACE

SID: Chen

Type: Access-allowed

Access Mask: "Read + Write + Execute"

Inherit Flag: "No_propagate + Directory_Inherit"

Limitations In Windows NT

- Only support 16 different access rights
- Inherit flag does not distinguish between types of objects with different access rights
 - Containers and Non-Containers
- Propagating access control changes to a tree of objects will be ambiguous
 - Propagated change of ACE conflicts with locally added ACE
- No mechanism for restricting the rights of a process other than disabling privileges.

Access Control In Windows 2000

- What's new in ACL of Windows 2000

Type	Inherit Flag	Access Mask	Object Type	Inherited Object Type
Specify this ACE is for ALLOW/DENY purposes	Specify how this ace should be inherited	<p>A mask to specify what kind of access rights this ACE is dealing with.</p> <p>e.g. Read, Write, Execute, Create, etc.</p>	<p>Identifies the type of object or property to which the ACE applies</p> <p><i>*property explained in next slide</i></p>	Controls which types of objects can inherit the ACE

New feature In Windows 2000 – Property Sets

- Property Sets:
 - What is a property?
 - Attributes of an object, e.g Name, age and weight, ect
- Global Unique Identifier(GUID)
 - Microsoft used term for Universally Unique Identifier(UUID)
 - Each Access control target(objects, properties) will be assigned a GUID
- Property Sets are useful:
 - Properties could be grouped into property sets, identified by **ONE GUID**
 - Only ACEs with no GUID or matching GUIDs are evaluated.

New feature In Windows 2000 – Inheritance Control

- Annotation
 - A tag specified by sub-objects dealing with changes of the ACEs pass down from parent-objects
- Dynamic Inheritance Control
 - Without Annotation
- Static Inheritance Control
 - With Annotation
- Comparison
 - Centralized management access control
 - Space and time cost.

New feature In Windows 2000 – Protection from untrusted code

- Motivation
 - Let user decides which subject have access to certain objects
 - Limiting the damage caused by misbehaving subjects
- Introduction of ***restricted context***
 - A restricted context is an access token with a restriction
 - A new field in **Access Token**
 - Apply deny-only attribute to some SIDs that should be restricted
 - When access secure objects, double check SID and the restricted SID list
 - E.g Browser creates restricted thread to show webpage content.

Wrap Up

Windows NT	Windows 2000
Only support 16 different access rights.	Extended the length of mask
Inheritance does not distinguish between types of objects	Object-specific ACE has the field “Inherited Object Type” to help differentiate that
Propagating access control changes to a tree of objects will be ambiguous	Using annotations and static inheritance to correctly propagate changed access control
No mechanism for restricting the rights of a process	Restricted context

Access control in Linux and Win 2000

	Linux	Windows 2000
Access rights	Read, write, execute	Support up to 32 different access rights
Inheritance	Mainly umask, but with setgid the objects inside can inherit	Support explicitly specified inheritance
ACE Types	Only have “allow”	Allow, deny, audit
Access control granularity	User level, controlled by uid	Thread level, controlled by restricted context in access token