

590.05 Lecture 4: OS Security

Xiaowei Yang

Previous lecture

- Security properties
 - Confidentiality
 - Integrity
 - Availability
 - Assurance
 - Authentication
 - Anonymous
- Tools for security
- Security principles

Today

- More security property: Accountability
- OS security
 - File system security
 - Buffer Overflow

Accountability and Freedom

Butler Lampson

Microsoft

September 26, 2005

Real-World Security

- It's about risk, locks, and **deterrence**.
 - Risk management: **cost of security < expected loss**
 - Perfect security costs way too much
 - Locks good enough that bad guys break in rarely
 - Bad guys get caught and punished enough to be **deterring**, so police / courts must be good enough.
 - Can recover from damage at an acceptable cost.
- Internet security similar, but
little accountability
 - Can't identify the bad guys, so can't deter them

How Much Security

- Security is costly—buy only what you need
 - You pay mainly in inconvenience
 - If there's no punishment, you pay a lot
- People *do* behave this way
- We don't *tell* them this—a big mistake
- *The best is the enemy of the good*
 - Perfect security is the worst enemy of real security
- **Feasible security**
 - Costs less than the value it protects
 - Simple enough for users to manage
 - Simple enough for vendors to implement

Causes of Security Problems

- Exploitable bugs
- Bad configuration
 - TCB: Everything that security depends on
Hardware, software, and **configuration**
 - Does formal policy say what I mean?
 - Can I understand it? Can I manage it?
- Why least privilege doesn't work
 - Too complicated, can't manage it


The unavoidable price of reliability is simplicity
—Hoare

Defensive strategies

- **Locks:** Control the bad guys
 - Coarse: Isolate—keep everybody out
 - Medium: Exclude—keep the bad guys out
 - Fine: Restrict—Keep them from doing damage
Recover—Undo the damage
- **Deterrence:** Catch bad guys, punish them
 - Auditing, police, courts or other penalties


The Access Control Model

1. **Isolation Boundary** to prevent attacks outside access-controlled channels
2. **Access Control** for channel traffic
3. **Policy management**




Isolation

- I am isolated if anything that goes wrong is my fault
 - Actually, my program's fault




- Attacks on:
 - Program
 - Isolation
 - Policy



Access Control Mechanisms: The Gold Standard

- **Authenticate** principals: Who made a request
 - Mainly people, but also channels, servers, programs (encryption implements channels, so key is a principal)
- **Authorize** access: Who is trusted with a resource
 - Group principals or resources, to simplify management
 - Can define by a property, e.g. “type-safe” or “safe for scripting”
- **Audit**: Who did what when?
 - *Lock* = *Authenticate* + *Authorize*
 - *Deter* = *Authenticate* + *Audit*



Making Isolation Work

- Isolation is imperfect: Can't get rid of bugs
 - TCB = 10-50 M lines of code
 - Customers want features more than correctness
- Instead, don't tickle them.
- How? Reject bad inputs
 - Code: don't run or restrict severely
 - Communication: reject or restrict severely
 - Especially web sites
 - Data: don't send; don't accept if complex

Accountability

- Can't identify bad guys, so can't deter them
- Fix? End nodes enforce accountability
 - Refuse messages that aren't accountable enough
 - or strongly isolate those messages
 - Senders are accountable if you can punish them
- *All trust is local*
- Need an ecosystem for
 - Senders becoming accountable
 - Receivers demanding accountability
 - Third party intermediaries
- To stop DDOS attacks, ISPs must play

Enforcing Accountability

- Not being accountable enough means end nodes will reject inputs
 - Application execution is restricted or prohibited
 - Communication is restricted or prohibited
 - Information is not shared or accepted
 - Access to devices or networks is restricted or prohibited

For Accountability To Work

- Senders must be able to make themselves accountable
 - This means pledging something of value
 - Friendship
 - Reputation
 - Money
 - ...
- Receivers must be able to check accountability
 - Specify what is accountable enough
 - Verify sender's evidence of accountability

Accountability vs. Access Control

- “In principle” there is no difference but
- Accountability is about punishment, not locks
 - Hence audit is critical
- Accountability is very coarse-grained

The Accountability Ecosystem

- Identity, reputation, and indirection services
- Mechanisms to establish trust relationships
 - Person to person and person to organization
- A flexible, simple user model for identity
- Stronger user authentication
 - Smart card, cell phone, biometrics
- Application identity: signing, reputation


Accountable Internet Access

- Just enough to block DDoS attacks
- Need ISPs to play. Why should they?
 - Servers demand it; clients don't get locked out
 - Regulation?
- A server asks its ISP to block some IP addresses
- ISPs propagate such requests to peers or clients
 - Probably must be based on IP address
 - Perhaps some signing scheme to traverse unreliable intermediaries?
- High priority packets can get through

Accountability vs. Freedom

- Partition world into two parts:
 - Green Safer/accountable
 - Red Less safe/unaccountable
- Two aspects, mostly orthogonal
 - User Experience
 - Isolation mechanism
 - Separate hardware with air gap
 - VM
 - Process isolation


Without R|G: Today



Entities


- Programs
- Network hosts
- Administrators

With R|G



Must Get Configuration Right

- Keep valuable stuff out of red
- Keep hostile agents out of green



Why R|G?

- Problems:
 - Any OS will always be exploitable
 - The richer the OS, the more bugs
 - Need internet access to get work done, have fun
 - The internet is full of bad guys
- Solution: Isolated work environments:
 - **Green**: important assets, only talk to good guys
 - **Don't tickle the bugs, by restricting inputs**
 - **Red**: less important assets, talk to anybody
 - **Blow away broken systems**
- Good guys: more trustworthy / accountable
 - Bad guys: less trustworthy or less accountable

Configuring Green

- Green = locked down = only whitelist inputs
- Requires professional management
 - Few users can make these decisions
 - Avoid “click OK to proceed”
- To escape, use Red
 - Today almost all machines are Red

R|G User Model Dilemma

- People don't want complete isolation
 - They want to:
 - Cut/paste, drag/drop
 - Share parts of the file system
 - Share the screen
 - Administer one machine, not multiple
 - ...
- But more integration can weaken isolation
 - Add bugs
 - Compromise security


Data Transfer

- Mediates data transfer between machines
 - Drag / drop, Cut / paste, Shared folders
- Problems
 - Red → Green : Malware entering
 - Green → Red : Information leaking
- Possible policy
 - Allowed transfers (configurable). Examples:
 - No transfer of “.exe” from R to G
 - Only transfer ASCII text from R to G
 - Non-spoofable user intent; warning dialogs
 - Auditing
 - Synchronous virus checker; third party hooks, ...

Where Should Email/IM Run?

- As productivity applications, they must be well integrated in the work environment (green)
- Threats—A tunnel from the bad guys
 - Executable attachments
 - Exploits of complicated data formats
- Choices
 - Run two copies, one in Green and one in Red
 - Run in Green and mitigate threats
 - Green platform does not execute arbitrary programs
 - Green apps are conservative in the file formats they accept
 - Route messages to appropriate machine

R|G and Enterprise Networks



- Red and green networks are defined as today:
 - IPSEC
 - Guest firewall
 - Proxy settings
 - ...
- The VMM can act as a router
 - E.g. red only talks to the proxy

Summary

- Security is about risk management
 - Cost of security < expected loss
- Security relies on deterrence more than locks
 - Deterrence requires the threat of punishment
 - This requires accountability
- Accountability needs an ecosystem
 - Senders becoming accountable
 - Receivers verifying accountability
- Accountability limits freedom
 - Beat this by partitioning: red | green
 - Don't tickle bugs in green, dispose of red

Today

- More security property: Accountability
- OS security
 - Buffer Overflow

Buffer Overflow Attacks

File system Security

General Principles

- Files and folders are managed by the operating system
- Applications, including shells, access files through an API
- Access control entry (ACE)
 - Allow/deny a certain type of access to a file/folder by user/group
- Access control list (ACL)
 - Collection of ACEs for a file/folder
- A file handle provides an opaque identifier for a file/folder
- File operations
 - Open file: returns file handle
 - Read/write/execute file
 - Close file: invalidates file handle
- Hierarchical file organization
 - Tree (Windows)
- DAG (Linux)

Discretionary Access Control (DAC)

- Users can protect what they own
 - The owner may grant access to others
 - The owner may define the type of access (read/write/execute) given to others
- DAC is the standard model used in operating systems
- Mandatory Access Control (MAC)
 - Alternative model not covered in this lecture
 - Multiple levels of security for users and documents
 - Read down and write up principles

Closed vs. Open Policy

Closed policy

- Also called “default secure”
- Give Tom read access to “foo”
- Give Bob r/w access to “bar”
- Tom: I would like to read “foo”
 - Access allowed
- Tom: I would like to read “bar”
 - Access denied

Open Policy

- Deny Tom read access to “foo”
- Deny Bob r/w access to “bar”
- Tom: I would like to read “foo”
 - Access denied
- Tom: I would like to read “bar”
 - Access allowed

Closed Policy with Negative Authorizations and Deny Priority

- Give Tom r/w access to “bar”
- Deny Tom write access to “bar”
- Tom: I would like to read “bar”
 - Access allowed
- Tom: I would like to write “bar”
 - Access denied
- Policy is used by Windows to manage access control to the file system

Access Control Entries and Lists

- An Access Control List (ACL) for a resource (e.g., a file or folder) is a sorted list of zero or more Access Control Entries (ACEs)
- An ACE refers specifies that a certain set of accesses (e.g., read, execute and write) to the resources is allowed or denied for a user or group
- Examples of ACEs for folder “Bob’s CS167 Grades”
 - Bob; Read; Allow
 - TAs; Read; Allow
 - TWD; Read, Write; Allow
 - Bob; Write; Deny
 - TAs; Write; Allow

Linux vs. Windows

- Linux
 - Allow-only ACEs
 - Access to file depends on ACL of file and of all its ancestor folders
 - Start at root of file system
 - Traverse path of folders
 - Each folder must have execute (cd) permission
 - Different paths to same file not equivalent
 - File's ACL must allow requested access
- Windows
 - Allow and deny ACEs
 - By default, deny ACEs precede allow ones
 - Access to file depends only on file's ACL
 - ACLs of ancestors ignored when access is requested

Linux File Access Control

- File Access Control for:
 - Files
 - Directories
 - Therefore...
 - `\dev\` : *devices*
 - `\mnt\` : *mounted file systems*
 - What else? *Sockets, pipes, symbolic links...*

Linux File System

- Tree of directories (folders)
- Each directory has links to zero or more files or directories
- Hard link
 - From a directory to a file
 - The same file can have hard links from multiple directories, each with its own filename, but all sharing owner, group, and permissions
 - File deleted when no more hard links to it
- Symbolic link (symlink)
 - From a directory to a target file or directory
 - Stores path to target, which is traversed for each access
 - The same file or directory can have multiple symlinks to it
 - Removal of symlink does not affect target
 - Removal of target invalidates (but not removes) symlinks to it
 - Analogue of Windows shortcut or Mac OS alias

Unix Permissions

- Standard for all UNIXes
- Every file is owned by a user and has an associated group
- Permissions often displayed in compact 10-character notation
- To see permissions, use `ls -l`

```
jk@sphere:~/test$ ls -l
total 0
-rw-r----- 1 jk ugrad 0 2005-10-13 07:18 file1
-rwxrwxrwx 1 jk ugrad 0 2005-10-13 07:18 file2
```

Permissions Examples (Regular Files)

-rw-r—r--	read/write for owner, read-only for everyone else
-rw-r-----	read/write for owner, read-only for group, forbidden to others
-rwx-----	read/write/execute for owner, forbidden to everyone else
-r--r--r--	read-only to everyone, including owner
-rwxrwxrwx	read/write/execute to everyone

Permissions for Directories

- Permissions bits interpreted differently for directories
- *Read* bit allows listing names of files in directory, but not their properties like size and permissions
- *Write* bit allows creating and deleting files within the directory
- *Execute* bit allows entering the directory and getting properties of files in the directory
- Lines for directories in `ls -l` output begin with `d`, as below:

```
jk@sphere:~/test$ ls -l
Total 4
drwxr-xr-x  2 jk ugrad 4096 2005-10-13 07:37 dir1
-rw-r--r--  1 jk ugrad     0 2005-10-13 07:18 file1
```

Permissions Examples (Directories)

drwxr-xr-x	all can enter and list the directory, only owner can add/delete files
drwxrwx---	full access to owner and group, forbidden to others
drwx--x---	full access to owner, group can access known filenames in directory, forbidden to others
-rwxrwxrwx	full access to everyone

File Sharing Challenge

- Creating and modifying groups requires root
- Given a directory with permissions `drwx-----x` and a file in it
 - Give permission to write the file to user1, user2, user3, ... without creating a new group
 - Selectively revoke a user
- Solution 1
 - Give file write permission for everyone
 - Create different random hard links: user1-23421, user2-56784, ...
- Problem! Selectively removing access: hard link can be copied
- Solution 2
 - Create random symbolic links
- Problem! Symbolic link tells where it points

Working Graphically with Permissions

- Several Linux GUIs exist for displaying and changing permissions
- In KDE's file manager Konqueror, right-click on a file and choose Properties, and click on the Permissions tab:
- Changes can be made here (more about changes later)



Special Permission Bits

- Three other permission bits exist
 - Set-user-ID (“suid” or “setuid”) bit
 - Set-group-ID (“sgid” or “setgid”) bit
 - Sticky bit

Set-user-ID

- Set-user-ID (“suid” or “setuid”) bit
 - On executable files, causes the program to run as file owner regardless of who runs it
 - Ignored for everything else
 - In 10-character display, replaces the 4th character (x or -) with s (or S if not also executable)
 - rwsr-xr-x: setuid, executable by all
 - rwxr-xr-x: executable by all, but not setuid
 - rwSr--r--: setuid, but not executable - not useful

Set-group-ID

- Set-group-ID (“sgid” or “setgid”) bit
 - On executable files, causes the program to run with the file’s group, regardless of whether the user who runs it is in that group
 - On directories, causes files created within the directory to have the same group as the directory, useful for directories shared by multiple users with different default groups
 - Ignored for everything else
 - In 10-character display, replaces 7th character (x or -) with s (or S if not also executable)

-rwxr-sr-X: setgid file, executable by all

drwxrwsr-x: setgid directory; files within will have group of directory

-rw-r-Sr--: setgid file, but not executable - not useful

Sticky Bit

- On directories, prevents users from deleting or renaming files they do not own
- Ignored for everything else
- In 10-character display, replaces 10th character (x or -) with t (or T if not also executable)

`drwxrwxrwt`: sticky bit set, full access for everyone

`drwxrwx--T`: sticky bit set, full access by user/group

`drwxr--r-T`: sticky, full owner access, others can read (*useless*)

Working Graphically with Special Bits

- Special permission bits can also be displayed and changed through a GUI
- In Konqueror's Permissions window, click Advanced Permissions:
- Changes can be made here (more about changes later)



Root

- “root” account is a super-user account, like Administrator on Windows
- Multiple roots possible
- File permissions do not restrict root
- This is *dangerous*, but necessary, and OK with good practices

Becoming Root

- **su**
 - Changes home directory, PATH, and shell to that of root, but doesn't touch most of environment and doesn't run login scripts
- **su -**
 - Logs in as root just as if root had done so normally
- **sudo <command>**
 - Run just one command as root
- **su [-] <user>**
 - Become another non-root user
 - Root does not require to enter password

Changing Permissions

- Permissions are changed with `chmod` or through a GUI like Konqueror
- Only the file owner or root can change permissions
- If a user owns a file, the user can use `chgrp` to set its group to any group of which the user is a member
- root can change file ownership with `chown` (and can optionally change group in the same command)
- `chown`, `chmod`, and `chgrp` can take the `-R` option to recur through subdirectories

Examples of Changing Permissions

chown -R root dir1	Changes ownership of dir1 and everything within it to root
chmod g+w,o-rwx file1 file2	Adds group write permission to file1 and file2, denying all access to others
chmod -R g=rwX dir1	Adds group read/write permission to dir1 and everything within it, and group execute permission on files or directories where someone has execute permission
chgrp testgrp file1	Sets file1's group to testgrp, if the user is a member of that group
chmod u+s file1	Sets the setuid bit on file1. (Doesn't change execute bit.)

Octal Notation

- Previous slide's syntax is nice for simple cases, but bad for complex changes
 - Alternative is octal notation, i.e., three or four digits from 0 to 7
- Digits from left (most significant) to right(least significant):
[special bits][user bits][group bits][other bits]
- Special bit digit =
$$(4 \text{ if setuid}) + (2 \text{ if setgid}) + (1 \text{ if sticky})$$
- All other digits =
$$(4 \text{ if readable}) + (2 \text{ if writable}) + (1 \text{ if executable})$$

Octal Notation Examples


644 or 0644	read/write for owner, read-only for everyone else
775 or 0775	read/write/execute for owner and group, read/execute for others
640 or 0640	read/write for owner, read-only for group, forbidden to others
2775	same as 775, plus setgid (useful for directories)
777 or 0777	read/write/execute to everyone (<i>dangerous!</i>)
1777	same as 777, plus sticky bit

Limitations of Unix Permissions

- Unix permissions are not perfect
 - Groups are restrictive
 - Limitations on file creation
- Linux optionally uses POSIX ACLs
 - Builds on top of traditional Unix permissions
 - Several users and groups can be named in ACLs, each with different permissions
 - Allows for finer-grained access control
- Each ACL is of the form *type:[name]:rwx*
 - Setuid, setgid, and sticky bits are outside the ACL system

Minimal ACLs

- In a file with minimal ACLs, *name* does not appear, and the ACLs with *type* “user” and “group” correspond to Unix user and group permissions, respectively.
 - When name is omitted from a “user” type ACL entry, it applies to the file owner.



ACL Commands

- ACLs are read with the `getfacl` command and set with the `setfacl` command.
- Changing the ACLs corresponding to Unix permissions shows up in `ls -l` output, and changing the Unix permissions with `chmod` changes those ACLs.
- Example of `getfacl`:

```
jimmy@techhouse:~/test$ ls -l
total 4
drwxr-x---  2 jimmy jimmy 4096 2005-12-02 04:13 dir
jimmy@techhouse:~/test$ getfacl dir
# file: dir
# owner: jimmy
# group: jimmy
user::rwx
group::r-x
other::---
```

More ACL Command Examples


```
jimmy@techhouse:~/test$ setfacl -m group::rwx dir
jimmy@techhouse:~/test$ ls -l
total 4
drwxrwx--- 2 jimmy jimmy 4096 2005-12-02 04:13 dir
```

```
jimmy@techhouse:~/test$ chmod 755 dir
jimmy@techhouse:~/test$ getfacl dir
# file: dir
# owner: jimmy
# group: jimmy
user::rwx
group::r-x
other::r-x
```

Extended ACLs

- ACLs that say more than Unix permissions are extended ACLs
 - Specific users and groups can be named and given permissions via ACLs, which fall under the group class (even for for ACLs naming users and not groups)
- With extended ACLs, mapping to and from Unix permissions is a bit complicated.
- User and other classes map directly to the corresponding Unix permission bits
- Group class contains named users and groups as well as owning group permissions. How to map?

Mask-type ACLs



- Unix group permissions now map to an ACL of *type* “mask”, which is an upper bound on permissions for all group class ACLs.
- All group class ACLs are logically and-ed with the mask before taking effect
 - $\text{rw--} \& \text{r-x--} = \text{r----x--}$
- The ACL of *type* “group” with no *name* still refers to the Unix owning group
- Mask ACLs are created automatically with the necessary bits such that they do not restrict the other ACLs at all, but this can be changed

Extended ACL Example

```
jimmy@techhouse:~/test$ ls -l
total 4
drwxr-xr-x  2 jimmy jimmy 4096 2005-12-02 04:13 dir
jimmy@techhouse:~/test$ setfacl -m user:joe:rwx dir
jimmy@techhouse:~/test$ getfacl dir
# file: dir
# owner: jimmy
# group: jimmy
user::rwx
user:joe:rwx
group::r-x
mask::rwx
other::r-x
```

```
jimmy@techhouse:~/test$ ls -l
total 8
drwxrwxr-x+ 2 jimmy jimmy 4096 2005-12-02 04:13 dir
```

Extended ACL Example Explained

- The preceding slide grants the named user `joe` read, write, and execute access to `dir`.
 - `dir` now has extended rather than minimal ACLs.
- The mask is set to `rwx`, the union of the two group class ACLs (named user `joe` and the owning group).
- In `ls -l` output, the group permission bits show the mask, not the owning group ACL
 - Effective owning group permissions are the logical and of the owning group ACL and the mask, which still equals `r-x`.
 - This could reduce the effective owning group permissions if the mask is changed to be more restrictive.
- The `+` in the `ls -l` output after the permission bits indicates that there are extended ACLs, which can be viewed with `getfacl`.

Default ACLs

- The kind of ACLs we've mentioned so far are access ACLs.
- A directory can have an additional set of ACLs, called default ACLs, which are inherited by files and subdirectories created within that directory.
 - Subdirectories inherit the parent directory's default ACLs as both their default and their access ACLs.
 - Files inherit the parent directory's default ACLs only as their access ACLs, since they have no default ACLs.
- The inherited permissions for the user, group, and other classes are logically and-ed with the traditional Unix permissions specified to the file creation procedure.

Default ACL Example

```
jimmy@techhouse:~/test$ setfacl -d -m group:webmaster:rwx
dir
jimmy@techhouse:~/test$ getfacl dir
# file: dir
# owner: jimmy
# group: jimmy
user::rwx
user:joe:rwx
group::r-x
mask::rwx
other::r-x
default:user::rwx
default:group::r-x
default:group:webmaster:rwx
default:mask::rwx
default:other::r-x
```

Note how this starts the default ACLs out as equal to the existing access ACLs plus the specified changes.

Default ACL Example Continued

```
jimmy@techhouse:~/test$ mkdir dir/subdir
jimmy@techhouse:~/test$ getfacl dir/subdir
# file: dir/subdir
# owner: jimmy
# group: jimmy
user::rwx
group::r-x
group:webmaster:rwx
mask::rwx
other::r-x
default:user::rwx
default:group::r-x
default:group:webmaster:rwx
default:mask::rwx
default:other::r-x
```


The default ACLs from the parent directory are both the access and default ACLs for this directory. Group webmaster has full access.

Default ACL Example Continued

```
jimmy@techhouse:~/test$ touch dir/file
jimmy@techhouse:~/test$ ls -l dir/file
-rw-rw-r-- 1 jimmy jimmy 0 2005-12-02 11:36 dir/file
jimmy@techhouse:~/test$ getfacl dir/file
# file: dir/file
# owner: jimmy
# group: jimmy
user::rw-
group::r-x                                #effective:r--
group:webmaster:rwx                         #effective:rwx-
mask::rw-
other::r--
```

The default ACLs from the parent directory are the basis for the access ACLs on this file, but since touch creates files without any execute bit set, the user and other classes, and the group class as well via the mask ACL, have their execute bits removed to match.

NTFS Permissions




Basic NTFS Permissions

NTFS Permission	Folders	Files
Read	Open files and subfolders	Open files
List Folder Contents	List contents of folder, traverse folder to open subfolders	Not applicable
Read and Execute	Not applicable	Open files, execute programs
Write	Create subfolders and add files	Modify files
Modify	All the above + delete	All the above
Full Control	All the above + change permissions and take ownership, delete subfolders	All the above + change permissions and take ownership


Multiple NTFS permissions

- NTFS permissions are cumulative
- File permissions override folder permissions
- Deny overrides Allow




NTFS: permission inheritance

Permission Inheritance



Block of Inheritance




NTFS File Permissions

- **Explicit:** set by the *owner* for each user/group.
- **Inherited:** dynamically inherited from the explicit permissions of ancestor folders.
- **Effective:** obtained by combining the explicit and inherited permission.

Determining effective permissions:


- By default, a user/group has no privileges.
- Explicit permissions override conflicting inherited permissions.
- Denied permissions override conflicting allowed permissions.




Access Control Algorithm

- The DACL of a file or folder is a sorted list of ACEs
 - Local ACEs precede inherited ACEs
 - ACEs inherited from folder F precede those inherited from parent of F
 - Among those with same source, Deny ACEs precede Allow ACEs
- Algorithm for granting access request (e.g., read and execute):
 - ACEs in the DACL are examined in order
 - Does the ACE refer to the user or a group containing the user?
 - If so, do any of the accesses in the ACE match those of the request?
 - If so, what type of ACE is it?
 - **Deny:** return `ACCESS_DENIED`
 - **Allow:** grant the specified accesses and if there are no remaining accesses to grant, return `ACCESS_ALLOWED`
 - If we reach the end of the DACL and there are remaining requested accesses that have not been granted yet, return `ACCESS_DENIED`

Example




NTFS move vs. copy in same volume




- If you **move** a file or a folder inside the same volume your permission will be the same of the **source** folder
- If you **copy** a file or a folder inside the same volume your permission will be the same of the **destination** folder

NTFS move vs. copy across volumes




- If you **copy** or **move** a file or a folder on different volumes your permission will be the same of the **destination** folder

Setting File Permissions in Win XP




- NTFS permissions in Windows XP Pro are disabled by default.
- Using **Folder Options...** from **Tools** menu inside **Windows Explorer** is possible to activate NTFS permission in windows by unchecking **Use simple file sharing**




Windows Tools


- Access control management tools provide detailed information and controls, across multiple dialogs.
- Focus on single file/folders.
- It is challenging for an inexperienced user, or a system administrator dealing with very large file structures, to gain a global view of permissions within the file system



Treemap Access Control Evaluator (TrACE)





TrACE Highlights



- At a glance, determine the explicit, inherited, and effective permissions of files and folders.
- Understand access control relationships between files and their ancestors
- Quickly evaluate large directory structures and find problem areas
- Layout based on treemaps

What is a Treemap?

- A visualization method to display large hierarchical data structures (trees)
- Layout based on nested rectangles.
- Treemaps were introduced by Ben Shneiderman in “Tree visualization with tree-maps: 2-d space-filling approach”; TOG 1991



Options

Permissions View Options

C:\Documents and S

Browse

Choose User/Group:

Alex Heitzmann

Style Options

Choose Color Scheme:

Full Spectrum

Zoom

In

Out

Legend

NoAccess
ReadOnly
WriteOnly
ReadWrite
ExecuteOnly
ExecuteAndRead
ExecuteAndWrite
ExecuteAndReadAndWrite

Inherited Permissions


Explicit Permissions

- rtWorkDocs
- TreeMapTest
 - Colors
 - Patterns
 - Shapes
 - Circle
 - Rectangle
 - Square
 - Triangle
 - Desktop.ini

C:\Documents and Settings\Alex Heitzmann\My Documents\My Projects\TreeMapTest\Shapes\Rectangle

Inherited	Privilege	Allow/Deny	Propagation Flags
False	Write	Deny	ContainerInherit, ObjectInherit
True	Write, Synchronize	Allow	ContainerInherit, ObjectInherit
True	Read, Synchronize	Allow	ContainerInherit, ObjectInherit

C:\Documents and Settings\Alex Heitzmann\My Documents



Options

Permissions View Options

C:\Documents and S [Browse]

Choose User/Group:

Alex Heitzmann []

Style Options

Choose Color Scheme:

Baseline []

Zoom

In [] Out

Legend

NoAccess
ReadOnly
WriteOnly
ReadWrite
ExecuteOnly
ExecuteAndRead
ExecuteAndWrite
ExecuteAndReadAndWrite

■ Inherited Permissions
 ● Explicit Permissions

Tree View


- rtWorkDocs
- TreeMapTest
 - Colors
 - Patterns
 - Shapes
 - Circle
 - Rectangle
 - Square
 - Triangle

D:\Desktop\ini

C:\Documents and Settings\Alex Heitzmann\My Documents\My Projects\TreeMapTest\Shapes\Rectangle

Inherited	Privilege	Allow/Deny	Propagation Flags
False	Write	Deny	ContainerInherit, ObjectInherit
True	Write, Synchronize	Allow	ContainerInherit, ObjectInherit
True	Read, Synchronize	Allow	ContainerInherit, ObjectInherit

C:\Documents and Settings\Alex Heitzmann\My Documents



Acknowledgment

- Much of these POSIX ACL slides are adapted (and some pictures are taken) from Andreas Grünbacher's paper *POSIX Access Control Lists on Linux*, available online at:

<http://www.suse.de/~agruen/acl/linux-acls/>

What is an Exploit?

- An **exploit** is any **input** (i.e., a piece of software, an argument string, or sequence of commands) that takes advantage of a bug, glitch or vulnerability in order to cause an attack
- An **attack** is an unintended or unanticipated behavior that occurs on computer software, hardware, or something electronic and that brings an advantage to the attacker


Buffer Overflow Attack

- One of the most common OS bugs is a **buffer overflow**
 - The developer fails to include code that checks whether an input string fits into its buffer array
 - An input to the running process exceeds the length of the buffer
 - The input string overwrites a portion of the memory of the process
 - Causes the application to behave improperly and unexpectedly
- Effect of a buffer overflow
 - The process can operate on malicious data or execute malicious code passed in by the attacker
 - If the process is executed as root, the malicious code will be executing with root privileges

Address Space

- Every program needs to access memory in order to run
- For simplicity sake, it would be nice to allow each process (i.e., each executing program) to act as if it owns all of memory
- The address space model is used to accomplish this
- Each process can allocate space anywhere it wants in memory
- Most kernels manage each process' allocation of memory through the virtual memory model
- How the memory is managed is irrelevant to the process


Virtual Memory



Mapping virtual addresses to real addresses

Unix Address Space

- **Text**: machine code of the program, compiled from the source code
- **Data**: static program variables initialized in the source code prior to execution
- **BSS** (block started by symbol): static variables that are uninitialized
- **Heap** : data dynamically generated during the execution of a process
- **Stack**: structure that grows downwards and keeps track of the activated method calls, their arguments and local variables



Vulnerabilities and Attack Method

- Vulnerability scenarios
 - The program has **root** privileges (**setuid**) and is launched from a shell
 - The program is part of a web application
- Typical attack method
 1. Find vulnerability
 2. Reverse engineer the program
 3. Build the exploit


Buffer Overflow Attack in a Nutshell

- First described in
 - Aleph One. Smashing The Stack For Fun And Profit. e-zine www.Phrack.org #49, 1996
- The attacker exploits an unchecked buffer to perform a buffer overflow attack
- The ultimate goal for the attacker is getting a shell that allows to execute arbitrary commands with high privileges
- Kinds of buffer overflow attacks:
 - Heap smashing
 - Stack smashing

Buffer Overflow

domain.c

```
Main(int argc, char *argv[ ])
/* get user_input */
{
    char var1[15];
    char command[20];
    strcpy(command, "whois ");
    strcat(command, argv[1]);
    strcpy(var1, argv[1]);
    printf(var1);
    system(command);
}
```




- Retrieves domain registration info
- e.g., domain brown.edu

strcpy() Vulnerability

domain.c

```
Main(int argc, char *argv[])
/*get user_input*/
{
    char var1[15];
    char command[20];
    strcpy(command, "whois ");
    strcat(command, argv[1]);
    strcpy(var1, argv[1]);
    printf(var1);
    system(command);
}
```



- `argv[1]` is the user input
- `strcpy(dest, src)` does not check buffer
- `strcat(d, s)` concatenates strings


strcpy() vs. strncpy()

- Function `strcpy()` copies the string in the second argument into the first argument
 - e.g., `strcpy(dest, src)`
 - If source string > destination string, the overflow characters may occupy the memory space used by other variables
 - The **null character** is appended at the end automatically
- Function `strncpy()` copies the string by specifying the number **n** of characters to copy
 - e.g., `strncpy(dest, src, n); dest[n] = '\0'`
 - If source string is longer than the destination string, the overflow characters are discarded automatically
 - You have to place the **null character** manually

Return Address Smashing

```
void fingerd (...) {  
    char buf[80];  
    ...  
    get(buf);  
    ...  
}
```

- The Unix `fingerd()` system call, which runs as root (it needs to access sensitive files), used to be vulnerable to buffer overflow
- Write malicious code into buffer and overwrite return address to point to the malicious code
- When return address is reached, it will now execute the malicious code with the full rights and privileges of root



Unix Shell Command Substitution

- The Unix shell enables a command argument to be obtained from the standard output of another
- This feature is called **command substitution**
- When parsing command line, the shell replaces the output of a command between back quotes with the output of the command
- Example:
 - File `name.txt` contains string `farasi`
 - The following two commands are equivalent
 - `finger `cat name.txt``
 - `finger farasi`

Shellcode Injection


- An exploit takes control of attacked computer so injects code to “spawn a shell” or “shellcode”
- A shellcode is:
 - Code assembled in the CPU’s native instruction set (e.g. x86 , x86-64, arm, sparc, risc, etc.)
 - Injected as a part of the buffer that is overflowed.
- We inject the code directly into the buffer that we send for the attack
- A buffer containing shellcode is a “payload”

Buffer Overflow Mitigation

- We know how a buffer overflow happens, but **why** does it happen?
- This problem could not occur in Java; it is a C problem
 - In Java, objects are allocated dynamically on the heap (except ints, etc.)
 - Also cannot do pointer arithmetic in Java
 - In C, however, you can declare things directly on the stack
- One solution is to make the buffer dynamically allocated
- Another (OS) problem is that **fingerd** had to run as root
 - Just get rid of **fingerd**'s need for root access (solution eventually used)
 - The program needed access to a file that had sensitive information in it
 - A new world-readable file was created with the information required by **fingerd**

Stack-based buffer overflow detection using a random canary

Normal (safe) stack configuration:



Buffer overflow attack attempt:



- The canary is placed in the stack prior to the return address, so that any attempt to over-write the return address also over-writes the canary.