FreeBSD Advanced Security Features

Robert N. M. Watson

Security Research
Computer Laboratory
University of Cambridge

14 September, 2007

Introduction

- Welcome!
  - Introduction to some of the advanced security features in the FreeBSD operating system
- Background
  - Introduce a series of access control and audit security features used to manage local security
  - Features appeared between FreeBSD 4.0 and FreeBSD 6.2, and build on the UNIX security model
  - To talk about new security features, we must understand the FreeBSD security architecture

Post-UNIX Security Features

- Securelevels
- Pluggable authentication modules (OpenPAM)
- Crypto library and tools (OpenSSL)
- Resource limits
- Jails, jail securelevels
- GBDE, GELI
- IPFW, PF, IPFilter
- KAME IPSEC, FAST_IPSEC
- **Access control lists (ACLs)**
- **Security event audit**
- **Mandatory access control (MAC)**
- 802.11 security

Brief History of the TrustedBSD Project

- TrustedBSD Project founded in April, 2000
  - Provide trusted OS extensions for FreeBSD
  - DARPA funding began in July, 2001
  - Continuing funding from a variety of government and industry sponsors
  - Work ranges from immediately practical to research
  - Many of these features are production-quality; some are still under development
  - Many features now also present in Mac OS X

FreeBSD Security Architecture

FreeBSD Security Architecture

- FreeBSD security architecture is the UNIX security architecture
  - Entirely trusted monolithic kernel
  - UNIX process model
  - Kernel UIDs/GIDs selected by user-space tools
  - Privileged root user
  - Various forms of access control (permissions, ...)
- New security features extend this model to address its weaknesses

Kernel and User Processes

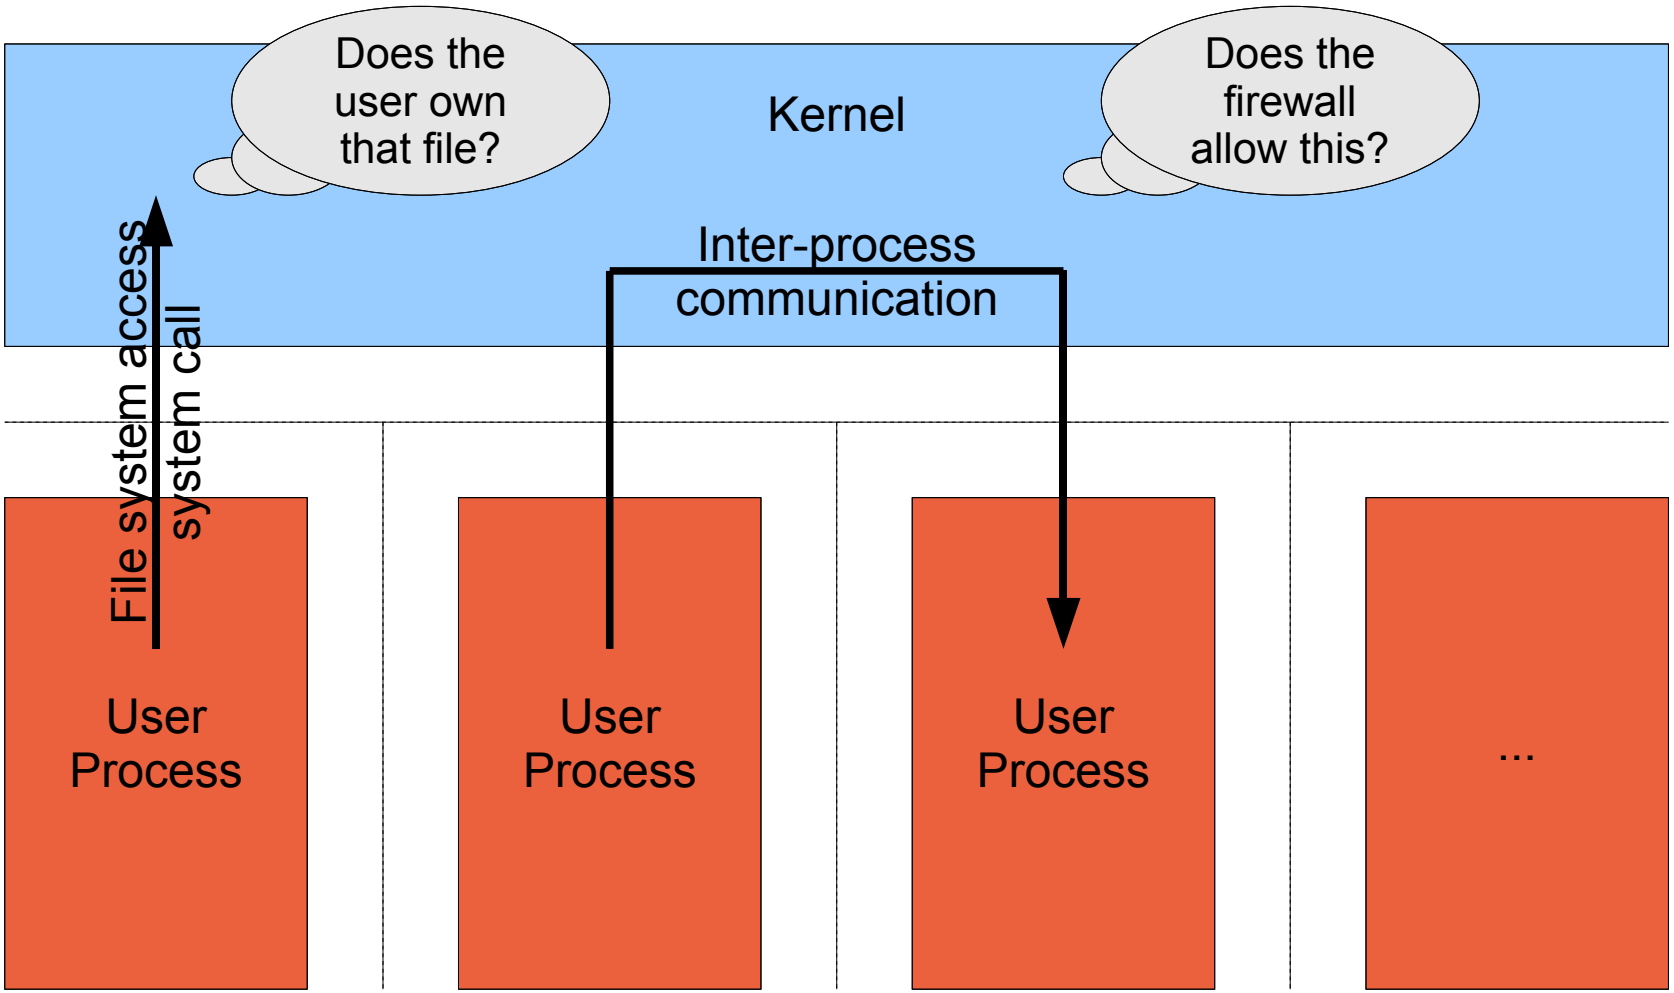

Security Architecture: Kernel Access Control Policy

- Objects owned by a user and group
- Mandatory inter-user protections
  - No inter-user process control (debugging, ...)
  - Only owner of an object can control its protections
  - Special protections for setuid, setgid processes
- Discretionary protections
  - File permissions and ACLs allow owner to grant specific rights to other users and groups
  - Used to protect both system and user data

Security Architecture: The User-space Security Model

- Low-level kernel primitives provide foundation:
  - Process isolation
  - Process credentials and privilege
  - Privilege escalation through setuid/setgid
  - Object ownership and access control
- Kernel knows nothing about password files, login, remote access, home directories, ...
  - All implemented in user-space software layer using kernel primitives

Security Architecture: Authentication and Remote Access

- Kernel provides low-level networking primitives
  - Concepts such as telnet, SSH in user-space
  - Map network I/O into simulated tty input
- User authentication entirely in user-space
  - Pluggable Authentication Modules (PAM) invoked by remote access daemons
  - Kernel UIDs and GIDs set by daemon at login
  - Cache consistency issues between /etc and kernel

Security Architecture: Conclusion

- Layered UNIX security architecture
  - Kernel provides low-level process, process credential, and system services
  - User-space libraries and tools implement users, authentication, remote access
- Security features we discuss extend this model
  - Increased functionality
  - Increased flexibility

Access Control Lists

Access Control Lists (ACLs)

- Extend UNIX file permissions
  - Allow flexible assignment of rights by and for users
  - Required by Orange Book C2, CC CAPP
- ACLs supported in most operating systems
  - POSIX.1e ACLs (Solaris, IRIX, FreeBSD, Linux)
  - NT ACLs (Windows, NFSv4, Mac OS X, ZFS)
- FreeBSD UFS implements POSIX.1e ACLs
  - NT ACL mapping provided by Samba

Configuring UFS2 ACLs

- UFS2 ACLs stored in extended attributes
- Compile UFS ACL support into kernel
  - options UFS_ACL
  - Enabled by default in GENERIC kernel
- ACLs must be administratively enabled for each file system they will be used with
  - tunefs -a enable
  - File system must be unmounted or mounted read-only (best done from single-user mode)

UNIX File Permissions

- Permission mask in file mode
  - Assigns rights to file ***owner***, ***group***, and ***other***
  - Possible rights: ***read***, ***write***, ***execute***
- Certain other special bits in file mode
  - ***setuid***: process takes on UID of file when executing
  - ***setgid***: process takes on GID of file when executing
  - ***sticky*** bit limits unlink rights in directory (/tmp)
- Expressiveness of file permissions very limited
  - Only administrator can modify group membership

POSIX.1e ACLs

- Allow file owner to assign rights for additional users and groups
  - UNIX permissions for ***owner***, ***group***, ***other***
  - POSIX.1e ACL entries assign rights for for ***additional users*** and ***additional groups***
  - Directories have an optional ***default*** ACL
    - Set ACLs on new files or sub-directories in subtree
- POSIX.1e provides a ***mask*** ACL entry to support file mode compatibility for applications

Example ACL

User	rw-
Group	r--
Other	---

User	[owner]	rw-
User	robert	rw-
Group	[owner]	r--
Group	www	r--
Mask		rw-
Other		---

Example ACL

- One file with only a basic ACL (UNIX permission mask)
- One file with an extended ACL
  - One additional user
  - One additional group
  - Mask granting at most read/write to groups and additional users

```
cinnamon% getfacl without_acl
#file:without_acl
#owner:0
#group:0
user::rw-
group::r--
other::---

cinnamon% getfacl with_acl
#file:with_acl
#owner:0
#group:0
user::rw-
user:robert:rw-
group::r
group:www:r--
mask::rw-
other::---
```

ACLs on Newly Created Files and Directories

- Directories have ***access*** and ***default*** ACLs
- If the parent directory has only a basic ACL, UNIX creation rules apply
- If the parent directory has a default ACL, special creation rules apply:
  - Access ACL of child will be default ACL of parent masked by requested creation mode **and** umask
  - New subdirectories inherit parents' default ACL

ACL Tools

- Modifications to existing commands
  - mount(8) – Show when ACLs are enabled
  - ls(1) – Show when an ACL is present with “+”
  - tar(1) – Back up and restore ACLs on files
- New ACL commands
  - getfacl(1) – Retrieve the ACL on one or more files
  - setfacl(1) – Set the ACL on one or more files

ACL Documentation

- Man pages
  - getfacl(1), setfacl(1)
- FreeBSD Handbook chapter
  - “File System Access Control Lists”

ACL Conclusion

- Access control lists add flexibility to UNIX file protection model
  - Users can assign rights to other users and groups
  - Avoid administrative intervention for user collaboration
  - Available in all FreeBSD versions with UFS2
  - Backwards compatible with UNIX permissions
  - Portable to other UNIX operating systems

Security Event Auditing

Security Event Auditing

- Auditing logs system security events
  - Secure, reliable, fine-grained, configurable
- A variety of uses including
  - Post-mortem analysis
  - Intrusion detection
  - Live system monitoring, debugging
- Orange Book, Common Criteria CAPP
- Found in most commercial operating systems

Audit Logs, Records, and Events

- Audit log files are called “trails”, contain records
- Audit records describe individual events
  - ***Attributable*** to authenticated user (or not)
  - ***Selected*** (configured to be audited)
- Most audit events fall into three classes
  - Access control
  - Authentication
  - Security management

What events can be audited?

- Access control examples
  - System calls checking for super user privilege
  - System calls with file system access control checks
  - Login access control decisions
- Authentication, Account Management
  - Password changes, successful authentication, failed authentication, user administration
- Audit administration events

FreeBSD Security Event Auditing Architecture

- Audit records describe security events
- Audit records managed by kernel audit engine
- Audit daemon manages trails, configuration
- Sun's BSM audit trail file format and API
- Administrators control log granularity
- Kernel, privileged processes submit records
- UNIX DAC permissions protect audit log

BSM Audit Record Format

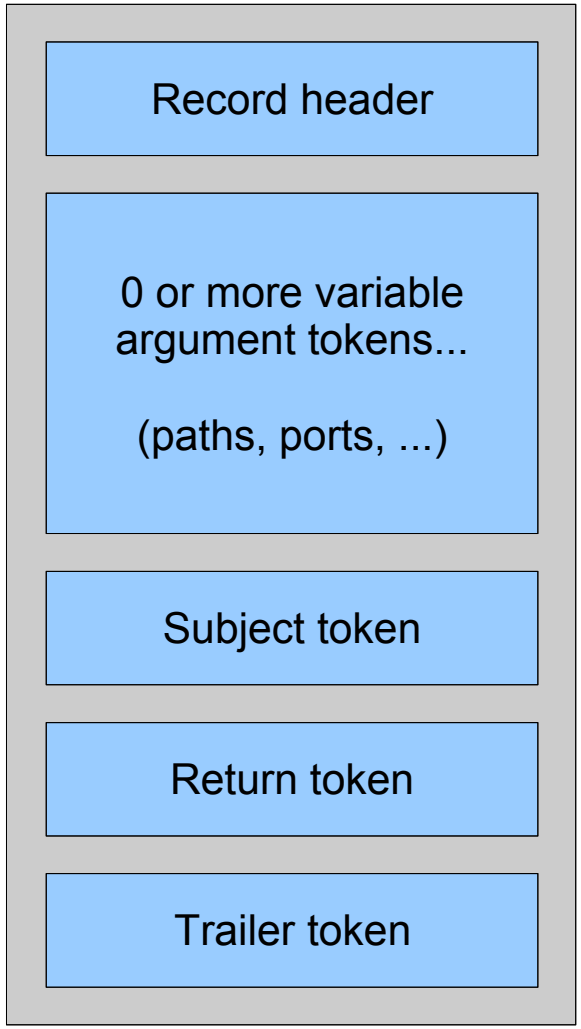


```
<record version="10" event="OpenSSH login"
modifier="0" time="Fri May 18 04:19:56 2007"
msec="274" >
<subject audit-uid="robert" uid="robert" gid="robert"
ruid="robert" rgid="robert" pid="44835" sid="44835"
tid="42666 24.114.252.226" />
<text>successful login robert</text>
<return errval="success" retval="0" />
</record>

<record version="10" event="execve(2)" modifier="0"
time="Fri May 18 07:04:15 2007" msec="933" >
<exec_args><arg>pine</arg></exec_args>
<path>/usr/local/bin/pine</path>
<attribute mode="555" uid="root" gid="wheel"
fsid="90" nodeid="71201" device="336464" />
<subject audit-uid="robert" uid="robert" gid="robert"
ruid="robert" rgid="robert" pid="51933" sid="51927"
tid="49811 24.114.252.226" />
<return errval="success" retval="0" />
</record>
```

Audit Selection

- Potential audit record volume is huge
  - Terabytes/hour on busy, fully audited system
- Two points for audit record selection
  - Audit pre-selection to limit audit records created
  - Audit post-selection, or reduction, to eliminate undesired records later
- FreeBSD support both models
  - Administrator can apply filters to users at login time
  - Administrator can use tools to reduce trails

Audit Trail Reduction

- Reduction selects records from audit trails
  - Long-term archiving
    - E.g., reducing an audit trail to just login/logout data
  - Immediate inspection or analysis
    - E.g., selecting specific records to print with praudit(8)
- auditreduce(8) operates on record streams
  - Input and output may be trail files or pipelines
  - Select by user ID, date or time, event type, or object
  - “grep” for audit records

Audit Pipes

- Historically, audit for post-mortem analysis
- Today, for intrusion detection / monitoring
- Audit pipes provide live record feed
  - Lossy queue
  - Discrete audit records
  - Independent streams

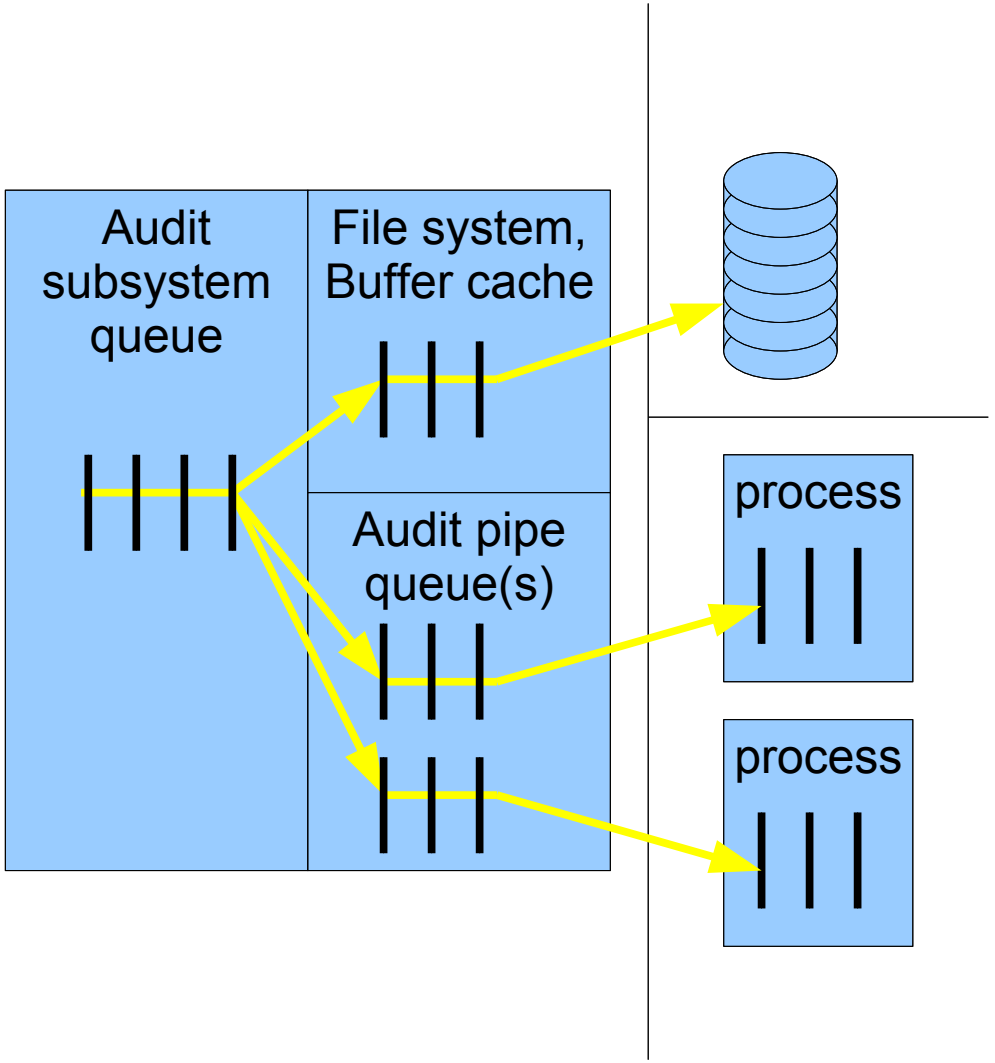

Audit Documentation

- Extensive man pages
  - audit(4), auditpipe(4)
  - audit(8), auditreduce(8), praudit(8), auditd(8)
  - audit.log(5), audit_control(5), audit_user(5), ...
- FreeBSD Handbook chapter
  - “Security Event Auditing”
- TrustedBSD audit implementation paper
  - “The FreeBSD Audit System”

Audit: Conclusion

- Powerful tool for tracking and monitoring use
  - Fine-grained, reliable, and secure logging
- Status
  - Experimental feature in FreeBSD 6.2
  - Production feature in FreeBSD 6.3 and later
- Google Summer of Code 2007
  - Distributed audit daemon
  - Graphical viewing and analysis tool

Mandatory Access Control (MAC)

Mandatory Access Control (MAC)

- Administrator defines mandatory rules under which users and processes interact
  - Contrast with Discretionary Access Control (DAC)
  - File ACLs protect files at ***discretion*** of the owner
- Historically, Multi-Level Security (MLS)
  - Data is labelled with sensitivity levels and compartments to indicate what protection required
- Recently, much more broad definition
  - Mandatory as opposed to a specific policy

TrustedBSD MAC Framework

- Kernel framework allows policy modules to modify kernel access control policy
  - Add new constraints
  - Track use of resources
  - Attach security labels to objects
- Two general common classes of policies
  - Ubiquitous information labelling policies
  - Hardening policies

MAC Policies for FreeBSD

- FreeBSD includes a number of policies
  - Labelled: Biba, MLS, LOMAC, partition
  - Hardening: portacl, seeotheruids, ugidfw
- Open source third party policies
  - Cryptographically signed binaries
  - SEBSD (SELinux FLASK/TE)
  - mac_privs privilege management policy
- Proprietary policies in FreeBSD-based products
- Mac OS X policy modules to be open source

User/Group File System Firewall

- Rule-based file system protection policy
  - ugidfw(8) command manages a rule list similar to that in a network firewall
- Administrator restrict access by user or group
  - Overrides file permissions and ACLs
  - No data or subject labelling required
  - Only allow ***read***, ***execute***, and ***stat*** by ***www*** user on objects owned by ***robert***

```
ugidfw set 100 subject uid www object uid robert mode rxs
.
.
.
```

MAC Documentation

- FreeBSD man pages
  - mac(4), getfmac(8), setfmac(8)
- FreeBSD Handbook chapter
  - “Mandatory Access Control (MAC)”
- TrustedBSD implementation papers
  - “The TrustedBSD MAC Framework: Extensible Access Control for FreeBSD 5.0”
  - “Design and Implementation of the TrustedBSD MAC Framework”

Conclusion

- Introduction to FreeBSD Security Architecture
- Several advanced FreeBSD security features
  - ACLs
  - Audit
  - MAC
- Further information can be found in:
  - The FreeBSD Handbook
  - http://www.TrustedBSD.org