

CNS Lecture 14

Writing secure code
Crypto APIs
Secure applications

594 paper due 12/1/06
Assignment 10 forensics



In the news



- Remember we said to avoid cross-site scripting use `HTMLEntities()` in PHP ... well today, buffer overflows in this function
- In case you hadn't noticed, about 9 out of 10 emails today are spam, sent by botnets around the world
- GNU RADIUS format string vulnerability
- Acer Notebooks ActiveX arbitrary command execution

CNS Lecture 14 - 2



You are here ...



- | Attacks & Defenses | Cryptography | Applied crypto |
|----------------------------|-----------------------|----------------|
| • Risk assessment ✓ | • Random numbers ✓ | • SSH ✓ |
| • Viruses ✓ | • Hash functions ✓ | • PGP ✓ |
| • Unix security ✓ | MDS, SHA, RIPEMD | • S/Mime ✓ |
| • authentication ✓ | • Classical + stego ✓ | • SSL ✓ |
| • Network security ✓ | • Number theory ✓ | • Kerberos ✓ |
| Firewalls, vpn, IPsec, IDS | • Symmetric key ✓ | • IPsec ✓ |
| • Forensics ✓ | DES, Rijndael, RC5 | • Crypto APIs |
| • Secure coding | • Public key ✓ | |
| | RSA, DSA, D-HECC | |

CNS Lecture 14 - 3



Software security

- Software engineering
 - Why do smart people write buggy software?
- Writing secure code
 - Secure design
 - Secure coding & tools
 - Testing for security
 - Defensive programming



CNS Lecture 14 - 4



First computer bug ?

Photo # NH 96566-KN First Computer "Bug", 1945

1942

9/9

0500 Anttan started
1000 - stopped - action ✓ (1.2700 7.000 500 015
13:45:00 AM - MC 2.1300 7.000 500 015
200 PPO - 2.1300 7.000 500 015
2000 2.1300 7.000 500 015
2000 2.1300 7.000 500 015
Relays are on 020 full speed and back
to relay - slow fast.

1100 Started Cassin Tape (Sine chart)
1525 Started Multi-Angle Test


1545 Relay #20 Panel F (MOT) in relay.

1640 First actual case of bug being found
1640 changed stand
1700 stand down.

CNS Lecture 14 - 5



Software bugs



- Flaws in software are expensive -- \$60 billion/yr (NIST)
- Well-known bugs
 - Therac-25 (1985-1987) - radiation therapy machine killed 3 patients
 - Denver baggage handling system - complex (300 computers), late and 50% cost overrun, abandoned
 - Ariane 5 (1996) - loss of guidance system, fit. pt to 16-bit int overflow
 - USS Yorktown (1998) - aircraft carrier dead in the water for hours due to zero divide
 - Mars Climate Orbiter (September 23rd, 1999) - \$125m loss due to failure to convert meters to feet
- Just examples of when things go bad on their own. What if you have an active adversary who is looking for and exploiting bugs? Making an "unlikely event" a certainty in the millions of copies of the software in use around the world.



CNS Lecture 14 - 6



Bug density?

- 10 bugs per 1,000 lines of code?
 - Million lines of code → 10,000 bugs
- Complexity breeds bugs
- Of course all of these bugs aren't exploitable
- But it's a bug hunters paradise
 - The bad guys are winning



Lines of code (millions)	
NT 3.1	6
NT3.5	10
NT 4.	16
Windows 2000	29
Windows XP	40
Windows Vista	50
RedHat 7.1	30
Linux 6.0 kernel	6
Debian 2.2	56
Debian 3.1	213
Mac OS X 10.4	86
OpenSSL	0.2
Apache	0.2

CNS Lecture 14 - 7

Why is software development so hard?


- Complex
 - No two parts the same
 - Growth adds to complexity in non-linear fashion
 - Interaction between data, function, user, devices
- Conformance to human requirements
- Changeable
 - Malleable
 - Easy to change, so let's change it
- Structure is impossible to visualize
- Human intensive
- Market pressures
 - Eliminate all bugs, never ship
 - Ship now, let customers debug



CNS Lecture 14 - 8

Risk Exposure


- Risk Exposure $RE = Prob(Loss) * Size(Loss)$
- "Loss" – financial; reputation; future prospects, ...
- Loss due to unacceptable dependability



CNS Lecture 14 - 9

© Boehm USC-CSE

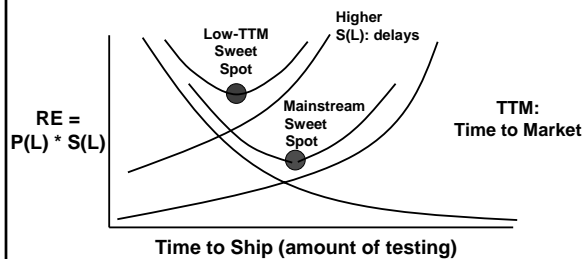
- Loss due to unacceptable dependability
- Loss due to market share erosion



CNS Lecture 14 - 10

© Boehm USC-CSE


Internet Startup



CNS Lecture 14 - 11

© Boehm USC-CSE

Safety-Critical or Secure System



CNS Lecture 14 - 12

© Boehm USC-CSE

Dilemma

Security conflicts with
 -Ease of use
 -Speed
 -Feature richness

Pick any two.

CNS Lecture 14 - 13

Software development

- **You hack up some code**
 - Single developer
 - “Toy” applications
 - Short lifespan
 - Single or few stakeholders
 - Architect = Developer = Manager = Tester = Customer = User
 - One-of-a-kind systems
 - Built from scratch
 - Minimal maintenance

- **Enterprise strength**
 - Teams of developers with multiple roles
 - Complex systems
 - Indefinite lifespan
 - Numerous stakeholders
 - Architect ≠ Developer ≠ Manager ≠ Tester ≠ Customer ≠ User
 - System families
 - Reuse to amortize costs
 - Maintenance accounts for over 60% of overall development costs

CNS Lecture 14 - 14

Software Development Lifecycle Waterfall Model

Lots of methodologies

- clean room
- PSP
- SDM
- pair programming
- agile
- object-oriented, reuse

CNS Lecture 14 - 15

Playing the software development game ...

- **The trusting model**
 - Design/code/test to provide proper function
 - Bugs are random and rare
- **The defensive model**
 - Design/code knowing you have an active adversary
 - Test not only that the right things happen, but also that wrong things don't happen
 - Bugs may be vulnerabilities and exploited

CNS Lecture 14 - 16

Computers at Risk

“The developers of secure software cannot adopt the various probabilistic measure of quality that developers of other software can. For many applications, it is quite reasonable to tolerate a flaw that is rarely exposed and to assume that its having occurred once does not increase the likelihood that it will occur again. It is also reasonable to assume that logically independent failures will be statistically independent and not happen in concert. In contrast, a security vulnerability, once discovered, will be rapidly disseminated among a community of attackers and can be expected to be exploited on a regular basis until it is fixed.”

CNS Lecture 14 - 17

Software engineering for security

- Secure design
- Secure implementation
- Security testing
- Secure deployment
- Root cause analysis for bugs
- Policy and training

Secure by Design
 Secure by Default
 Secure in Deployment

CNS Lecture 14 - 18

Secure design

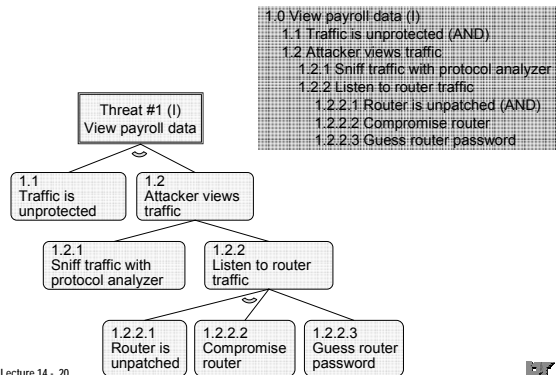
- **Assess the risk**
 - Detail your assets
 - Know the threats and your attackers
 - Mitigate the threats
 - Costs: time & money ... acceptable risk
- **Formal threat modeling**
 - Attack trees
 - Reduce the attack surface
- **Apply secure design principles**
- **Think like a bad guy**

CNS Lecture 14 - 19



Threat Modeling Process

Identify the Threats by Using Attack Trees



CNS Lecture 14 - 20



Design principles

- **Principle of least privilege**
 - Give only those privileges needed to complete a task
- **Principle of fail-safe defaults**
 - Access should be denied unless it is specifically permitted
- **Principle of economy of mechanism**
 - Security mechanisms should be as simple as possible
- **Principle of complete mediation**
 - All accesses to objects must be mediated
- **Principle of open design**
 - Security should not depend on secrecy of design or implementation
- **Principle of separation of privilege**
 - Don't grant permission based on a single condition (su: password+wheel grp)
- **Principle of least common mechanism**
 - Mechanisms used to access resources should not be shared
- **Principle of psychological acceptability**
 - Security mechanisms should not make resource access more difficult

CNS Lecture 14 - 21



How are the bad guys finding bugs?



- **Open source** → **code review**
 - grep for str*()
 - Analyze input and packet parsing
- **Blackbox testing**
 - "random" alteration of input - big username, missing fields ...
 - Cause a crash
 - Analyze dump
 - experiment
 - Disassembly (IDA Pro)
 - Look for str*(), network code
- **Bug announcements, or diff patched vs old images**
 - Attack unpatched versions
- **Motivation?**
 - Attribution
 - Financial gain

CNS Lecture 14 - 22



Secure implementation

- **Trained developers**
- **Language choice** - Java, C++, C#, C, VB, perl, php, ...
- **Trusted APIs (safestring, banned APIs)**
- **Development tools**
 - Source code control
 - Compile/link/debug
 - Compiler warnings
- **Peer code review**
 - Specifically "security reviews"
 - Use checklists
- **Static analysis**
 - Microsoft PREFIX and PRefast
 - RATS, ITS4, Klocwork, Coverity, ESC/Java

```
perl -T
Perl has option for reporting use of
"tainted" variable at runtime
system("mail " . $form_data("email"));
```

CNS Lecture 14 - 23



Code reviews

- **Probably your best return on investment for security**
- **Will increase coding "cost", but worth it for production software**
- **Identify critical components for review**
 - Processes that run with elevated privilege
 - Processes that talk to the net
 - Processes with "user input"
 - Software that parses packets/input
- **Peer review source code and "patches" (person or team)**
- **Have a checklist of things to look for**
 - Use of tainted data, banned functions, false assumptions
 - Failure to check return code, signed/unsigned, toctou
 - Update the list as bugs are found
- **Sign off on review**
- **Engineers need to be trained to recognize flaws/vulnerabilities**
- **What about legacy code? Or outside libraries (e.g. OpenSSL)**

CNS Lecture 14 - 24



Sample code

```

/* really bad telephone number lookup code, expects "name=foo" on stdin
static char cmd[128];
static char format[] = "grep %s phone.list\n";
main()
{
    char buff[256];
    gets(buff);
    sprintf(cmd,format,buff+5);
    printf(cmd); // debug, remove later
    system(cmd);
}

-- more fragments
p = malloc(sizeof(buff));
strcpy(p,userdata);
strncpy(p,userdata,sizeof(p)); // better ?
strncpy(p,userdata,sizeof(buff)); // better?
buff[sizeof(buff)] = 0;

--
char digest[16], tmp[16];
--
MD5_final(digest,&context);
strncpy(tmp,digest,sizeof(digest));

```

CNS Lecture 14 - 25



More broken code snippets

```

char *p = NULL;
if (argc == 2) p = argv[1]; //suggest if (2 == argc) - why
*p = 'A';

while(a=5){
    --
}

int fcn(){
    char *p;
    int i;
    p = malloc(1024);
    --
    return (i+5);
}

```

CNS Lecture 14 - 26



Static analysis

- Automated source code analysis
- Looking for

- Potential NULL pointer dereferences
- Access beyond an allocated area (e.g. array or dynamically allocated buffer); otherwise known as a buffer overflow
- Writes to potentially read-only memory
- Reads of potentially uninitialized objects
- Resource leaks (e.g. memory leaks and file descriptor leaks)
- Use of memory that has already been deallocated
- Out of scope memory usage (e.g. returning the address of an automatic variable from a subroutine)
- Failure to set a return value from a subroutine
- Buffer and array underflows
- if (x=3) ...

- List of warnings, many false positives
- 10x slower than compile

Filename	Report file location path
Filelist	Report path: /usr/src/linux-2.6.10-1.1.10
Files	Files: 1772
Functions	Functions: 2048
Lines	Lines: 30434
Errors and warnings	Errors: 107
Warnings	Warnings: 24
Average	
Lines per function	14.8
Errors per line	0.0035
Warnings per line	0.00078
Errors per function	0.052
Warnings per error	0.15
Lines per warning	1215



CNS Lecture 14 - 27



Coverity's static analysis of Linux kernel

- Growth in linux kernel in drivers
- Complexity grows but not % of security flaws

2001 2.4.1 kernel 1.6 MLOC 1000 flaws
2004 2.6.4 kernel 5 MLOC 900 flaws



CNS Lecture 14 - 28



Coverity's analysis of open source software

Code Base	Lines of Code	Number of Errors	Analysis Time (secs)	Defects/Linearity
Antenna	87,282	106	6	1.207
Avahi	107,828	81	11	0.750
Edklib	1,767,801	143	108	0.241
Firebird	288,701	149	19	0.480
Firefox	803,818	108	24	0.330
FreeBSD	1,587,168	420	107	0.417
Gtk	301,801	113	18	0.362
Gtk+	680,881	142	80	0.202
IceWM	1,804,014	876	112	0.488
kernel	37,247	12	1	0.324
kernel	21,882	28	4	0.123
Linux	3,717,481	1080	204	0.289
Minicom	484,044	216	38	0.446
MySQL	807,488	136	48	0.371
NetBSD	1,152,748	142	16	0.123
OpenLDAP	264,350	148	33	0.487
OpenSSL	166,101	48	19	0.288
OpenSSH	68,810	7	4	0.101
Perl	478,708	89	23	0.188
PHP	458,817	266	88	0.474
Python	876,342	293	38	0.342
Python	89,814	16	4	0.188
Python	258,272	86	18	0.372
Scheme	201,048	278	24	0.486
Scout	83,818	48	4	0.578
SQLite	80,117	21	6	0.278
SQLite	151,841	43	8	0.283
TCL	181,048	69	11	0.372
TeX	303,248	73	88	0.241
TeX	2,014,880	1681	224	0.774
TeX	278,248	347	81	0.482
TeX	178,748	8	4	0.041


- Less than 1 bug per 1,000 lines of code

CNS Lecture 14 - 29



Testing for security

- Regression testing, unit testing, integration testing, system testing
- Most testing is aimed at testing conformance/quality, that code does what it's supposed to do
- Security testing is different, making sure bad things don't happen
 - if you knew to test for it, you wouldn't have coded it wrong...
 - Independent test group




While most functions of bugs are the result of missing functionality, most security bugs are the result of extra functionality.

CNS Lecture 14 - 30



Secure testing

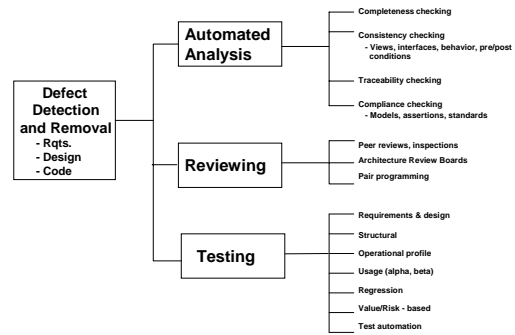
- **Run time tests**
 - Memory leak tester (purify)
 - Fuzz testing, random user/packet input (udpsic)
 - Smarter fuzz testing -- codenomicon
 - Lots of web app testing products
- **Penetration testing**
 - Think/attack like a bad guy
 - Monitor bugtraq/CERT, because the bad guys are pen testing!



CNS Lecture 14 - 31



Software Defect Detection Opportunity Tree



CNS Lecture 14 - 32

© Boehm USC-CSE



Defense in Depth

- Training, code review, testing won't find all the bugs (prevention)
- Buffer overflow defenses
 - Stackguard/propolice
 - Hardware memory protection (RO/NX)
 - Address randomization (text/heap/stack/data) ... see PaX
- Fail gracefully, fault tolerance
- Log events and error conditions
- Watch for crashing daemons/servers



CNS Lecture 14 - 33



Security training

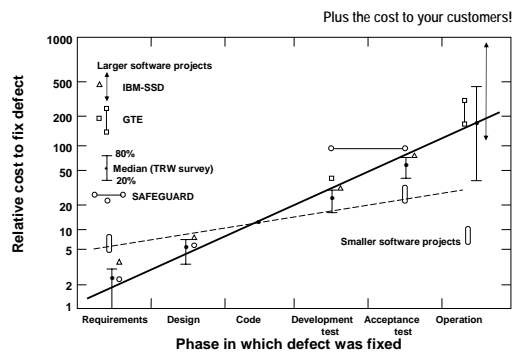
- Educate your designers, developers, testers, and users
- How can you do a risk assessment, if you don't know the risks?
- How can you do code reviews looking for security flaws, if you don't know what security flaws look like?
- Yearly re-train?
- Metrics? Bugs/1000 lines of code



CNS Lecture 14 - 34



Factor-of-100 Growth in Software Cost-to-Fix



CNS Lecture 14 - 35

© Boehm USC-CSE



Fixing bugs

- Root cause analysis
- Each defect found can trigger five analyses:
 - Debugging: eliminating the defect
 - Regression: ensuring that the fix doesn't create new defects
 - Similarity: looking for similar defects elsewhere
 - Insertion: catching future similar defects earlier
 - Prevention: finding ways to avoid such defects
- Example, strcpy() cause for buffer overflows
 - ban strcpy() and provide safe string libs
 - Training
 - Static analysis tools to detect



CNS Lecture 14 - 36



Software Defect Reduction Top-10 List

from CeBASE <http://www.cebase.org>

1. Finding and fixing a software problem after delivery is often 100 times more expensive than finding and fixing it during the requirements and design phase.
2. About 40-50% of the effort on current software projects is spent on avoidable rework.
3. About 80% of the avoidable rework comes from 20% of the defects.
4. About 80% of the defects come from 20% of the modules and about half the modules are defect free.
5. About 90% of the downtime comes from at most 10% of the defects.
6. Peer reviews catch 60% of the defects.
7. Perspective-based reviews catch 35% more defects than non-directed reviews.
8. Disciplined personal practices can reduce defect introduction rates by up to 75%.
9. All other things being equal, it costs 50% more per source instruction to develop high-dependability software products than to develop low-dependability software products. However, the investment is more than worth it if significant operations and maintenance costs are involved.
10. About 40-50% of user programs have nontrivial defects.

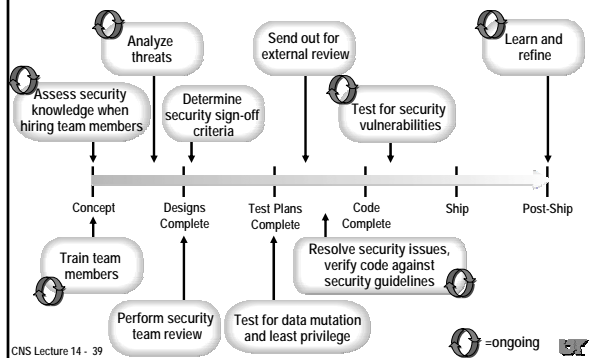
CNS Lecture 14 - 37

Microsoft secure coding

- Security push – audit legacy code for security flaws
- For Vista, Security Development Life Cycle
 - Design – threat modeling, least privilege, reviews
 - Implementation – banned APIs, tools, review, bluehat
 - Deployment – secure by default, root cause analysis of bugs
- Security training – from one course to a dozen
- Address Space Layout Randomization (ASLR)
 - Randomize text/data/heap
 - Stackguard (!GS)
 - Support RO/NX
- TPMs and Bitlocker
- Upper management support

CNS Lecture 14 - 38


Secure Product Development Timeline



CNS Lecture 14 - 39

Design security in from the beginning

- Good design and code reviews are more effective than testing
- Should the software vendor be held accountable for loss due to software flaws?



CNS Lecture 14 - 40

Adding security in

- Our ncp labs (assignments 4, 7, 8)
 - Design problems
 - coding problems
- TCP protocol
 - Assumed “cooperating” processes
 - SYN flooding, malformed packets ...
- NTP protocol
- SNMP protocol (Simple Network Management Protocol)
 - Or Security Not My Problem
 - Clear text read/write group (v1)
 - Added notion of views (ACLs/ least privilege) (v2c)
 - Added authenticity/privacy (v3)

CNS Lecture 14 - 41

SNMP

- Simple Network Management Protocol
- Used to manage network devices (routers, switches, toasters,...)
- Network manager and data base (MIB) and network agents
 - ASN.1/BER data encoding
 - Object ID (OID) (long numeric tag) provides unique variable names (tree)
- Simple datagram protocol (UDP port 161)
 - GET GET-NEXT GET-RESPONSE SET TRAP (port 162)
- Version 1, simple (no) security (community string)
 - SET’s disabled
 - Security for SNMP v1: Security Not My Problem
- Numerous implementation flaws discovered with Finland/Codenomicon smart fuzz tester
 - Buffer overflows community string
 - ASN.1/BER bugs (Tag Length Value TLV)

CNS Lecture 14 - 42

SNMP v3 security (RFC 3414)

Threats

- Impersonation
- Modification
- Replay/re-order/delay
- disclosure

Countermeasures (v3)

- **Timeliness**
 - Loosely synchronized monotonically increasing time indicators
- **Authentication**
 - MD5 HMAC (96 bit)
 - 128 bit (16-byte) key
 - SHA optional (crypto agile)
- **Privacy**
 - DES/CBC (AES & D-H option)
 - 8 byte key + 8 byte IV (from key)
 - Export restrictions ☹
- **Additional control thru router filters**

Key update:

1. administrator picks newkey
2. mgt. station generates random # rand and calculates $\text{delta} = \text{newkey} \oplus \text{hash}(\text{rand}, \text{oldkey})$
3. send client delta, rand
4. client recovers $\text{newkey} = \text{delta} \oplus \text{hash}(\text{rand}, \text{oldkey})$
no Perfect Forward Security ☹

CNS Lecture 14 - 43



Crypto APIs



- functions for developing crypto applications
- provide data privacy and integrity and authenticity
- portable -- architecture and OS
- services
 - establish context (keying, algorithm negotiation)
 - encrypt/decrypt and authenticate
 - certificate mgt. (fetch, CRL)
 - sign/verify
 - encode/decode (PKCS, ASN)
 - protocols (LDAP, SSL/TLS, key update)
- multiple algorithms -- symmetric, asymmetric, hash, random, primality, big numbers, compress, encode

CNS Lecture 14 - 44



Example C APIs

roll your own

- build from algorithms/shareware
- GNU's multiprecision lib
- encoding/portability
- Schneier PC disks

shareware packages

- OpenSSL or Young's libdes/SSLey or RSA's BSAFE
- Blaze's cryptolib
 - DSA, RSA, D-H, ElGamal, Rabin
 - DES/3DES, MD5, SHA
 - bigmath, truerand, primality, quantize
- elliptic curve routines

other: crypto++, JAVA crypto

CNS Lecture 14 - 45



Crypto++

- Dal's C++ class library
- block ciphers (incl. AES)
- hashes/MACs (SHA/MD5/Tiger/RIPEM/Panama)
- RSA/DSA/ECC/D-H
- primes/PKCSs/compress/encode
- benchmarks
- class hierarchy with abstract base classes
- filter/pipeline metaphor

```
class MD5 : public IteratedHash<word32, false, 64>
{ ... }

MD5 md;
SecByteBlock digest(md.DigestSize());

md.Update(testSet[i].input, testSet[i].inputLen);
md.Final(digest);
```

base classes

```
BlockTransformation
BufferedTransformation
BufferedTransformation:Err
Exception
FixedBlockSize
HashModule
PK_Signer:KeyTooShort
PK_Encryptor
PK_FixedLengthCryptoSystem
PK_FixedLengthDecryptor
PK_FixedLengthEncryptor
PK_Precomputation
PK_SignatureSystem
PK_SignatureSystemWithRecovery
PK_Signer
PK_SignerWithRecovery
PK_SimpleKeyAgreementDomain
PK_Verifier
PK_VerifierWithRecovery
PK_WithPrecomputation
RandomAccessStreamCipher
RandomNumberGenerator
StreamCipher
```

CNS Lecture 14 - 46



Java JCE



Java Security

Exportable

- Crippled crypto or signed applications

Provider-based architecture (CSFs)

hashes, block ciphers, PRNG

- DES, AES, Blowfish, SHA, MD5, HMAC

public key (RSA/DSA), SSL, and D-H

certificate mgt (keytool)

- public key generation
- signing (CA)
- signing JAR's

also Cryptix API



```
import java.io.*;
import java.net.*;
import java.security.*;
import javax.crypto.*;

...
KeyGenerator generator =
    KeyGenerator.getInstance("DES");
generator.init(new SecureRandom());
SecretKey key = generator.generateKey();
Cipher encipher =
    Cipher.getInstance("DES/ECB/PKCS5Padding");
Cipher decipher =
    Cipher.getInstance("DES/ECB/PKCS5Padding");
encipher.init(Cipher.ENCRYPT_MODE, key);
decipher.init(Cipher.DECRYPT_MODE, key);
...
ciphertext = encipher.doFinal(cleartxt);
cleartxt = decipher.doFinal(ciphertext);
...
Signature dsa =
    Signature.getInstance("SHAwithDSA");
PrivateKey priv = pair.getPrivate();
dsa.initSign(priv);
dsa.update(data);
byte[] sig = dsa.sign();
```

Cryptographic Service Provider

CNS Lecture 14 - 47



OpenSSL API openssl.org



widely used reference implementation

Our basic C crypto toolkit

- Hashing (MD*, SHA, RIPEMD)
- Random numbers and prime numbers
- Big number library
- Symmetric key crypto (RC4, blowfish, AES, DES)
- Public key crypto (RSA, DSA, D-H, ECC)
- Support for crypto hardware accelerators (engines)
- Crypto-agile wrappers (EVP)

SSL/TLS

- API for creating SSL network connection (asnmt 9)
- Commands for key/certificate management (your own CA)
 - genrsa -- key generation
 - verify -- verify cert
 - pkcs12 -- convert cert encodings
 - x509 -- cert mgt
 - req -- generate cert or CSR

Keep current:

- New Features
- Bug Fixes

CNS Lecture 14 - 48



Example CAPI

Kerberos

- Client context established (tickets, keys) KDC
- server/principal must be registered
- `krb_mk_safe()`, `krb_mk_priv()`
- user provides transport

DCE

- context established (tickets, keys) KDC
- server/principal must be registered
- `rpc_binding_set_auth_info()`
- RPC mechanism provides transport




Figure 14.1 Overview of Kerberos

CNS Lecture 14 - 49

Kerberos v4 API example

- client sends safe and private message
- server decodes and prints
- server must be registered with Kerberos and server key in server's `/etc/krbtab` (HOST and SERVICE in example C code)



```
krb_mk_req(authent, service, instance, realm, checksum)
krb_rd_req(authent, service, instance, from_addr, ad, fn )
krb_get_cred(service, instance, realm, credentials)
krb_mk_priv(in, out, in_length, schedule, key, sender, receiver)
krb_rd_priv(in, in_length, schedule, key, sender, receiver, msg_data)
krb_mk_safe(in, out, in_length, key, sender, receiver)
krb_rd_safe(in, length, key, sender, receiver, msg_data)
```

CNS Lecture 14 - 50

Kerberos v4 client

```
main()
{
    KTEXT_ST k; /* Kerberos data */
    KTEXT ktxt = &k;
    CREDENTIALS c; /* ticket & session key */
    CREDENTIALS *cred = &c;
    des_key_schedule sched; /* session key schedule */
    if ((serv = getservbyname(SERVICE, "udp")) == NULL) {
        fprintf(stderr, "service unknown: %s/udp\n", SERVICE);
        exit(1);
    }
    if ((host = gethostbyname(HOST)) == (struct hostent *) 0) {
        fprintf(stderr, "%s: unknown host\n", HOST);
        exit(1);
    }
    bzero((char *)s_sock, sizeof(s_sock));
    bcopy(host->h_addr, (char *)s_sock.sin_addr, host->h_length);
    s_sock.sin_family = AF_INET;
    s_sock.sin_port = serv->s_port;
    gethostname(hostname, sizeof(hostname));
    host = gethostbyname(hostname);
    if ((sock = socket(AF_INET, SOCK_DGRAM, 0)) < 0) {
        perror("opening datagram socket");
        exit(1);
    }
}
```

CNS Lecture 14 - 51

Client continued ...

```
bzero((char *)c_sock, sizeof(c_sock));
bcopy(host->h_addr, (char *)c_sock.sin_addr, host->h_length);
c_sock.sin_family = AF_INET;
if (bind(sock, (struct sockaddr *)&c_sock, sizeof(c_sock)) < 0) {
    perror("binding datagram socket");
    exit(1);
}
/* Get local realm, not needed, just an example */
if ((l = krb_get_lrealm(c_realm, l)) != KSUCCESS) {
    fprintf(stderr, "can't find local Kerberos realm\n");
    exit(1);
}
printf("Local Kerberos realm is %s\n", c_realm);
/* Get Kerberos realm of host */
s_realm = krb_realmofhost(HOST);

/* Get ticket for server from TGS, create krb_mk_req message */
if ((i = krb_mk_req(ktxt, SERVICE, HOST, s_realm, cksam)) != KSUCCESS) {
    fprintf(stderr, "%s\n", krb_err_txt[i]);
    exit(1);
}
printf("Got credentials for %s.\n", SERVICE);
```

CNS Lecture 14 - 52

Client cont.

```
/* Send authentication info to server */
i = sendto(sock, (char *)ktxt->dat, ktxt->length, flags, (struct sockaddr *)&s_sock,
sizeof(s_sock));
if (i < 0)
    perror("sending datagram message");
printk("Sent authentication data: %d bytes\n", i);
/* PREPARE KRB_MK_SAFE MESSAGE */
/* Get my address */
bzero((char *)c_sock, sizeof(c_sock));
i = sizeof(c_sock);
getsockname(sock, (struct sockaddr *)&c_sock, &i);
/* Get session key from my collection of tickets */
i = krb_get_cred(SERVICE, HOST, s_realm, cred);
/* Make the safe message */
len = krb_mk_safe(MSG, ktxt->dat, strlen(MSG)+1, cred->session, &c_sock, &s_sock);
/* Send it */
i = sendto(sock, (char *)ktxt->dat, (int) len, flags, (struct sockaddr *)&s_sock,
sizeof(s_sock));
/* PREPARE KRB_MK_PRIV MESSAGE */
/* Get key schedule for session key */
key_sched(cred->session, sched);
/* Make the encrypted message */
len = krb_mk_priv(MSG, ktxt->dat, strlen(MSG)+1, sched, cred->session, &c_sock, &s_sock);
/* Send it */
i = sendto(sock, (char *)ktxt->dat, (int) len, flags, (struct sockaddr *)&s_sock,
sizeof(s_sock));
}
```

CNS Lecture 14 - 53

Kerberos v4 server

```
main()
{
    KTEXT_ST k;
    KTEXT ktxt = &k; /* Kerberos data */
    AUTH_DAT ad; /* authentication data */
    struct sockaddr_in c_sock; /* client's address */
    MSG_DAT msg_data; /* session key schedule */
    des_key_schedule sched; /* session key schedule */
    /* Set up server address */
    bzero((char *)s_sock, sizeof(s_sock));
    s_sock.sin_family = AF_INET;
    serv = getservbyname(SERVICE, "udp");
    s_sock.sin_port = serv->s_port;
    gethostname(hostname, sizeof(hostname));
    host = gethostbyname(hostname);
    bcopy(host->h_addr, (char *)s_sock.sin_addr, host->h_length);
    sock = socket(AF_INET, SOCK_DGRAM, 0);
    bind(sock, (struct sockaddr *)&s_sock, sizeof(s_sock));
    /* GET KRB_MK_REQ MESSAGE */
    i = read(sock, (char *)ktxt->dat, MAX_KTXT_LEN);
    ktxt->length = i;
    /* Check authentication info */
    i = krb_rd_req(ktxt, SERVICE, HOST, any, &ad, "");
    if (i != KSUCCESS) {
        ...
    }
}
```

Receive ticket from client

Verify ticket and extract session key using server key in `/etc/krbtab`

CNS Lecture 14 - 54

Server cont.

```

/* GET KRB_MK_SAFE MESSAGE */
i = sizeof(c_sock);
i = recvfrom(sock, (char *)ktxt->dat, MAX_KTXT_LEN, flags,
             (struct sockaddr *)&c_sock, &i);
/* Verify the checksummed message */
i = krb_rd_safe(ktxt->dat, i, (ad.session) &c_sock, &s_sock, &msg_data);
if (i != KSUCCESS) {
    ...
printf("Safe message is: %s\n", msg_data.app_data);
/* NOW GET DECRYPTED MESSAGE */
key_sched(ad.session) sched);
i = sizeof(c_sock);
i = recvfrom(sock, (char *)ktxt->dat, MAX_KTXT_LEN, flags,
             (struct sockaddr *)&c_sock, &i);
i = krb_rd_priv(ktxt->dat, i, sched, (ad.session) &c_sock, &s_sock, &msg_data);
if (i != KSUCCESS) {
    ...
printf("Decrypted message is: %s\n", msg_data.app_data);
}

```

CNS Lecture 14 - 55



kerberizing

- you can add Kerberos calls to your own client/servers
- need Kerberos data base, authenticator, ticket-granting server, and administrative programs
- can use klogin, but better if you have kerberized BSD utilities
- Kerberos calls added to login, r-utilities, NFS


CNS Lecture 14 - 56



GSS

Generic Security Service

- IETF RFC 1508 1509, APt RFC 2743 2744
- establishes a security context and provides security services
- application goes through 4 phases
 1. establish identity (authentication)
 2. negotiate shared security context
 3. exchange messages -- privacy, integrity
 4. destroy context
- example implementations in Kerberos v5 and DCE distributions
- API is independent of OS and network protocols
- application is responsible for message transport or file I/O
- used in globus (grid security)
- Java class GSSUtil
- primary functions
 - gss_acquire_cred()
 - gss_init_sec_context()
 - gss_seal() gss_unseal()
 - gss_sign() gss_verify()



CNS Lecture 14 - 57



Example CAPI

PKCS 11 cryptoki/ FORTEZZA

- crypto token/smartcard API
- routines for
 - card discovery
 - PIN verify
 - encrypt/decrypt
 - hash/sign/verify
 - wrap/unwrap key
 - random generation
 - management functions



CNS Lecture 14 - 58



Example CAPI

Microsoft CAPI

- 23 crypto services
- context and key generation
- key exchange
- encryption/signing

```

CryptAcquireContext() CryptGenKey() CryptGetKey()
CryptEncrypt() CryptDecrypt() CryptHashData()
CryptSignHash() CryptVerify() CryptGenRandom()

```

Other proprietary CAPIs - SpyruS or Cyrix

• actual algorithms (RSA, DES, MD5) provided by cryptographic service provider (CSP) can replace underlying CSP.
 • gov'ts or vendors might define CSP's (e.g., FORTEZZA).
 • RSA CSP is provided.
 • Microsoft must sign the CSP.
 • supposedly easy to switch from strong to weak crypto by swapping CSP.
 • used in Since Internet Explorer 3.0
 • samples and manual available

CNS Lecture 14 - 59



Where to encrypt?


- link layer
 - encrypting modem, net board (wireless)
 - transparent, fast
 - protects only one link (pt-to-pt)
 - info may be exposed in OS

network/transport layer

- IPsec IPv6(v4), VPN
- transparent
- selectable (policy)
- appl./host/net keying

application layer

- custom applications (FGP, ssh)
- CAPTs can help (openSSL)
- intrusive, but flexible
- key for every logical circuit



CNS Lecture 14 - 60



Secure applications

App's we've already seen

- ssh
- PGP
- SSL/netescape
- S/MIME
- characteristics
 - key mgmt/schedule
 - ciphers/hashes
 - public key/symmetric key
 - random numbers
 - Protocols and encodings
 - Crypto agile

CNS Lecture 14 - 61



nautilus

nautilus sinks clipper ships ...

- encrypted Internet/modem phone
- UNIX, Windows
- various audio encodings
- D-H or shared secret
- 3DES, Blowfish, IDEA
- random bits from microphone noise



CNS Lecture 14 - 62



speakefreely



- Use microphone/speaker for secure voice over IP
- Uses GSM compression, SGI's AtoD PCM, VAT/RTP options
- Crypto
 - AES/Blowfish/IDEA/DES to encrypt audio (CBC within a "block")
 - Session key can be provided by user as ASCII passphrase or by a key file (e.g. dd count=10 of=rand.dat if=/dev/random)
 - sfmike -e will generate a session key (128 bit encoded into 32 ASCII characters that Bob could PGP email to Alice, or read over phone?)
 - For "conference call" -Z option will create session keys for each user and PGP/PGP encrypt the key with their public key!

```
makeSessionKey() in mike .c creates random session key with
getpid() getpid() clock() time() gettimeofday() gethostname()
getuid() getuid() getdomainname() and /dev/random
all hashed with MD5
```

CNS Lecture 14 - 63



zfone

- Zimmerman's voice-over-IP secure phone (IETF draft)
- Uses ideas from PGPfone
- Doesn't need PKI
 - Does peer-to-peer key establishment
- Header extensions to RTP to support Diffie-Hellman
 - Establishes ephemeral session key/IV for Secure RTP
 - Displays authentication string for each user to verify authenticity and prevent man-in-the-middle attacks
 - Can include pre-shared secret and running shared secret
- Secure RTP uses
 - HMAC-SHA1
 - AES 128 (counter mode)
 - 112-bit session salt key
 - 2⁴⁸ key derivation rate

CNS Lecture 14 - 64



Disk encryption cfs

- des or pgp -c is cumbersome
- Blaze AT&T
- most BSD-based UNIX and linux
- local NFS server on loopback interface
- encrypted directories
- backups etc. are unaffected
- DES/3DES/MacGuffin
- /crypt is mounted
- command to attach/detach (with DES key)
- if CFS unavailable, ccat can be used

```
startup cfs
cfsd
mount -o port=3049,intr localhost:null /crypt

Create encrypted directory
mkdir /usr/dunigan/secrets
Key: xxxxxxxxxxxx

attach /usr/dunigan/secrets tom
Key: xxxxxxxxxx (now have /crypt/tom)

do normal file operations cat, vi, etc.

detach tom

/usr/dunigan/secrets/80b6e856778
contents are gibberish (encrypted)
```

CNS Lecture 14 - 65



cfs encryption

- need random access for file
 - ECB works but crypto weak
 - CBC can do random reads, but need to re-encrypt all to do block update
- cfs computes two DES keys from passphrase
- uses first key to create 512KB pseudo-random bit mask using DES-OFB -- stored for life of attach
- file block is XOR'd with mask corresponding to byte offset in file mod size of mask (512KB) and then encrypted with DES-ECB
- optionally file IV's generated from inode number and stored in file's group-ld field
- 4 times slower than regular NFS I/O benchmarks, but only about 1.5 slower for "normal" applications (compile)


Are plaintext versions lying around? swap space, backup files (vi), /dev/mem ...
 wipe utility to over-write plaintext?
 Is the passphrase lying in memory/swap space?
 Unattended terminal session? -- timeout?

CNS Lecture 14 - 66



Windows Encrypting File System (EFS)


- Encrypt folders (DES-X or 3DES)
- Provides key recovery, you need a certificate




CNS Lecture 14 - 67

Windows EFS

Windows generates a random number for the file encrypting key (FEK). User's public key is used to encrypt the FEK. Encrypted key is stored in extension attribute to the file. FEK also encrypted under recovery agent's public key.




"wipe" service provided to scrub clear text file blocks.



CNS Lecture 14 - 68

BitLocker vista


- Drive encryption optionally using TPM, optional user password
- Solve the stolen laptop problem
 - Removing the drive or booting with thumb drive and becoming Administrator still doesn't give you access to the bitlocked disk
- Each sector encrypted with AES CBC
 - no authentication? Because it would require more bits per sector
 - Since CBC any change will effect the rest of the sector
 - Uses sector number as part of encryption
 - IV is encryption of sector key and sector number
- Uses "Elephant diffuser"
 - Provide poor man's authentication
 - Fast (32-bit word XOR and rotate)
- Performance (pentium 4)
 - AES 20 cycles/byte
 - Diffuser 10 cycles/byte
 - 5% slower



CNS Lecture 14 - 69

MAC disk encryption

- Encrypt a disk "volume" using AES
 - Key can be added to your keychain for one-time signon
- FileVault to AES-encrypt your home directory
 - Key recovery option with Master password
 - Disable auto-login (duhhh)
 - Backups unencrypted (unless you backup the vault)



CNS Lecture 14 - 70

Next time

Review

CNS Lecture 14 - 71