

More Security

Security Issues - NT vs UNIX passwords
Network Security
Database Security
Encryption
SSL

Security Issues

It is easy to run a secure computer system. You merely have to disconnect all dial-up connections and permit only direct-wired terminals, put the machine and its terminals in a shielded room, and post a guard at the door.

F.T. Grampp & R.H. Morris

What is Computer Security?

- Security is keeping unauthorized people or programs from doing things that you do not want them to do on your system.
- But, we must define what exactly we are trying to protect
 - CPU cycles
 - Files, storage devices, data
 - Telecommunications
 - Identity and access privileges

Security Policies

- A security policy is a set of decisions that define an organization's attitude and actions toward the issue of security.
- The limits of acceptable behaviour and the appropriate responses to unacceptable behaviour must be defined.
- Policies must take into account the nature of the organization, what is to be protected, from whom, and at what cost.

Password Compromising

- The simplest way into a system is to appear as if you are a legitimate user and just log in!
- A successful login is defined as the pairing of a login id with a password in a reasonable number of tries.
- Early login systems stored passwords in plain text files but attempted to hide the location of the password file – not a great strategy!

Password Security

- Obviously, passwords must be stored:
 - File containing plain text of the passwords
 - File containing encrypted versions of the passwords
 - File containing a unique obscured representation of the passwords
 - Password is used as a parameter in an irreversible mathematical function that calculates a hash value

Windows NT Password Security

- The password file (*\\WINNT\\SYSTEM32\\CONFIG\\SAM*) contains hash values generated by the Windows/NT hash function with the user's password as an input parameter.
- The default control access lists allow read access to the password database.

The UNIX Hashing Function

- DES-like algorithm is used to calculate the hash value.
- The password is used as the DES key (eight 7-bit chars make a 56 bit DES key) to encrypt a block of binary zeros.
- The result of the encryption is the hash value.
- The password is not encrypted – it is the key used to perform the encryption!

The UNIX Hashing Function

- Two random characters (the *salt*) are introduced into the algorithm to ensure that two equal passwords result in two different hash values.
- The *salt* is a 12-bit random number that is concatenated with a user's password before it is hashed.
- The *salt* is stored with the hashed result.

The Power of Salt!

- Try the following:

```
perl
print crypt("fred","am");
^D
amLH9TiZZksc perl
print crypt("fred","an");
^D
anvepwCPZQ2Z6
```
- Same password but different hash value!

The NT Hashing Function

- Uses the Internet standard MD4 hashing algorithm to generate a 16-byte hash of the Unicode encoded password.
- LanManager hashing function
 - Uses a 14-byte password (112 bits) which is split into two 56-bit entities.
 - These two 56-bit entities are used to DES-encrypt a fixed (known) 8-byte magic number.
 - The result of these two encryptions are concatenated to form a 16-byte hash value.

The NT Hashing Function

- NT stores both the LanManager (DES) and NT (MD4) hash values in the SAM (Security Accounts Manager) part of the registry.
- These two 16-byte hash values are further obfuscated by DES-encrypting them, but the keys used for this encryption are known and so it is easy to recover the hash values.

Which is more secure?

- The random element in the UNIX hashing algorithm makes it difficult for a “cracker” to perform password attacks because the cracking program has to introduce the salt as well.
- Both NT and UNIX (shadow) password files are readable only by system administrators (root) and so crack programs have to have root privilege!

An NT Password Attack

- If the NT SAM database is obtained (installation may leave a copy of the password file in */WINNT/REPAIR*) then an NT cracking program can be run.
- NT cracking program as not as sophisticated as their UNIX counterparts.
 - But because the NT hashing functions do not use a salt, the hashing/matching process is much faster – rumor says 5 minutes to match a 860,000 word list against a SAM DB of 1,000 users.

Introduction to Network Security

<http://www.interhack.net/pubs/network-security/>

Risk Management: The Game of Security

- There are two extremes: absolute security and absolute access.
- Every organization needs to decide for itself where between the two extremes of total security and total access they need to be.
- A policy needs to articulate this, and then define how that will be enforced with practices.

Lessons Learned

Did anyone ever do backups?

- Operational requirements should dictate the backup policy, and this should be closely coordinated with a disaster recovery plan.

Don't put data where it doesn't need to be

- Information that doesn't need to be accessible from the outside world sometimes is, and this can needlessly increase the severity of a break-in dramatically.

Avoid systems with single points of failure

- Any security system that can be broken by breaking through any one component isn't really very strong.
- In security, a degree of redundancy is good

Stay current with relevant operating system patches

- Watch the vendors' security advisories.
- Exploiting old bugs is still one of the most common (and most effective!) means of breaking into systems.

Watch for relevant security advisories

- Keep a close watch on groups like CERT and CIAC.



CERT (www.cert.org)

- The CERT Coordination Center (CERT/CC) is a center of Internet security expertise.
- It is located at the Software Engineering Institute, a federally funded research and development center operated by Carnegie Mellon University.

CIAC (<http://ciac.llnl.gov/ciac>)

- Computer Incident Advisory Capability
- CIAC provides on-call technical assistance and information to Department of Energy (DOE) sites faced with computer security incidents.
- This central incident handling capability is one component of all encompassing service provided to the DOE community.

CIAC (<http://ciac.llnl.gov/ciac>)

- CIAC is a founding member of FIRST, the Forum of Incident Response and Security Teams, a global organization established to foster cooperation and coordination among computer security teams worldwide.
- To protect and ensure authenticity all of CIAC electronic publications, bulletins and advisories are signed with a PGP encryption key.

Have someone on staff be familiar with security practices

- Having at least one person who is charged with keeping abreast of security developments is a good idea.
- They should be familiar with the *dos and don'ts* of security, from the such sources as the *Site Security Handbook* (RFC1244).
 - <http://www.net.ohio-state.edu/rfc1244/intro.html> or <http://www.landfield.com/rfcs/rfc1244.html>

Site Security Handbook

- This handbook is a guide to setting computer security policies and procedures for sites that have systems on the Internet.
- It lists issues and factors that a site must consider when setting their own policies.
- It makes some recommendations and gives discussions of relevant areas.

Firewalls

- Connecting an organization to the Internet provides a two-way flow of traffic.
- To provide some level of separation between an organization's intranet and the Internet, firewalls have been employed.
- A firewall is simply a group of components that collectively form a barrier between two networks.

Bastion host

- A general-purpose computer used to control access between the internal (private) network (intranet) and the Internet (or any other untrusted network).
- Typically, these are hosts running a flavor of the Unix operating system that has been customized in order to reduce its functionality to only what is necessary in order to support its functions.

Bastion host

- Many of the general-purpose features have been turned off, and in many cases, completely removed, in order to improve the security of the machine.

Router

- A special purpose computer for connecting networks together.
- Routers also handle certain functions, such as routing , or managing the traffic on the networks they connect.

Access Control List (ACL)

- Many routers now have the ability to selectively perform their duties, based on a number of facts about a packet that comes to it.
- This includes things like origination address, destination address, destination service port, and so on.
- These can be employed to limit the sorts of packets that are allowed to come in and go out of a given network.

Demilitarized Zone (DMZ)

- The DMZ is a critical part of a firewall: it is a network that is neither part of the untrusted network, nor part of the trusted network.
- But, this is a network that connects the untrusted to the trusted.
- The importance of a DMZ is tremendous: someone who breaks into your network from the Internet should have to get through several layers in order to successfully do so.

Proxy

- This is the process of having one host act in behalf of another.

Database Security

From *Building Secure Software*
by John Viega and Gary McGraw, 2002.

DB Security Woes

- Most databases ignore security issues.
- The biggest concern with databases is whether the connection between the database and the user is encrypted.
 - Strong encryption for connections is support by an add-on to Oracle and MySQL supports SSL connections.

Malicious Data

- Malicious data from untrusted users can severely damage a database.
- Authentication of users and data is essential.

Authentication

- Databases provide password authentication and access control to tables.
- Protects against other users of the database but not against your application's own users.

Statistical Attack

- An attacker can make inferences based on available information that reveals facts that should not have been revealed.
- This type of attack is difficult to defend against.

Using Stats to Violate Privacy

- You let someone perform queries against your customer database but you do not want to violate customer privacy.
 - You do not let this user see personal data but they can still use queries to generate data that violates privacy.
 - E.g. A query that only returns one customer may give you enough information to identify that person.

Approaches

- Filter out queries that are not aggregate queries.
 - Aggregate queries operate on a set of rows and returns only a single result.
 - AVG, COUNT, MAX, MIN, SUM

More problems...

- But if COUNT ever returns 1 we have the same problem.
 - Restrict aggregates to a specified number of records, i.e. must be > 10 .
- Must also guard against the complement of the query.

Even more problems...

- If a hacker can identify any unique entries then they can pose a query that exposes other information.
 - The attacker identifies a query that can be answered on an interesting field, e.g. income.
 - A query is formed from the original query, logically OR-ed with the information that the attacker is really interested in.

Example

- Want to find the income of a quarantined patient in Guelph - we know that there is only one such person.
- But our aggregate queries are restricted to numbers of rows > 100 .
- We know that Toronto has hundreds of people in quarantine.

Queries

- `SELECT COUNT (*) FROM qpatients WHERE city = "Toronto";`
 - Result = 99
- `SELECT AVG(income) FROM qpatients WHERE city="Toronto" OR city="Guelph";`
 - Result = \$60,500

Queries

- `SELECT AVG(income) FROM qpatients WHERE city="Toronto";`
 - Result = \$60,000
 - $(99 * 60000 + y) / 100 = 60500$
 - Target salary = \$110,000

And in conclusion...

- There is no good, general-purpose strategy for defeating a statistical attack.
- Defense Strategies:
 - Log queries and perform a manual audit periodically to see if there are suspicious queries.
 - Modify the database to preserve the statistical properties but affect the ability to track individuals.
 - Add random noise to the statistical results.

Access Control

- Components for access control in SQL databases are objects, actions, and privileges.
- Objects are tables, views, columns.
- Default: only the owner (creator) of an object has access to it.
- The owner can grant privileges to other users.

Access Control

- Privileges are composed of
 - The entity granting the privilege.
 - The entity to which the privilege is granted.
 - The object for which the permission is granted.
 - The action granted.
 - Whether the privilege can be granted to others.

Access Control

- GRANT action(s) ON object TO user(s);
 - GRANT SELECT ON patients TO Doc1;
 - GRANT SELECT, INSERT, DELETE, UPDATE ON patients TO Doc1;
 - GRANT ALL ON patients TO Doc1;

Access Control

- GRANT SELECT ON stats TO PUBLIC;
- GRANT SELECT ON stats TO Doc1 WITH GRANT OPTION;
- REVOKE ALL ON stats FROM PUBLIC;
- REVOKE GRANT OPTION FOR SELECT ON stats FROM Doc1;

Access Control

- Views are virtual tables that are created from one or more real tables in the DB.
- Views can be used to enforce access policies.
- Users can be assigned views.
- The downside to views is that they can be slow and should not be used as an excuse to allow clients to connect directly to your database.

Field Protection

- Some fields in the database need to be protected from everyone, e.g. passwords, credit card numbers.
- Solution: encryption.
- Caution: do not use built-in encryption.
 - There are password-decoding utilities for Oracle databases in circulation.

Field Protection

- The Safe Solution: do your own encryption before storage and your own decryption after retrieval.
- Downsides:
 - Cannot storage binary strings so must encode the encrypted data.
 - No efficient search mechanism.
 - Performance hit.
 - Expanding data - may not be able to store encrypted data in existing fields

Encryption – An Introduction

Definitions

- Encryption is the transformation of data in ways that are hard to copy or reverse by unauthorized persons.
- The principle properties of encryption techniques are
 - Confidentiality
 - Authenticity
 - Integrity

Confidentiality

- Confidentiality is the assurance that people who are not party to the encryption technique cannot read the encrypted message.
- This confidentiality is usually guaranteed by the fact that the sender and receiver of the encrypted message share some sort of decoding information that *is known only to them*.

Authenticity

- An encrypted message is authentic if the receiver is *reasonably sure* that the incoming message is indeed from the appropriate sender.
- If Alice is sending Bob a message encrypted using information that only they share, then Bob should not have to worry about forged messages because they should not decrypt properly.

Integrity

- Integrity refers to the fact that an encrypted message should not be modifiable except by the authorized sender.
- Since a forger should not be able to read the message, they should not be able to change any of the contents.
- Random changes to the encrypted message will usually render it *unreadable*.

Private Key Encryption

- This is the traditional method of cryptography.
- The sender uses a key to encrypt a message using some function T .
- Early examples of this type of system
 - Polybios square
 - CAESAR

Problems with Private Keys

- Private key systems rely on the sender and receiver *both having the key*.
 - Keys have to be delivered to both parties – this is both cumbersome, time-consuming, and possibly insecure or dangerous.
- Even without the key, simple systems are breakable by *brute force* – if the algorithm is known then try all key combinations.

Modern Private Key Encryption

- In 1977, the Data Encryption Standard (DES) was introduced by the National Bureau of Standards in the U.S.
- It was the first, modern, commercial encryption standard.
- Its main attributes include
 - Extremely fast
 - Behaves according to the avalanche effect

Avalanche Effect

- A small change in the key will yield very large differences in the produced ciphertext.
- This property makes the algorithm extremely difficult and costly to break.
- The Electronic Frontier Foundation has created a \$220,000 machine to crack DES-encrypted messages!
- The newest version of DES (triple-DES) should still provide more than enough security.

DES – The Algorithm

- DES encrypts groups of 64 message bits or 16 hexadecimal numbers.
- DES uses keys which are also apparently 16 hex numbers long, however, every 8th key bit is ignored so the effective key size is 56 bits.
- If the message to be encrypted is not a multiple of 64 bits than the message is padded.

Block Cipher

- There can be different padding schemes, e.g. padding with 0's.
- DES is a *block cipher*.
 - It operates on plaintext blocks of a given size and returns ciphertext blocks of the same size.
- DES results in a permutation among the 2^{64} possible arrangements of 64 bits.
- Each block is divided into two blocks of 32 bits – a left half block **L** and right half **R**.

Cracking DES

- As long ago as 1975, Diffie and Hellman predicted that DES was vulnerable to a brute force attack.
 - Brute force works by working its way through as many as possible of the 2^{56} possible keys until a *sensible* plaintext is retrieved.
- In 1998, John Gilmore of the EFF built a machine (*Deep Crack*) that could go through the entire 56-bit key space in an average of 4.5 days.

Triple-DES

- DES with two 56-bit keys!
 - The first key encrypts the plaintext message.
 - The second key is used to DES-decrypt this encrypted message – but of course this will not work – it just *scrambles* the message even more.
 - The scrambled message is encrypted again with the first key to yield the final ciphertext.
- This increases the key space to 2^{112} .

Introduction to SSL

The Secure Sockets Layer Protocol

- <http://developer.netscape.com/docs/manuals/security/slin/index.htm>
- <http://www.rsa.com>

Secure Sockets Layer Protocol

- The Secure Sockets Layer (SSL) protocol was originally developed by Netscape.
- SSL is an accepted Internet standard for authenticated and encrypted communication between clients and servers.
- The new Internet Engineering Task Force (IETF) standard called Transport Layer Security (TLS) is based on SSL.

TCP/IP and SSL

- The SSL protocol runs above TCP/IP and below application-layer protocols such as HTTP or IMAP.
- SSL uses TCP/IP on behalf of the other protocols, and allows
 - an SSL-enabled server to authenticate itself to an SSL-enabled client
 - the client to authenticate itself to the server
 - both machines to establish an encrypted connection.

SSL Server Authentication

- Allows a user to confirm a server's identity.
- SSL-enabled client software can use public-key cryptography to check that a server's certificate and public ID are valid and have been issued by a certificate authority (CA) listed in the client's list of trusted CAs.
- This confirmation might be important if the user, wants to check the receiving server's identity (e.g. credit card transactions).

SSL Client Authentication

- Allows a server to confirm a user's identity.
- SSL-enabled server software can check that a client's certificate and public ID are valid and have been issued by a certificate authority (CA) listed in the server's list of trusted CAs (same as server authentication).
- This confirmation might be important if the server, wants to check the recipient's identity (e.g. banking transactions).

Encrypted SSL Connection

- Requires all information sent between a client and a server to be encrypted by the sending software and decrypted by the receiving software, thus providing a high degree of confidentiality, i.e. a private transaction.
- All data sent over an encrypted SSL connection is checked for tampering -- automatically determining whether the data has been altered in transit.

Ciphers Used with SSL

- The SSL protocol supports the use of a variety of different cryptographic algorithms, or **ciphers**, for use in operations such as authenticating the server and client to each other, transmitting certificates, and establishing session keys.

Cipher Suites

- Clients and servers may support different **cipher suites** depending on factors such as
 - the version of SSL they support
 - company policies regarding acceptable encryption strength
 - government restrictions on export of SSL-enabled software

Cipher Suite Usage

- The SSL handshake protocol determines how the server and client negotiate which cipher suites they will use to authenticate each other, to transmit certificates, and to establish session keys.

Cipher Suites

- **DES**
 - Data Encryption Standard, an encryption algorithm used by the U.S. Government.
- **DSA**
 - Digital Signature Algorithm, part of the digital authentication standard used by the U.S. Government.
- **KEA**
 - Key Exchange Algorithm, an algorithm used for key exchange by the U.S. Government.

Cipher Suites

- **MD5**
 - Message Digest algorithm developed by Rivest.
- **RC2 and RC4**
 - Rivest encryption ciphers developed for RSA Data Security.
- **RSA**
 - A public-key algorithm for both encryption and authentication. Developed by Rivest, Shamir, and Adleman.

Cipher Suites

- **RSA key exchange**
 - A key-exchange algorithm for SSL based on the RSA algorithm.
- **SHA-1**
 - Secure Hash Algorithm, a hash function used by the U.S. Government.

Cipher Suites

- **SKIPJACK**
 - A classified symmetric-key algorithm implemented in FORTEZZA-compliant hardware used by the U.S. Government.
- **Triple-DES**
 - DES applied three times.

RSA

- Key-exchange algorithms like KEA and RSA govern the way in which the server and client determine the symmetric keys they will both use during an SSL session.
- The most commonly used SSL cipher suites use RSA key exchange.

Choice of Suite

- The SSL 2.0 and SSL 3.0 protocols support overlapping sets of cipher suites.
- Administrators can enable or disable any of the supported cipher suites for both clients and servers.
- When a client and server exchange information during the SSL handshake, they identify the strongest enabled cipher suites they have in common and use those for the SSL session.

Choice of Suite

- Decisions about which cipher suites a particular organization decides to enable depend on trade-offs among the sensitivity of the data involved, the speed of the cipher, and the applicability of export rules.

Strength of Ciphers

- Some organizations may want to disable the weaker ciphers to prevent SSL connections with weaker encryption.
- Disabling support for all 40-bit ciphers effectively restricts access to network browsers that are available only in the United States.

Why??

- Due to U.S. government restrictions on products that support anything stronger than 40-bit encryption
- Unless the server involved has a special Global Server ID that permits the international client to step up to stronger encryption.

But...

- 40-bit ciphers can be broken relatively quickly.
- If one is concerned about eavesdropping and if your user community can legally use stronger ciphers then the 40-bit ciphers should be disabled.

The Strongest Cipher Suite

- Permitted for deployments within the United States only.
- Appropriate for banks and other institutions that handle highly sensitive data.
- **Triple DES, which supports 168-bit encryption, with SHA-1 message authentication.**

Triple DES

- Triple DES is the strongest cipher supported by SSL, but it is not as fast as RC4.
- Triple DES uses a key three times as long as the key for standard DES.
- Because the key size is so large, there are more possible keys than for any other cipher, approximately $3.7 * 10^{50}$.