Probable Plaintext Cryptanalysis of IPSEC

Steven M. Bellovin

smb@research.att.com

908-582-5886

AT&T Labs Research

Murray Hill, NJ 07974

Attacking DES

- DES is a strong cipher, but its key size is too short to protect against brute-force cryptanalysis.
- Most published designs for DES-crackers assume a block of known plaintext.
- Often, we don't know any full blocks of plaintext, but we may know enough.

IPSEC

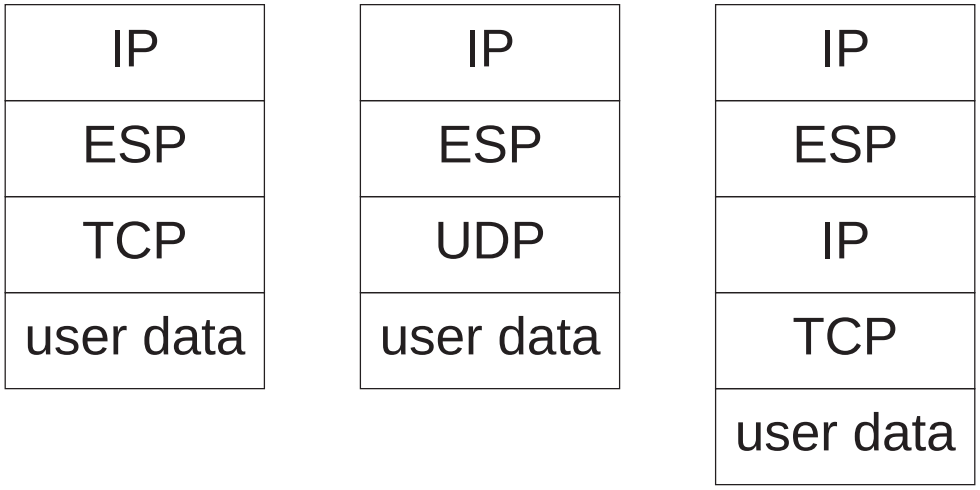

The encryption header contains a security association identifier, a replay prevention counter, authentication data, length information, padding, and the “next protocol” indicator.

Single Packet Attacks

- Pick out one packet.
- From that packet's characteristics, guess the values of many bits in a single block. Feed that into a key search engine.
- Test candidate keys on other blocks or (ultimately) by human analysis.

Example: IP Header

Replay Counter An early draft said that the replay counter started at 1.

Probable plaintext available: 30-31 bits.

Version/header len/TOS/precedence Almost always 4510_{16}.

Probable plaintext available: 16 bits.

Packet length We know that 30-40% of packets are TCP ACK packets, exactly 40 bytes long. By looking at the received length of the packet, we can often pick these out, giving us the IP length.

Probable plaintext available: 16 bits.

Src/Dst Addresses If host-to-host tunnel mode is used, the encrypted addresses will match the outside addresses exactly.

Probable plaintext available: 64 bits.

In firewall-to-firewall mode or host-to-firewall mode, the high-order bits will often match or otherwise be known from other sources.

Probable plaintext available: 32-48 bits.

Double Packet Attacks

- Pick out two packets from the same stream.
- Do trial decryption on both packets.
- See if the fields match.
- Statistically, adding a small value to a counter changes very few bits.

☞ More or less doubles cost per key recovered.

Example: TCP Header

Port numbers In the same connection, port numbers will match exactly.
Probable plaintext available: 32 bits.

Seq/Ack Typically, these fields change in only a few bit positions on successive packets. The length of the packet may be a clue.
Probable plaintext available: 18-30 bits.

Flags/Window/Urgent Flags are generally constant (5018_{16}); window size of the sender will remain constant; urgent pointer is usually 0.
Probable plaintext available: 48 bits.

Other fields that match closely between packets are the IP packet ID, TTL, protocol, and checksum fields, the ESP replay counter, and the UDP port number fields.

Total Probable Plaintext

	Single	Double
IP	54–58	127
TCP	88	124
UDP	28	46

If tunnel mode is used, the IP header will have 32 or 64 more bits of probable plaintext.

For double packet attacks, the ESP header has at least 30-31 more bits of probable plaintext if the IV is known.

The Role of Traffic Analysis

- Can be used to find packet pairs for double-packet analysis.
- TCP open sequence is very characteristic; initial packets have more known plaintext (sequence and acknowledgment fields).
- Packet timings and lengths give away clues to port numbers.
- Closely spaced packets are often from the same TCP stream (packet trains).
- Related conversations, such as multiple image fetches for a Web page, often have common addresses, source port numbers, etc.
- Per-connection keying makes some of this easy.

Defenses

- Avoid host-to-host tunnel mode.
- For firewall-to-firewall mode, use secret internal addresses.
- Use host-pair or firewall-pair keying, not per-connection keying.
  ☞But this has other cryptographic weaknesses!
- Compression:
  - Abbreviate sequence and acknowledgment fields, as in VJ header compression.
  - With per-connection keying, eliminate port numbers as well.
  - Possibly use data compression with a keyed dictionary.
- Don't use ciphers as weak as DES. . .

World War II Analogs

- Cryptanalysts often tried likely texts for message starts as plaintext for the bombes.
- Parallel solution of the same message in Enigma and simpler systems provided other cribs.
- Repeated encryption of the initial rotor settings provided unknown values that nevertheless had to match upon decryption.