

Security and Cryptography 2015

Mirosław Kutylowski

grading criteria: 50% exam, 50% assignments

skills to be learned: developing end-to-end security systems, they must be flawless!

rules: do not memorize the standards, they come and go. Only the skills are important

exam date: TBA (quickly to enable internships in February)

I. EXAMPLE TO LEARN FROM: PKI FAILURE

reasons for PKI failure (According to Schneier):

1. whom we trust and for what? why CA should be trusted??
2. who is using my key? (private key - there are really no clones??)
3. how secure is the verifying computer? (no cryptography can help if the verifier software is cheating)
4. who is the signer? (ambiguity unless there is a trustworthy ID registry)
5. is CA an authority? (really not an authority for data contained in the certificate. Certificate based on fake documents...)
6. is the user part of the security design? (no, the user is free to behave in a stupid way)
7. separation CA and RA brings new threats
8. How did CA verify the certificate holder? (certificate issued for ..., but how to know that this was really this person)
9. How secure are the certification practices? (revocation, etc)
10. the customers wish to run single-sign-on

Answers: yes, but

- tradition in Nordic countries
 - honest system participants
 - the best one can do
 - ...
- no solutions from crypto community

MAJOR PROBLEM: how to design/buy sound systems?

II. COMMON CRITERIA FRAMEWORK

<http://www.commoncriteriaportal.org/>

Idea:

- standardize the process of
 - designing products (Security Target ST),
 - designing requirements (Protection Profile, PP)
 - evaluation of products (licensed labs checking conformance of implementation with the documentation)
- international agreement of bodies from some countries (USA, France, UK, Germany, India, Turkey, Sweden, Spain, Australia, Canada, Malaysia, Netherlands, Korea, New Zealand, Italy, Turkey) but Israel only “consuming”, no Poland, China, Singapore,
- idea: ease the process,
- support for certification industry

Value:

- CC certification does not mean a product is secure
- it only says that it has been developed according to PP
- assurance level concerns only the stated requirements , e.g. trivial requirements \Rightarrow high EAL level (common mistake in public procurement: EAL level ... without specifying PP)
- clean up the zoo of different assumptions, descriptions, ...

Example for PP: BAC (Basic Access Control)

- encryption primitive $EM(K, S) = \text{Enc}(KB_{\text{Enc}}, S) \parallel \text{MAC}(KB_{\text{Mac}}, \text{Enc}(KB_{\text{Enc}}, S), S)$ where $K = \{KB_{\text{Enc}}, KB_{\text{Mac}}\}$
- steps:
 1. The MRTD chip sends the nonce r_{PICC} to the terminal
 2. The terminal sends the encrypted challenge $e_{\text{PCD}} = EM(K, r_{\text{PCD}}, r_{\text{PICC}}, K_{\text{PCD}})$ to the MRTD chip, where r_{PICC} is the MRTD chip’s nonce, r_{PCD} is the terminal’s randomly chosen nonce, and K_{PCD} is keying material for the generation of the session keys.
 3. The MRTD chip decrypts and verifies r_{PICC} , responds with $e_{\text{PICC}} = EM(K, r_{\text{PICC}}, r_{\text{PCD}}, K_{\text{PICC}})$
 4. The terminal decrypts and verifies r_{PCD}

- 5. both sides derive $K_{\text{Enc}}, K_{\text{Mac}}$ from master key $K_{\text{PICC}} \text{ XOR } K_{\text{PCD}}$ and sequence number derived from randoms (key derivation function)
- K derived from information available on the machine readable zone (optical)
- implementation: biometric passports.
- simple system. Really?

Common Criteria Protection Profile Machine Readable Travel Document with ICAO Application, Basic Access Control BSI-CC-PP-0055

1. Introduction

1.1 PP reference

1 Title: Protection Profile - Machine Readable Travel Document with ICAO Application and Basic Access Control (MRTD-PP)

Sponsor: Bundesamt für Sicherheit in der Informationstechnik CC Version: 3.1 (Revision 2)

Assurance Level: The minimum assurance level for this PP is EAL4 augmented.

General Status: Final

Version Number: 1.10

Registration: BSI-CC-PP-0055

Keywords: ICAO, machine readable travel document, basic access control

1.2 TOE Overview

- Target of Evaluation
- "is aimed at potential consumers who are looking through lists of evaluated TOEs/Products to find TOEs that may meet their security needs, and are supported by their hardware, software and firmware"
- important sections:
 - Usage and major security features of the TOE
 - TOE type
 - Required non-TOE hardware/software/firmware
- Definition, Type

which parts, which general purpose, which functionalities are present and which are missing, eg. ATM card with no contactless payments
- Usage and security features

crucial properties of the system (high level) and security features from the point of view of the security effect and not how it is achieved

- life cycle
the product in the whole life cycle including manufacturer and destroying
- Required non-TOE hardware/software/firmware: other components that can be crucial for evaluation

2. Conformance Claim

- CC Conformance Claim: version of CC
- PP claim: other PP taken into account in a plug-and-play way
- Package claim: which EAL package level

EAL packages:

- The CC formalizes assurance into 6 categories (the so-called "assurance classes" which are further subdivided into 27 sub-categories (the so-called "assurance families"). In each assurance family, the CC allows grading of an evaluation with respect to that assurance family.
- assurance classes:
 - development:
 - ADV_ARC - 1 1 1 1 1 1 architecture requirements
 - ADV_FSP 1 2 3 4 5 5 6 functional specifications
 - ADV_IMP - - - 1 1 2 2 implementation representation
 - ADV_INT - - - - 2 3 3 “is designed and structured such that the likelihood of flaws is reduced and that maintenance can be more readily performed without the introduction of flaws”?
 - ADV_SPM - - - - - 1 1 security policy modeling
 - ADV_TDS - 1 2 3 4 5 6 TOE design
 - guidance documents
 - AGD_OPE 1 1 1 1 1 1 1 Operational user guidance
 - AGD_PRE 1 1 1 1 1 1 1 Preparative procedures
 - life-cycle support
 - ALC_CMC 1 2 3 4 4 5 5 CM capabilities
 - ALC_CMS 1 2 3 4 5 5 5 CM scope
 - ALC_DEL - 1 1 1 1 1 1 1 Delivery
 - ALC_DVS - - 1 1 1 2 2 Development security

- ALC_FLR - - - - - Flaw remediation
 - ALC_LCD - - 1 1 1 1 2 Life-cycle definition
 - ALC_TAT - - - 1 2 3 3 Tools and techniques
- security target evaluation
- ASE_CCL 1 1 1 1 1 1 1 Conformance claims
 - ASE_ECD 1 1 1 1 1 1 1 Extended components definition
 - ASE_INT 1 1 1 1 1 1 1 ST introduction
 - ASE_OBJ 1 2 2 2 2 2 2 Security objectives
 - ASE_REQ 1 2 2 2 2 2 2 Security requirements
 - ASE_SPD - 1 1 1 1 1 1 Security problem definition
 - ASE_TSS - 1 1 1 1 1 1 TOE summary specification
- tests
- ATE_COV 1 2 2 2 3 3 Coverage
 - ATE_DPT 1 1 3 3 4 Depth
 - ATE_FUN 1 1 1 1 2 2 Functional tests
 - ATE_IND 1 2 2 2 2 3 Independent testing
- vulnerability assesment
- AVA_VAN 1 2 2 3 4 5 5 Vulnerability analysis
- for example, a product could score in the assurance family developer test coverage (ATE_COV):
 - 0: It is not known whether the developer has performed tests on the product;
 - 1: The developer has performed some tests on some interfaces of the product;
 - 2: The developer has performed some tests on all interfaces of the product;
 - 3: The developer has performed a very large amount of tests on all interfaces of the product
 - example more formal: ALC_FLR
 - ALC_FLR.1:
 - The flaw remediation procedures documentation shall describe the procedures used to track all reported security flaws in each release of the TOE.
 - The flaw remediation procedures shall require that a description of the nature and effect of each security flaw be provided, as well as the status of finding a correction to that flaw.

- The flaw remediation procedures shall require that corrective actions be identified for each of the security flaws.
 - The flaw remediation procedures documentation shall describe the methods used to provide flaw information, corrections and guidance on corrective actions to TOE users.
- ALC_FLR.2:
 - first four like before
 - The flaw remediation procedures shall describe a means by which the developer receives from TOE users reports and enquiries of suspected security flaws in the TOE.
 - The procedures for processing reported security flaws shall ensure that any reported flaws are remediated and the remediation procedures issued to TOE users.
 - The procedures for processing reported security flaws shall provide safeguards that any corrections to these security flaws do not introduce any new flaws.
 - The flaw remediation guidance shall describe a means by which TOE users report to the developer any suspected security flaws in the TOE.
- ALC_FLR.3:
 - first 5 as before
 - The flaw remediation procedures shall include a procedure requiring timely response and the automatic distribution of security flaw reports and the associated corrections to registered users who might be affected by the security flaw.
 - next 3 as before
 - The flaw remediation guidance shall describe a means by which TOE users may register with the developer, to be eligible to receive security flaw reports and corrections.
 - The flaw remediation guidance shall identify the specific points of contact for all reports and enquiries about security issues involving the TOE.
- 7 predefined ratings, called evaluation assurance levels or EALs. called EAL1 to EAL7, with EAL1 the lowest and EAL7 the highest
- Each EAL can be seen as a set of 27 numbers, one for each assurance family. EAL1 assigns a rating of 1 to 13 of the assurance families, and 0 to the other 14 assurance families, while EAL2 assigns the rating 2 to 7 assurance families, the rating 1 to 11 assurance families, and 0 to the other 9 assurance families
- monotonic: EAL $n+1$ gives at least the same assurance level as EAL n in each assurance families
- levels:
 - EAL1: Functionally Tested:
 - correct operation, no serious threats

- minimal effort from the manufacturer
- EAL2: Structurally Tested
 - delivery of design information and test results,
 - effort on the part of the developer than is consistent with good commercial practice.
- EAL3: Methodically Tested and Checked
 - maximum assurance from positive security engineering at the design stage without substantial alteration of existing sound development practices.
 - developers or users require a moderate level of independently assured security, and require a thorough investigation of the TOE and its development without substantial re-engineering.
- EAL4: Methodically Designed, Tested and Reviewed
 - maximum assurance from positive security engineering based on good commercial development practices which, though rigorous, do not require substantial specialist knowledge, skills, and other resources.
 - the highest level at which it is likely to be economically feasible to retrofit to an existing product line.
- EAL5: Semiformally Designed and Tested
- EAL6: Semiformally Verified Design and Tested
- EAL7: Formally Verified Design and Tested

CEM -Common Evaluation Methodology

- given CC documentation, EAL classification etc, perform a check
- idea: evaluation by non-experts, semi-automated, mainly paper work
- mapping:
 - assurance class \Rightarrow activity
 - assurance component \Rightarrow sub-activity
 - evaluator action element \Rightarrow action
 - developer action element \Rightarrow work-unit
 - content and presentation of evidence element \Rightarrow work unit
- responsibilities:
 - sponsor: requesting and supporting an evaluation. different agreements for the evaluation (e.g. commissioning the evaluation), providing evaluation evidence.

- developer: produces TOE, providing the evidence required for the evaluation on behalf of the sponsor.
 - evaluator: performs the evaluation tasks required in the context of an evaluation, performs the evaluation sub-activities and provides the results of the evaluation assessment to the evaluation authority.
 - evaluation authority: establishes and maintains the scheme, monitors the evaluation conducted by the evaluator, issues certification/validation reports as well as certificates based on the evaluation results
- verdicts: pass, fail, inconclusive
 - parts:
 - evaluation input task (are all documents available to perform evaluation?)
 - evaluation sub-activities
 - evaluation output task (describe the Observation Report (OR) and the Evaluation Technical Report (ETR)).
 - demonstration of the technical competence task

3 Security Problem Definition

- **Object Security Problem (OSP):** "The security problem definition defines the security problem that is to be addressed.
 - axiomatic. deriving the security problem definition outside the CC scope
 - the usefulness of the results of an evaluation strongly depends on the security problem definition.
 - spend significant resources and use well-defined processes and analyses to derive a good security problem definition.
 - good example:

Secure signature-creation devices must, by appropriate technical and operational means, ensure at the least that:

 - 1) The signature-creation-data used for signature-creation can practically occur only once, and that their secrecy is reasonably assured;
 - 2) The signature-creation-data used for signature-creation cannot, with reasonable assurance, be derived and the signature is protected against forgery using currently available technology;
 - 3) The signature-creation-data used for signature-creation can be reliably protected by the legitimate signatory against the use of others
 - **assets:** entities that someone places value upon. Examples of assets include: - contents of a file or a server; - the authenticity of votes cast in an election; - the availability of an electronic commerce process; - the ability to use an expensive printer; - access to a classified facility.
- no threat no asset

- **Threats:** threats to assets
- **Assumptions:** assumptions are acceptable, where certain properties of the TOE environment are already known,
 - but not when they are derived from specific properties of the TOE

4. Security objectives

- "The security objectives are a concise and abstract statement of the intended solution to the problem defined by the security problem definition. Their role:
 - a high-level, natural language solution of the problem;
 - divide this solution into partwise solutions, each addressing a part of the problem;
 - demonstrate that these partwise solutions form a complete solution to the problem.
- bridge between the security problem and Security Functional Requirements (SFR)
- **mapping objectives to threats:** table, each threat should be covered, each objective has to respond to some threat

answers to questions:

 - what is really needed?
 - have we forgot about something?
- **rationale:** verifiable explanation why the mapping is sound

5. Extended Component Definition

- In many cases the security requirements (see the next section) in an ST are based on components in CC Part 2 or CC Part 3.
- in some cases, there may be requirements in an ST that are not based on components in CC Part 2 or CC Part 3.
- in this case new components (extended components) need to be defined

6.1 SFR (Security Functional requirements)

- *The SFRs are a translation of the security objectives for the TOE. They are usually at a more detailed level of abstraction, but they have to be a complete translation (the security objectives must be completely addressed) and be independent of any specific technical solution (implementation). The CC requires this translation into a standardised language for several reasons: - to provide an exact description of what is to be evaluated. As security objectives for the TOE are usually formulated in natural language, translation into a standardised language enforces a more exact description of the functionality of the TOE. - to allow comparison between two STs. As different ST authors may use different terminology in describing their security objectives, the standardised language enforces using the same terminology and concepts. This allows easy comparison.*
- predefined classes:

- Logging and audit class FAU
 - Identification and authentication class FIA
 - Cryptographic operation class FCS
 - Access control families FDP_ACC, FDP_ACF
 - Information flow control families FDP_IFC, FDP_IFF
 - Management functions class FMT
 - (Technical) protection of user data families FDP_RIP, FDP_ITT, FDP_ROL
 - (Technical) protection of TSF data class FPT
 - Protection of (user) data during communication with external entities families FDP_ETC, FDP_ITC, FDP_UCT, FDP_UIT, FDP_DAU, classes FCO and FTP
- There is no translation required in the CC for the security objectives for the operational environment, because the operational environment is not evaluated
 - customizing SFRs: refinement (more requirements), selection (options), assignment (values), iterations (the same component may appear at different places with different roles)
 - rules:
 - check dependencies between SFR - In the CC Part 2 language, an SFR can have a dependency on other SFRs. This signifies that if an ST uses that SFR, it generally needs to use those other SFRs as well. This makes it much harder for the ST writer to overlook including necessary SFRs and thereby improves the completeness of the ST.
 - security objectives must follow from SFR's - Security Requirements Rationale section (Sect.6.3) in PP
 - if possible, use only standard SFR's

6.2 Security Assurance Requirements

- The SARs are a description of how the TOE is to be evaluated. This description uses a standardised language (to provide exact description, to allow comparison between two PP).
-

III. EIDAS REGULATION

goals:

- interoperability, comparable levels of trust
- merging national systems into pan-European one
- trust services, in particular: identification, authentication, signature, electronic seal, time-stamping, delivery, Web authentication
- supervision
- information about

- focused on public administration systems. However, the rules for all trust services except for closed systems (not available to anyone).

tools:

- common legal framework
- supervision system
- obligatory exchange of information about security problems
- common understanding of assurance levels

technical concept:

- Member State provides an online system enabling identification and authentication with means from the member state used abroad
- notification scheme for national systems
- if notified (some formal and technical conditions must be fulfilled), then every member state must admit it in own country within 12 months

Identification and authentication:

- eID cards – Member States are free to introduce any solution, the Regulation attempts to change it and build a common framework from a zoo of solutions
- breakthrough claimed, but likely to fail

Signature:

- electronic seal with the same conditions as electronic signature,
- the seal is aimed for legal persons
- weakening conditions for qualified electronic signatures: admitting server signatures and delegating usage of private keys

new:

- electronic registered delivery service
- Webpage authentication

Example of requirements (electronic seal):

Definition:

“electronic seal creation device” means configured software or hardware used to create an electronic seal;

“qualified electronic seal creation device” means an electronic seal creation device that meets mutatis mutandis the requirements laid down in Annex II;

Art. 36

An advanced electronic seal shall meet the following requirements:

- (a) it is uniquely linked to the creator of the seal;
- (b) it is capable of identifying the creator of the seal;
- (c) it is created using electronic seal creation data that the creator of the seal can, with a high level of confidence under its control, use for electronic seal creation; and
- (d) it is linked to the data to which it relates in such a way that any subsequent change in the data is detectable.

Annex II:

- (a) the confidentiality of the electronic signature creation data used for electronic signature creation is reasonably assured;
 - (b) the electronic signature creation data used for electronic signature creation can practically occur only once;
 - (c) the electronic signature creation data used for electronic signature creation cannot, with reasonable assurance, be derived and the electronic signature is reliably protected against forgery using currently available technology;
 - (d) the electronic signature creation data used for electronic signature creation can be reliably protected by the legitimate signatory against use by others.
2. Qualified electronic signature creation devices shall not alter the data to be signed or prevent such data from being presented to the signatory prior to signing.
 3. Generating or managing electronic signature creation data on behalf of the signatory may only be done by a qualified trust service provider.
 4. Without prejudice to point (d) of point 1, qualified trust service providers managing electronic signature creation data on behalf of the signatory may duplicate the electronic signature creation data only for back-up purposes provided the following requirements are met:
 - (a) the security of the duplicated datasets must be at the same level as for the original datasets;
 - (b) the number of duplicated datasets shall not exceed the minimum needed to ensure continuity of the service.

Art. 30

1. Conformity of qualified electronic signature creation devices with the requirements laid down in Annex II shall be certified by appropriate public or private bodies designated by Member States.

notification system:

An electronic identification scheme eligible for notification if:

- (a) issued by the notifying state
- (b) at least one service available in this state;
- (c) at least assurance level low;
- (d) ensured that the person identification data is given to the right person
- (e) ...
- (f) availability of authentication online, for interaction with foreign systems (free of charge for public services), no specific disproportionate technical requirements
- (g) description of that scheme published 6 months in advance
- (h) meets the requirements from the implementing act

Assurance levels:

- regulation, Sept. 2015, implementation of eIDAS
- reliability and quality of
 - enrolment
 - electronic identification means management
 - authentication
 - management and organization
- authentication factors
 - possession based
 - knowledge based
 - inherent (physical properties)
- enrolment: (for all levels):
 1. Ensure the applicant is aware of the terms and conditions related to the use of the electronic identification means.
 2. Ensure the applicant is aware of recommended security precautions related to the electronic identification means.
 3. Collect the relevant identity data required for identity proofing and verification.
- identity proofing and verification (for natural persons):

low:

1. The person can be assumed to be in possession of evidence recognised by the Member State in which the application for the electronic identity means is being made and representing the claimed identity.
2. The evidence can be assumed to be genuine, or to exist according to an authoritative source and the evidence appears to be valid.
3. It is known by an authoritative source that the claimed identity exists and it may be assumed that the person claiming the identity is one and the same.

substantial: low plus:

1. The person has been verified to be in possession of evidence recognised by the Member State in which the application for the electronic identity means is being made and representing the claimed identity

and

the evidence is checked to determine that it is genuine; or, according to an authoritative source, it is known to exist and relates to a real person

and

steps have been taken to minimise the risk that the person's identity is not the claimed identity, taking into account for instance the risk of lost, stolen, suspended, revoked or expired evidence; or

2. options related to other trustful sources

high: substantial plus

(a) Where the person has been verified to be in possession of photo or biometric identification evidence recognised by the Member State in which the application for the electronic identity means is being made and that evidence represents the claimed identity, the evidence is checked to determine that it is valid according to an authoritative source; and the applicant is identified as the claimed identity through comparison of one or more physical characteristic of the person with an authoritative source; or

- electronic identification means management:

low:

1. The electronic identification means utilises at least one authentication factor.
2. The electronic identification means is designed so that the issuer takes reasonable steps to check that it is used only under the control or possession of the person to whom it belongs.

substantial:

1. The electronic identification means utilises at least two authentication factors from different categories.
2. The electronic identification means is designed so that it can be assumed to be used only if under the control or possession of the person to whom it belongs.

high:

1. The electronic identification means protects against duplication and tampering as well as against attackers with high attack potential
2. The electronic identification means is designed so that it can be reliably protected by the person to whom it belongs against use by others.

- Issuance , delivery and activation:

low:

After issuance, the electronic identification means is delivered via a mechanism by which it can be assumed to reach only the intended person.

substantial:

After issuance, the electronic identification means is delivered via a mechanism by which it can be assumed that it is delivered only into the possession of the person to whom it belongs.

high:

The activation process verifies that the electronic identification means was delivered only into the possession of the person to whom it belongs.

- suspension, revocation and reactivation:

all levels:

1. It is possible to suspend and/or revoke an electronic identification means in a timely and effective manner.
2. The existence of measures taken to prevent unauthorised suspension, revocation and/or reactivation.
3. Reactivation shall take place only if the same assurance requirements as established before the suspension or revocation continue to be met.

- authentication mechanism:

substantial:

1. The release of person identification data is preceded by reliable verification of the electronic identification means and its validity.
2. Where person identification data is stored as part of the authentication mechanism, that information is secured in order to protect against loss and against compromise, including analysis offline.
3. The authentication mechanism implements security controls for the verification of the electronic identification means, so that it is highly unlikely that activities such as guessing, eavesdropping, replay or manipulation of communication by an attacker with enhanced-**basic attack potential** can subvert the authentication mechanisms.

high:

... by an attacker with **high attack potential** can subvert the authentication mechanisms.

- audit:

low:

The existence of periodical internal audits scoped to include all parts relevant to the supply of the provided services to ensure compliance with relevant policy.

substantial:

The existence of periodical **independent** internal or external audits

high:

1. The existence of periodical **independent external audits** scoped to include all parts relevant to the supply of the provided services to ensure compliance with relevant policy.
2. Where a scheme is directly managed by a government body, it is audited in accordance with the national law.

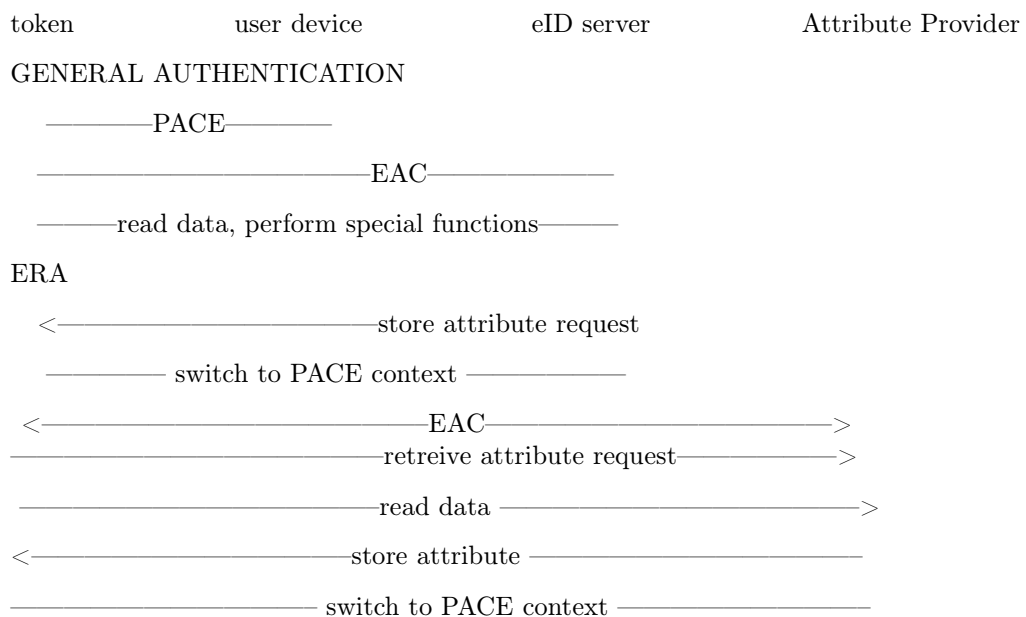
IV. eIDAS TOKEN SPECIFICATION, BSI

- Technical guideline, security mechanisms for electronic travel documents, not focused on readers
- cryptographic mechanisms:
 - Password Authenticated Connection Establishment (PACE)
 - Terminal Authentication Version 2 (TA2).
 - Chip Authentication Version 3 (CA3)
 - Restricted Identification (RI)
 - Pseudonymous Signatures (PS)
- procedures
 - General Authentication Procedure (GAP)

- Enhanced Role Authentication (ERA)
- PIN Management
- terminal types:
 - inspection system
 - authentication terminal - government or private, terminal rights to be checked, GAP must be used
 - attribute terminal- extension of Authentication Terminal, ERA must be used
 - signature management terminal - key creation, signature creation
 - signature terminal - GAP must be used
 - privileged terminals: category: inspection terminals and some authentication terminals explicitly authorized. Signature terminals are never privileged
- user credentials:
 - MRZ-Password
 - CAN
 - PIN - always blocking (RC reaches 0 then blocked)
 - PUK - blocking or non-blocking
- password blocking: RC=0 password blocked, RC=1 - password suspended and the correct CAN must be entered during the same session to resume the password. Resume is volatile.
- switching session context: a stack of protocols, when terminating a protocol we return to the context on the top of the stack
- password authentication:
 - PACE- global passwords, VERIFY-application local
 - Inspection terminal SHALL use CAN or MRZ
 - authentication terminal SHALL use PIN, but the CAN can be allowed by the terminal
 - signature terminal: PIN, CAN or PUK
- Extended Access Control:
 - 1. Terminal Authentication v2: terminal SHALL generate ephemeral keys used later for Chip Authentication, only standard parameters, ephemeral keys authenticated, result: read/write access granted
 - 2. Passive Authentication: terminal reads and verifies Security Objects, compares the data obtained before PACE

- 3. Chip Authentication v2 or v3: afterwards secure channel restarted
- General Authentication Procedure:
 - i. password verification - PACE
 - ii. EAC
 - iii. read/write data
- Enhanced Role Authentication – authentication terminal with proper rules can proceed as follows:
 - i. authentication terminal sends an ATTRIBUTE REQUEST to eIDAS token. token makes a link between the request and the terminal's sector
 - ii. restore session context of PACE, store context of Chip authentication
 - iii. EAC with attribute provider
 - iv. proceed attribute request, write the resulting attributes to the eIDAS token, the access rights restricted to terminals with proper rights
 - v. restore session context: PACE, then Chip Authentication
 - vi. terminal may read the stored attributes
- online authentication:
 - eID server: remote part of authentication terminal
 - user device: interacts with user, eIDAS token and eID server, but not authorized to read eIDAS token data, access rights only after authentication with the eID server

Protocol chart:



— switch to e-ID session context —
—————read attribute—————>

- unauthenticated terminals:
 - i. password verification based on PACE:
 - terminal does not show its type
 - can choose password type
 - after authentication secure messaging
 - ii. authentication with CAN resumes PIN
 - iii. updating retry counter

- authenticated terminals: after terminal authentication the terminal becomes authenticated

Cryptographic building blocks:

- hash $H(m)$
- compression function for public key: $Comp(PK)$
- projected representation of a public key $\Pi(PK)$
- symmetric key algorithms:
 - deriving key for encryption $K_{Enc} = KDF_{Enc}(K, [r])$
 - $K_{Mac} = KDF_{Mac}(K, [r])$
 - $K_{\pi} = KDF_{\pi}(\pi)$
 - encryption and decryption
 - MAC
- asymmetric algorithms:
 - domain parameters
 - keys (page 19):
 - eIDAS: ephemeral on both sides
 - Chip authentication: static on side of the chip
 - Chip authentication version 3: ephemeral on both sides based on static Chip's key

- Restricted Identification: token uses a static key, sector public key, sector specific identifier
- KA - key agreement (like DH)
- signatures, mapping to RSA and ECDSA described

Pseudonymous signature:

- used for anonymous signature and for Chip Authentication v3
- keys:
 - domain parameters D_M and a pair of global keys (PK_M, SK_M)
 - public key PK_{ICC} for a group of eIDAS tokens, the private key SK_{ICC} known to the issuer of eIDAS tokens (called manager)
 - for a token the manager chooses $SK_{ICC,2}$ at random, then computes $SK_{ICC,1}$ such that $SK_{ICC} = SK_{ICC,1} + SK_M \cdot SK_{ICC,2}$
 - a sector (domain) holds private key SK_{sector} and public key PK_{sector} .
 - a sector has revocation private key $SK_{revocation}$ and public key $PK_{revocation}$
 - sector specific identifiers $I_{ICC,1}^{sector}$ and $I_{ICC,2}^{sector}$ of the eIDAS token in the sector
- signing: with keys $SK_{ICC,1}$, $SK_{ICC,2}$ and $I_{ICC,1}^{sector}$ and $I_{ICC,2}^{sector}$ for PK_{sector} and message m
 - i. choose K_1, K_2 at random
 - ii. compute
 - $Q_1 = g^{K_1} \cdot (PK_M)^{K_2}$
 - $A_1 = (PK_{sector})^{K_1}$
 - $A_2 = (PK_{sector})^{K_2}$
 - iii. $c = \text{Hash}(Q_1, I_{ICC,1}^{sector}, A_1, I_{ICC,2}^{sector}, A_2, PK_{sector}, m)$ (variant parameters and Π omitted here)
 - iv. compute
 - $s_1 = K_1 - c \cdot SK_{ICC,1}$
 - $s_2 = K_2 - c \cdot SK_{ICC,2}$
 - v. output (c, s_1, s_2)
- verification:

compute

 - $Q_1 = (PK_{ICC})^c \cdot g^{s_1} \cdot (PK_M)^{s_2}$
 - $A_1 = (I_{ICC,1}^{sector})^c \cdot (PK_{sector})^{s_1}$

- $A_2 = (I_{\text{ICC},2}^{\text{sector}})^c \cdot (\text{PK}_{\text{sector}})^{s_2}$
- recompute c and check against the c from the signature
- why it works?

$$\begin{aligned} (\text{PK}_{\text{ICC}})^c \cdot g^{s_1} \cdot (\text{PK}_M)^{s_2} &= (\text{PK}_{\text{ICC}})^c \cdot g^{K_1} \cdot (\text{PK}_M)^{K_2} \cdot g^{-c\text{SK}_{\text{ICC},1}} \cdot (\text{PK}_M)^{c\text{SK}_{\text{ICC},2}} \\ &= (\text{PK}_{\text{ICC}})^c \cdot g^{K_1} \cdot (\text{PK}_M)^{K_2} \cdot g^{-c\text{SK}_{\text{ICC},1}} \cdot (g)^{-c\text{SK}_M \cdot \text{SK}_{\text{ICC},2}} \\ &= (\text{PK}_{\text{ICC}})^c \cdot g^{K_1} \cdot (\text{PK}_M)^{K_2} \cdot g^{-c \cdot \text{SK}_{\text{ICC}}} = g^{K_1} \cdot (\text{PK}_M)^{K_2} = Q_1 \end{aligned}$$
- there is a version without A_1, A_2 and the pseudonyms $I_{\text{ICC},1}^{\text{sector}}, I_{\text{ICC},2}^{\text{sector}}$

PACE (Password Authenticated Connection Establishment)

- ICAO Doc 9303: Basic Access Control/PACE and EAC v1 (=Chip Authentication v1+ Terminal Authentication v1) MUST be used
- password based authentication protocol
- password on the side of the token: stored, on the terminal: input by the user
- steps:
 - i. token chooses s at random
 - ii. token computes $z = \text{Enc}(K_\pi, s)$, where $K_\pi = \text{KDF}(\pi)$ and sends z to the reader together with the parameters D_{PICC}
 - iii. the reader recovers s
 - iv. the reader and the token compute $D_{\text{Mapped}} = \text{Map}(D_{\text{PICC}}, s)$ (mapping function)
 - v. the reader and the token perform anonymous Diffie-Hellman key agreement based on the ephemeral domain parameters (ephemeral values based on D_{Mapped} as a generator), shared secret K obtained
 - vi. they create session keys $K_{\text{Mac}} = \text{KDF}_{\text{Mac}}(K)$ and $K_{\text{Enc}} = \text{KDF}_{\text{Enc}}(K)$
 - vii. exchange and verification of tokens: $T_{\text{PCD}} = \text{MAC}(K_{\text{MAC}}, \text{ephemeral key of PICC})$
 $T_{\text{PICC}} = \text{MAC}(K_{\text{MAC}}, \text{ephemeral key of PCD})$
 - viii. Secure Messaging restarted

Terminal authentication v2

- Chip Authentication MUST be performed after Terminal Authentication (condition repeated in the description of CHA v2 only)
- simple challenge-response algorithm, undeniable, resistant to replay
- ephemeral public key for ChA as a side effect

- steps:
 - i. the terminal send the certificate chain to eIDAS token, it has to confirm the key PK_{PCD}
 - ii. the token checks PK_{PCD}
 - iii. the terminal creates ephemeral pair of keys, sends the compressed version of $\widetilde{PK_{PCD}^{CA}}$ to the token
 - iv. token replies with a random nonce r_{PICC}
 - v. the terminal signs with SK_{PCD} the following data: r_{PICC} , compressed version of $\widetilde{PK_{PCD}^{CA}}$
 - vi. the token checks the signature

Chip authentication v2

- static DH authentication with the ephemeral key of the terminal
- steps:
 - i. the token sends its public key PK_{PICC}
 - ii. the terminal sends ephemeral public key from TA (uncompressed)
 - iii. static DH key agreement with SK_{PICC} and ephemeral public key on side of the token, and PK_{PICC} and ephemeral secret key on side of the terminal, master key K generated
 - iv. token chooses r_{PICC} , computes $K_{Enc} = KDF_{Enc}(K, r_{PICC})$, $K_{Mac} = KDF_{Mac}(K, r_{PICC})$
 - v. token computes the tag $T_{PICC} = MAC(K_{Mac}, \text{ephemeral public key of the terminal})$
 - vi. the terminal checks the tag
 - vii. secure messaging restarted using K_{Enc} and K_{Mac}

Chip authentication v3

- alternative to Chip authentication v2 and RI
- claimed: “message-deniable strong authentication”, “pseudonymity without using the same key on several chips”, “possibility of whitelisting eIDAS tokens”
- scheme:
 - i. phase 0: terminal authentication, ephemeral key for terminal in phase 1 chosen and signed
 - ii. phase 1: key agreement like DH with ephemeral keys on both sides, restarting secure messaging with new keys

iii. phase 2:

- static keys on the side of the chip: $SK_{ICC,1}$, $SK_{ICC,2}$, PK_{ICC} and the parameters
 - terminal sends PK_{sector} to the chip, the chip compares it with the “compressed” version received during Terminal Authentication
 - chip reconstructs $I_{ICC,1}^{sector} = (PK_{sector})^{SK_{ICC,1}}$ and $I_{ICC,2}^{sector} = (PK_{sector})^{SK_{ICC,2}}$
 - chip creates pseudonymous signature using $I_{ICC,1}$, $I_{ICC,2}$ as pseudonym and the secret keys $SK_{ICC,1}$, $SK_{ICC,2}$ over the ephemeral key given by the terminal
- If PACE GM used before ChA v3 then one can reuse the ephemeral key from the terminal
 - checking the key PK_M is obligatory (otherwise it would be easy to forge the token)

Restricted Identification

- optional
- depending on the version, deanonymization might be possible or not (depending on PK_{sector})
- executed after Terminal Authentication and Chip Authentication (not specified which version, but with v3 it does not makes sense)
- sector specific identifier computed as $\text{Hash}(\text{key computed via DH from } PK_{sector} \text{ and } SK_{ID})$
- blacklisting impossible in case of group key compromise (from ChA v2)

Pseudonymous Signature as replacement of RI

- whitelisting possible in case of group key compromise (claimed as new but possible for RI)
- the second part from ChA v3, the key PK_{sector} used as sector public key

PSA - Pseudonymous Signature Authentication

- the sector public key = the ephemeral public key from ephemeral DH key agreement (now DH explicitly mentioned)

PSM - Pseudonymous Signature of a Message

- TA and ChA must be executed before
- message to be signed comes from the terminal
- public key unspecified

PSC - Pseudonymous Signature of Credentials

- used in combination with ERA
- Attribute Terminal involved, but eIDAS token creates the signature himself (after breaking group key one can also create the PSC)

- public key unspecified
- terminal rights to get the attributes are to be checked

PROBLEMS:

- security properties not stated, they can be derived via tedious analysis
- lack of security proofs
- underspecified (details may turn the token to be insecure)
- powerful adversary able to break into the token may crate fake ID's, unless whitelist approach used

V. STANDARDS VERSUS SECURITY

Bleichenbacherr's RSA signature forgery based on implementation error

The attack works for PKCS-1 padding.

The PKCS-1 padding consists of a byte of 0, then 1, then a string of 0xFF bytes, then a byte of zero, then the "payload" which is the hash+ASN.1 data.

Graphically:

```
00 01 FF FF FF ... FF 00 ASN.1 HASH
```

The signature verifier first applies the RSA public exponent to reveal this PKCS-1 padded data, checks and removes the PKCS-1 padding, then compares the hash with its own hash value computed over the signed data.

The error that Bleichenbacher exploits is if the implementation does not check that the hash+ASN.1 data is right-justified within the PKCS-1 padding. Some implementations remove the PKCS-1 padding by looking for the high bytes of 0 and 1, then the 0xFF bytes, then the zero byte; and then they start parsing the ASN.1 data and hash. The ASN.1 data encodes the length of the hash within it, so this tells them how big the hash value is. These broken implementations go ahead and use the hash, without verifying that there is no more data after it. Failing to add this extra check makes implementations vulnerable to a signature forgery, as follows.

An attacker forges the RSA signature for an exponent of 3 by constructing a value which is a perfect cube. Then he can use its cube root as the RSA signature. He starts by putting the ASN.1+hash in the middle of

the data field instead of at the right side as it should be. Graphically:

```
00 01 FF FF ... FF 00 ASN.1 HASH GARBAGE
```

This gives him complete freedom to put anything he wants to the right of the hash. This gives him enough flexibility that he can arrange for the value to be a perfect cube.

There are some other variations of this attack, for example some implementations of the signature verification algorithm neglect to check if the field "parameters" inside "digestAlgorithm" field is NULL. In such a case an attacker may put some GARBAGE here, making the attack still possible even if the algorithm verifies that the HASH is right-justified.

Chosen Ciphertext Attacks Against Protocols Based on RSA Encryption Standard PKCS-1 – the Million Message Attack.

An attacker has access to an oracle that, for any chosen ciphertext, indicates whether the corresponding plaintext $C^d \bmod n$ has the correct format according to the encryption standard.

Then after querying the oracle with about 2^{20} adaptively chosen ciphertexts the attacker is able to calculate the plaintext corresponding to the original ciphertext.

The oracle might be for example a HSM that receives ciphertexts resulting from the key-wrap algorithm.

The attack exploits such vulnerabilities like

- different error messages returned by the attacked device when decryption fails on different stages of the decryption algorithm
- different timings of execution of the decryption algorithm when the PKCS-1 encryption padding is correct and when it is incorrect.

If a device supports the PKCS-1 encryption padding and the implementation of the PKCS-1 decryption on the device is vulnerable, then the million message attack works also when

- the ciphertext is calculated according to a padding different than PKCS-1
- the “ciphertext” is the plaintext for which we want to obtain a signature (dangerous for a situation when the same key is used for decryption and for signatures, and decryption is not PIN protected).

VI. FORMAL SECURITY PROOFS

Security model e.g. for PACE

Background (Hanzlik, MK)

data confidentiality: nobody can understand any data from the communication between an eID and a terminal, except for this eID and this terminal. By “data” we mean:

- workload data to be transmitted via the channel established according to the protocol,
- partner specific data (such as partner identity) - if sending them (explicitly or implicitly) results from the protocol execution.

data integrity: a third person cannot manipulate without detection the data exchanged between the eID and the terminal. This concerns in particular manipulating identity data.

session integrity: if a party A accepts at some moment a session executed presumably with a single partner, then indeed this interaction of A emerged in interaction with a single partner.

partner authentication: if a partner A accepts a session as a session with party C, then A indeed has been talking with C until this moment (maybe with somebody playing man-in-the-middle, but only passively). Partner authentication might be mutual or one-sided. In case of PACE, there is one-sided authentication of an eID.

owner’s consent: eID is used only when the user agrees and with the terminal chosen by the user.

proof non-transferability: a party A interacting with a party B cannot prove against a third party C that it is interacting with B, and cannot authenticate in this way the data received from B. This should be understood that executing the protocol does not provide additional cryptographic evidence over the data mentioned in the data confidentiality condition.

Case study: KEA

- Diffie-Hellman based key-exchange protocol, mutual authentication for the parties
- developed by NSA, declassified in 1998, no security analysis
- attacked in 2005, Lauter, Mityagin, extension KEA+ proposed, security proven by reduction proofs
- naive protocol:
 - party A chooses x at random and sends to B:
 - g^x and $\text{sign}_A(g^x, B)$
 - party B chooses y at random and sends to BA:
 - g^y and $\text{sign}_B(g^y, A)$
 - both: verify the signature, compute $g^{x \cdot y}$ as for DH protocol
 - attack:
 - if ephemeral x of A from communication between A and B revealed, then ...
 - the adversary resends g^x and $\text{sign}_A(g^x, B)$ to C and can impersonate A as he can compute the session key
- KEA:
 - A and B hold, respectively, the keys: private a and b , and public keys g^a and g^b
 - A and B select ephemeral secret keys x and y at random and exchange g^x and g^y
 - each party computes $g^{a \cdot y}$ and $g^{b \cdot x}$ (static DH protocol)
 - session key computed as $F(g^{a \cdot y} \text{ xor } g^{b \cdot x})$ (just like Blake-Wilson, D. Johnson, and A. Menezes: $\text{Hash}(g^{a \cdot y}, g^{b \cdot x})$)
- Unknown Key Share (UKS) – a formal attack on KEA:
 - Mallet registers the same key g^a as Alice
 - Alice starts a session with Bob but session intercepted by Mallet
 - Mallet starts a session with Bob as Mallet
 - Mallet forwards the values g^x and g^y
 - therefore Alice and Bob compute the same session key

- Mallet corrupts one session and get a session key for the second one - contradicting AKE security
- KEA+
 - session key computed as $F(g^{a \cdot y}, g^{b \cdot x}, A, B)$
- KEA+C
 - keys as for KEA
 - A chooses x at random and sends g^x
 - B chooses y at random, computes $L = \text{Hash}(g^{a \cdot y}, g^{b \cdot x}, A, B)$
 - B responds with g^y and $\text{MAC}_L(0)$
 - A computes L , checks $\text{MAC}_L(0)$ and responds with $\text{MAC}_L(1)$
 - B checks $\text{MAC}_L(1)$
- security properties:
 - AKE (Authenticated Key Exchange)-
 - the adversary controls all communication
 - the adversary can corrupt some of the parties.
 - the adversary must select an uncorrupted session called a test session and then he is given a challenge, which is either the session key of the test session or a randomly selected key.
 - the adversary wins if can distinguish between these 2 cases.
 - PFS (Perfect Forward Security):
 - AKE experiment
 - the adversary can corrupt a party A (reveal the long-term secret key),
 - test session: a session of A occurred before corrupting A
 - KCI (Key Compromise Impersonation)
 - the adversary gets a long-term secret key of A
 - attempt to impersonate as other party to A
 - of course, the adversary can impersonate A to anyone
 - advantage of the adversary A running algorithm \mathcal{A} :

$$|\Pr(\mathcal{A}(\text{data}, \text{real key}) = 1) - \Pr(\mathcal{A}(\text{data}, \text{random key}))|$$
 the advantage should be “negligibly small”

- reduction proofs:
 - assume that there is an adversary A breaking scheme U
 - choose a cryptographic assumption P
 - from a case p for P construct a case u for U
 - show how to run A on u
 - the environment need not to behave exactly as the scheme U
 - the difference between real U and the simulated one should be impossible to detect by A
 - breaking u should lead to breaking p with a fair probability
 - finally: compute the advantage of the resulting adversary breaking p
- modelling via oracles:
 - atomic actions that can be initiated by the adversary
 - all interactions with the system defined by the oracles
 - specification of adversary's power
- typical oracles:
 - Reveal: reveal ephemeral key
 - Reveal: reveal session key
 - Corrupt: reveal long-time key
 - Execute(A, B): make A and B execute the protocol
 - Send: send a message to A and get its reaction (if any) – the messages may come from the protocol, but might be faulty
 - Test: a session ends after key establishment, no workload communication (this can be added with the tested key), must concern a *fresh session*
 - *fresh session*: exclude situation where for instance via corruptions it is possible to break the session
- AKE for KEA+:
 - reduction via Gap Diffie-Hellman (CDH under assumption that DDH easy)
 - ROM for hash function
 - ways to distinguish between the random from real key: hash value must be asked
 - possibilities for the real key K to appear in the experiment:
 1. Forging: enforce Hash on the tuple $(\text{CDH}(A, Y), \text{CDH}(B, X), A, B)$

- 2. Key-replication attack. succeed to create another session with the same “signature” ($\text{CDH}(A, Y), \text{CDH}(B, X), A, B$) and so the same secret key
- key replication: impossible, since $A' = A$ and $\text{CDH}(A', Y') = \text{CDH}(A, Y)$ implies $Y = Y'$. Similarly $X = X'$ and the sessions are identical
- forging: case of a single session:
 - adversary observes a single session between honest A and B
 - problem GDH for (X_0, Y_0)
 - the long term key of A chosen as X_0 , the response of B chosen as Y_0 , the rest executed as in the scheme description
 - learning the key requires asking hash oracle about $(\text{CDH}(X_0, Y_0), g^{b \cdot x}, A, B)$
- forging the in general case: problem since A involved in many interactions but we do not know the secret key. Idea: replace with a random key
 - all users initialized according to the scheme, except for A
 - Hash simulated by HSim
 - sessions not involving A executed according to the protocol (and HSim)
 - a session (A, C, role) :
 - C public key of C
 - if A initiator, then it chooses x at random, sends g^x , gets reply Y , session key $\text{HSpec}(1, Y, C^x, A, C)$
 - if A responder, then it waits for X , chooses y at random, sends g^y , gets reply Y , session key $\text{HSpec}(2, X, C^y, C, A)$
 - a session (C, A, role) :
 - as in the scheme description
 - except for test session where Y_0 sent and the session key not computed
- reveal and corrupt key: as described by the scheme
- $\text{HSim}(Z_1, Z_2, B, C)$ – random oracle on valid signatures
 - if asked before, then repeat the answer
 - check all previous $\text{HSpec}(i, Y, Z, B, C) = v$ and check if $Z = Z_{3-i}$ and $\text{DDH}(X_0, Y, Z_i) = \text{true}$. If yes, then return v .
 - if not found then return random w and remember it
- $\text{HSpec}(i, Y, Z, B, C)$ - random oracle for cases when adversary does not know the secret key of A . For input (Z_1, Z_2, B, C) , where $Z_i = \text{CDH}(X_0, Y)$ and $Z_{3-i} = Z$

VII. CATACRYPT

catastrophy cryptography

- what happens if assumptions broken (e.g. DL solvable for some group)?
- "post-quantum crypto"

reality:

- post-quantum is at early stage, no industrial products, logistically impossible to replace
- no plans, scenarios, ...
- catastrophe is already there

TLS and DH real security

mistakes:

- risk of common (standard) groups
- cryptanalysis: most efficient number field sieve (NFS):
 - complexity subexponential (for \mathbb{Z}_p it is
$$\exp(1.93 + o(1))(\log p)^{1/3}(\log \log p)^{2/3}$$
)
 - most time precomputation independent from the target number y (where $\log y$ to be computed in a given group)
 - the time dependant from y can be optimized to subexponential but much lower
 - 512-bit groups can be broken, MitM attack can be mounted
- standard safe primes – seem to be ok, but attacker can amortize the cost over many attacks
- TLS with DH: frequently "export-grade" DH with 512 bit primes, about 5% of servers support DHE_EXPORT, most servers (90% and more) use a few primes of a given length, after a precomputation breaking for a given prime: reported as 90 sec
- TLS: client wants DHE, server offers DHE_EXPORT, but one can manipulate the messages exchanged, so that the client treats the (p_{512}, g, g^b) as a response to DHE – it is not an implementation bug!
 - handshake time is a problem, but some protocols allow. sending TLS warning alert that reset the countdown
 - ephemeral key hashing

- sometimes non safe prime used ($\frac{p-1}{2}$ composite), Pohling-Hellman method can be used
 - DH-768 breakable on academic level, DH-1024 on the state level
 - recommendations:
 - avoid fixed prime groups
 - transition to EC
 - deliberately do not downgrade security even if seems to be ok
 - follow the progress in computer algebra
-

VIII. HARDWARE TROJANS

methods of testing:

- functional tests
- internal tests circuitry
- optical inspection (destructive) - can detect modifications on layout level

Idea: change properties that are not visible under microscope: increase aging effects, manipulate transistors so that the output is fixed

Dopant Trojans

CMOS inverter: (image Wikipedia)



Figure 1.

where: A is the source, Vdd positive supply , Vss is ground
 upper transistor: PMOS (allows current flow at low voltage)

lower transistor: NMOS (allows current flow at high voltage)

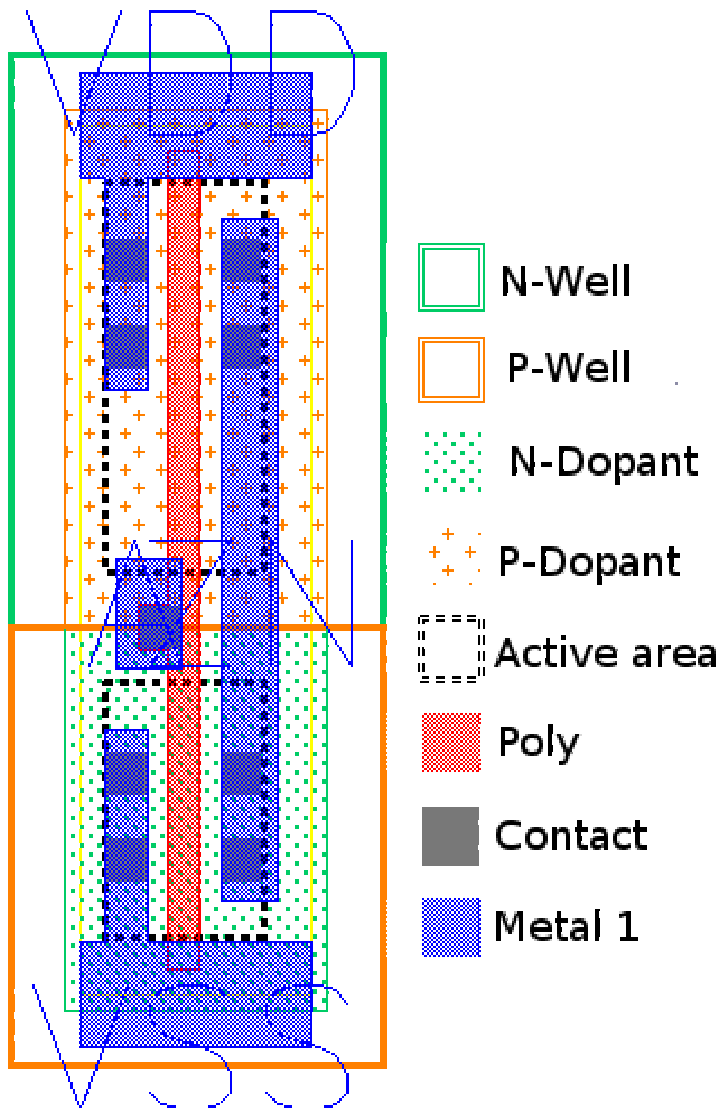
how it works:

- if voltage is low then the lower transistor is in high resistance state and the current from Q flows to Vdd (high voltage)
- if voltage is high then the upper transistor is in high resistance state and the current from Q flows to Vss while Vdd has low voltage

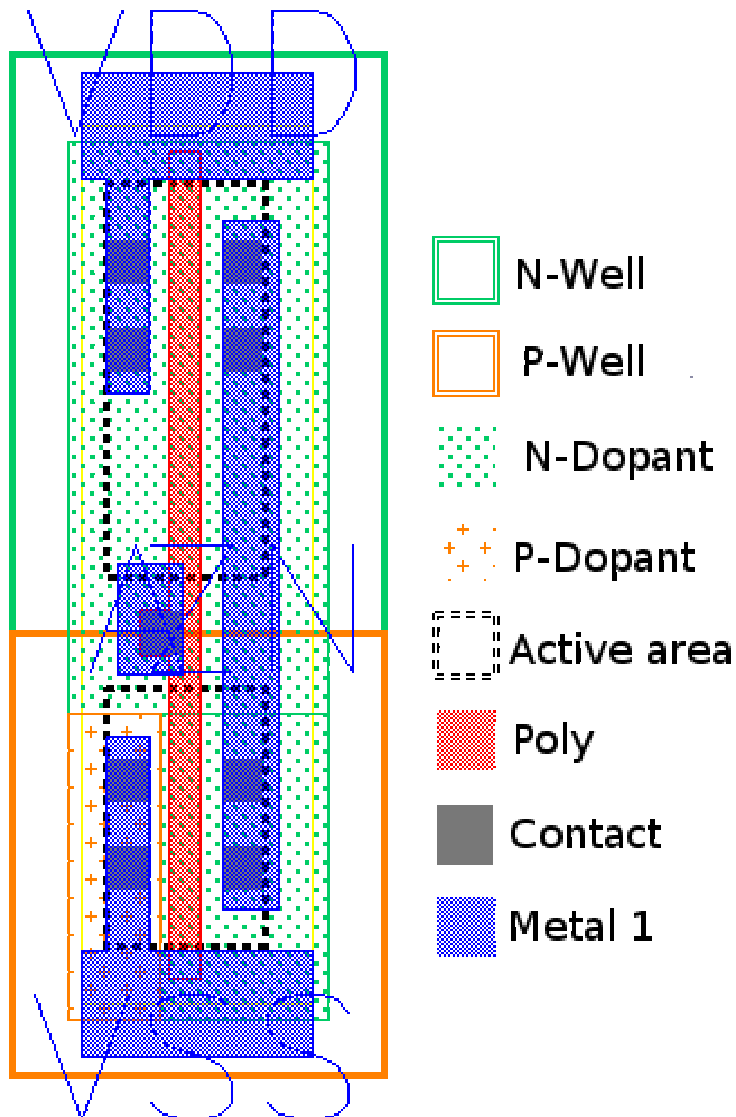
PMOS: in dopant area “holes” (positive) playing the role of conductor, low voltage creates depletion area, high voltage attracts them

NMOS: in dopant area electrons (negative) playing the role of conductor, high voltage pushes the electrons out

CMOS inverter in the “bird eye perspective”:



Trojan design:



- whatever happens the VDD is connected to the output

Trojan TRNG

TRNG consists of

- entropy source (physical)
- self test circuit (OHT - inline health test)
- deterministic RNG, Intel version:
 - conditioner (computes seeds to rate matcher) and rate matcher (computes 128 bit numbers)
 - derivation, internal state (K, c) :

1. $c := c + 1, r := \text{AES}_K(c)$

2. $c := c + 1, x := \text{AES}_K(c)$

3. $c := c + 1, y := \text{AES}_K(c)$

4. $K := K \oplus x$

5. $c := c \oplus y$

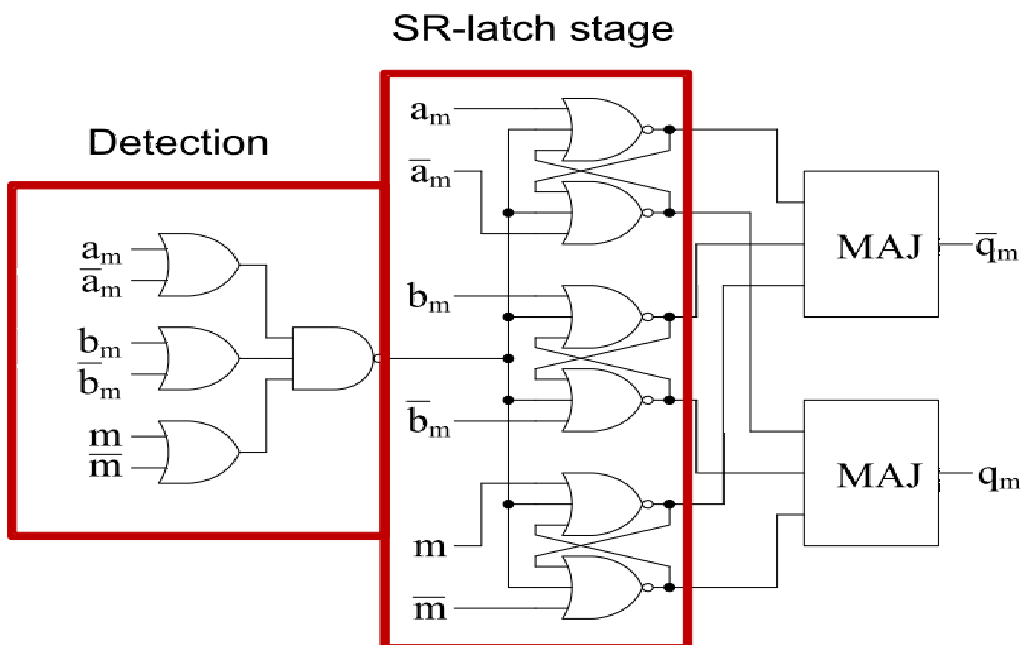
- attack: fix K by applying Trojan transistors, if K is known, then it is easy to find internal state c from r and then the consecutive random numbers r
- problem with OHT: tests with some values have to create known outputs (32 CRC from the last 4 outputs), knowing the test one can find K by exhaustive search

Side channel Trojan:

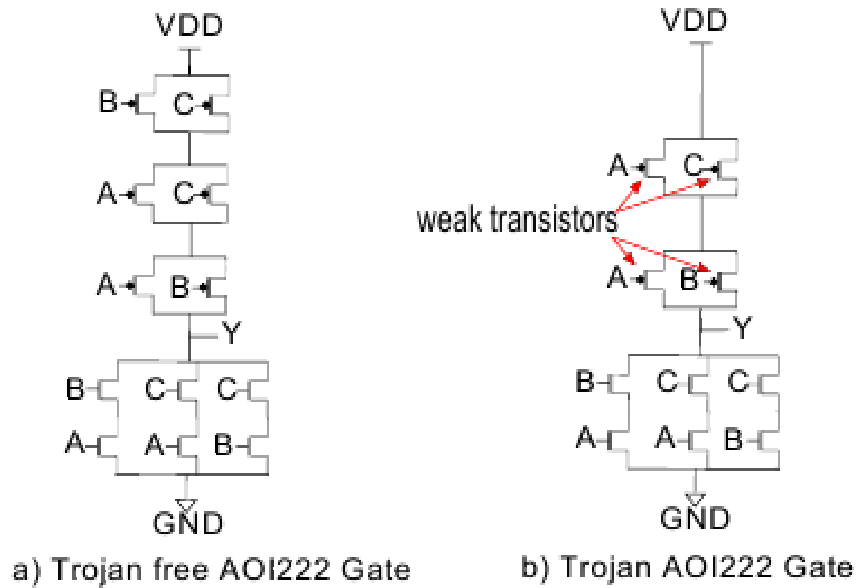
- side channel resistant logic: Masked Dual Rail Logic
 - i. for each a both a and negation of a computed
 - ii. precharge: each phase preceded by charging all gates
 - iii. masking operations by random numbers:

computing $a \wedge b$:

- input $a \oplus m, a \oplus \bar{m}, b \oplus m, b \oplus \bar{m}, m, \bar{m}$
- detection, SR-latch stage and majority gat



attacking not-majority gate:



Idea: instead of cutting output a low voltage

- the same behavior except for $A = 0$ and $B, C = 1$, where good output but high power consumption due to connection between VDD and VSS

Defense methods:

- problem: Trojan may be triggered by some particular event, detection becomes harder
- problem: Trojan may work in very particular physical conditions, e.g. temperature, voltage
- on-chip checks: detection of unexpected behavior, e.g. delay characteristics: workload path and a shadow path that provides result after fixed time, + comparison
- methods to enable activation in certain areas only
- inserting PUFs, (either randomize as much as possible - noise over trojan information) or keep deterministic

VIII. HOW TO CREATE A SYSTEM IN THE WORST POSSIBLE WAY?

Example: CHIP AUTHENTICATION PROGRAMME

“optimisation is the process of taking something that works and replacing it with something that almost works but is cheaper”

CAP - idea:

- cheap, Chip&PiN device
- keyboard, display, chip reader
- protecting PIN (it does not go to the PC)

- CAP device not personalized, one can use own card with CAP in a bank, or borrow from somebody
- recommended to use own CAP device
- optimized to be as much as possible based on EMV (standard for electronic purse)

used in UK

operation modes:

- identify: returns one-time code (like RSA-token) (based on symmetric key)
- respond: responds to a challenge using symmetric key
- sign: just as respond, however takes account number and value to generate the response

Protocol overview:

1. select application of the card (CAP has some fixed identifiers)
2. read records: account number, certificates , ... but important: CDOL1, CDOL2 (card object lists) and CAP (bit filter defining the protocol execution)
3. PIN verification
4. ciphertext generation: GENERATE AC command, response: Authorisation Request Cryptogram (ARQC), then the reader asks for Application Authentication Cryptogram (AAC) indicating cancelling the transaction (according to EMV)

challenge: AA (Authorized Amount), UN (Unpredictable number)

- for identify both are 0
- for respond: AA=0, UN=challenge
- for sign: AA=transaction value, UN= destination account

Response:

- based on the following data: ATC (application transaction counter), CID (Cryptogram Identification Data), IAD (Issuer Application Data - contains result of PIN verification), AC (Application Cryptogram - MAC (3DES CBC MAC) of the rest)
- CAP filter used to determine which bits to take
- NatWest: 5 least significant bits of ATC, 20 least significant bits of MAC, 1 bit from IAD
- Barclay: top bit of CID, 8 least significant bits of ATC, 17 least significant bits of MAC
- HBOS: top bit of CID, 7 least significant bits of ATC, 17 bits of MAC (not in one block), 1 bit from IAD

Verification: recomputed with the secret key shared with the card

Application:

- bank decides how to use (mode + semantic field)

- NatWest: respond mode, 8 bits of challenge, 4 random, 4 =last 4 digits of destination account, not used for login, transaction value not authenticated
- Barclay: identify necessary for login, for transaction: sign with destination account and transaction value (no freshness from bank, only ACT against replay – but might be played later)

Serious mistakes:

- checking PIN, result available on the device (mugging threat) – this concerns also cards of other banks
- the same PIN for ATM and online authentication – some keys on the CAP clean and some used - after stealing it one has 3 trials, 24 permutations on 4 keys, pbb to guess PIN to ATM becomes $\frac{1}{8}$
- CAP has no secret, infected PC may emulate CAP
- GSM in CAP to transmit secrets
- complicated instruction manual, the user may insert something else than intended account number
- overloading: sign with transaction with 0 value is valid for response (for a random account-nounce)
- NatWest: nonce as 4 digits in respond challenge, Chip&PIN terminal requests a number of responses from the card, later number of challenges from the online bank, there would be a match due to birthday paradox
(there are info indicating the attack: the number of requests, the change of transaction counter)

critical mistake: MITM regarding PIN verification

- PIN verification result never explicitly stated. Info to the bank contained in TVR (terminal verification results) and IAD (Issuer Application Data)
- TVR states possible failure conditions for authentication, in success not indicated which method used
 - bit8=1: cardholder verification was not successful
 - bit7=1: unrecognized CVM
 - bit6=1: PIN Try limit exceeded
 - bit5=1: PIN required and PIN pad not present
 - bit4=1: PIN required, PIN pad present and PIN not entered
 - bit3=1: online PIN entered

- IAD may contain info on whether the PIN has been verified, but cannot be read by the terminal (proprietary format), So terminal can have a different picture of the situation
 - bit4=1: Issuer Authentication performed and failed
 - bit3=1: offline PIN performed
 - bit2=1: offline PIN verification performed and failed
 - bit1=1: unable to go online

- attack:
 1. tricking the terminal by sending 0x9000 to **Verify** without sending PIN to the card
 2. card thinks that the terminal is not supporting PIN and skip PIN or uses signature
 3. card does not increase PIN retry counter
 4. issuer thinks that the terminal was not supporting PIN and accepts

- practical case (as described in 2015 paper after 2011 case in Belgium)
 - credit cards stolen, used in Belgium, police used intersection analysis (card usage, SIM cards in the proximity) to identify the criminals
 - "minimal effort design", just to work. Implementation of the attack with MiTM
 - hardware: FUN chip attached to the original chip, wires connected (contacts of the FUN with contacts of the original chip), the card has traces of manipulation. thickness: .82 mm (instead of .76mm)
 - functional: data embossed on the card does not match the data from the chip, accepts any PIN, some wrong responses

What went wrong:

- no evaluation, no public certification report
- no reaction to S&P paper from 2011
- specification EMV: thousands of pages
- certification costs
- designing a solution: chaos, no sufficiently detailed documentation and regime
- CC very likely to fail:
 - asset: PIN, password, protected against use on a PC
 - no methodology to answer the question: **what are side-effects of protecting one asset**

- important: security is **not monotonic**: improving situation with respect to one threat may worsen situation to another one. **Not reflected by CC framework.**
- **optimization is necessary, but may lead to situation that is worse than the original one**

(other solution: a shadow PIN for the case of mugging)

IX. MODELLING UNSECURITY

attacks:

- hit-and-run
- hit-and-stay
- insider

life-cycle of a solution based on key secrecy:

1. T_0 : key created
2. T_1 : forensics-based key non-compromised
3. T_2 : key compromise (known to the adversary)
4. T_3 : forensics-based key already compromised
5. T_4 : key ceased from operation

$[T_1, T_4]$ is a gray period. Should be as short as possible

Example countermeasure: DSAS framework

archive for signatures,

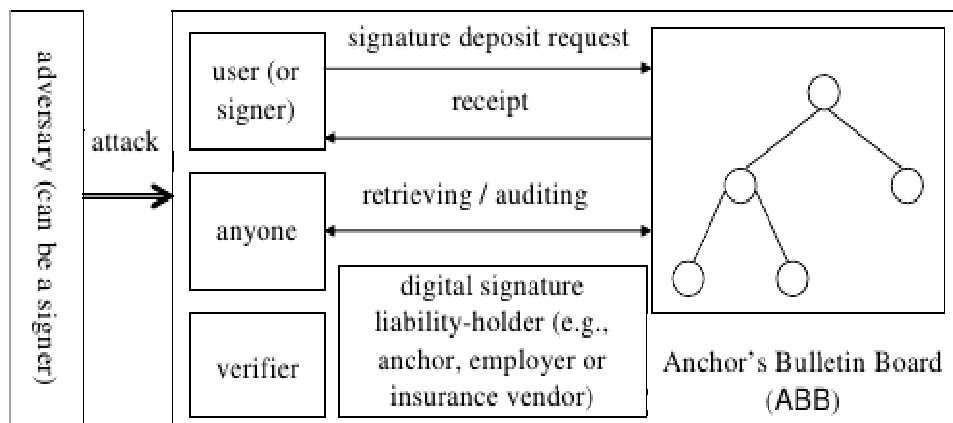
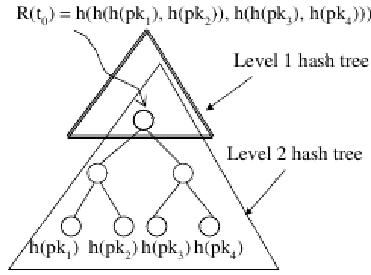
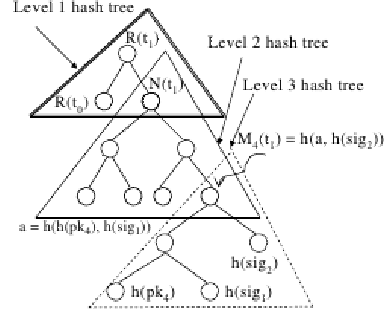


ABB:

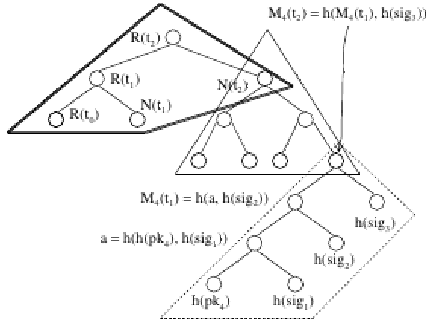
- 3 levels of trees
- bottom level: signatures corresponding to one user in a separate tree, leftmost leaf holds public key, the root is a leaf in the level 2 tree
- middle level trees: binary tree for all user, the root is a leaf in top level tree
- top tree has nonleaf nodes corresponding to old roots and old leaves of trees of level 2



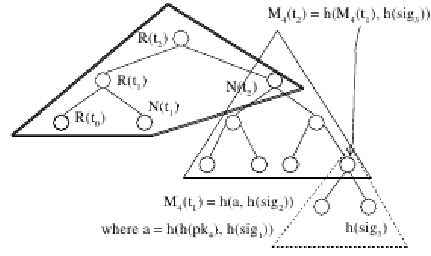
(a) ABB_0 at time t_0 (signatures-preserved case and signatures-compressed case): initialization with four user public keys



(b) ABB_1 at time t_1 (signatures-preserved case and signatures-compressed case): user pk_4 deposited two signatures



(c) ABB_2 at time t_2 (signatures-preserved case): user pk_4 deposited one new signature



(d) ABB_2 at time t_2 (signatures-compressed case): user pk_4 deposited one new signature

X. CRYPTOGRAPHIC FORENSICS

how to detect that a card has been cloned:

fail stop signatures: work if the clone created by cryptanalysis and deriving secret key and not by card inspection

- key generation by a trusted party: p, q chosen as for RSA, a - odd integer such that $\frac{p-1}{2a}$ is a prime and $q-1$ and a are coprime
- user chooses secret keys $sk_1, sk_2 \in Z_n^2$, public keys: $pk_1 = \alpha^a, pk_2 = sk_1^a \bmod n^2$
- signing m : $s := sk_1 \cdot sk_2^m$
- verification: $s^a := pk_1 \cdot pk_2^m$

- fail-stop idea: there are many solutions for sk_1 and sk_2 (namely a for each)
 a different solution yields $s'^a = s^a \bmod n$, then $s'^a = s^a \bmod q$, $s' = s \bmod q$ since a is invertible modulo q

ctrl-signatures:

actors:

Inspection Authority: IA has a long period secret key k_{master} . For a user U , IA determines the control key $c_U := Hash_1(U, k_{master})$, and a pair of inspection keys: the private key $i_U = Hash_2(U, k_{master})$, and the public key $I_U = g^{i_U}$.

Card Issuer: for a user U , Card Issuer obtains the keys c_U and I_U from IA and installs them in the SSCD issued for U .

Signatories: the SSCD of a user U holds the preinstalled keys c_U and I_U , as well as the private signature key x_U created at random by the SSCD, and the public key $X_U = g^{x_U}$. (Note that the SSCD does not hold the key i_U .)

Certification Authority: CA has standard keys for issuing certificates for the public keys of the users, just as in PKI built according to the X.509 framework.

footprints:

Generating $f_U(k)$ - a hidden footprint for k and user U .

input: I_U, k

$f := Hash_3(I_U^k)$;

output d least significant bits of f

For the inspection procedure carried out by Inspection Authority there is an alternative way for computing $f_U(k)$ (this is essential, since parameter k is present only on the SSCD):

Alternative generation of $f_U(k)$.

input: $i_U, r = g^k$

$f := Hash_3(r^{i_U})$;

output d least significant bits of f

Creating the i th signature by SSCD of user U for message M .

input: a message M

“choose k at random so that $f_U(k) = \rho_U^i$ ”

proceed with the signing algorithm Sign with

the first signature component $r = g^k$

Inspection

Below we describe inspection of the signature list created by a user U .

1. User U presents a list S_1, S_2, \dots, S_t of allegedly all signatures created with S of U , where the signatures appear on the list in the order in which they have been created. (If the signing time is included in the signatures, it is not necessary to specify the order of creating signatures.)
2. Apart from the regular verification of each signature S_i , the Inspection Authority checks all footprints. Namely, for each signature $S_j = (r_j, s_j)$, $j \leq t$, IA computes the footprint $\omega_j := f_U(r_j)$.
3. If $(\omega_1, \omega_2, \dots, \omega_t) = (\rho_U^1, \rho_U^2, \dots, \rho_U^t)$, then inspection result is positive.

XI. COMMUNICATION SECURITY – SSL/TLS

Padding attack (Serge Vaudenay)

Scenario:

- for encryption the plaintext should have the length as a multiple of b
- pad the plaintext with n occurrences of n , always pad something
- the resulting padded plaintext x_1, \dots, x_N encrypt in CBC mode with IV (fixed or random) and a block cipher:

$$y_1 = \text{Enc}(\text{IV} \oplus x_1), \quad y_i = \text{Enc}(y_{i-1} \oplus x_i)$$

- CBC:
 - efficiency
 - confidentiality limits: if IV fixed one can check that two plaintexts have the same prefix of a given size
 - CBC-MAC has security flaws: m_1 and m_2 augment by extra blocks: due to birthday paradox we might create the same MAC

attack:

- manipulate the ciphertext
- destination node decrypts, can see incorrect padding
- decision: what to do if padding incorrect?
 - reject: creates padding oracle
 - proceed: enables manipulation of the data

last word oracle:

- goal: compute $\text{Dec}(y)$
- create an input for padding oracle:
 - $r = r_1 \dots r_b$ chosen at random, $c := r \parallel y$
 - oracle call: if $O(c) = \text{valid}$, then $y_b = r_b \oplus 1$ whp
 - recognizing other cases:
 1. pick r_1, r_2, \dots, r_b at random, take $i = 0$
 2. put $r = r_1 r_2 \dots r_{b-1} (r_b \oplus i)$
 3. run padding oracle on $r \parallel y$, if result "invalid" then increment i and goto (2)
 4. $r_b := r_b \oplus i$
 5. for $j = b$ to 2:
 - $r := r_1 \dots r_{b-j} (r_{b-j+1} \oplus 1) r_{b-j} \dots r_b$
 - ask padding oracle for $r \parallel y$, if "invalid" then output $(r_{b-j+1} \oplus j) \dots (r_b \oplus j)$ and halt
 6. output $r_b \oplus 1$

block decryption oracle

let $a_1 \dots a_b$ be the plaintext of y

decryption:

- get a_b via the last word oracle
- proceed step by step learning a_{j-1} once a_j, \dots, a_b are already known
 1. set $r_k := a_k \oplus (b - j + 2)$ for $k = j, \dots, b$
 2. set r_1, \dots, r_{j-1} at random, $j := 0$
 3. $r := r_1 \dots r_{j-2} (r_{j-1} \oplus i) r_j \dots r_b$

4. if $O(r|y) = 0$, then $i := i + 1$ and goto 3
5. output $r_{j-1} \oplus i \oplus (b - j + 2)$

decryption oracle

- block by block
- the only problem with the first block if IV is secret

bomb oracles:

- padding oracle in SSL/TLS breaks the connection if padding error, so can be used only once
- bomb oracle: try a longer part at once

other paddings:

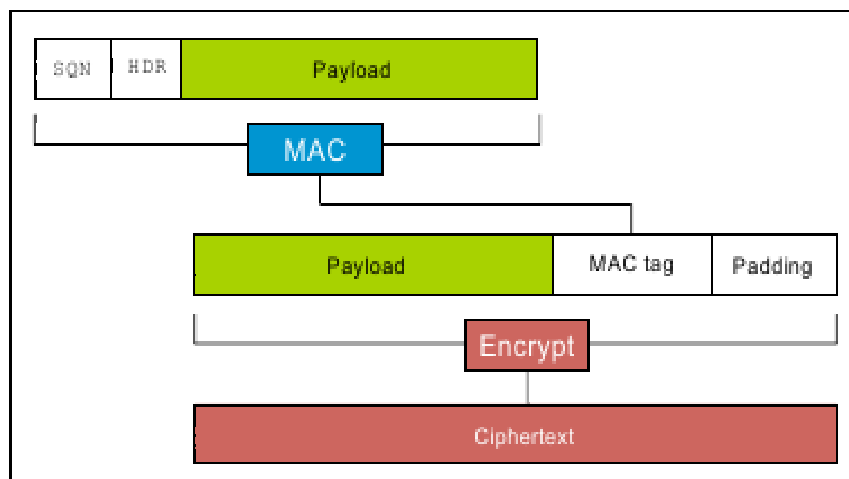
- $00\dots 0n$ instead of $nn\dots n$ - also vulnerable
- $12\dots n$ instead of $nn\dots n$ - also vulnerable
- $\langle \text{random} \rangle n$ instead of $nn\dots n$ - last word only, possible detection of padding length 1 (if encrypted twice with fixed IV)

Applications for (old) versions of SSL/TLS, ...

- MAC applied before padding, so padding oracle techniques can be applied
- wrong MAC and wrong padding create the same error message - from SSL v3.0, debatable whether it is impossible to recognize situation via side channel (response time)
- TLS attempts to hide the plaintext length by variable padding
- checking the length of padding: take the last block y , send $r|y$ where the last word of r is $n \oplus 1$. acceptance means that the padding is of length n
- checking longer paddings: send ry_1y_2 where y_1y_2 are the last blocks
- IPSEC: discards message with a wrong padding, no error message
- WTLS: decryption-failed message in clear (!) session not interrupted
- SSH: MAC after padding (+)

Lucky Thirteen

- concerns DTLS (similar to TLS for UDP connections)
- MAC-Encode-Encrypt paradigm (MEE), MAC is HMAC based



- 8-byte SQN, 5-byte HDR (2 byte version field, 1 byte type field, 2 byte length field)
- size of the MAC: 16 bytes (HMAC-MD5), 20 bytes (HMAC-SHA1), 32 bytes (HMAC-SHA-256)
- padding: $p + 1$ copies of p , at least one byte must be added
- after receiving: checking the details: padding, MAC, (underflow possible if padding manipulated and removing blindly)
- HMAC of M :

$$T := H((K_a \oplus \text{opad}) || H((K_a \oplus \text{ipad}) || M))$$
to M append the length field encoded

- **Distinguishing attack:**

- M_0 : 32 arbitrary bytes followed by 256 copies of 0xFF
- M_1 : 287 bytes followed by 0x00
- both 288 bytes, 18 plaintext blocks
- encoded $M_d || T || \text{pad}$, we aim to guess d
- C – the ciphertext
- create a ciphertext C' by truncating all parts corresponding to $T || \text{pad}$
- give $\text{HDR} || C'$ for decryption
- if M_0 : the 256 copies of 0xFF interpreted as padding and removed, remaining 32 bytes as short message and MAC, calculating MAC: 4 hash computed, then typically error returned to the attacker
- if M_1 : 8 hash evaluations

Plaintext recovery attacks

- C^* – the block of ciphertext to be broken, C' – the ciphertext block preceding it
- we look for P^* , where $P^* = \text{Dec}(C^*) \oplus C'$
- assume CBC with known IV, $b = 16$ (as for AES). $t = 20$ (as for HMAC-SHA-1)
- let Δ be a block of 16 bytes, consider

$$C^{\text{att}}(\Delta) = \text{HDR} || C_0 || C_1 || C_2 || C' \oplus \Delta || C^*$$

4 non-IV blocks in plaintext, the last:

$$P_4 = \text{Dec}(C^*) \oplus (C' \oplus \Delta) = P^* \oplus \Delta$$

- case 1: P_4 ends with 0x00 byte:
 - 1 byte of padding is removed, the next 20 bytes interpreted as MAC, 43 bytes left - say R . MAC computed on $\text{SQN} || \text{HDR} || R$ of 56 bytes
- case 2: P_4 ends with padding pattern of ≥ 2 bytes:
 - at least 2 bytes of padding removed, 20 bytes interpreted as MAC, at most 42 bytes left, MAC over at least $42 + 13 = 55$ bytes
- case 3: P_4 ends with no valid padding:
 - according to RFC of TLS 1.1, 1.2 treated as with no padding, 20 bytes treated as MAC, verification of MAC over $44 + 13 = 57$ bytes
 - MAC computed to avoid other timing attack!
- time: case 1 and 3: 5 evaluations of SHA-1, case 2: 4 evaluations of SHA-1, detection of case 2 possible in LAN
- in case 2: most probable is the padding 0x01 0x01, all other paddings have probability about $\approx \frac{1}{256}$ of probability of 0x01 0x01, so we may assume that $P_4 = P^* \oplus \Delta$ ends with 0x01 0x01. Then we derive the last two bytes of P^* .

repeat the attack with Δ' that has the same last two bytes to check if the padding has the length bigger than 2.

- after recovery of the last two bytes the rest recovered byte by byte from right to left:
 - the original padding attack
 - e.g. to find 3rd rightmost byte set the last two bytes Δ so that P_4 ends with 0x02 0x02, then try different values for the Δ_{13} so that Case 2 occurs (meaning that P_4 ends with 3 bytes 0x02)
 - average time: $14 \cdot 2^7$ trials
- practical issues:
 - for TLS after each trial connection broken, so multi-session scenario
 - timing difference small, so necessary to gather statistical data

- complexity in fact lower, since the plaintexts not from full domain : e.g. http user-name and password are encoded Base64
- partial knowledge may speed up the recovery of the last 2 bytes
- less efficient configuration of the lengths for HMAC-MD5 and HMAC-SHA-256

BEAST

attack, phase 0:

1. P to be recovered (e.g. a password, cookie, etc), requires ability to force Alice to put secret bits on certain positions
2. force Alice to send $0\dots 0P_0$ (requires malware on Alice computer)
3. eavesdrop and get $C_p = \text{Enc}(C_{p-1} \oplus 0\dots 0P_0)$
4. guess a byte g
5. force Alice to send the plaintext $C_{i-1} \oplus C_{p-1} \oplus 0\dots 0g$
6. Alice sends $C_i = \text{Enc}(C_{i-1} \oplus C_{i-1} \oplus C_{p-1} \oplus 0\dots 0g) = \text{Enc}(C_{p-1} \oplus 0\dots 0g)$
7. if $C_i = C_p$ then $P_0 = g$

attack phase 1:

1. P_0 already known
2. force Alice to send $0\dots 0P_0P_1$ and proceed as in phase 0

last phase: we get the test for the whole $P_0\dots P_{15}$

protection: browser must be carefully designed and do not admit injecting plaintexts (SOP- Same Origin Protection). Some products do not implement it.

CRIME (2012)

- based on compression algorithm used by some (more advanced) versions of TLS
- compression: LZ77 and then Huffman encoding, LZ77- sliding window approach: instead of a string put a reference to a previous occurrence of the same substring
- idea of recovering cookie:

```
POST / HTTP/1.1
Host: example.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:14.0) Gecko/20100101 Firefox/14.0.1
Cookie: secretcookie=7xc89f94wa96fd7cb4cb0031ba249ca2
Accept-Language: en-US,en;q=0.8

(... body of the request ...)
```

Listing 1: *HTTP request of the client*

modified POST:

```
POST /secretcookie=0 HTTP/1.1
Host: example.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:14.0) Gecko/20100101 Firefox/14.0.1
Cookie: secretcookie=7xc89f94wa96fd7cb4cb0031ba249ca2
Accept-Language: en-US,en;q=0.8
```

(... body of the request ...)

Listing 2: HTTP request modified by the attacker

LZ77 compresses the 2nd occurrence of secretcookie= or secretcookie=0. We try all secretcookie=i to find out the case when compression is easier (secretcookie=7) when the first character recovered the attacker repeats the attack for the second character (trying all “secretcookie=7i” in the preamble)

TIME

- again based on compression but now on the server side (from the client to the server compression might be disabled and CRIME fails)
- works if the server includes the client’s request in the response (most do!)
- works even if SOP is enabled. SOP does not control data with the tag `img`, so the attacker can manipulate length
- attacker requires malicious Javascript on the client’s browser
- attacker tries to get the secret value sent from the server to the client
- mechanism:
 - as in CRIME, the request sends “secretvalue=x” where x varies
 - the response is compressed, so it takes either “secretvalue=” or “secretvalue=x”
 - the length manipulated so that either two or one packets – connection specific data must be used: Maximum Transmission Unit
 - RTT (round trip time) measured
- independent on the browser, it is not an implementation attack!
- countermeasure: restrict displaying images

BREACH

Browser Reconnaissance and Exfiltration via Adaptive Compression of Hypertext

- attack against HTTP compression and not TLS compression as in case of CRIME
- a victim visits attacker-controlled website (phishing etc).
- force victim’s computer to send multiple requests to the target website.

- check sizes of responses

```
GET /product?id=12345&user=CSRFtoken=<guess> HTTP/1.1
Host: example.com
```

Listing 4: *Compromised HTTP request*

```
<form target="https://example.com:443/products/catalogue.aspx?id=12345&user=CSRFtoken=<guess>" >
...
<td nowrap id="tdErrLgf">
<a href="logoff.aspx?CSRFtoken=4bd634cda846fd7cb4cb0031ba249ca2">Log Off</a></td>
```

Listing 5: *HTTP response*

- requirements: application supports http compression, user's input in the response, sensitive data in the response
- countermeasures:
 - disabling compression
 - hiding length
 - no secrets in the same response as the user's data
 - masking secret: instead of S send $R||S \oplus R$ for random R (fresh in each response)
 - trace behaviour of requests and warn the user

POODLE (2014)

in SSL v.3.0 using technique from BEAST:

- encrypted POST request:

```
POST /path Cookie: name=value... <r\n\r\n> body ||20-byte MAC||padding
```
- manipulations such that:
 - the padding fills the entire block (encrypted to C_n)
 - the last unknown byte of the cookie appears as the last byte in an earlier block encrypted into C_i
- attack: replace C_n by C_i and forward to the server
usually reject
accept if $\text{Dec}_K(C_i)[15] \oplus C_{n-1}[15] = 15$, thereby $P_i[15] = 15 \oplus C_{n-1}[15] \oplus C_{i-1}[15]$
proceed in this way byte by byte
- downgrade dance: provoke lower level of protection by creating errors say in TLS 1.0, and create connection with SSL v3.0
- the attack does not work with weak (!) RC4 because of no padding

Weaknesses of RC4

- known weaknesses:
 - the first 257 bytes of encryption strongly biased, ≈ 200 bytes can be recovered if ≈ 232 encryptions of the same plaintext available
simply gather statistics as in case of Caesar cipher
 - at some positions (multiplies of 256) if a zero occurs then the next position more likely to contain a zero
- broadcast attack: force the user to encrypt the same secret repeatedly and close to the beginning
- countermeasure: no secrets in the initial part!

TLS 1.2

differences with TLS 1.1 and TLS 1.0 (Edukacja runs with TLS 1.0):

- explicit IV instead of implicit IV
- IDEA and DES 64bit removed
- MD5/SHA-1 PRF 65 is replaced with a suite specified hash function – SHA-256 for all TLS 1.2 suites, but in the future also SHA-3, ...
- digitally-signed element includes the hash algorithm used
- Verify_data length is no longer fixed length \Rightarrow TLS 1.2 can define SHA-256 based cipher suites
- new encryption modes allowed: CCM, GCM

CCM encryption mode

Prerequisites: block cipher algorithm; key K ; counter generation function; formatting function; MAC length $Tlen$

Input: nonce N ; payload P of $Plen$ bits; valid associated data A

Computation: Steps:

1. formatting applied to (N, A, P) , result: blocks B_0, \dots, B_r
2. $Y_0 := \text{Enc}_K(B_0)$
3. for $i = 1$ to r : $Y_i := \text{Enc}_K(B_i \oplus Y_{i-1})$
4. $T := \text{MSB}_{Tlen}(Y_r)$
5. generate the counter blocks $\text{Ctr}_0, \text{Ctr}_1, \dots, \text{Ctr}_m$ for $m = Plen/128$
6. for $j = 0$ to m : $S_j := \text{Enc}_K(\text{Ctr}_j)$
7. $S := S_1 || \dots || S_m$
8. $C := (P \oplus \text{MSB}_{Plen}(S)) || (T \oplus \text{MSB}_{Plen}(S))$

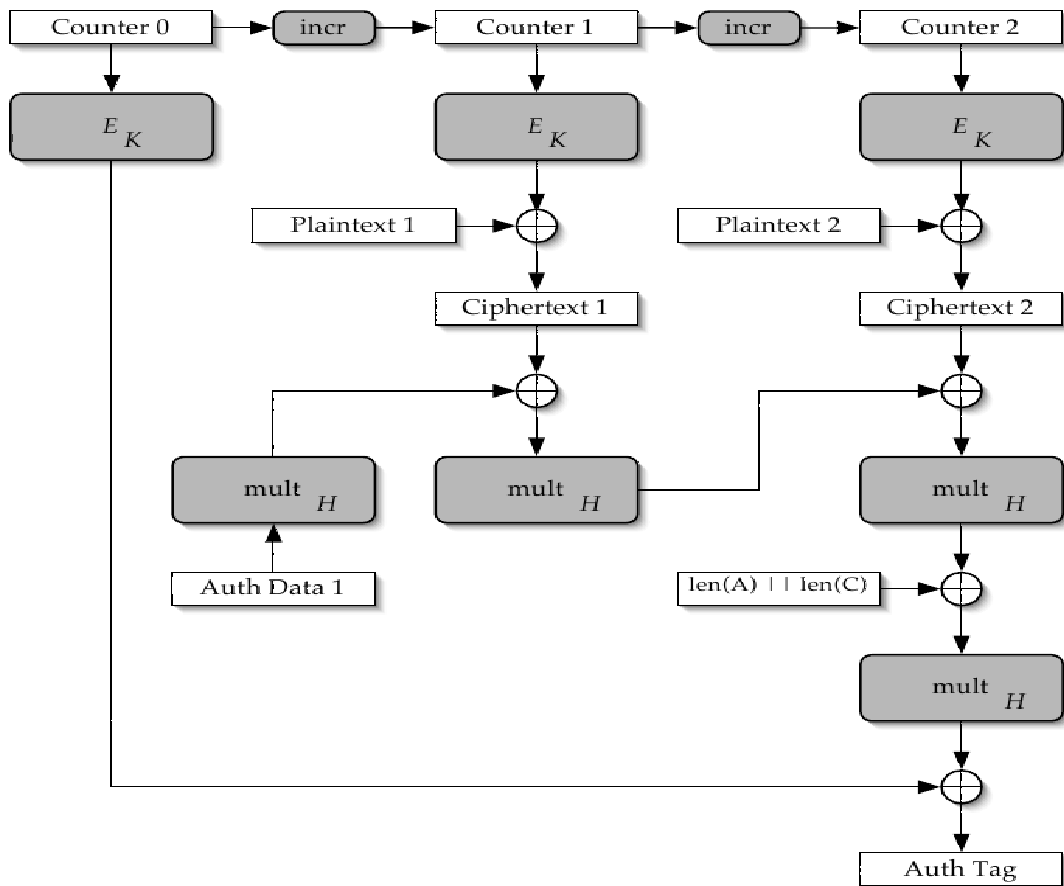
Decryption:

1. return INVALID, if $Clen < Tlen$
2. generate the counter blocks $Ctr_0, Ctr_1, \dots, Ctr_m$ for $m = Plen/128$
3. for $j = 0$ to m : $S_j := Enc_K(Ctr_j)$
4. $S := S_1 || \dots || S_m$
5. $P := MSB_{Clen}(C) \oplus MSB_{Plen}(S)$
6. $T := LSB_{Tlen}(C) \oplus MSB_{Tlen}(S_0)$
7. If N , A or P invalid, then return INVALID, else reconstruct B_0, \dots, B_r
8. recompute Y_0, \dots, Y_r
9. if $T \neq MSB_{Tlen}(Y_r)$, then return INVALID, else return P .

GCM (The Galois/Counter Mode)

Computation: Steps:

1. $H := Enc_K(0^{128})$
2. $Y_0 := IV || 0^{31}1$ if length of IV should be 96
or $Y_0 := GHASH(H, \{\}, IV)$
3. $Y_i := incr(Y_{i-1})$ for $i = 1, \dots, n$ (counter computation)
4. $C_i := P_i \oplus Enc_K(Y_i)$ for $i = 1, \dots, n-1$ (counter based encryption)
5. $C_n^* := P_n \oplus MSB_u(Enc_K(Y_n))$ (the last block need not to be full)
6. $T := MSB_t(GHASH(H, A, C)) \oplus Enc_K(Y_0)$



Details of computation of the tag

$\text{GHASH}(H, A, C) = X_{m+n+1}$ where m is the length of authenticating information A , and:

X_i equals:

$$\begin{array}{ll} 0 & \text{for } i = 0 \\ (X_{i-1} \oplus A_i) \cdot H & \text{for } i = 1, \dots, m-1 \\ ((X_{i-1} \oplus (A_m^* || 0^{128-v})) \cdot H & \text{for } i = m \\ (X_{i-1} \oplus C_i) \cdot H & \text{for } i = m+1, \dots, m+n-1 \\ ((X_{m+n-1} \oplus (C_m^* || 0^{128-u})) \cdot H & \text{for } i = m+n \\ ((X_{m+n} \oplus (\text{len}(A) || \text{len}(C))) \cdot H & \text{for } i = m+n+1 \end{array}$$

Decryption:

1. $H := \text{Enc}_K(0^{128})$
2. $Y_0 := \text{IV} || 0^{31}1$ if length of IV should be 96
or $Y_0 := \text{GHASH}(H, \{\}, \text{IV})$
3. $T' := \text{MSB}_t(\text{GHASH}(H, A, C)) \oplus \text{Enc}_K(Y_0)$, is $T = T'$?
4. $Y_i := \text{incr}(Y_{i-1})$ for $i = 1, \dots, n$
5. $P_i := C_i \oplus \text{Enc}_K(Y_i)$ for $i = 1, \dots, n$
6. $P_n^* := C_n^* \oplus \text{MSB}_u(\text{Enc}_K(Y_n))$

XII. CERTIFICATES and – SSL/TLS

“Certified Lies”

- rogue certificates + MitM attack: the user believes that is directed elsewhere
- no control over root CA’s worldwide, indicated either by operating system or the browser
- compelled assistance from CA’s ?

ROGUE Certificates and MD5

- target: create a certificate (webserver, client) that has not been issued by CA
- not forging a signature but:
 - i. find two messages that $\text{Hash}(M_0) = \text{Hash}(M_1)$ and M_0 as well as M_1 have some common prefix that you expect in a certificate (e.g. the CA name)
 - ii. submit a request corresponding to M_0 , get a certificate with the signature over $\text{Hash}(M_0)$
 - iii. copy the signature to a certificate based on M_1

- problems: some data in M_0 are to be guessed : sequential number, validity period, other are known in advance: distinguished name, ...

legitimate website certificate		rogue CA certificate
serial number		serial number
issuing CA		issuing CA
validity period		validity period
domain name	chosen prefixes	rogue CA name
		1024 bit RSA public key
		extensions
		"CA=true"
		tumor
2048 RSA public key	collision bits	
extension "CA=false"	identical suffix	

Table.

- finding M_0 and M_1 has to be fast (otherwise the guess about the serial number and validity will fail) - e.g. a day over the weekend
- attack on MD5, general picture:

message A		message B
prefix P		prefix P'
padding S_r		padding S'_r
birthday blocks S_b		birthday blocks S'_b
near-collision block $S_{c,1}$		near-collision block $S'_{c,1}$
near-collision block $S_{c,2}$		near-collision block $S'_{c,2}$
...		...
near-collision block $S_{c,r}$	←collision→	near-collision block $S'_{c,r}$
suffix		suffix

Table.

- identical prefix, birthday bits, near collision blocks:
 - birthday bits: 96, end at the block boundary, RSA bit in certificate, tumor (ignored part by almost all software) in rogue

birthday bits make the difference of intermediate hash value fall into a good class
 - then 3 near-collision 512-bit blocks. website $208 + 96 + 3 \cdot 512 = 1840$ bits of RSA modulus. rogue certificate: tumor
 - after collision bits, $2048 - 1840 = 208$ bits needed to complete the RSA modulus of the webpage.

– continued so that two prime factors:

→ B denotes the fixed 1840-bit part of the RSA modulus followed by 208 bits

- select a random 224-bit integer q until $B \bmod q < 2^{208}$, continue until both q and $\lfloor B/q \rfloor$ are prime
- (purely esthetic reasons: smallest fact is more than 67-digit prime)
- ... one can create RSA signature for the webpage for the certificate request

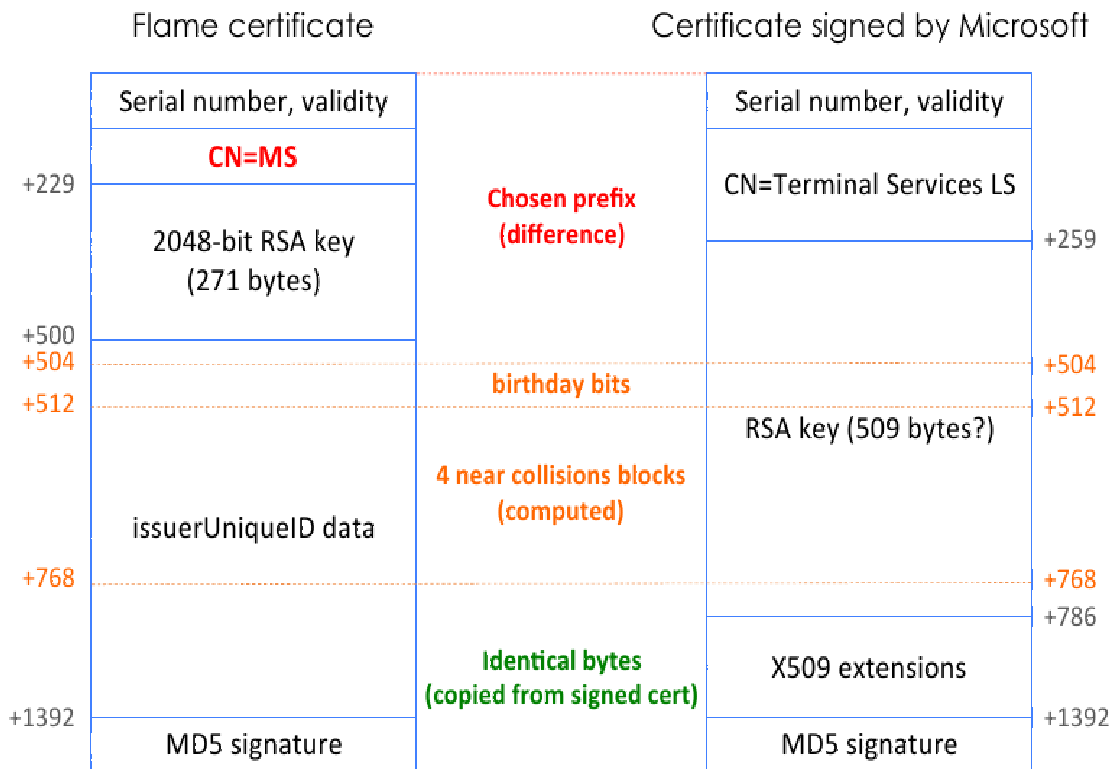
- attack complexity (number of hash block evaluations) for chosen prefix MD5: 2^{49} at 2007, 2^{39} in 2009 (computation costly in 2009), not much motivation for more work - remove MD5 certificates! (For a collision: 2^{16})

for SHA-1 still 2^{77} in 2012 (for a collision: 2^{65})

- history:
 - attack found
 - real collision computed as a proof-of-concept
 - CA informed and given time
 - publication
 - code available

FLAME

- malware discovered 2012, 20MB, sophisticated code, mainly in Middle East, government servers,
- draft of the attack:
 - client attempts to resolve a computer name on the network, in particular make WPAD (Web Proxy Auto-Discovery Protocol) requests
 - Flame claims to be WPAD server, provides wpad.dat configuration file
 - victim that gets wpad.dat sets its proxy server to Flame computer (later no sniffing necessary!)
 - Windows updates provided (but must be signed) – main problem!
 - signatures obtained for terminal Services, certificates issued by Microsoft LSRA PA. No Extended Key Usage restrictions – allows code signing, but Microsoft Hydra X.509 extension – cannot be used for code-signing on Vista and Windows 7
 - till 2012 still signatures with MD5 hash used
 - MD5 collision necessary to remove extension



MD5 attack draft

MD5:

- padding to the length $448 \bmod 512$ with $10\dots$, then the length of the message as 64 bit
- partition into 512 bit blocks
- IHV_{*i*} after block *i*, consist of four 32-bit numbers a_i, b_i, c_i, d_i . the initial values a_0, b_0, c_0, d_0 are fixed
- $IHV_i = MD5Compress(IHV_{i-1}, M_i)$
- output IHV_{*N*} (reformatted)
- compression function:
 - steps 0,...,63 (4 rounds of 16 steps)
 - each step involves modular addition, left rotation, nonlinear function f_t , involves addition constant AC_t and rotation constant RC_t
 - nonlinear function: $f_t(x, y, z) =$

$$\begin{aligned}
 F(x, y, z) &= (x \wedge y) \oplus (\bar{x} \wedge z) && \text{for } 0 < t < 16 \\
 G(x, y, z) &= (z \wedge x) \oplus (\bar{z} \wedge y) && \text{for } 16 \leq t < 32 \\
 H(x, y, z) &= x \oplus y \oplus z && \text{for } 32 \leq t < 48 \\
 I(x, y, z) &= y \oplus (x \vee \bar{z}) && \text{for } 48 \leq t < 64
 \end{aligned}$$
 - w_t composed of message blocks: repetitions occur (important for security):

W_t equals:

$$\begin{array}{ll} m_t & \text{for } 0 \leq t < 16 \\ m_{(1+5t) \bmod 16} & \text{for } 16 \leq t < 32 \\ m_{(5+3t) \bmod 16} & \text{for } 32 \leq t < 48 \\ m_{(7t) \bmod 16} & \text{for } 48 \leq t < 64 \end{array}$$

- rolling notation: the values kept are $Q_t, Q_{t-1}, Q_{t-2}, Q_{t-3}$, computed Q_{t+1} and retained $Q_{t+1}, Q_t, Q_{t-1}, Q_{t-2}$. Computation:

$$\begin{aligned} F_t &= f_t(Q_t, Q_{t-1}, Q_{t-2}), \\ T_t &= F_t + Q_{t-3} + AC_t + W_t \\ R_t &= RL(T_t, RC_t) \\ Q_{t+1} &= Q_t + R_t \end{aligned}$$

- final step: (quite important! enables elimination of differences step by step)

$$\text{MD5Compress}(\text{IHV}, B) = (a + Q_{61}, b + Q_{64}, c + Q_{63}, d + Q_{62})$$

Notation:

- primes for second copy
- $\delta X = X - X'$
- ΔX is δX in BSDR notation: a 32-bit word X is defined as $(k_i)_{i=0}^{31}$, where $X = \sum_{i=1}^{31} 2^i \cdot k_i$, and $i \in \{-1, 0, 1\}$
- many such representations, taken the one with maximal number of zeroes

Birthday part

- constructed near-collision blocks (based on Xiaoyun Wang observation): they cannot erase a difference in a , only identical differences can be removed from c and d parts
- assumption: before using near-collision block we adjust so that the differences are $\delta \text{IHV}_n = (0, \delta b, \delta c, \delta d)$
- birthday search: we aim to have a property on 64 bits,
- average number of calls to MD5 compression function $\approx \sqrt{\pi} 2^{32}$
- in the original paper: more conditions (relationship between δb and δc) - time-space trade-off – so that the time for the birthday part and near collision part are of the same order
- the search for collision: usual technique of a pseudorandom walk and storing only characteristic points on a walk
- a collision found in this way is not always useful as the prefix might be the same, number of pairs of near-collision blocks must be small, etc

Near-collision blocks

- basic differential path for a near-collision block:

Table 2 Family of partial differential paths using $\delta m_{11} = \pm 2^{p-10 \bmod 32}$, where $s_0, \dots, s_{w'} \in \{-1, 0, +1\}$ and $w' = \min(w, 31 - p)$ for a fixed $w \geq 0$

t	δQ_t	δF_t	δW_t	δT_t	δR_t	RC_t
31	$\mp 2^{p-10 \bmod 32}$					
32	0					
33	0					
34	0	0	$\pm 2^{p-10 \bmod 32}$	0	0	16
35 – 60	0	0	0	0	0	.
61	0	0	$\pm 2^{p-10 \bmod 32}$	$\pm 2^{p-10 \bmod 32}$	$\pm 2^p$	10
62	$\pm 2^p$	0	0	0	0	15
63	$\pm 2^p$	0	0	0	0	21
64	$\pm 2^p$ $+ \sum_{\lambda=0}^{w'} s_\lambda 2^{p+21+\lambda \bmod 32}$					

Note: Interesting values for the parameter w are between 2 and 5.

- overview:
 - only δm_{11} not equal to 0. It occurs only at steps 34 (to erase a difference δQ_{31}), however it occurs at step 61 as well
 - then the difference propagates to Q_{62}, Q_{63} (affecting the change on c and d) as well as Q_{64} (affecting b), difference obtained

$$\pm \left(0, 2^p + \sum_{\lambda=0}^{w'} s_\lambda \cdot 2^{p+21+\lambda \bmod 32}, 2^p, 2^p \right)$$

- different characteristics for different w' - for large w more differences might be removed but also pbb lower
- starting differences: $\delta c = \sum_i k_i \cdot 2^i$ and $\delta b - \delta c = \sum_i l_i \cdot 2^i$ (expressed as NAFs)
- let $k_i \neq 0$: use differential path with $m_{11} = k_i \cdot 2^{i-10 \bmod 32}$ to eliminate the difference $k_i \cdot 2^i$ in c and d . A side effect: change of δb by

$$k_i 2^i + \sum_{\lambda=i+21}^{i+21+w'} l_\lambda \cdot 2^{\lambda \bmod 32}$$

in this expression we carefully choose w to get coordinates l_λ

- finally $\delta c=0$, but there are some differences in b , say we get $\delta \hat{b} = \sum_{\lambda=0}^{31} e_\lambda l_\lambda 2^\lambda$, where $e_\lambda = 0$ (if the coordinate has been nullified) or $e_\lambda = 1$, the weight of $\delta \hat{b}$ is not higher than the weight of δb
- the differences from δb eliminated as follows:
 - let $\text{NAF}(\delta \hat{b}) = \sum_{\lambda=0}^{31} \hat{l}_\lambda \cdot 2^\lambda$. Choose j such that $\hat{l}_j \in \{-1, 1\}$ and $j - 21 \bmod 32$ is minimal
 - Then the difference $\sum_{i=j}^{j+w'} \hat{l}_i 2^i$ with $w' = \min(w, 31 - (j - 12 \bmod 32))$ can be eliminated from $\delta \hat{b}$ with $m_{11} = 2^{j-31 \bmod 32}$

- side effect: a new difference $2^{j-21 \bmod 32}$ in b , c and d
- the side effect eliminated using $\delta m_{11} = 2^{31} \bmod 32$
- result: a new difference vector $(0, \delta \bar{b}, 0, 0)$ with weight of $\text{NAF}(\delta \bar{b})$ smaller than the weight of $\text{NAF}(\delta b)$
- construction of near-collision blocks:
 - the key is so-called *differential path* for MD5Compress, for IHV , IHV' and δb
 - $\delta F_t = f_t(Q'_t, Q'_{t-1}, Q'_{t-2}) - f_t(Q_t, Q_{t-1}, Q_{t-2})$
 $\delta T_t = \delta F_t + \delta Q_{t-3} + \delta W_t$
 $\delta R_t = \text{RL}(T'_t, \text{RC}_t) - \text{RL}(T_t, \text{RC}_t)$
 $\delta Q_{t+1} = \delta Q_t + \delta R_t$
 - δF and δR cannot be uniquely determined given $(\delta Q_t, \delta Q_{t-1}, \delta Q_{t-2})$ and δT_t
 - differentials: corresponding to each step in the rolling notation. We start from $\text{IHV} = (Q_{-3}, Q_0, Q_{-1}, Q_{-2})$ and $\text{IHV}' = (Q'_{-3}, Q'_0, Q'_{-1}, Q'_{-2})$ and δB and leading to $(\delta Q_{61}, \delta Q_{62}, \delta Q_{63}, \delta Q_{64})$
 - bitconditions:

Table 5 Boolean function bitconditions

$q_t[i]$	Condition on $(Q_t[i], Q'_t[i])$	Direct/indirect	Direction
0	$Q_t[i] = Q'_t[i] = 0$	Direct	
1	$Q_t[i] = Q'_t[i] = 1$	Direct	
~	$Q_t[i] = Q'_t[i] = Q_{t-1}[i]$	Indirect	Backward
v	$Q_t[i] = Q'_t[i] = Q_{t+1}[i]$	Indirect	Forward
!	$Q_t[i] = Q'_t[i] = \overline{Q_{t-1}[i]}$	Indirect	Backward
y	$Q_t[i] = Q'_t[i] = \overline{Q_{t+1}[i]}$	Indirect	Forward
m	$Q_t[i] = Q'_t[i] = Q_{t-2}[i]$	Indirect	Backward
w	$Q_t[i] = Q'_t[i] = Q_{t+2}[i]$	Indirect	Forward
#	$Q_t[i] = Q'_t[i] = \overline{Q_{t-2}[i]}$	Indirect	Backward
h	$Q_t[i] = Q'_t[i] = \overline{Q_{t+2}[i]}$	Indirect	Forward
?	$Q_t[i] = Q'_t[i] \wedge (Q_t[i] = 1 \vee Q_{t-2}[i] = 0)$	Indirect	Backward
q	$Q_t[i] = Q'_t[i] \wedge (Q_{t+2}[i] = 1 \vee Q_t[i] = 0)$	Indirect	Forward

- constructing differential paths:
 - i. for steps 1-11: forward
 - ii. for steps 64-16 backwards
 - iii. then try to fill the gap between 11 and 16
- very subtle case specific techniques

XIII. CACHE ATTACKS

idea:

- applies to multiprocess architectures, with strict separation between processes offered by the system: hypervisor and virtualization, sandboxing, ...
- trying to get secrets from one processes by another process with no privileges
- despite separation protection the processes share cache
- there is a strict control over the cache content but **cache hits and cache misses** might be detected by **timing for the attacker's process** (and not of the victim process)
- the timing for cache access should somehow depend on the sensitive information to be retrieved
- difficulty: other than in the classical cryptanalysis – access to plaintext or ciphertext might be impossible (they belong to the victim process) - the attacker can only predict something

cache:

- cache is necessary: gap between CPU speed and latency of memory access, innermost cache access $\approx 0.3\text{ns}$, main memory access $\approx 50\text{ns}$ to 150ns
- set-associative memory cache:
 - cache line of B byte
 - S cache sets, each consisting of W cache lines
 - when a cache miss occurs, then a memory block is copied into one of cache lines evicting its previous contents
 - a memory block of address a can be cached only into the cache set with the index i such that $i = \lfloor a/B \rfloor$ — **this is crucial for the attack**
- cache levels: slight complication to the attacks but differences of timing enable to recognize the situation

CASE STUDY: AES encryption

AES software implementation:

- particularly vulnerable because of its design
- AES defined in algebraic terms, but lookup table typically the fastest
- key expansion: round zero: simply the key bytes directly, other rounds: key expansion reversible (details irrelevant for the attack)

- fast implementation based on tables T_0, T_1, T_2, T_3 and $T_0^{(10)}, T_1^{(10)}, T_2^{(10)}, T_3^{(10)}$ for the last round (with no MixColumns)

- round operation

$$\left(x_0^{(r+1)}, x_1^{(r+1)}, x_2^{(r+1)}, x_3^{(r+1)}\right) := T_0(x_0^r) \oplus T_1(x_5^r) \oplus T_2(x_{10}^r) \oplus T_3(x_{15}^r) \oplus K_0^{(r+1)}$$

$$\left(x_4^{(r+1)}, x_5^{(r+1)}, x_6^{(r+1)}, x_7^{(r+1)}\right) := T_0(x_4^r) \oplus T_1(x_9^r) \oplus T_2(x_{14}^r) \oplus T_3(x_3^r) \oplus K_1^{(r+1)}$$

$$\left(x_8^{(r+1)}, x_9^{(r+1)}, x_{10}^{(r+1)}, x_{11}^{(r+1)}\right) := T_0(x_8^r) \oplus T_1(x_{13}^r) \oplus T_2(x_2^r) \oplus T_3(x_7^r) \oplus K_2^{(r+1)}$$

$$\left(x_{12}^{(r+1)}, x_{13}^{(r+1)}, x_{14}^{(r+1)}, x_{15}^{(r+1)}\right) := T_0(x_{12}^r) \oplus T_1(x_1^r) \oplus T_2(x_6^r) \oplus T_3(x_{11}^r) \oplus K_3^{(r+1)}$$

attack notation:

- $\delta = B/\text{entrysize}$ of lookup table, typically: $\text{entrysize}=4\text{bytes}$, $\delta = 16$
- for a byte y let $\langle y \rangle = \lfloor y/\delta \rfloor$, it indicates a memory block of y in T_l
- if $\langle y \rangle = \langle z \rangle$ then request to the same memory block of the lookup table
- $Q_k(p, l, y) = 1$ iff AES encryption of plaintext p under key K accesses memory block of index y in T_l at least once in 10 rounds
- $M_k(p, l, y)$ a measurement that has expected value bigger in case when $Q_k(p, l, y) = 1$ then in case when $Q_k(p, l, y) = 0$

“synchronous attack”

- plaintext random but known, one can trigger encryption (e.g. for VPN with unknown key, dm-crypt of Linux)
- phase 1: measurements, phase 2: analysis
- from experiments: AES key recovered using 65 ms of measurements (800 writes) and 3 sec analysis
- **round-one attack:** the first round attacked

i. accessed indices are simply $x_i^{(0)} = p_i \oplus k_i$ for $i = 0, \dots, 15$

ii. finding information $\langle k_i \rangle$ of k_i – test candidates \bar{k}_i

iii. if $\langle k_i \rangle = \langle \bar{k}_i \rangle$ and $\langle y \rangle = \langle p_i \oplus \bar{k}_i \rangle$ then $Q_k(p, l, y) = 1$ for the lookup $T_l(x_i^{(0)})$

iv. if $\langle k_i \rangle \neq \langle \bar{k}_i \rangle$ then there is no lookup in block y for T_l during the first round, but

- there are $4 \cdot 9 - 1 = 35$ other accesses affected by other plaintext bits

- probability that none of them accesses block y for T_l is

$$\left(1 - \frac{\delta}{256}\right)^{35} \approx 0.104 \text{ for } \delta = 16$$

v. few dozens of samples required to find a right candidate for $\langle \bar{k}_i \rangle$

vi. together we determine $\log(256/\delta) = 4$ bits of each byte of the key

- vii. no more possible for the first round, not enough to start brute force (still 64 bits to be found!)
- viii. in reality more samples needed due to noise in measurements $M_k(p, l, y)$ and not $Q_k(p, l, y)$

– **two-round attack:** the second round attack because of the missing bits

i. exploiting equations derived from Rijndael specification:

$$x_2^{(1)} = s(p_0 \oplus k_0) \oplus s(p_5 \oplus k_5) \oplus 2 \bullet s(p_{10} \oplus k_{10}) \oplus 3 \bullet s(p_{15} \oplus k_{15}) \oplus s(k_{15}) \oplus k_2$$

$$x_5^{(1)} = s(p_4 \oplus k_4) \oplus 2 \bullet s(p_9 \oplus k_9) \oplus 3 \bullet s(p_{14} \oplus k_{14}) \oplus s(p_3 \oplus k_3) \oplus s(k_{14}) \oplus k_1 \oplus k_5$$

$$x_8^{(1)} = \dots$$

$$x_{15}^{(1)} = \dots$$

where s stands for the Rijndael Sbox, and \bullet means multiplication in the field with 256 elements

ii. lookup for $T_2(x_2^{(1)})$:

- $\langle k_0 \rangle, \langle k_5 \rangle, \langle k_{10} \rangle, \langle k_{15} \rangle, \langle k_2 \rangle$ already known
- low level bits of $\langle k_2 \rangle$ influence only low bits of $x_2^{(1)}$ so not important for cache access pattern
- the upper bits of $x_2^{(1)}$ can be determined after guessing low bits of k_0, k_5, k_{10}, k_{15} : there are δ^4 possibilities ($=16^4$)
- a correct guess yields a lookup in the right place
- an incorrect guess: some $k_i \neq \bar{k}_i$ so

$$x_2^{(1)} \oplus \bar{x}_2^{(1)} = c \bullet s(p_i \oplus k_i) \oplus c \bullet s(p_i \oplus \bar{k}_i) \oplus \dots$$

(for c depending on i) where ... depends on different random plaintext bits and therefore random

differential properties of AES studied for AES competition:

$$\Pr [c \bullet s(p_i \oplus k_i) \oplus c \bullet s(p_i \oplus \bar{k}_i) \neq z] > 1 - \left(1 - \frac{\delta}{256}\right)^3$$

so the false positive for lookup:

- $\left(1 - \frac{\delta}{256}\right)^3$ for computing $T_2(x_2^{(1)})$
- $\left(1 - \frac{\delta}{256}\right)$ for computing each of the remaining T_2
- together $\left(1 - \frac{\delta}{256}\right)^{38}$

- this yields about 2056 samples necessary to eliminate all wrong candidates
- it has to be repeated 3 more times to get other nibbles of key bytes

iii. optimization: guess $\Delta = k_i \oplus k_j$ and take $p_i \oplus p_j = \Delta$, then i.e. $s(p_0 \oplus k_0) \oplus s(p_5 \oplus k_5)$ cancels out and we have to guess less bits (4 instead of 8)

- **similar attack: last round** - created ciphertext must be known to the attacker, otherwise similar. subkey from the last round learnt, but keyschedule is reversible
- **measurement: Evict+Time**
 - i. procedure:
 1. trigger encryption of p
 2. evict: access memory addresses so that one cache set overwritten completely
 3. trigger encryption of p
 - ii. in the evicted cache set one cache line from T_l
 - iii. measure time: if long then cache miss and the encryption refers to known δ positions
 - iv. practical problem: triggering may invoke other activities and timing is not precise
- **measurement: Prime+Probe**
 - i. procedure
 1. (prime) read A : a contiguous memory of the size of the cache – results in overwriting the entire cache
 2. trigger an encryption of p (partial eviction at places where lookup used)
 3. (probe:) read memory addresses of A that correspond to $M_k(p, l, y)$
 - ii. easier: timing for probe suffice to check if encryption used a given cache set
- **complications in practice:**
 - i. address of lookup tables in the memory - unknown so how they are loaded to the cache unknown – offset can be found by considering all offsets and then statistics for each offset (experiments show good results even on noisy environment)
 - ii. hardware prefetcher may disturb the effects. Solution: read and write the addresses of A according to a pseudorandom permutation
- **practical experiments:** e.g. Athlon 64, no knowledge of addresses mapping, 8000 encryptions with Prime & Probe
Linux dm-crypt (disk, filesystem, file encryption): with knowledge of addressing, 800 encryptions (65 ms), 3 seconds analysis, full AES key
- **extensions of the attack:**
 - on some platforms timing shows also position of the cache line (better resolution for one-round attack)
 - remote attacks (VPN, IPSec): with requests that trigger immediate response (situation yet unclear about practicality)

“asynchronous attack”

- no knowledge of plaintext, no knowledge of ciphertext
- one-round attack
- based on frequency F of bytes in e.g. English texts, frequency score for each of $\frac{256}{\delta}$ blocks of length δ
- F is nonuniform: most bytes have high nibble equal to 6 (lowercase characters “a” through “o”)
- find j such that j is particularly frequent indicates $j = 6 \oplus \langle k_i \rangle$ and shows $\langle k_i \rangle$
- complication: this frequency concerns at the same time k_0, k_5, k_{10}, k_{15} affecting T_0 so we learn 4 nibbles but not their actual allocation to k_0, k_5, k_{10}, k_{15}
roughly number of bits learnt: $4 \cdot (4 \cdot 4 - \log 4!) \approx 4 \cdot (16 - 3.17) \approx 51$ bits
- experiment: OpenSSL, measurements 1 minute, 45.27 bits of information on the 128-bit key gathered

Countermeasures

- “no reliable and practical countermeasure” so far
- implementation based on no-lookup but algebraic algorithm (slow!!!) or bitslice implementation (sometimes possible and nearly as efficient as lookup)
- alternative lookup tables: if smaller than less data leaks (but for cryptanalysis bigger Sboxes increase security)
- data-independent access to memory blocks - every lookup causes a redundant read in all memory blocks, generally: oblivious computation possible theoretically but overhead makes it impractical
- masking operations: \approx “we are not aware of any method that helps to resist our attack”
- cache state normalization: load all lookup tables - equires deep changes in OS and reduces efficiency, even then LRU cache policy may leak information which part has been used!
- process blocking: again, deep changes in OS
- disable cache sharing: deep degradation of performance
- “no-fill” mode during encryption:
 - preload lookup tables
 - activate “no-fill”
 - encrypt
 - deactivate “no-fill”

the first two steps critical and no other process is allowed to run
possible only in privileged mode, cost of operation prohibitive

- dynamic table storage: e.g. many copies of each table, or permute tables
details architecture dependent and might be costly
- hiding timing information: adding random values to timing makes the statistical analysis harder but still feasible
- try to protect some rounds (the first 2 and the last one) with any mean – but may be there are other attack techniques...
- cryptographic services at system level: good but are unflexible
- sensitive status for user processes: erasing all data when interrupt
- specialized hardware support: seems to be the best choice

but the problem is not limited to AES or crypto – many sensitive data operations are not cryptographic and a coprocessor does not help