FFI Forsvarets forskningsinstitutt

Public Key Infrastructure – scaling perspectives

Finseskolen 2012

Anders Fongen, PhD
Norwegian Defence Research Establishment
anders.fongen@ffi.no

Outline of presentation

- Short intro to PKI architecture and services
- Optimization opportunities
- Traffic estimates
- Estimates of a multi-tiered PKI
- Conclusion

PKI architecture and services

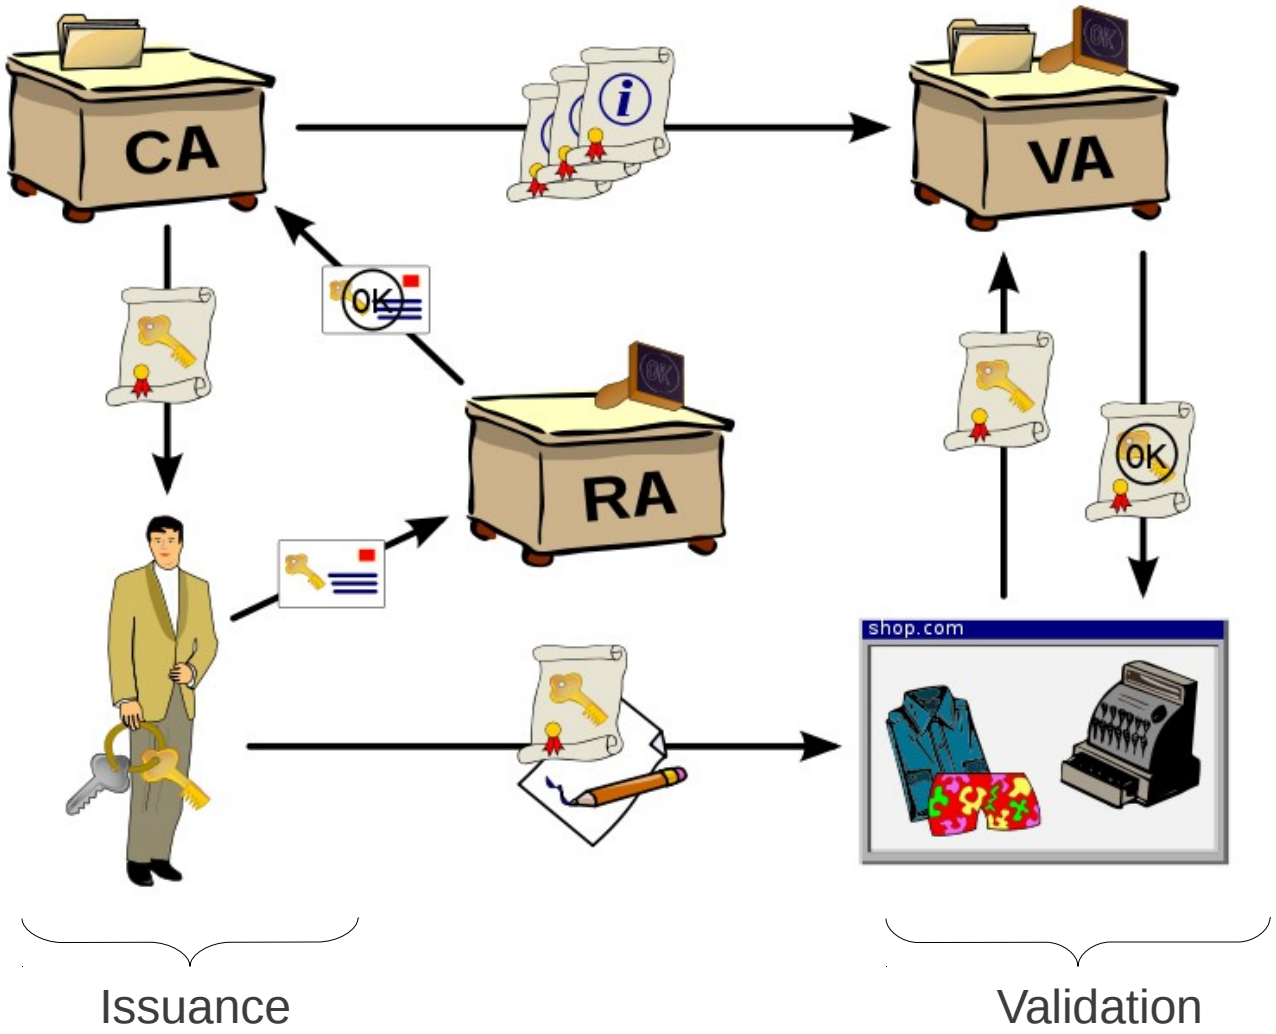

Source: Wikipedia Commons

Optimization principles

- Identity bottlenecks
  - Certificate issuance are unlikely candidates (too seldom)
  - Validation and CRL distribution generate large data volumes
  - Network near the client is likely to be poorest
- Reduce message size
  - Redundancy elimination: compression, normalization
- Reduce message frequency
  - Relaxed consistency, caching
- Exploit topological properties
  - Aggregated multicast, cooperative caching, overlay networks
  - Cross-layer techniques
- Identify consequences for “COTS deployment”

Optimization opportunities

- Delta CRLs (contain only recent revocations)
  - Not well supported by COTS software
- Push-based distribution of CRLs
  - Employs multicast middleware
- Exclude certificate from signature structure
  - Receiver obtains it on-demand
- Cache validation results for a while (trust has a lifetime)
  - A “freshness cache” is required
  - Hit/miss rate will be analyzed shortly

Online validation

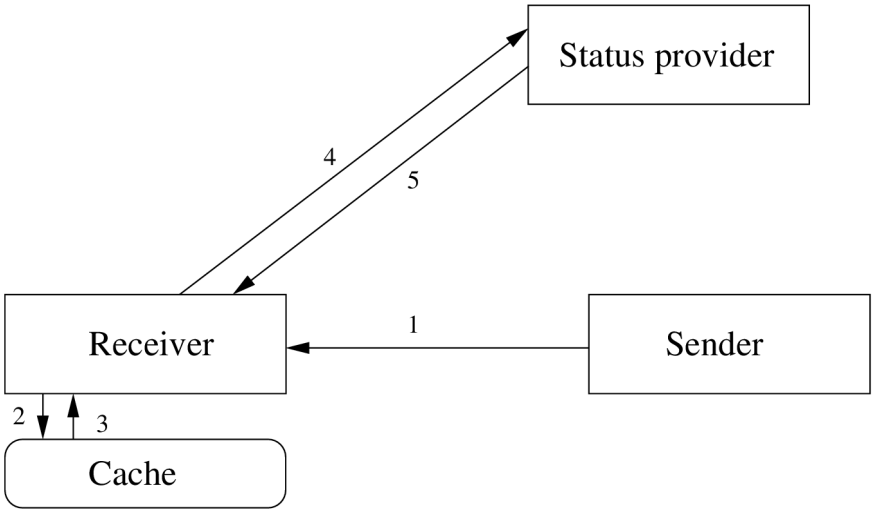

- The validation result is assumed to be “cacheable”
- Used for subsequent validation of certificates already “seen”

Chosen parameter values for the analysis

Parameter name	value
Number of users	1000
Number of certificates (N)	10000
Revocation latency (t_e)	4 hours
Messages received per day (λ_{tot})	300
Revocation rate (r)	10 % per year

Caching of recent validation results

- In the beginning the cache is empty
  - most received certificates will need an “ordinary” validation
- The cache fills up gradually
  - as more certificates turn up for the first time
- The miss rate reaches an non-zero asymptote in a “stable state”
  - due to the freshness requirement and cache entry expiration

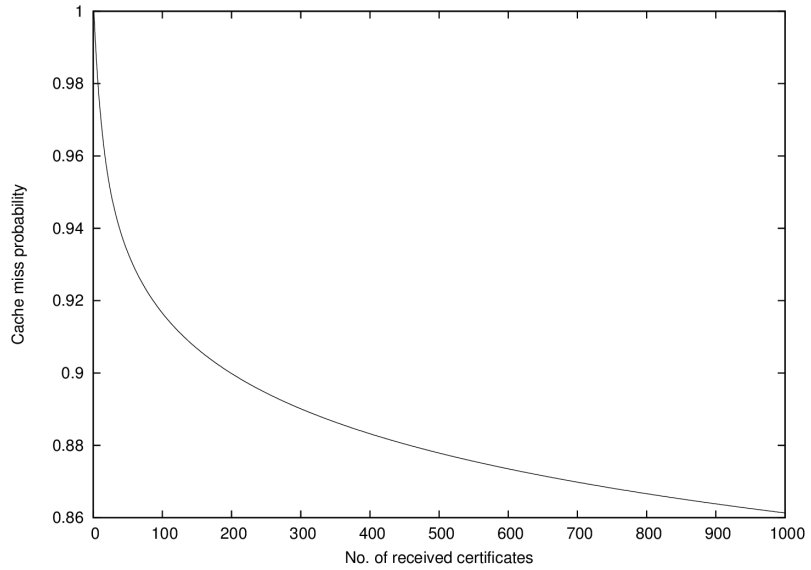

lifetime cert population

p_{miss}(4 h, 10000) = 0.86

p_{miss}(4 h, 10) = 0.15

p_{miss}(2 h, 10) = 0.25

p_{miss}(6 mon, 10000) = 0.47

p_{miss}(6 mon, 10) ≈ 0

300 messages per day, 1000 messages received

Size of revocation lists

- Certificates have an estimated revocation rate *r*, e.g. 10 % per year
- Certificates are issued and revoked at a uniform rate
  - Average age is half their lifetime
- With a lifetime of *x* years, there should be *r*x/2* certificates on the revocation list.

- *Well, not exactly*, but this holds as an approximation
- With x = 1 year, r= 0.1/year and a population of 10000 certs, the CRL size is:
  - 10000 * 0.05 * 36 bytes = *19 kB*
- CRLs can grow big and potentially create huge traffic peaks.

Multi-tiered PKI

- Several CAs in a hierarchy
  - Trust anchor at the top
  - Issuing of End Entity (EE) certificates takes place at the bottom
- Relying party must validate the entire *certificate chain*
  - Provided either in signatures or by a *Certificate Store* (CS)
- Investigated configurations:
  - Validation based on CRLs
  - Validation based on VA
  - Short lived certificates (no revocation)
  - COTS/non-COTS configuration

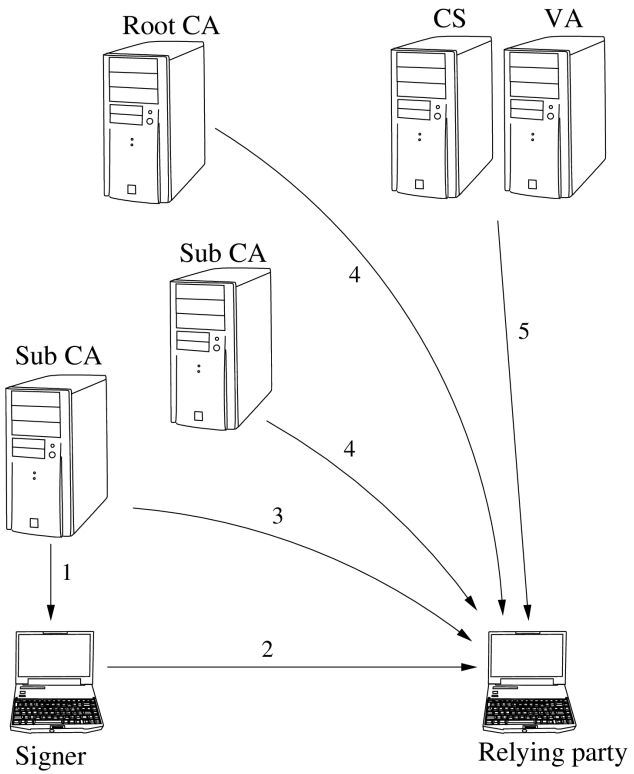

Validation based on revocation lists

- Validating party must obtain list regarding EE cert and 2 CA cert
  - 23 kB (during the interval t_e)
- Certificates included in signature (“COTS compliant”)
  - 6 kB * 50 messages (during t_e)
  - **Total 323 kB**
- Certificates retrieved on-demand from a cert store (CS)
  - Cached for their remaining lifetimes (average 6 months)
  - p_{miss}(6 mon, 10000) = 0.47 (for EE certs in a freshness cache)
  - For CA certs $p_{miss} \approx 0$
  - Retrieval operation from a CS takes 2 kB (est.)
  - 2 kB * 50 * 0.47 = 47 kB (during t_e)
  - **Total 70 kB**

Validation based on VAs (status providers)

- Validation results are cached for the duration of t_e
- We assume 10 CA certificates per tier
  - $p_{miss}(4\ h, 10) = 0.15$
  - $p_{miss}(4\ h, 10000) = 0.86$
- Invocation of a VA service makes 3 kB of network traffic
- Traffic related to validation is
  - 50 * 3 kB*(0.86+0.15+0.15) = 174 kB
- Certificates included in signature (“COTS compliant”)
  - 6 kB * 50 + 174 kB = **474 kB**
- Certificates retrieved on-demand from a cert store (CS)
  - 2 kB * 50 * 0.47 + 174 kB = **221 kB**

Validation of short-lived certificates

- Issued with a lifetime of t_e
- Never revoked, always valid
- Retrieved by the signer/sender every t_e
- EE certificate included in every signature (not cached)
- CA cert validation retrieved from VA by relying party (and cached)
  - $p_{miss}(2\ h, 10) = 0.25$
  - Retrieval of validation result (PoV) takes 1.5 kB
- Traffic related to validation of CA certificates is
  - 50 * (1.5 kB + 1.5 kB) * 0.25 = 38 kB
- Traffic related to relying party (during t_e)
  - 2 kB + 50 * 2 kB + 38 kB = **140 kB**

Summary table of validation alternatives

Table shows client-side traffic related to signature validation of 50 messages over 4 hours

Validation approach	*COTS compliant*	*COTS non-compliant*	*Sensitive to*
Revocation lists	323 kB	70 kB	Size of cert population
Online status provider	474 kB	221 kB	Message volume, loss of connections
Short lived certificates		140 kB	Message volume

Does not say anything about the *traffic distribution* within those 4 hours

Advantages of short-lived certificates

- Straightforward semantics
  - “Is it valid?” is not complementary to “Is it revoked”?
  - Shifts “burden of proof” from receiver to sender, where it belongs
- Revocation lists not needed
  - Validation is always conclusive (never “maybe”)
- Scales better
  - Validation cost unaffected by the size of the certificate population
  - Avoids the traffic peaks associated with CRL distribution
- Applies well to a cross-domain operation
  - Only certificates need to cross, not revocation lists
- CA private key is less exposed than the VA private key
  - The VA holds a key of unlimited trust right behind a public service point. The CA does not have a public service point.

An “optimized” PKI arrangement

(1) sub CA certificate issuance

(2) EE certificate issuance

(3) PoV (regarding CA) issuance

(4) EE certificate included in message signature

(5) PoV distribution network (high hit rate)

An alternative approach is to issue a common PoV for all (max 10) CAs, resulting in a larger, but singleton data structure for distribution.

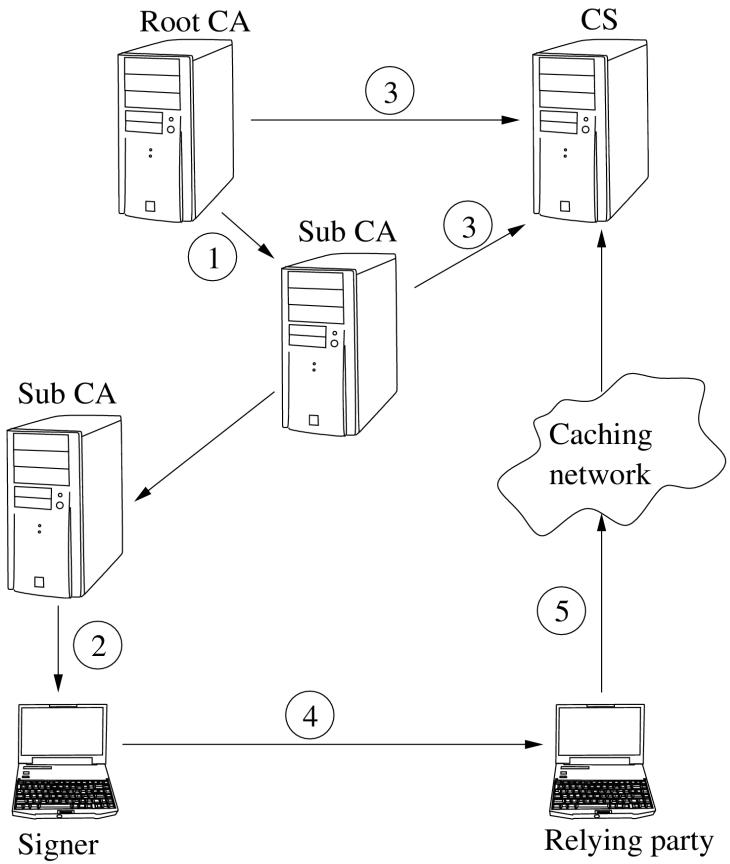

Conclusion

- The use of short-lived certificate offers the following advantages:
  - Scalability (insensitive to size of certificate population)
  - Low traffic volume
  - Even traffic rates (no peaks)
  - Low dependency on connectivity
    - e.g. Push or prefetch PoV when connected
  - Reduced number of trust anchors (eliminates the VA)
  - Improved facilities for cross domain operation