Addressing the Costs of Revocation Information Delivery for Public Key Infrastructures

{ **Cable**Labs® ◦ Open**CA**}

Massimiliano Pala <m.pala@cablelabs.com>

Problem Statement

- One of the most Important step in validating certificates is to check for the availability of revocation information
- Today, the deployment of revocation infrastructures, especially for large PKIs (i.e., the population of active certificates is quite large), comes with large price tags attached to it
  - OCSP does not scale well for very large PKIs
  - Costs are directly proportional to the size of the certificate population
- Different techniques have been deployed to be able to deploy OCSP responders, especially for high-frequency environments
  - Even when using pre-computed responses, the deployment costs are associated to the entire certificates' population

General Considerations

- Today, the larger the PKI is, the higher the costs of providing a good revocation infrastructure is
  - Pre-Signing one response for each certificate for each validity interval
- OCSP not optimized for the common case. Specifically, for non-revoked certificates, we notice that there is actually “no revocation information”
  - CSPs manage hundreds if not thousands of different PKIs, each of which requires the distribution of revocation information. The population of active certificates, in this case, can easily reach the hundreds of millions

What can we do to optimize the system for the most common case (valid) ?

General Considerations (cont.)

- Asking for the status of a certificate, in practice, means asking for the status of the full chain (up to the Root)
  - What can we do to limit the number of needed round trips ?
- OCSP allows for multiple single responses to be added to the basic response from a single responder
  - Section 4.2.2.3 of RFC 6960 allows for extra responses to be added, however, at the same time, it suggests that this option SHALL NOT be used (only for improving the pre-generation performance or cache efficiency as per RFC5019, Section 2.2.1)

How can we provide full-chain information that the client can trust ?

General Considerations (cont.)

- IETF mostly focused its attention on the Internet PKI where the trust model that mostly uses certificates for server-side authentication and other type of credentials for client-side authentication
- Other New and Existing environments (e.g., IoT, Industrial, Cable, etc.) implement a different trust model where both sides use certificates for authentication and the population of active certificates can easily grow into the millions
  - Client-Side revocation checking, specifically for large PKI environments, is often quite expensive to provide
  - Current deployments prefer to use cloud-provided access controls instead of leveraging revocation infrastructures

Current Approaches

- Today, revocation information is mainly distributed via three different paths
  - Certificates Revocation Lists (CRLs)
  - Online Certificate Status Protocol Responses (OCSP)
  - ~~Other Proprietary Methods~~
- We focus our discussion on the first two methods

Method #1: CRL

- Certificate Revocation Lists provide an authenticated list of revoked certificates identified by their serial numbers
- CRLs grow due to revocation events and shrink due to expiration events, therefore the size of a CRL is mostly unpredictable and can grow beyond acceptable sizes
- The `issuerDistributionPoints` could be used to partition the space and limit the worst-case scenario
  - the partition of the certificate space has to be decided in advance
  - could be avoided with the introduction of a scope extension in the CRLs specifying that it applies to a range of certificate serial numbers

Method #2: OCSP

- OCSP responses provide the revocation information related to one (or more) certificates as requested by the client
- Although OCSP responders can provide responses for certificates the client did not request the status for (i.e., for full chain checking), no implementation, today, uses this feature
  - Trust Model Issue – why should I trust a CA to report its own revocation status or the revocation status of CAs up the chain?
- OCSP responders pre-compute responses and serve them from CDNs
  - A large number of signatures must be generated at each validity period
  - Can we reduce the number of signatures (and associated costs) for each CA ?

Method #3: Range Responses

- There is NO revocation information to provide in the ranges between revocation entries (i.e., only revoked serial numbers have valid revocation information associated with them)
  - The full range does not have any revocation info
- Many certificates issued by a PKI are seldom validated – in a pre-computed OCSP responses environment, some of the pre-computation is not relevant, but still needed
  - Not all the population of certificates is validated every day
- By providing responses that group together multiple certificates we can reduce the number of required signatures for each validity period
  - reduces the number of “wasted” signatures

Method #3: Range Responses (cont.)

Num of Rev = 3; Num of Certs = 11.5+ M

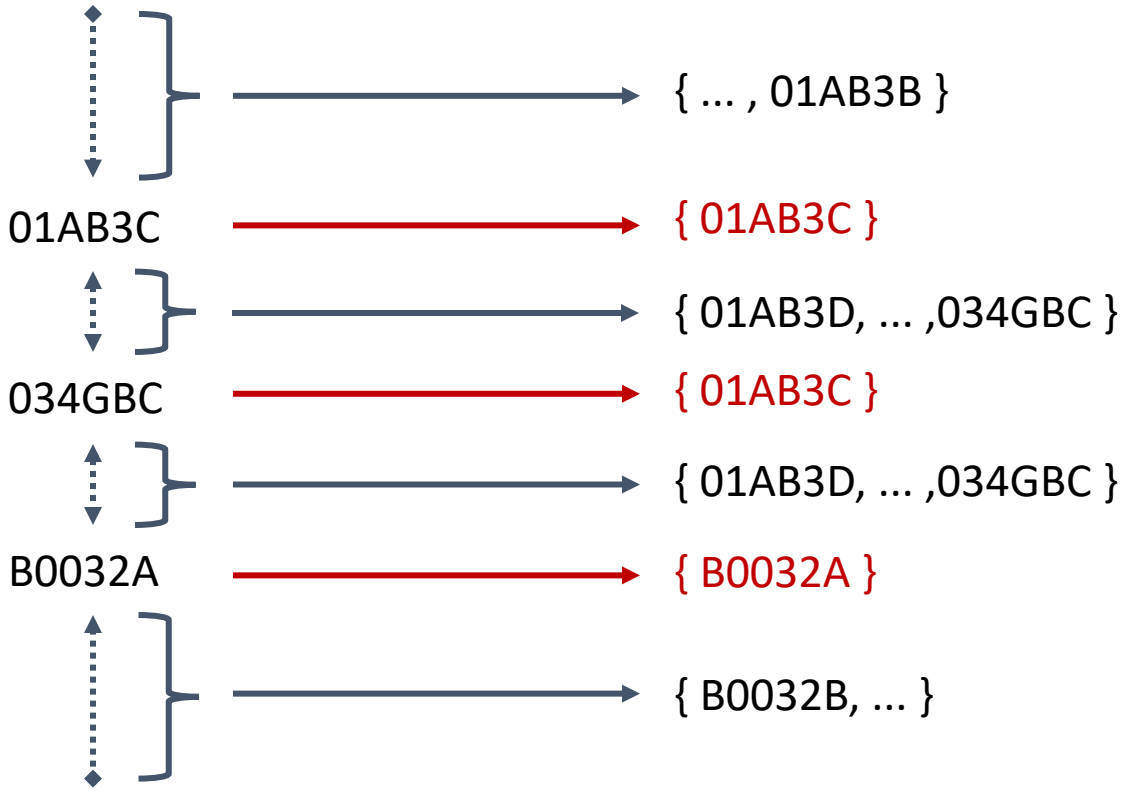

Responses = 2x + 1 (x = Number of Revocations)

Method #3: Range Respons… (cont.)

Num of Rev = 3

01A…

B0032A…

7 Responses

vs.

11,535,147+ Responses
(0xB0032A+)

Re…ses = 2x + 1 (x = Number of Revocations)

Current Status of Discussion

- Conversation has happened on IETF SecDispatch, LAMPS, and PKIX
- Mixed Reaction to the problem statement and proposal for a solution
  - There has been agreement that the limitations are there
  - Mixed reaction to the proposed range-based responses approach
  - Positive feedback on providing full chain responses by encapsulating responses from responders up in the chain
  - Some requests for gathering $$$ figures on the costs of revocation have been made to the list – we are contacting some of the largest providers to get some of these numbers
  - So far, we have received interests from some of the major PKI providers – DigiCert in particular supports the idea and its deployment.

Summary

- Revocation Information distribution is one of the pillar of trust infrastructures (current trust models)
- The OCSP protocol provides a good replacement for CRLs, however it is not optimized for the average case and costs depend on the size of the active certificate population
- With the introduction of New Environments and Scenarios (e.g., IoT) the size of new Trust Infrastructures will increase with the introduction of client-side authentication
- In order to reduce the costs of providing revocation information for large infrastructures, we propose to update the protocol to be able to provide “ranged” responses

Future Work

- What is the best path forward ?
  - Our initial work seems promising (a good fit for our industry and IoT)
  - More work is needed to gather revocation costs related to the Internet PKI (i.e., what is the average ratio revoked/total population) and be able to answer the inquiries from the ML(s)
  - BoF ? Individually Sponsored ? Existing WG ?

- Call for participation
  - Are there other CSPs that are interested (please contact us)

- ***References***

  - *Range Responses proposal for Revocation Information (placeholder – soon to be populated with an initial proposal) -* https://datatracker.ietf.org/doc/draft-pala-ocspv2/