Workgroup: SIDROPS
Internet-Draft:
draft-li-sidrops-roa-granularity-problem-
statement-00
Published: 23 October 2023
Intended Status: Best Current Practice
Expires: 25 April 2024
Authors: Y. Li J. Yao D. Ma
CNIC, CAS CNNIC ZDNS

On the Operational Granularity of RPKI ROA Management: Problem Statement and Requirements

Abstract

When using the Resource Public Key Infrastructure (RPKI) to perform route origin validation (ROV) with route origin authorizations (ROAs), there have been security and usability issues identified and reported. This memo revisits these issues from the perspective of the operational granularity of ROA management, demonstrates problems and their root cause with the existing ROA encoding scheme, summarizes design requirements to address them, and evaluates three potential solutions. Though neither of existing solutions satisfies all requirements, a hybrid solution composed of two existing schemes is recommended to use in ROA management.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 25 April 2024.

Table of Contents

1. Introduction

In the RPKI, an ROA is an attestation that the owner of an IP address block has authorized an autonomous system (AS) to originate routes in BGP for one or more prefixes within this address block[RFC6482]. ROAs are used to identify route hijacking with RPKI-ROV[RFC6483][RFC6811].

However, there are quite a few concerns about ROAs, in terms of security, reachability and scalability. These issues have been noticed from way back, but were attributed to the partial deployment of RPKI or human errors in configuring ROAs all the while. Departure from this widely accepted understanding, this memo explores the root cause of these issues in a new perspective, demonstrates problems with the existing ROA encoding scheme, and reveals that the root cause of them is the coarse-grained management of ROAs. For this reason, some security or usability issues will retain even when the RPKI finishes its deployment and all operators are careful enough.

Further, this memo summarizes the requirements in designing an efficient and concrete solution to resolve security and usability issues with ROAs, evaluates two possible solutions with these design principles and suggests the future direction for improvement.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Problem Statement

With the existing ROA encoding scheme, an IP prefix, as the root prefix, along with a maxLength parameter specifies the authorization of an AS originating a set of IP prefixes covered by the root prefix. These IP prefixes form a full binary tree rooted at the root prefix, and are referred to as a prefix block throughout this memo. In case that the maxLength parameter is missing, only the ROA prefix is authorized. Below, the problems with this scheme are demonstrated from the perspective of ROA management, and will retain even when the RPKI deployment is complete and there is no human errors.

3.1. Scattered Authorization Issuance

Authorization issuance is the most common and important operation in managing ROAs. As the network management becomes more and more flexible, route origin authorizations issued to an AS is always scattered, for business or engineering purposes, such as resource reservation, resource allocation, route aggregation, traffic engineering, to name only a few. In this case, it's difficult to determine a proper maxLength parameter in ROA encoding. On one hand, a too-short maxLength parameter can render the legitimate BGP routes that carry the IP prefixes covered by but not included in this ROA "invalid"; on the other hand, a too-long maxLength parameter can lead to forged-origin sub-prefix hijack[RFC9319].

Take the example shown in Figure 1 for instance. The resource holder owns an IP address block starting from 202.127.16.0 to

202.127.31.255, and issues an AS to originate 4 IP prefixes, A (202.127.16.0/20), B (202.127.16.0/21), C (202.127.24.0/21) and E (202.127.20.0/22), within this address block. In case that the maxLength parameter is set to 21, a BGP route carrying the prefix E will be identified as "invalid" in RPKI ROV and thus becomes unreachable. While if the maxLength parameter is set to 22, three other prefixes D, F, G will be included in the authorization. In this case, the attacker can use these more specific prefixes to hijack part of the traffic destined for the prefixes A, B or C, leading to security risks.

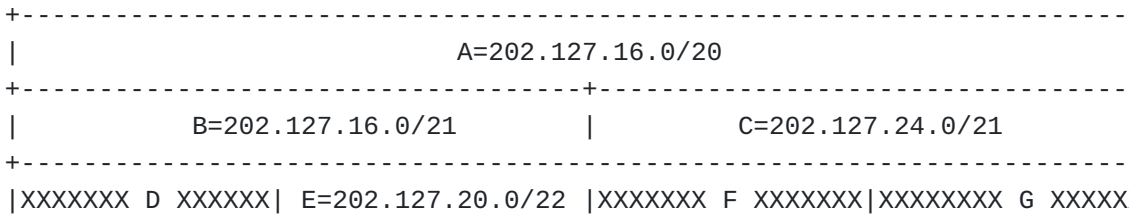

Figure 1: an example of scattered authorization issuance.

This dilemma is a natural consequence of the fact that the authorization is managed at a granularity of the bundle of several IP prefixes. This operational granularity is too coarse and is the root cause of the tradeoff to make between security and reachability. Worse yet, this granularity issue is essentially resulted from the fundamental encoding scheme taken by ROAs, and is thus hard to resolve completely within the existing framework.

A statistical study on 10 BGP RIBs, collected from the same collector at Jan 1st every year since 2014 till now, and with corresponding ROA records demonstrates how the potential security and reachability issues resulted by the dilemma in ROA encoding would be. In a RIB, a secure route prefix is the one that can be validated via ROAs; while, correspondingly, a route prefix that will be identified as "invalid" via ROAs is called unauthorized route prefix. Besides, the authorized but not routed prefixes may be utilized by the attacker to hijack all or part of the traffic destined for other routed prefixes. These affected prefixes are thus called insecure route prefixes.

As clearly demonstrated in Table 1, the fast growth of RPKI deployment does help increase the number of route prefixes protected by ROAs. Correspondingly, the increasing speed of the number of unauthorized route prefixes has slowed down obviously since 2019. However, the number of insecure route prefixes increases as well, with a similar trend as the increasing of secure route prefixes. In a word, the promotion of RPKI deployment can help protect more route prefixes but will also bring in more security risks due to improper configuration of ROAs.

3.2. Partial Authorization Revocation

In comparison with the authorization issuance, authorization revocation is an infrequent operation in managing ROAs, but it is important to keep the authorization status up-to-date. In case that the authorization revocation is not performed timely and properly, the IP prefixes that are no longer authorized to be originated by an AS will remain authorized, opening an attack vector to route hijacking as a result.

With the existing framework of managing ROAs, there is no "update" operation defined. Consequently, the operational granularity in authorization revocation is the whole ROA that previously issued. In case that only a part of the authorization is required to revoke, two steps are mandatory. First, all previously issued ROAs that contain this part should be withdrawal, this adds a burden to the network operator because he has to remember all previously issued ROAs and what IP prefixes are authorized in each of them. Second, all the IP prefixes from the other parts of original ROAs should be authorized again, triggering the case of "scattered authorization issuance".

Take the example shown in Figure 2 for instance. The original ROA contains 7 IP prefixes, where the authorization of two prefixes C (202.127.24.0/21) and F (202.127.24.0/22) is required to revoke. Then, the whole ROA should be withdrawal and the authorization of the rest 5 prefixes (A, B, D, E, G) should be issued again.

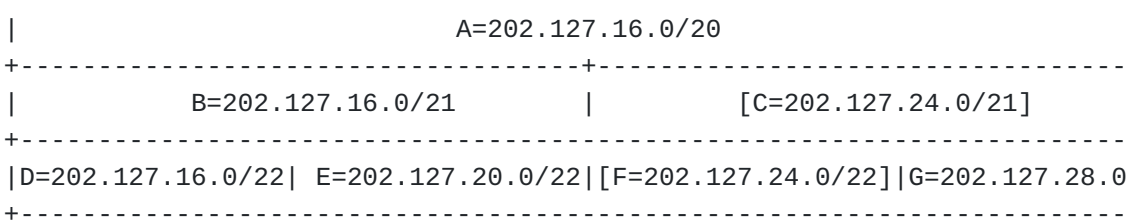

Figure 2: an example of partial authorization revocation

A statistical study on 10 months of RPKI data records summarizes the instances of authorization revocations every month. The results are shown in Table 2. There are 13~147 partial authorization revocations everyday on average, which have taken a non-negligible part (8.2%~49.8%) of total instances of revocations.

3.3. Multi-Authorization Resolution

Generally, different ASes may issue their ROAs independently without any negotiation. Thus, it's possible that their ROAs may overlap, leading to multi-authorization where the same set of IP prefixes are authorized to be originated from different ASes.

In some cases, multi-authorization is useful, such as for the purpose of supporting MOAS or for the mitigation of DDoS attacks[RFC9319]; but it may also be extremely harmful in some other cases, such as the case that one AS allocates a sub IP address space to its customer and their ROAs overlap. In this case, the attacker can use those IP prefixes authorized to but won't be used by the provider (which has been allocated to the customer) for route hijacking.

For example, suppose the provider owns an IP address block starting from 202.127.16.0/20, and is authorized to originate any IP prefix within this address block that is with a length not longer than 22. Then, it allocates a sub address block starting from 202.127.24.0/21 to its customer, which is authorized to originate any IP prefix within this sub address block that is with a length not longer than 23. As a result, their ROAs overlap, where the intersection contains 3 IP prefixes, C (202.127.24.0/21), F (202.127.24.0/22) and G (202.127.24.0/22). Generally, the provider will not announce BGP routes with these IP prefixes any more, since they have been allocated to the customer, but they are still included in the provider's ROA. This opens a near-permanent attack vector to route hijacking, where the attacker can use the IP prefixes in the intersection to hijack part of the traffic destined for the IP prefix A (202.127.16.0/20) owned by the provider, or that destined for the IP prefix C (202.127.24.0/21) owned by the customer.

To resolve this security risk, the provider should revoke the authorization of IP prefixes in the intersection between its ROA and its customer's ROA, triggering partial authorization revocation as a result.

A statistical study on 10 BGP RIBs and corresponding ROA records collected at the same day was conducted to evaluate the multi-authorization between two ASes, and the results are shown in Table 3. The case of possible resource allocation accounts for 48.8% ~ 76.3% of all instances in a day, where there is a potential security risk as discussed above.

3.4. A Brief Summary

As the growth of the Internet and the RPKI deployment, and in regard to flexible network operations, ROA management has to deal with scattered authorization issuance, partial authorization revocation and multi-authorization resolution. However, with the existing scheme, the operational granularity of issuing route origin authorizations is a bundle of IP prefixes that share the same prefix length, while the operational granularity of authorization revocation is a whole ROA. This sharply complicates the management

of ROAs in the mentioned scenarios, resulting in non-negligible security and usability issues.

The logical relation among the problems stated above are shown in Figure 3. More specifically, it opens attack vector to route hijacking that if the partial authorization revocation or the multi-authorization resolution does not complete timely or properly. Actually, the key operation of the multi-authorization resolution is to perform partial authorization revocation, where the key operation turns to be scattered authorization issuance finally. To complete this kind of issuance, there is a bad dilemma in determining the maxLength parameter to balance the security and usability. Besides, to properly perform partial authorization revocation, the network operator has to record all previously issued ROAs and what IP prefixes they contain. This can also be treated as a usability issue.

Accordingly, all the problems stated in this memo are essentially caused by the coarse-grained management of route origin authorizations with the existing scheme, and thus has nothing to do with the RPKI deployment and human errors. In another word, these security and usability issues will retain even when the RPKI finishes its deployment and all network operators are careful enough.

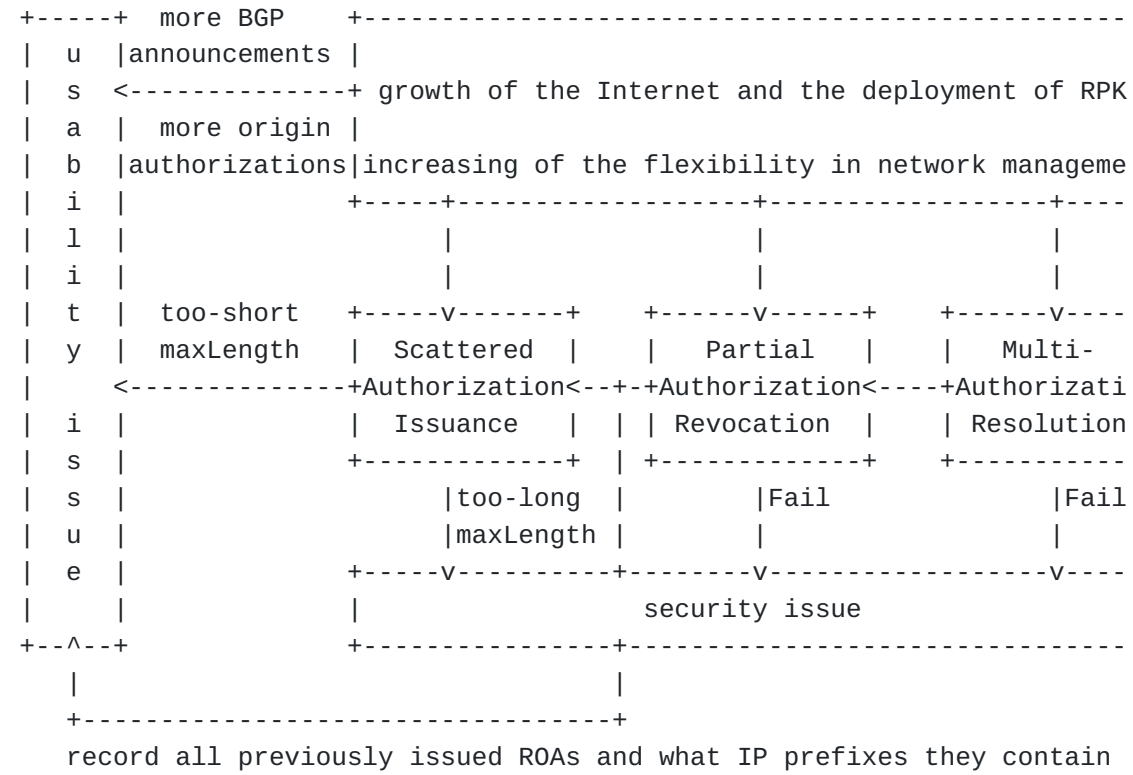

Figure 3: logical relation among stated problems.

4. Requirements for Encoding and Managing ROAs

In order to address the security and usability issues demonstrated above, a new scheme for encoding and managing ROAs is required. This memo summarizes the requirements and four design principles in response to security, reachability and scalability issues.

4.1. Fine Granularity

Given that the coarse granularity in authorization management is the root cause of the dilemma between security and reachability, the most important design principle is to break this limitation with a solution that achieves fine-grained authorization management, especially that enables the management at the prefix granularity.

4.2. Incremental Update

To facilitate authorization revocation and thus multi-authorization resolution, the support of incremental update of authorizations is important. This allows an ROA to be updated on demand without excessive withdrawal or redundant authorizations.

4.3. High Efficiency and Scalability

At last, the encoding efficiency as well as the scalability of the whole scheme should be taken into consideration, in regard to the fast growth of global routing tables and RPKI deployment. In another word, the management of ROAs should not add too much burden to network operators or BGP routers, when there are more and more origin authorizations and BGP routes.

5. Evaluation of Potential Solutions

In the current literature, there are 3 potential solutions that might be able to satisfy above requirements to some extent. The first two of them are extensions over the existing scheme, while the last one introduces a novel encoding scheme.

5.1. Single-prefix ROA

A straightforward solution is to eliminate the maxLength parameter for exact authorization issuance, by organizing ROAs as several single-prefix authorizations. In this way, authorizations are managed prefix by prefix and thus efficient efficient incremental updates can be achieved, since the only update operation toward an ROA, namely a single-prefix authorization, is to revoke it.

The only but a big drawback of this solution is the encoding efficiency and its scalability. Every authorized prefix has to be managed as separate authorizations, adding up the management burden.

In addition, the encoding of the authorization of a big prefix block (such as the AS0 ROAs) produces a huge number of entries, which may increase exponentially.

5.2. Minimal ROA

On basis of the single-prefix ROA, all authorized prefixes that form a full binary tree are aggregated into one ROA, with the maxLength parameter added back to specify the longest length of the IP prefixes contained in this authorization. This is called minimal ROA[RFC9319] and it minimizes the use of maxLength parameter while keeping the feature of exact authorization issuance. Essentially, it improves the encoding efficiency and scalability over the single-prefix ROA, but lacks in the support of efficient incremental updates.

5.3. Hanging ROA

Hanging ROA proposes to use a bitmap to encode a set of authorized IP prefixes. This solution has a very good nature to support authorization management at the prefix granularity. Besides, it can enable flexible and efficient incremental updates with the help of bit-masking operations. However, this solution may become inefficient in dealing with a big prefix block, where either a rather long bitmap or a large number of bitmaps are required, restricting its scalability.

5.4. Requirements Satisfaction

Figure 4 compares the 3 schemes in view of their satisfaction to every design requirements. In terms of encoding efficiency and scalability, two scenarios are evaluated separately, where the differences are the number and the volume of authorizations to encode. In the first case, there are only a few big prefix blocks, each containing 64 or more IP prefixes, to encode; while there are a large number of small prefix blocks to encode in the second case.

			encoding efficiency and scalabil	
	operational granularity	support of incremental updates	case1: a few big prefix blocks	case2: lots of s prefix blocks
Single-prefix ROA	prefix	Yes	Bad	Bad
Minimal ROA	a set of prefixes	No	Good	Bad
Hanging ROA	prefix	Yes	Bad	Good

Figure 4: summary of 3 potential solutions.

In summary, both the single-prefix ROA and the hanging ROA manage authorizations at the prefix granularity, and thus support efficient incremental updates. In regard to the encoding efficiency and scalability, the minimal ROA and the hanging ROA are better than the single-prefix ROA, but they are good at different scenarios. unfortunately, neither of the 3 potential solutions can perfectly satisfy all design requirements.

5.5. Performance Evaluation

10 sets of authorized route prefixes are generated with 10 BGP RIBs and their corresponding ROA records collected at the same day, and are used to evaluate the encoding efficiency and scalability of the 3 schemes. In encoding each of the 10 sets, the number of encoded blocks in IPv4 and IPv6 are measured respectively.

The results are recorded in Table 4. Clearly, the hanging ROA has the highest encoding efficiency, and reduces the encoding cost in IPv4 by 30.4% ~ 66.9% and 27.1% ~ 65.1% in comparison with the single-prefix ROA and the minimal ROA respectively, while the reductions in IPv6 are 1.5% ~ 47.1% and 1.5% ~ 45.9% respectively.

Table 5 shows the detailed results about the minimal ROA and the hanging ROA with different settings of defining what are big prefix blocks. Since in both scheme, AS0 ROAs are maintained in the same form, they are eliminated from the results. It is clear that the minimal ROA is more efficient to encode big prefix blocks than the hanging ROA; while the hanging ROA outperforms the minimal ROA in encoding small prefix blocks. Besides, there is no prefix block that contains more than 63 prefixes, namely the difference between the

prefix length and the maxLength is 6 or larger, in both IPv4 and IPv6. This reflects the trend that more and more network operators prefer to issue exact authorizations.

6. Recommendations

For authorization issuance, it is recommended to authorize exactly what are about to use in BGP with a hybrid encoding scheme. More specifically, hanging ROA should always be adopted to ensure efficient fine-grained management, unless when there is a need to authorize a big prefix block that contains more than 63 prefixes. Then, minimal ROA should be adopted alternatively.

Once the authorization status changes, the corresponding updates toward previously issued ROAs should be performed timely and properly, especially those for authorization revocation.

In case of resource allocation, the customer should obtain its authorizations as soon as possible to avoid reachability issues, while the provider should make careful checks and timely operations (revocations) to ensure that there is no unnecessary intersection between its ROA and its customer's ROA.

7. Security Considerations

This document analyzes and reveals the root cause of potential security risks under RPKI-ROV.

The recommended hybrid ROA encoding scheme MAY improve the security of using RPKI-ROV, and MUST NOT bring in additional security issues.

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, DOI 10.17487/

RFC6482, February 2012, <https://www.rfc-editor.org/info/rfc6482>.

[RFC6483] Huston, G. and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)", RFC 6483, DOI 10.17487/RFC6483, February 2012, <https://www.rfc-editor.org/info/rfc6483>.

[RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2012, <https://www.rfc-editor.org/info/rfc6811>.

[RFC9319] Gilad, Y., Goldberg, S., Sriram, K., Snijders, J., and B. Maddison, "The Use of maxLength in the Resource Public Key Infrastructure (RPKI)", RFC 9319, DOI 10.17487/RFC9319, January 2012, <https://www.rfc-editor.org/info/rfc9319>.

9.2. Informative References

[RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", RFC 3410, DOI 10.17487/RFC3410, December 2002, <https://www.rfc-editor.org/info/rfc3410>.

Appendix A. Statistic Results

A.1. Scattered Authorization Issuance

	2014-07-21	2015-07-21	2016-07-21	2017-07-21	2018-07-21	2019-07-21	2020-07-21	2021-07-21	2022-07-21	2023-07-21
secure route prefix	15K	22K	28K	38K	52K	101K	166K	256K	335K	412K
insecure route prefix	9K	11K	13K	16K	18K	37K	49K	71K	95K	127K
unauthorized route prefix	4K	3K	3K	6K	12K	7K	8K	10K	8K	9K
total	546K	596K	661K	724K	802K	893K	967K	1065K	1132K	1178K

Table 1: Statistics on different kinds of route prefixes.

A.2. Partial Authorization Revocation

	2022-10	2022-11	2022-12	2023-01	2023-02	2023-03	2023-04	2023-05	2023-06	2023-07
total instances of complete revocation	3235	3268	3696	51137	9051	6715	5644	10428	5958	4727
total instances of partial revocation	3207	409	979	4583	2670	2795	1624	1759	2969	1293

Table 2: Statistics of authorization revocation instances.

A.3. Multi-Authorization Resolution

	2014-07-21	2015-07-21	2016-07-21	2017-07-21	2018-07-21	2019-07-21	2020-07-21	2021-07-21	2022-07-21	2023-07-21
possible MOAS	38 (7%)	70 (9%)	70 (5%)	87 (4%)	170 (4%)	551 (7%)	643 (6%)	1701 (6%)	2410 (6%)	2615 (6%)
resource allocation	439 (76%)	580 (74%)	951 (72%)	1664 (67%)	2124 (49%)	4923 (63%)	6708 (64%)	19052 (69%)	24995 (67%)	29739 (65%)
other cases	97 (17%)	135 (17%)	299 (23%)	732 (29%)	2056 (47%)	2316 (30%)	3183 (30%)	6746 (25%)	10098 (27%)	13572 (30%)

Table 3: Statistics of multi-authorization instances

Appendix B. Performance Evaluation

	IPv4			IPv6		
	single prefix ROA	minimal ROA	hanging ROA	single prefix ROA	minimal ROA	hanging ROA
2014-07-21	19289	17888	7394	1017	995	855
2015-07-21	28512	26512	10562	1834	1804	1459
2016-07-21	32523	30188	13069	2665	2627	2016
2017-07-21	42553	39755	16366	4003	3919	2826
2018-07-21	58925	54877	21815	5835	5693	3938
2019-07-21	92366	86835	32095	10382	10228	6606
2020-07-21	163924	155150	54198	19790	19456	11059
2021-07-21	240474	222451	82579	34886	34059	18439
2022-07-21	323632	298512	113996	55957	54443	30281
2023-07-21	403635	370282	144493	76952	74137	41851

Table 4: Encoding cost with 3 schemes on 10 sets of authorizations.

threshold of block size	IPv4				IPv6			
	Big prefix block		Small prefix block		Big prefix block		Small prefix block	
	minimal ROA	hanging ROA	minimal ROA	hanging ROA	minimal ROA	hanging ROA	minimal ROA	hanging ROA
1	11995	10952	358287	136908	878	1088	73259	4083
3	1894	2136	368388	142816	193	363	73944	4150
7	208	526	370074	144094	21	166	74116	4169
15	48	318	370234	144221	10	84	74127	4176
31	6	25	370276	144479	0	0	74137	4185
63	0	0	370282	144493	0	0	74137	4185

Table 5: Encoding cost with different settings of block size threshold.

Authors' Addresses

Yanbiao Li
CNIC, CAS
No.2 Dongsheng South Rd, Haidian District
Beijing
100083
China

Email: lybmath@cnic.cn

Jiankang Yao
CNNIC
No.4 South 4th Street, Zhongguancun
Beijing
100190
China

Email: yaojk@cnnic.cn

Di Ma
ZDNS
No.4 South 4th Street, Zhongguancun
Beijing

100190
China

Email: madi@zdns.cn