Issuance Policies of Route Origin Authorization with a Single Prefix and Multiple Prefixes: A Comparative Analysis

Zetong Lai[1], Zhiwei Yan[2], Guanggang Geng*[3], Hidenori Nakazato[4]
Department of Cyber Security, Jinan University, Guangzhou, PR China[1,3]
National Engineering Laboratory for Naming and Addressing, China Internet Network Information Center, Beijing, PR China[2]
Faculty of Science and Engineering, Waseda University, Tokyo, Japan[4]

***Abstract*—Resource Public Key Infrastructure (RPKI) is a solution to mitigate the security issues faced by inter-domain routing. Within the RPKI framework, Route Origin Authorization (ROA) plays a crucial role as an RPKI object. ROA allows address space holders to place a single IP address prefix or multiple IP address prefixes in it. However, this feature has introduced security risks during the global deployment of RPKI. In this study, we analyze the current status of ROA issuance and discuss the impact of using two ROA issuance policies on RPKI security and synchronization efficiency. Based on the aforementioned work, recommendations are proposed for the utilization of ROA issuance policies.**

***Keywords*—BGP; RPKI; route origin authorization; inter-domain routing security; computer network protocols; routing**

I. Introduction

The Border Gateway Protocol (BGP) [1] is one of the most vital protocols on the Internet, responsible for the exchange routing and reachability information among autonomous systems (AS) on the Internet. However, the design of BGP neglected security considerations and the decentralized nature of the Internet, consequently giving rise to numerous security issues. Among these, BGP route hijacking poses the most severe risk, capable of triggering a cascade of catastrophic consequences such as data breaches, network outages, and malicious attacks [2]. To mitigate the issue of BGP route hijacking, the Internet Engineering Task Force (IETF) Secure Inter-Domain Routing (SIDR) working group has devised RPKI and consistently refined it.

RPKI is rooted in the concept of cryptographically verifying BGP update messages [3]. RPKI utilizes digital signatures to authorize and allocate Internet Number Resources (INR) [4], and verifies BGP update messages by using cryptographical RKPI objects. Much research has been conducted on enhancing the RPKI during the process of global deployment. In terms of the trust model, in 2016, Hari et al. [5] proposed a basic framework for decentralized internet infrastructure based on blockchain. This framework abstracts the allocation of IP address prefixes and the mapping relationship of IP address prefixes and AS Numbers (ASN) as transactions on the blockchain. By leveraging the distributed and tamper-resistant properties of the blockchain [6], preventing malicious operations and reducing the centralization of authority in the existing RPKI trust model. In terms of potential attack risks, Hlavacek et al. [7] explored the dependency of RPKI on DNS components and proposed that disruptions to DNS resolvers can lead to RPKI failures. Additionally, Hlavacek et al. [8] introduced a downgrade attack on RPKI and analyzed the potential damage caused by such attacks in existing RPKI deployment environments, providing defense recommendations based on these analyses. In terms of ROA security, Gilad et al. [9] conducted research on the improper use of the maxLength field in ROA, which poses security risks to RPKI, and provided configuration recommendations for the maxLength field.

This study focuses on ROA security. ROA is the most prevalent object in RPKI. The eContent structure of ROA includes a version field, an asID field, and an ipAddrBlocks field [10]. The version field defaults to zero. The asID field contains a single AS number, authorized by address space owners as the origin for IP address prefixes. The ipAddrBlocks field contains a list of one or more IP address prefixes that will be announced, allowing address space owners to place one or more IP address prefixes in ROA. However, when placing multiple IP address prefixes in ROAs, there is a security issue where INRs are unexpectedly validated as invalid, thereby diminishing the reliability of RPKI. In this study, we found this security issue arises only when ROA overclaims. Through further analysis, we attributed this security issue to the fate-sharing nature of ROA with multiple prefixes. In contrast, the absence of the fate-sharing nature in ROA with a single prefix avoids this security issue. Additionally, we identified two scenarios triggering this security issue through experiments. Then we analyzed the current ROA situation and found that many address space holders choose to use the issuance policy of ROA with multiple prefixes. This choice poses security risks to the current RPKI production environments. But compared to ROA with a single prefix, ROA with multiple prefixes offers the advantage of reducing ROA data volume. Requiring using the issuance policy of ROA with a single prefix in the RPKI production environment would impact the synchronization efficiency of RPKI. To evaluate this impact, we conducted experiments to compare the synchronization times under two different ROA issuance policies. The experimental results indicate that the increased synchronization time resulting from using the issuance policy of ROA with a single prefix is acceptable. Through the aforementioned works, we provided recommendations for using the issuance policy of ROA with a single prefix as the preferred option, and promoted the formulation of IETF Request for Comments (RFC) 9455 [11], enhancing the security of the RPKI.

This paper is organized as follows: Section II introduces the overview of the RPKI as the foundation for understanding this paper, Section III presents the analysis of the current ROA situation, Section IV describes security issues arising from ROA with multiple prefixes overclaiming, Section V shows our evaluation of the impact on synchronization efficiency in the current RPKI production environment when using the issuance policy of ROA with a single prefix, Section VI concludes our work.

II. Overview of RPKI

The RPKI system is primarily comprised of a certificate issuance system, a certificate storage system, and a certificate synchronization and verification mechanism. As illustrated in Fig. 1, the certificate issuance system allocates INRs through issuing certificates, followed by storing certificates in the certificate storage system. RPKI Relying Party(RP) synchronizes and verifies RPKI certificates and signature objects, and then provides the verification result to BGP routers for filtering purposes.

A. Issuance System

The certificate issuance system adopts a hierarchical certificate model that aligns with the allocation architecture of INR. At the top level, the Internet Assigned Numbers Authority (IANA) allocates INRs to the RIRs, which manage and allocate address spaces within their respective regions. RIRs allocate their INRs to the National Internet Registries (NIR), the Local Internet Registries (LIR), or the Internet Service Providers (ISP), who allocate INRs downstream to smaller network operators.

RPKI employs five independent RIRs as the trust anchors (TAs), which are AfriNIC, APNIC, ARIN, LACNIC, and RIPE NCC. RPKI is deployed through either the hosted model or the delegated model [12]. With the hosted model, RIR bears the responsibility of maintaining RPKI and providing CA service. This allows address space holders to focus on creating and maintaining ROAs. With the delegated model, address space holders are obliged to operate their CAs to create and maintain ROAs. Such a model affords address space holders autonomy in managing their IP address resources, reducing their reliance on RIR.

CA is an entity that is responsible for issuing CA certificate and end-entity (EE) certificate. The allocation of INR between CAs requires the parent CA to generate and sign a CA certificate for the child CA. After establishing a relationship between the parent CA and the child CA, the child CA is required to periodically request the parent CA to update the CA certificate to maintain the validity of the certificate chain. Krill [13], which is a widely used CA software, implements this mechanism by setting the request periodic interval to ten minutes. When CA allocates IP address prefixes to AS, CA needs to generate an EE certificate for AS. Once generated, the EE certificate is required to sign the ROA content that has been encapsulated using the Cryptographic Message Syntax (CMS) format [14]. The EE certificate and ROA have a one-to-one correspondence relationship. To simplify ROA issuance and revocation processes, the EE certificate is embedded in the corresponding ROA.

B. Storage System

RPKI storage system is comprised of multiple repository publication points, CAs store their CA certificates, ROAs, and Certificate Revocation Lists (CRLs) in their respective repository publication points. The repository publication point establishes a manifest [15] based on the stored files. Manifest is beneficial for detecting replay attacks and unauthorized in-flight modification or deletion of signed objects. Upon authorization of INRs is modified by CA, a real-time message will be promptly dispatched to notify its repository publication point to update RPKI objects.

In the RPKI storage system, the repository publication points are interconnected via two fields in the CA certificate, namely Subject Information Access (SIA) and Authority Information Access (AIA) [16]. The SIA field records the repository publication point address of CA, thereby facilitating the search for certificates issued by CA. Meanwhile, the AIA field records the repository publication point address of the parent CA, thereby enabling the retrieval of certificates issued by the parent CA. By utilizing the two aforementioned fields, it is theoretically feasible to systematically traverse the entire RPKI repository system.

The storage system supports data synchronization by means of both the RPKI Repository Delta Protocol (RRDP) [17] and rsync [18]. Considering the broad support for rsync across multiple operating systems, the SIDR working group chose to utilize rsync as the synchronization protocol for RPKI during its initial design. This choice promotes the widespread adoption and deployment of RPKI. Although rsync has implemented the incremental synchronization mechanism to reduce synchronized data, this approach is in high demand on computational resources. Hence, the SIDR working group devised RRDP as a substitute for rsync [19]. By utilizing storing space to decrease the demand for computational resources, RRDP requires that every repository publication point maintains updated files, documenting all modified operations (for example, updated manifests and CRLs, newly issued certificates, or ROAs) along with their corresponding timestamps in the repository publication point.

C. Synchronization and Verification Mechanism

RP is a critical component in the RPKI synchronization and validation mechanism. RP uses Trust Anchor Locators (TALs) to retrieve the CA certificates and public keys of each TA. The corresponding repository publication point address is obtained from the SIA field in the CA certificate. Afterward, RP synchronizes RPKI objects from the repository publication points of TAs by using either RRDP or rsync, with RRDP being the preferred synchronization option, and continues to synchronize repository publication points of the child CAs downwards. After synchronizing RPKI objects to the local cache, RP validates them by verifying each object along the certificate chain from top to bottom. Following this, RP parses the mapping relationships of IP address prefixes and ASN recorded in valid ROAs to generate a route filtering table. By default, common RP software typically synchronizes and validates at intervals of one hour or less [20].

The BGP router in AS utilizes the RPKI to Router (RTR) [21] protocol to regularly fetch the route filtering table from

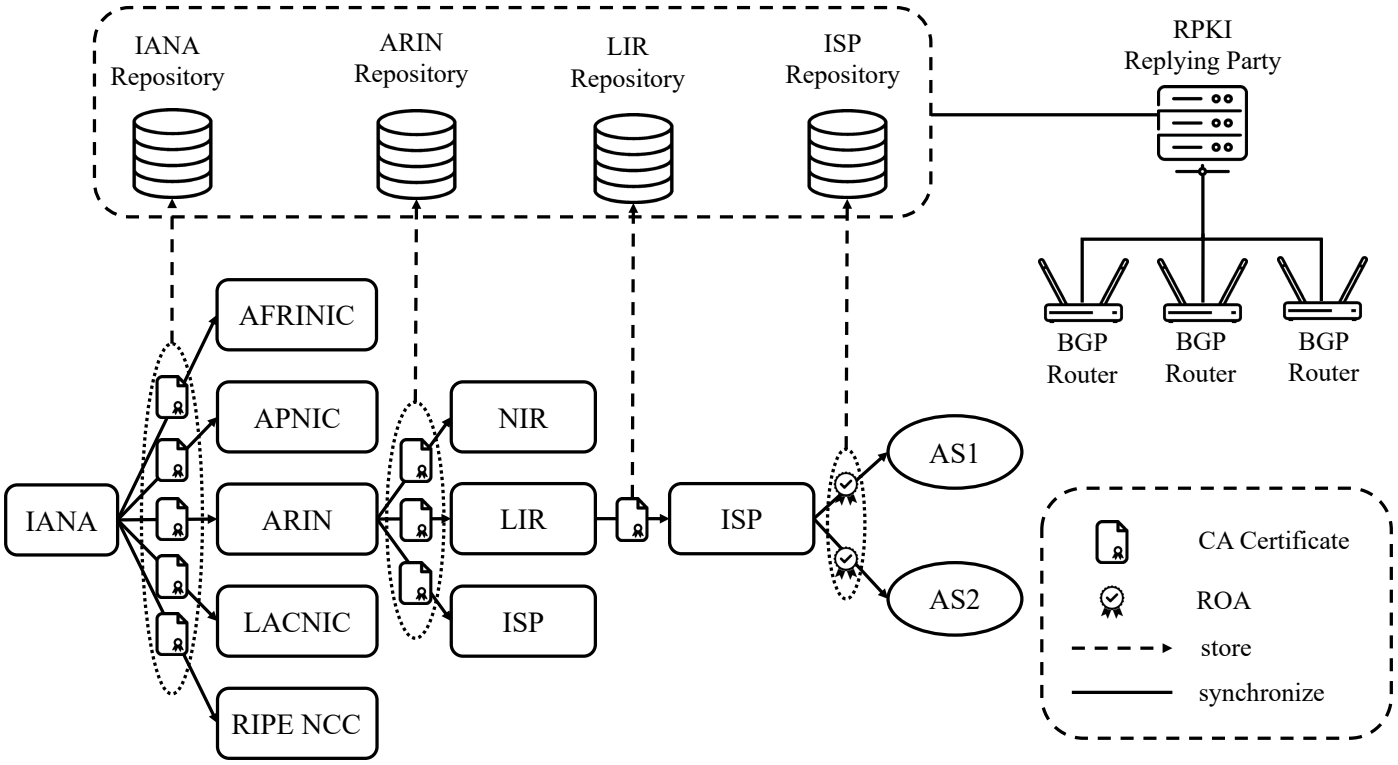

Fig. 1. RPKI system. The certificate issuance system is displayed below, the certificate storage system is shown above, and the certificate synchronization and verification mechanism is presented on the right side.

TABLE I. THE RELATIONSHIP BETWEEN VRPS AND THE VALIDITY OF ROUTES

IP address prefix of route	VRP match ASN of route	VRP mismatch ASN of route
Not covered by VRP	NotFound	NotFound
Covered by VRP	Valid	Invalid

RP. BGP router utilizes the route filtering table to perform route origin validation (ROV) [22] on the received BGP announcements, thereby sieving out invalid BGP routes. The relationship between the Validated ROA Payloads (VRPs) [23] in the route filtering table and the validity of the routes in BGP announcements is shown in Table I.

Covered by the IP address prefix of VRP refers to the length of the IP address prefix in VRP is shorter than that in route, and all the bits specified by the IP address prefix length of VRP are identical between VRP and route. Valid routes are accepted by the BGP router while invalid ones are rejected. The routes with the verification status of NotFound are accepted by default. BGP router administrator retains the ability to adjust the acceptance of routes with the verification status of NotFound in accordance with individual needs and preferences.

III. ROA ANALYSIS

In this section, we made the following analysis to elucidate the current ROA situation. The data utilized for this analysis was provided by RIPE NCC and Internet Multifeed Co. [24], up until February 25th, 2023.

TABLE II. THE NUMBER OF GLOBAL ROA

ROA type	Quantity
Total ROA	139484
ROA with a single prefix	110944
ROA with multiple prefixes	28540

TABLE III. THE NUMBER OF GLOBAL IP ADDRESS PREFIX

ROA type	Quantity
Total ROA	404101
ROA with multiple prefixes	293157

As shown in Table II, approximately 139484 ROA objects were globally issued. Further analysis reveals that around 110944 (79.54% of all ROA objects) ROAs contain a single IP address prefix, while the remaining 28540 (20.46% of all ROA objects) ROAs contain multiple IP address prefixes. Calculating the number of IP address prefixes within all ROAs with multiple IP prefixes, the statistical results are presented in Table III. Among 28,540 ROAs contain two or more IP address prefixes with a total of 293,157 IP address prefixes. Notably, despite the greater number of ROAs with a single prefix, the IP address prefixes contained in ROAs with multiple prefixes constitute 72.55% of the total IP address prefixes.

TABLE IV. THE AVERAGE SIZE OF EACH ROA IN FIVE TYPES OF ROAS

The number of IP address prefix in ROA	average size (bytes)
1	1999
2-10	1915
11-50	2157
51-100	2785
>100	5677

TABLE V. THE NUMBER OF ROA AND IP PREFIX ADDRESSES ISSUED WITH TWO POLICIES AMONG FIVE RIRS

RIR	ROAs with a single prefix	ROAs with multiple prefixes	IP address prefixes in ROAs with multiple prefixes
AfriNIC	2999	319	1562
ARIN	55943	2629	16166
APNIC	16543	6810	88166
LACNIC	15318	2081	16398
RIPE NCC	20141	16701	170865

Additionally, ROAs with multiple prefixes have been further categorized into four types based on the number of IP address prefixes contained in them: those containing 2-10, 11-50, 51-100, or more than 100 IP address prefixes. These categories of ROAs were analyzed alongside ROAs with a single prefix.

Table IV demonstrates that ROAs containing more than 100 IP address prefixes are, on average, only 2.8 times larger than ROAs containing one or two to ten IP address prefixes. It illustrates the effective reduction of both the quantity and size of ROA achieved by placing multiple IP address prefixes into one ROA.

Furthermore, Table V shows an analysis of ROA data in five RIRs. The quantity of ROAs with a single prefix is more than ROAs with multiple prefixes within each RIR. However, different RIRs have different ROA issuance policies. In AfriNIC and ARIN, the majority of IP address prefixes are issued via ROAs with a single prefix. The situation is reversed while in APNIC and RIPE NCC. Especially in RIPE NCC, the number of ROAs containing two or more IP address prefixes closely approximates ROAs containing only one single IP address prefix. In LACNIC, the number of IP address prefixes in both types of ROA is almost evenly divided.

IV. SECURITY RISK OF OVERCLAIMING

This section introduces the existing mitigation measures for the security risk of overclaiming and their shortcomings, then outlines two scenarios that using the issuance policy of ROA with multiple prefixes leads to INRs being unexpectedly validated as invalid due to overclaiming, and finally describes the adverse effects on routing security, and proposes mitigation strategies.

A. Shortcomings of Existing Mitigation Measure

The initial version of the certificate validation procedure requires that any certificate containing INR not held in the issuing certificate will be verified as invalid. The certificate signed by an invalid certificate is also verified as invalid. When the parent CA transfers or reclaims INRs, the CA certificate of the child CA will not refresh at once, causing the child CA to overclaim the transferred or reclaimed INRs. Consequently, any CA certificates or ROAs issued by the child CA will be verified as invalid before the CA certificate of the child CA is updated, irrespective of whether they contain transferred or reclaimed INRs.

To mitigate potential adverse effects on routing security, the IETF SIDR working group modified the certificate verification algorithm [25]. By using the modified algorithm, certificates and ROAs that do not contain transferred or reclaimed INRs are verified as valid. This modification effectively mitigates the issue of downstream certificate becoming invalid due to the issuing certificate being overclaimed. With the modified algorithm, utilizing the issuance policy of ROA with a single prefix, ROA overclaiming would only affect itself. However, when utilizing the issuance policy of ROA with multiple prefixes, even if ROA overclaims only one transferred or reclaimed INR, all INRs contained in the ROA will be verified as invalid due to the fate-sharing nature. This will cause the BGP router to filter routes inaccurately.

B. Parent CA and Child CA Deploy Repository Publication Points on Different Servers

As illustrated in Fig. 2, the parent CA initially allocated 192.168.1.0/24 and 192.168.2. 0/24 to the child CA. The child CA allocated IP address prefixes to AS65000 and AS65001 by issuing two ROAs: one containing 192.168.1.128/25 and 192.168.2.128/25, authorizing AS65000 as the origin; the other containing 192.168.2.0/25, authorizing AS65001 as the origin. The parent CA and child CA deployed repository publication points on different servers.

After a period of operation, the parent CA reclaimed 192.168.1.0/24 from the child CA. Subsequently, the child CA sent a request to the parent CA to update its CA certificate. Upon receiving this request, the parent CA notified the repository publication point I to update the CA certificate of the child CA and returned a response to notify the child CA that the update had been completed. The child CA received the response and notified the repository publication point II to update ROAs.

If the repository publication point is working, as the response to the update notification, the repository publication point II will revoke ROAs and generate ROAs that do not contain any IP address prefixes in the range of 192.168.1.0/24. However, due to a malfunction in the publication program, the repository publication point II could not update the ROAs. When RP attempted to synchronize data from the repository publication point II, it discovered that the RRDP service provided by the publication program was not working. Therefore, RP switched to utilizing rsync to synchronize data from the repository publication point II. During validating the RPKI objects, RP identified that 192.168.1.128/25 contained in 65000.roa was not held in the CA certificate of the child CA. Consequently, 65000.roa was validated as invalid. As a result, the route that announced AS65000 as the origin of 192.168.2.128/25 would be validated as NotFound or invalid.

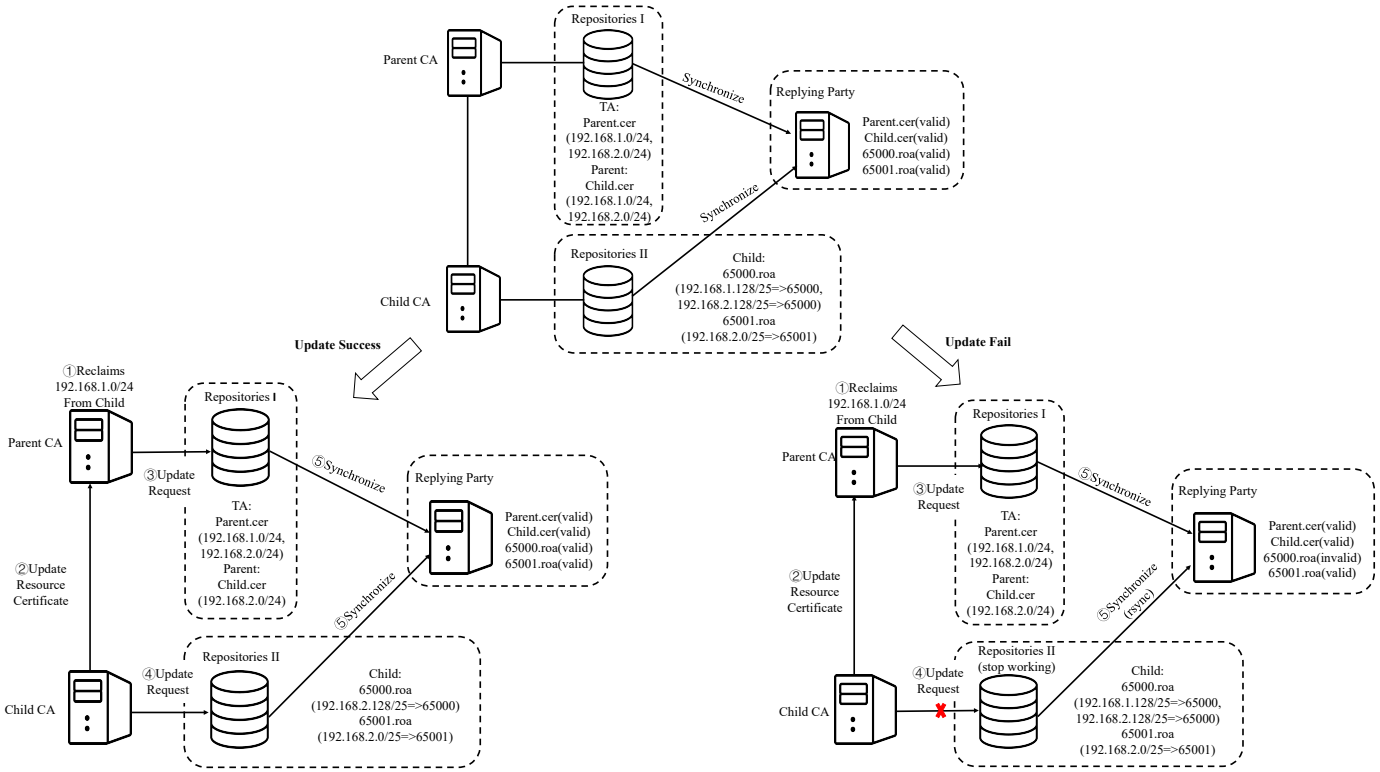

Fig. 2. When the repository publication points deploy on different servers, different results caused by publication program is working or malfunctioned during updating.

C. Update Latency Between Parent CA and Child CA

As shown in Fig. 3, the initial state mirrors that of Fig. 2, but both the parent CA and the child CA deployed the repository publication points on the same server. After a period of operation, 192.168.1.0/24 held by the parent CA was reclaimed. Following updating the CA certificate of the parent CA, 192.168.1.0/24 was not contained in the CA certificate. However, because the child CA has not updated its CA certificate, the CA certificate still contained 192.168.1.0/24 and the 65000.roa issued by it also contained 192.168.1.0/24. Until the child CA periodically sends the certificate update request, the parent CA updates the CA certificate of the child CA and the child CA updates the ROAs issued by it.

During update latency, if RP synchronizes the data from the repository publication points, the 65000.roa will be validated as invalid due to containing 192.168.1.128/25, which is not held by the CA certificate of the parent CA. This would cause the route that announced AS65000 as the origin of 192.168.2.128/25 to be validated as NotFound or invalid.

D. Security Risk and Mitigation Strategies

In the scenarios described in Sections IV.B and IV.C, the invalidation of 65000.roa would result in the absence of the VRP "192.168.2.128/24=>65000" from the VRP set acquired by the BGP router from the RP. When the BGP router receives a BGP announcement "192.168.2.128/24 originate from AS65000", if there exists a VRP whose prefix covers 192.168.2.128/24 in the VRP set, such as "192.168.0.0/16=>65002", the BGP router will validate this BGP announcement as invalid and rejected it. When traffic with a destination address within the 192.168.2.128/24 range passes through the BGP router, it will be forwarded to AS65002. Such route leakage will lead to severe performance degradation or even network outage [26]. If there is no existing VRP whose prefix covers 192.168.2.128/24 in the VRP set, the BGP router will validate this BGP announcement as NotFound and retain it. In this scenario, the 192.168.2.128/24 has lost the protection of RPKI, allowing malicious AS to launch BGP hijacking by crafting specific BGP announcements to steal traffic.

Both scenarios can be mitigated by eliminating the fate-sharing nature by adopting the issuance policy of ROA with a single prefix. Overclaiming triggered by the scenario described in section IV.B is rarely, because it is caused by software malfunctions. The scenario described in section IV.C may occur each time INRs from the child CA are reclaimed. In this scenario, except for adopting the issuance policy of ROA with a single prefix, the risk of overclaiming can be mitigated by promptly notifying the administrators of the child CA to manually update the CA certificate. However, because this requirement is difficult to accomplish, the existing CA software provides the periodical certificate update service. In addition, when the resources are reclaimed due to expiration without the awareness of administrators, manual and prompt update of CA certificates is impossible. Evidently, adopting the issuance policy of ROA with a single prefix emerges as the simplest and most efficacious method.

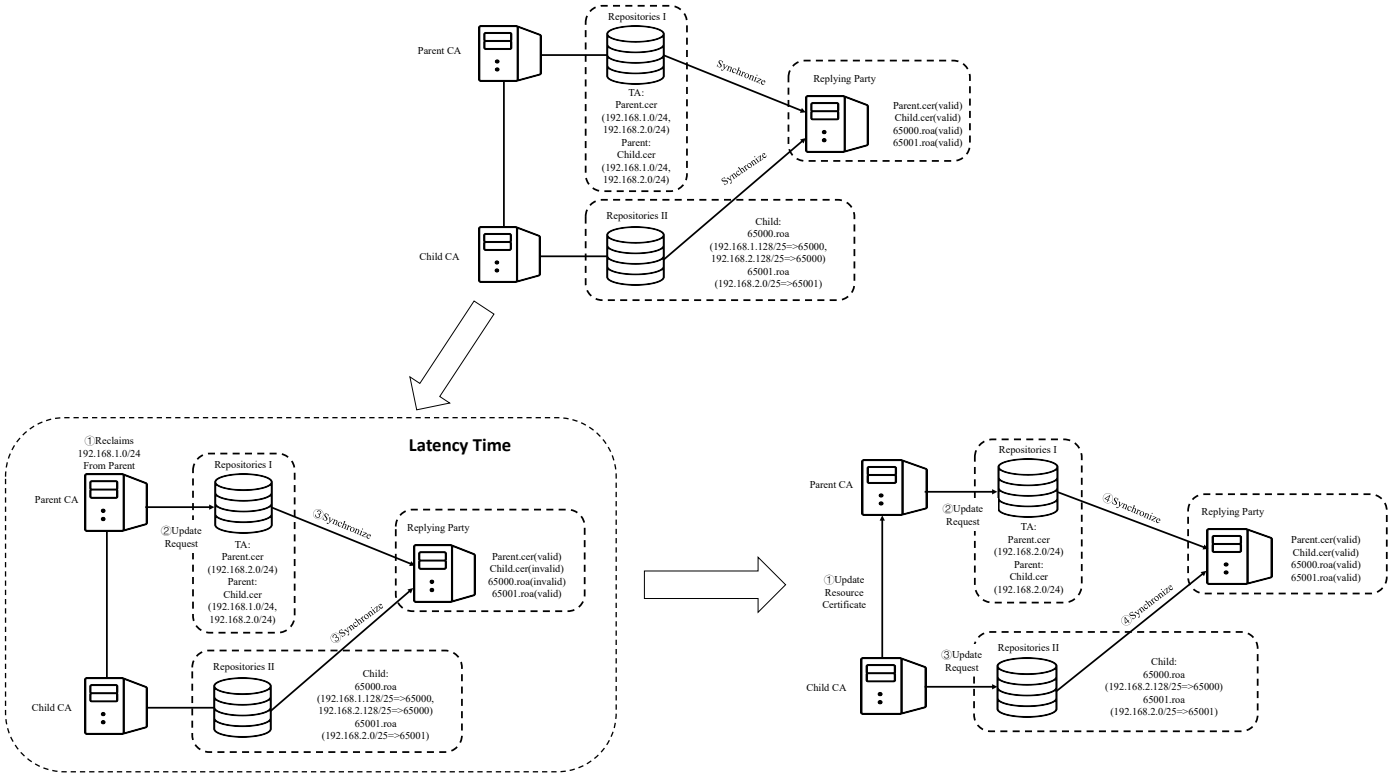

Fig. 3. When the resource of the parent CA change, the change of ROAs verification status during and after latency time for update.

V. Synchronization Efficiency

As shown in Table IV, it can be concluded that when the INR quantity is the same, using the issuance policy of ROA with multiple prefixes can significantly reduce the data size of ROA. The data size will affect the efficiency of the process of RP synchronizing data from the repository publication point. This section analyzes the impact of using two different ROA issuance policies on the efficiency of initial synchronization and incremental synchronization to discuss the feasibility of requiring the use of the issuance policy of ROA with a single prefix in the RPKI production environment.

A. Initial Synchronization

The initial synchronization refers to the synchronization that takes place when the local cache of RP is empty. In the course of RP operation, the transmission data volume during initial synchronization is the largest. The experiments were conducted to compare the synchronization efficiency of using two extreme ROA issuance policies. One policy involves placing only one IP address prefix in an ROA, while the other policy involves placing all IP address prefixes originating from the same AS in an ROA.

Two IP address prefix distribution schemes were considered for the experiments: the randomized distribution of IP address prefixes and the distribution of IP address prefixes from five currently operational RIRs. The randomized IP address prefix distribution is discreteness, but different address space holders have distinct tendencies of issuance policy in the current production environment as mentioned in section III. The randomized IP address prefix distribution is unable to simulate these tendencies. The distribution of IP address prefixes from five currently operational RIRs provides both discreteness and reflects the distinct tendencies of issuance policy of different address space holders in the current production environment. Utilizing the distribution of IP address prefixes from five currently operational RIRs as a sample makes experimental data more practical and representative of the current production environment. By using this sample, the impact of synchronization efficiency can be evaluated in the current production environment when all ROAs with multiple prefixes are transformed into ROAs with a single prefix.

Due to the potential for interference when using public IP address for experiments, the decision was made to choose the largest available private IP address prefix 10.0.0.0/8 in IPv4 for experiments. Similarly, the decision was determined to choose the testing IP address prefix of 2001:db8::/32 in IPv6 as advised by Krill for experiments. For IPv4 address prefixes, a right-shift operation was applied to the IP addresses by 8 bits, and the IP addresses prefix length was increased by 8. This benefited to map the modified IP address prefixes to 10.0.0.0/8 (e.g., 165.98.219.0/24 was modified to 10.165.98.219/32). For IPv6 address prefixes, a right-shift operation was applied to the IP addresses by 32 bits, and the IP addresses prefix length was increased by 32. By doing so, the modified IP address prefixes were able to map to 2001:db8::/32 (e.g., 2407:9e40::/32 was modified to 2001:db8:2407:9e40::/64). By selecting these two IP addresses and adopting these mapping operations, the IP address prefixes can be retained as much as possible.

The number of lost INRs does not exceed 0.46% of the

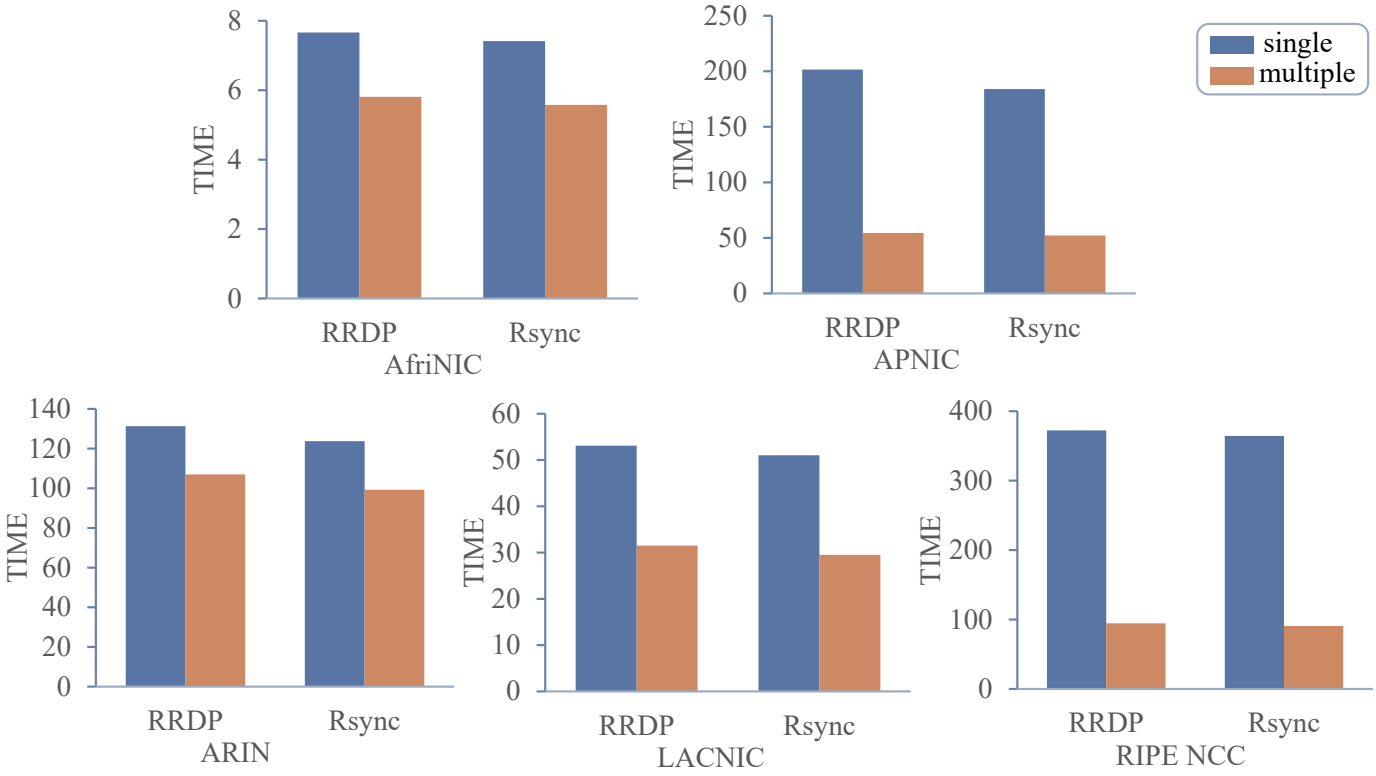

Fig. 4. The experiment results. The five histograms display the synchronization time of samples from five RIRs, using RRDP and Rsync under two different ROA issuance policies. The synchronization time is shown in seconds.

total quantity during the mapping process. In the lost INRs, the length of IPv4 prefix exceeds 24, and the length of IPv6 prefix exceeds 96. Approximately 38.5% of lost INRs are placed in ROAs with a single prefix in the current production environment. The remaining lost INRs are placed in ROAs with multiple prefixes in the current production environment, with an average of 3.7 INRs contained per ROA. When using the issuance policy of ROA with a single prefix, the lost INRs lead to a reduction of approximately 0.46% in synchronization time. While using the issuance policy of ROA with multiple prefixes, the lost INRs lead to a reduction of approximately 0.25% in synchronization time. These tiny errors do not impact the experimental conclusions.

The experiments used two servers, each equipped with 8 cores and 8GB of RAM. Krill software was selected to run the CA, and Routinator [27] was chosen to run the RP. One server was dedicated to running Krill to issue RPKI objects, while another server was utilized to run Routinator to synchronize data. TA issued multiple child CAs through the hosted model based on the number of IP address prefixes in each sample and allocated 10.0.0.0/8 and 2001:db8::/32 to each child CA. Each child CA managed a similar number of INRs. Following each completion of the initial synchronization, the local cache of RP was cleared and the next initial synchronization started. This step was repeated 30 times in each experiment sample of each ROA issuing policy. The synchronization efficiency comparison between two ROA issuing policies is predicated on the mean synchronization time of these 30 tests.

The experimental results are illustrated in Fig. 4, indicating that using the issuance policy of ROA with multiple prefixes leads to an enhancement in synchronization efficiency. The improvement of synchronization efficiency of two RIRs, AfriNIC and ARIN, is not significant. In contrast, the other two RIRs, APNIC and RIPE, show a marked improvement. It should be pointed out that even in RIPE NCC, which issues the greatest number of IP address prefixes, using the issuance policy of ROA with a single prefix does not lead to the initial synchronization time exceeding 7 minutes. The synchronization time of no more than 7 minutes for the initial deployment of RP is acceptable.

B. Incremental Synchronization

The synchronization except for initial synchronization was defined as incremental synchronization. Following the initial synchronization, RP periodically synchronizes updated files from the repository publication point at intervals no longer than one hour. Unlike the case in initial synchronization, using the issuance policy of ROA with multiple prefixes does not necessarily result in decreasing the transmission data volume in incremental synchronization.

In two distinct scenarios, using the issuance policy of ROA with a single prefix potentially decreases the transmission data volume. One situation is that only deletions are made to IP address prefixes in the incremental synchronization interval. In this scenario, adopting the issuance policy of ROA with a single prefix needs not to synchronize ROAs, while adopting the issuance policy of ROA with multiple prefixes needs to synchronize an entire ROA. Another scenario is when there

is a set of a large number of IP address prefixes originating from the same AS and a few operations (additions or deletions of IP address prefixes) made to this set in the incremental synchronization interval. When using the issuance policy of ROA with a single prefix, the ROAs containing the added IP address prefixes are required to be retransmitted. While using the issuance policy of ROA with multiple prefixes, all the IP address prefixes in the set are contained in an ROA. This ROA is required to be retransmitted. In this scenario, the size of the ROA with multiple prefixes may be larger than the total size of the few retransmission ROAs with a single prefix.

The above situations are not uncommon in production environments, thus utilizing ROA with multiple prefixes does not significantly enhance the efficiency of incremental synchronization.

VI. Conclusions

According to sections III, IV, and V, both the issuance policy of ROA with a single prefix and multiple prefixes possess distinct merits. The former provides greater flexibility and avoids the risk of overclaiming, thereby ensuring stable and valid route announcements. The latter reduces the quantity and size of ROAs, thereby augmenting synchronization efficiency.

Despite the obvious impact on the efficiency of initial synchronization caused by using the issuance policy of ROA with a single prefix in APNIC and RPIE NCC, it is worth noting that RP requires only one initial synchronization. Incremental synchronizations are frequent but the transmission data volume is small, hence exerting an inconspicuous influence on synchronization efficiency. Therefore, it is feasible to use the issuance policy of ROA with a single prefix in existing production environments. Above all, the fundamental purpose of RPKI is to ensure the security of BGP and its ability to provide the BGP router with accurate guidance regarding route filtering is vital. The validity of ROA assumes a pivotal role in this regard. It is deemed acceptable to compromise a certain degree of efficiency in order to ensure the validity of ROA.

In the current RPKI deployment environment, placing only one IP address prefix in ROA should be the preferred option in general situations. If the address space holder insists on placing multiple IP address prefixes into one ROA, the stability of INRs should be evaluated. The INRs that will not be revoked for a long time should be placed in the ROAs with multiple prefixes, while the unstable INRs should be individually placed in the ROAs with a single prefix. However, evaluating the stability of INRs cannot entirely avoid the security risks of overclaiming. Address space holders need to be aware of and assume the security risks by using the ROA with multiple prefixes.

Certainly, like RFC 9455 is a best current practice, the preferred option of ROA issuance policy may change with the ongoing refinement of RPKI. For instance, designing a mechanism that the parent CA proactive notifies the child CA when it reclaims INRs from the child CA can avoid the security risks of overclaiming in the scenario described in Section IV.C. However, the design, standardization, and deployment of such a mechanism take a considerable amount of time. Or using post-quantum cryptography [28]–[30] to protect the security of RPKI. Prior to the deployment of other effective mitigation measures, it is recommended to use the issuance policy of ROA with a single prefix.

Acknowledgment

We would like to thank our supervisors, Zhiwei Yan and Guanggang Geng, for their unwavering guidance and support throughout the course of this research. Thanks to the IETF SIDR Operations working group, which provided valuable expertise and assistance throughout the course of this research; and thanks to the RIPE NCC for sharing the ROA statistic.

References

[1] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," IETF, RFC 4271, January 2006.

[2] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," Proceedings of the IEEE, vol. 98, no. 1, pp. 100-122, 2009.

[3] G. Huston, M. Rossi, and G. Armitage, "Securing BGP — A Literature Survey," IEEE Communications Surveys & Tutorials, vol. 13, no. 2, pp. 199-222, 2010.

[4] Q. Xing, B. Wang, and X. Wang, "BGPcoin: Blockchain-Based Internet Number Resource Authority and BGP Security Solution," Symmetry, vol. 10, no. 9, 2018. [Online]. Available: https://www.mdpi.com/2073-8994/10/9/408

[5] A. Hari, and V. Lakshman, "The Internet Blockchain: A Distributed, Tamper-Resistant Transaction Framework for the Internet," in Proceedings of the 15th ACM Workshop on Hot Topics in Networks. Atlanta, GA, USA: ACM, 2016, pp. 204-210.

[6] M. Iansiti, and K. R. Lakhani, "The Truth About Blockchain," Harvard business review, vol. 95, no. 1, pp. 118-127, 2017.

[7] T. Hlavacek, P. Jeitner, D. Mirdita, H. Shulman, and M. Waidner, "Behind the Scenes of RPKI," in Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. Los Angeles, CA, USA: ACM, 2022, pp. 1413-1426.

[8] T. Hlavacek, P. Jeitner, D. Mirdita, H. Shulman, and M. Waidner, "Stalloris: RPKI Downgrade Attack," in 31st USENIX Security Symposium. Boston, MA, USA: Usenix Association, 2022, pp. 4455-4471.

[9] Y. Gilad, O. Sagga, and S. Goldberg, "MaxLength Considered Harmful to the RPKI," in Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies. New York, NY, USA: ACM, 2017, pp. 101-107.

[10] M. Lepinski, S. Kent, and D. Kong, "A Profile for Route Origin Authorizations (ROAs)," IETF, RFC 6482, February 2012.

[11] Z. Yan, R. Bush, G. Geng, T. de Kock, and J. Yao, "Avoiding Route Origin Authorizations (ROAs) Containing Multiple IP Prefixes," IETF, RFC 9455, August 2023.

[12] D. Mirdita, H. Shulman, and M. Waidner, "Poster: RPKI Kill Switch," in Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. New York, NY, USA: ACM, 2022, pp. 3423-3425.

[13] NLnetLabs, "RPKI Certificate Authority and Publication Server," 2023. [Online]. Available: https://github.com/NLnetLabs/krill

[14] R. Housley, "Cryptographic Message Syntax (CMS)," IETF, RFC 3852, July 2004.

[15] T. Hlavacek, P. Jeitner, D. Mirdita, H. Shulman, and M. Waidner, "Beyond Limits: How to Disable Validators in Secure Networks," in Proceedings of the ACM SIGCOMM 2023 Conference. New York, NY, USA: ACM, 2023, pp. 950-966.

[16] G. Huston, G. Michaelson, and R. Loomans, "A Profile for X.509 PKIX Resource Certificates," IETF, RFC 6487, February 2012.

[17] T. Bruijnzeels, O. Muravskiy, B. Weber, and R. Austein, "The RPKI Repository Delta Protocol (RRDP)," IETF, RFC 8182, July 2017.

[18] A. Tridgell, P. Mackerras, and W. Davison, "rsync protocol man page." [Online]. Available: https://linux.die.net/man/1/rsync

[19] A. Durand, "Resource public key infrastructure (RPKI) technical analysis," ICANN, Sep. 2020. [Online]. Available: https://icann-hamster.nl/ham/icann/octo/pub/octo-014-en.pdf

[20] J. Kristoff, R. Bush, C. Kanich, G. Michaelson, A. Phokeer, T. C. Schmidt, and M. Wählisch, "On Measuring RPKI Relying Parties," in Proceedings of the ACM Internet Measurement Conference. New York, NY, USA: ACM, 2020, pp. 484-491.

[21] R. Bush, and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1," IETF, RFC 8210, September 2017.

[22] G. Huston, and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)," IETF, RFC 6483, February 2012.

[23] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, "BGP Prefix Origin Validation," IETF, RFC 6811, January 2013.

[24] RIPE NCC and Internet Multifeed Co, "Index of /rpki," 2023. [Online]. Available: https://ftp.ripe.net/rpki/

[25] G. Huston, G. Michaelson, C. Martinez, T. Bruijnzeels, A. Newton, and D. Shaw, "Resource Public Key Infrastructure (RPKI) Validation Reconsidered," IETF, RFC 8360, April 2018.

[26] BGPMON, "Massive route leak causes Internet slowdown," 2015. [Online]. Available: https://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/

[27] NLnetLabs, "An RPKI Validator and RTR server," 2023. [Online]. Available: https://github.com/NLnetLabs/routinator

[28] M. Anastasova, R. Azarderakhsh, M. Mozaffari Kermani, and L. Beshaj, "Time-Efficient Finite Field Microarchitecture Design for Curve448 and Ed448 on Cortex-M4," International Conference of Information Security and Cryptology, 2022, pp. 292-314.

[29] M. Anastasova, R. Azarderakhsh, and M. Mozaffari Kermani, "Fast Strategies for the Implementation of SIKE Round 3 on ARM Cortex-M4," IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 68, no. 10, pp. 4129-4141, 2021.

[30] P. Sanal, E. Karagoz, H. Seo, R. Azarderakhsh, and M. Mozaffari-Kermani, "Kyber on ARM64: Compact Implementations of Kyber on 64-Bit ARM Cortex-A Processors," International Conference on Security and Privacy in Communication Systems, 2021, pp. 424-440.