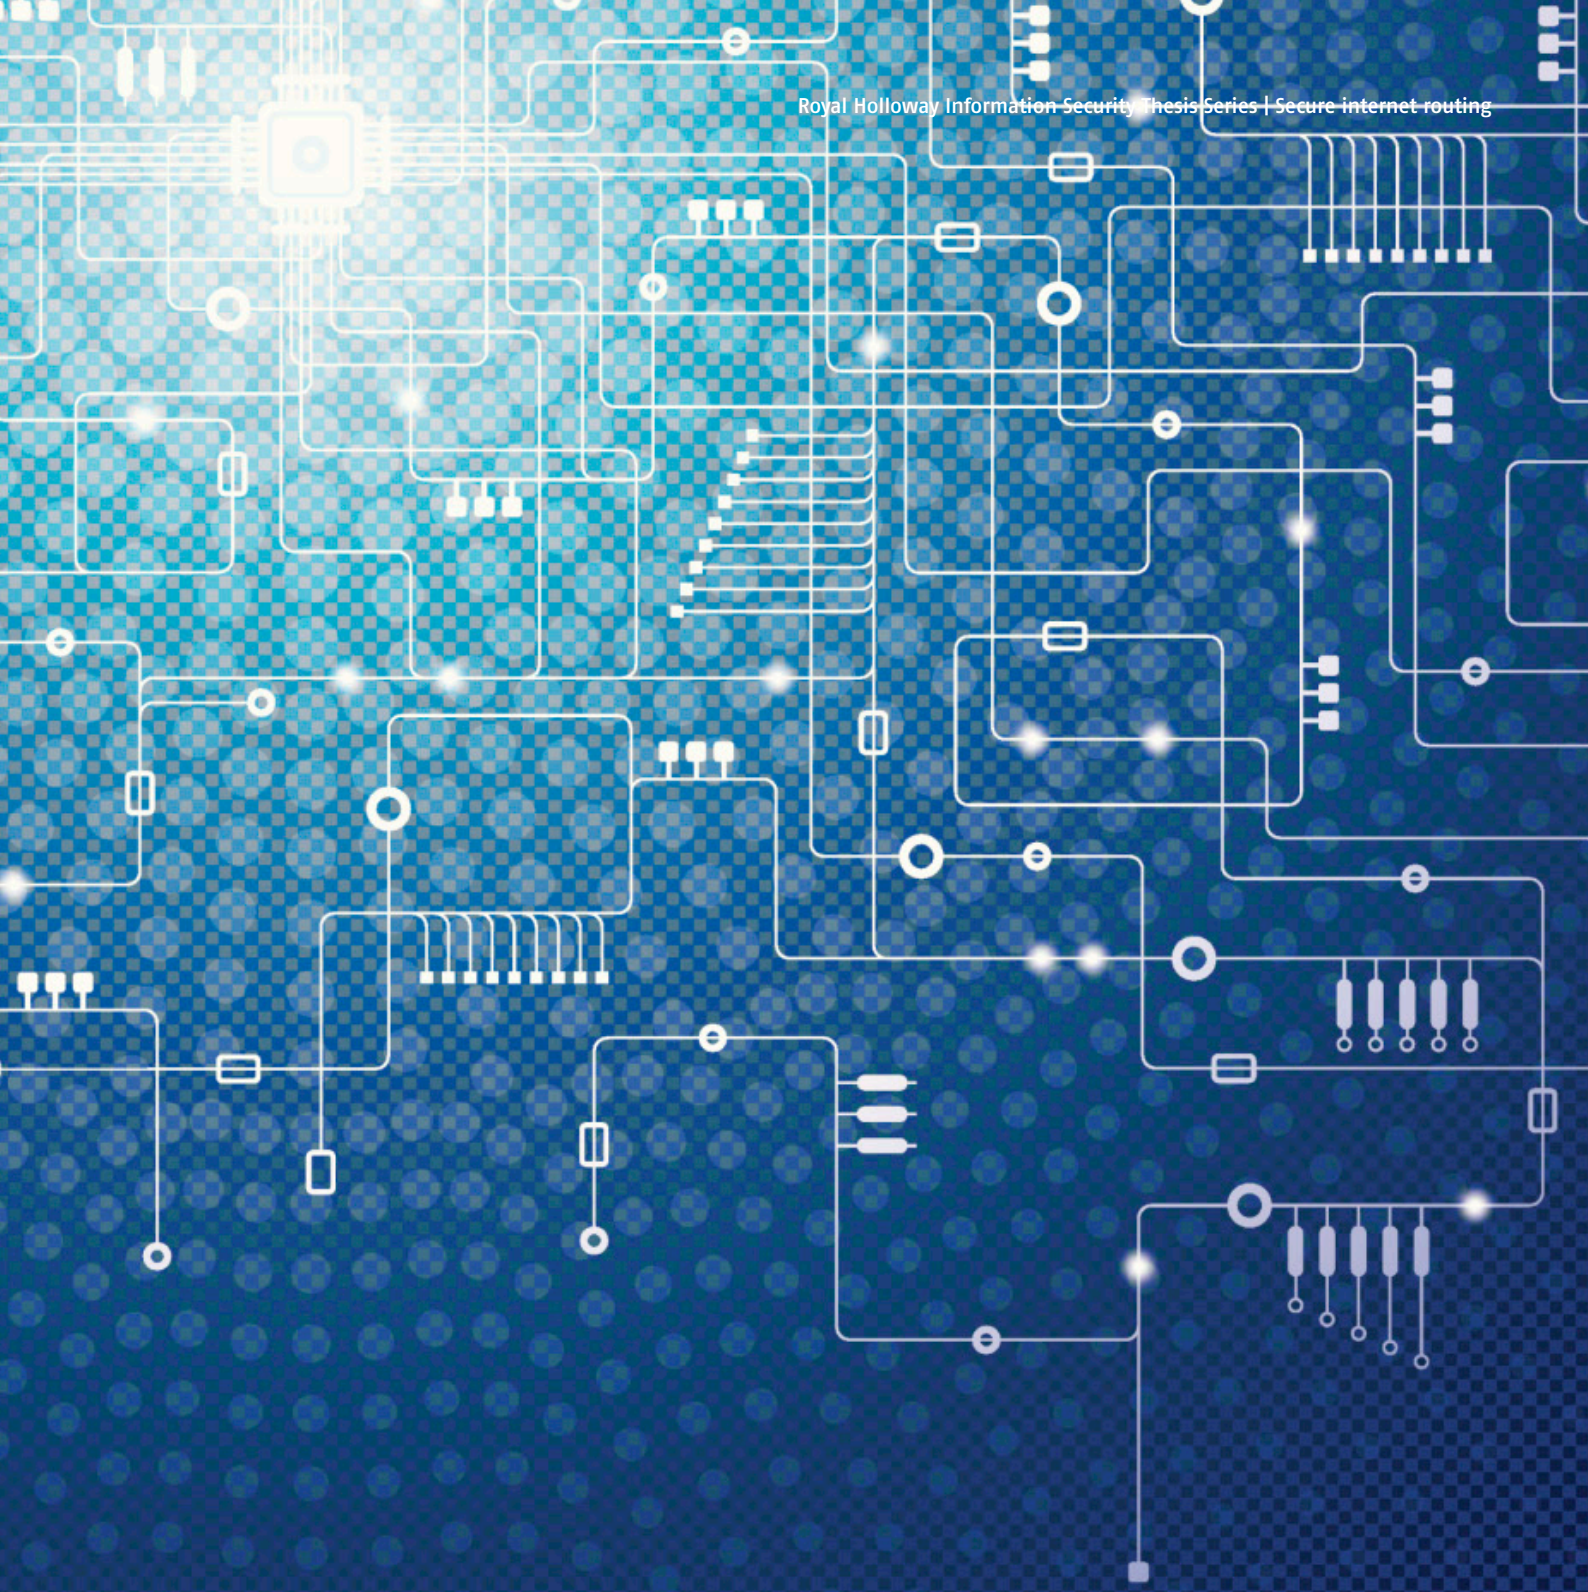

Secure internet routing

Can Resource Public Key Infrastructure provide framework for true security?

by Amreesh Phokeer, MSc (Royal Holloway) and Keith Martin, ISG, Royal Holloway

THINKSTOCK

Avoid losing your way on the information superhighway

Internet routing security has been a hot topic for many years. Can a new series of protocols released by the Internet Engineering Task Force (IETF) to operate via a framework known as the Resource Public Key Infrastructure (RPKI) provide the solution to internet users' security fears?

by Amreesh Phokeer and Keith Martin

Imagine you are driving your car from town A to town Z. Your journey is, of course, not a straight line and you need to go through towns C, K and R before reaching Z. How misleading would it be if the directions to the intermediate towns were wrongly indicated? Worse still, imagine that vandals had introduced fake road signs that force you along unnecessary detours through dangerous areas, or simply lead you to the middle of nowhere…

The internet is no different. Just as on the highway, internet traffic (data packets) requires proper and continuous directional information in order to choose the best route to its intended destination. Making sure that this information cannot be manipulated is crucial for the future of our activities on the internet.

After several years of unsuccessful standardisation work, the Internet Engineering Task Force (IETF) recently released a series of protocols that should enable interdomain routing to become more secure and robust through a framework known as the Resource Public Key Infrastructure (RPKI). In this article, we introduce RPKI and identify some of its deployment challenges.

The challenge of providing secure internet routing

Internet routing security has been in the limelight for more than a decade. All internet-related activities rely on a robust routing system to ensure that data packets are delivered to the right destinations. Disruptions to the routing system, either through accident or malice, can have unwanted consequences for the day-to-day life of internet users.

Human errors and router misconfigurations are particularly common sources of errors. A well-known incident was the 'Pakistan Telecom – YouTube hijack' in February 2008 (see box and Figure 1, page 3), when Pakistan Telecom knowingly advertised a prefix belonging to YouTube as its own. This announcement was propagated to its upstream providers, which resulted in the hijack of YouTube traffic.

As the internet keeps expanding and its use becomes ubiquitous, concerns have been raised over whether current routing systems are sufficiently reliable and secure. In recent years, many initiatives have been launched to try to better secure the current internet routing system.

Concerns have been raised over whether current routing systems are sufficiently reliable and secure

However, deploying new security mechanisms on such a large scale is not easy. Many of these initiatives were never deployed widely because of the excessive technical implementation overheads they required.

There are further complications. Internet routing works through disparate networks called autonomous systems (ASs). Each AS functions independently

The Pakistan Telecom – YouTube incident: what happened

On 24 February 2008, the YouTube website was unreachable for almost two hours. The BBC reported that, in an attempt to ban YouTube in Pakistan, Pakistan Telecom (AS 17557) had "hijacked" the address space of YouTube and propagated it to the wider internet through a Hong-Kong based provider, PCCW (AS 3491). The technique, called "sub-prefix hijacking", had been used previously in another famous AS 7007 attack.
Route selection is based on several criteria, one of which is the principle of "longest prefix match", which means BGP will always prefer a route where a more specific prefix is advertised. YouTube's network (208.65.152.0/22) is normally advertised by AS 36561, but on Sunday, 24 February 2008, AS 17557 (Pakistan Telecom) deliberately decided to redirect all traffic going to YouTube to a "black hole". To achieve this hijack, at 18:47 (UTC), AS 17557 started to announce a subnet of YouTube's network (208.65.153.0/24). As the more specific prefix /24 will always be chosen over the less specific prefix /22, all traffic to YouTube within Pakistan Telecom's internal network was effectively redirected to a black hole. Unfortunately, it made one mistake: it forgot to add an outgoing filter on its announcement and the hijacked route was propagated to its upstream provider, AS 3491 (PCCW Global).
AS 3491 was not doing any incoming filtering to check whether the route announced by AS 17557 was genuine. Therefore BGP routers around the world were notified about the (hijacked) route, with the result that all YouTube traffic was redirected to Pakistan Telecom.
At 20:07 (UTC), YouTube decided to advertise the same subnet (/24) as a countermeasure, in order to attract traffic back to its network. This worked and some parts of the world could then reach YouTube's network properly at AS 36561.
However, end-users who were closer to Pakistan Telecom were still being denied access. At 20:18, as a remedy, YouTube decided to advertise more specific prefixes – 208.65.153.128/25 and 208.65.153.0/25 – to attract all traffic to its network. The result was seen at 20:51.
Finally, at 21:01, AS 3491 (PCCW Global) filtered out the hijacked prefixes announced by AS 17557 (Pakistan Telecom), restoring the situation to normal.

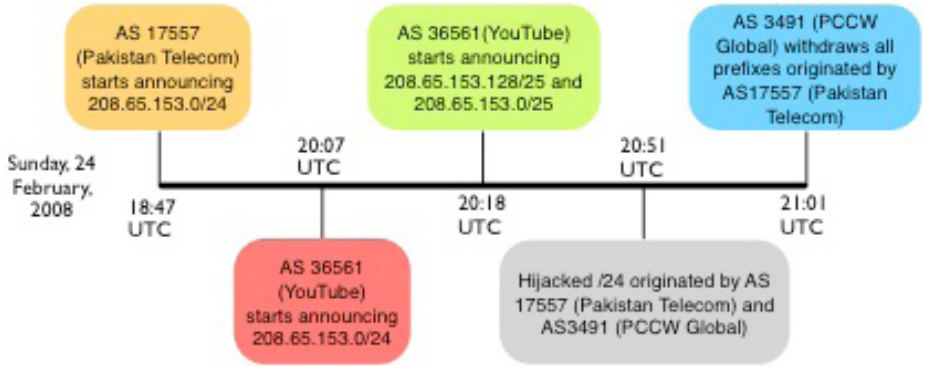

Figure 1: The hijack event timeline as reported by the RIPE NICC's routing information service

and is tied to other ASs through business and commercial relationships. Routing policies – the rules that determine how packets are transmitted from one network to another – are based on those potentially conflicting AS relationships. Therefore, competition and commercial interest are also in the picture, making the problem of secure routing even more complex.

How does internet routing work?

On the internet, data packets are routed by network devices called routers, whose role is to select the best possible path for a data packet to its destination based on a set of information it receives from other routers. Returning to our highway analogy, we normally choose the shortest path to our destination, but sometimes the shortest path is not necessarily the cheapest, especially if it involves a toll road.

Additional criteria such as this are taken into account when a router determines the optimal path to a destination address. Essentially, routers are rather like roundabouts with signposts (Figure 2). On the internet, those signposts are

updated dynamically using of a specially crafted protocol called the Border Gateway Protocol (BGP).

Figure 2: A roundabout signpost

Internet routers rarely work in a standalone mode and are grouped together to form networks (domains), within which packets are transmitted from one point to another. These are the autonomous systems mentioned earlier and each is identified by an AS number. Using the road analogy once again, an AS represents a city such as London, which has an internal network plus many junctions that allow traffic to flow in or out.

Those junctions consist of border routers and their role is to allow traffic to flow from one network to another (see Figure 3). In order for border routers to be aware which way to send traffic, they need to constantly send updated information to their peers using the BGP protocol. A typical BGP exchange between neighbouring routers in Figure 3 might be as follows:

Router B: *Hello Router A, if you want to send traffic to network Madrid, use path B->G->H.*
Router A: *Hello Router B, thanks for the update.*
Router F: *Hello Router E, if you want to send traffic to network Cape Town, use path E->F->H->I->J->K.*
Router E: *Hello Router F, thanks for the update.*

In reality, Router A and Router E will also receive update messages from other neighbouring routers, such as routers C and D, respectively, and the best path to the destination will be chosen based on several criteria, such as the number of hops to a destination, cost of transit, bandwidth and peer agreements.

How can internet routing be secured?

At the time the BGP protocol was conceived, security and trust in interdomain routing was not a big issue. The internet has therefore been functioning since its inception with this simple routing protocol at its heart. Unfortunately, BGP has two main vulnerabilities.

First, BGP does not have any strong mechanism to prevent a network from announcing an arbitrary route. In other words, on the internet today, any network

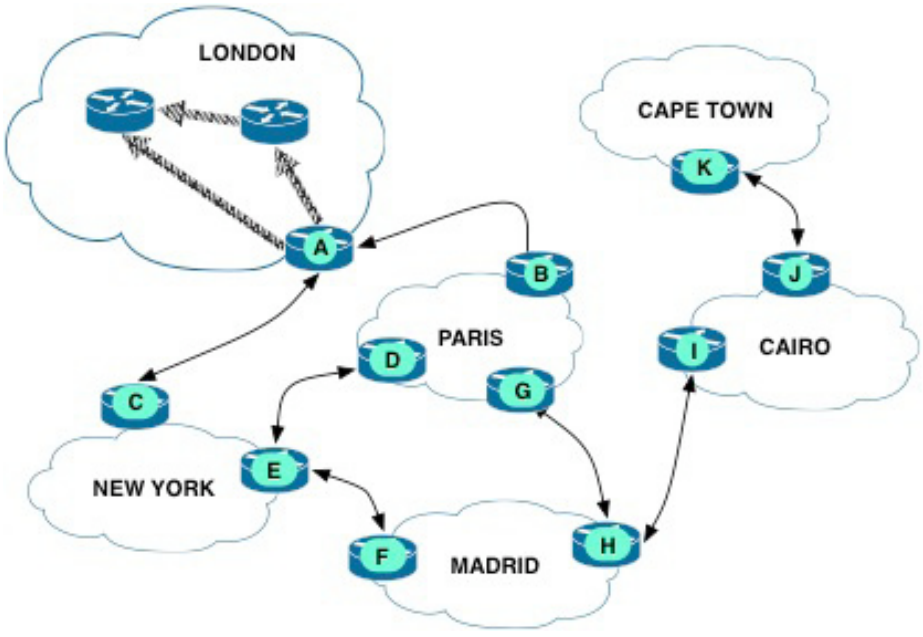

Figure 3: A typical network

could potentially announce: “I am Google, please send all Google traffic to my network.” This is called a BGP origination hijack. The second main issue is path validation. A router has no means to validate the information it receives about a route to a certain destination. As a relying party router, how can I be sure that the path that I am receiving has not been tampered with by an adversary?

If we want a reliable system to validate this kind of information, we need to make use of a strong security mechanism, such as digital signatures. And if we want a reliable system based on digital signatures, we need a supporting public key infrastructure (PKI) and certification process. The IETF recently proposed RPKI as such a framework. RPKI itself will not solve all interdomain routing problems, but it will, hopefully, provide the much-needed building blocks upon which internet routing security can be built.

If we want a reliable system to validate this information, we need to make use of a strong security mechanism

How to protect against false origin attacks

To protect against BGP origination attacks, we need a way to validate whether an AS that is originating an IP prefix on the internet has the right to do so. In other words, we need to have a secure way to certify that an AS is indeed the holder of the IP address space it is advertising to other networks.

In RPKI, this is done by using dedicated end-entity (EE) certificates to generate cryptographically signed route filters, called route origin authorisations (ROAs). EE certificates are generated by resource certificate authorities, which are usually run by resource holders such as the Internet Assigned Numbers Authority (IANA), regional internet registries (RIRs), local internet registries (LIRs) or internet service providers (ISPs), depending on location in the hierarchy (Figure 4).

An ROA is a structured signed object whose role is to attest that an AS has been authorised, by the holder of a resource, to advertise the address space to other BGP users. ROAs by themselves do not contain any routing validation information; they only represent the authority of an AS over a prefix. The validity of an ROA is tied to the EE certificate that it includes. The EE certificate has a validity period determined by the resource holder. The ROA is signed by the private key corresponding to the public key in the EE certificate. A simplified

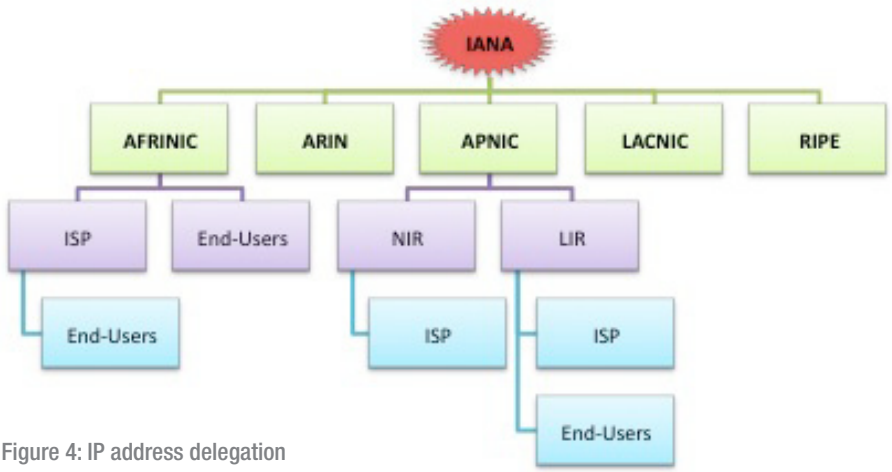

Figure 4: IP address delegation

A network is made up of IP addresses (an IP prefix) and those prefixes are allocated by the IANA (Internet Assigned Numbers Authority) to RIRs (regional internet registries). The RIRs, in turn, assign IP prefixes to local or national internet registries, ISPs or end-users such as big companies. Therefore, in the RPKI, the IANA naturally takes the role of the trust anchor. The five RIRs would run their CAs and would delegate signing authority to their national or local internet registries (NIRs, LIRs) to which their IP resources have been allocated. Similarly, if an ISP received resources from an LIR, the LIR would delegate the relevant signing authority to the ISP.

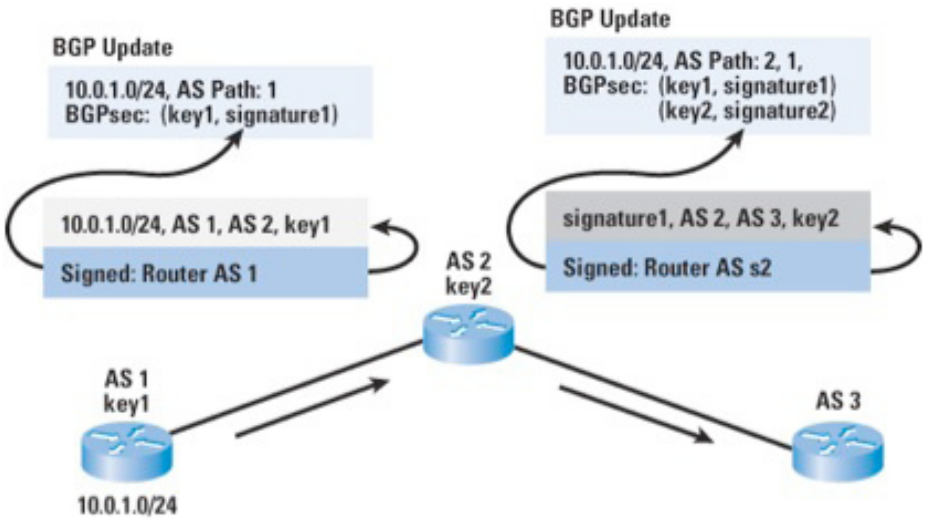

Figure 5: A chain of interlocking signatures Source: Cisco

view of the content of an ROA (without the EE certificate) is shown in Figure 6. The main information here is the field “as_id”, which contains the origin AS number authorised to advertise the prefix in the “prefixes” field. It also contains information from the EE signing certificate, such as its validity period.

With RPKI, when a relying party router receives a route advertisement from a peer, it will now be able to check whether there is a valid ROA by verifying the digital signature of the EE certificate of the ROA. The route is considered valid if this signature can be verified.

How to protect against path manipulation

We thus have a solution to ascertain that an origin AS has the right to advertise a prefix. Origin validation can help BGP users detect misconfigurations or malicious announcements, but it cannot prevent a malicious attacker from faking the path information in a BGP update message. How can we now make sure that the origin AS received in a BGP update message is indeed coming from the AS that is claimed?

This question is essentially one of making sure that the sequence of AS numbers received (the AS path) actually reflects a real path to the destination. RPKI addresses this by using router certificates. Such a certificate is issued to a router and attests that it is the rightful holder of an AS number. The router, equipped with such a certificate and a corresponding private signature key, is therefore able to sign outgoing update messages.

The idea is that the AS path is recursively signed before being sent to the next hop. This creates an interlocking chain of signatures, each of which can be validated individually. Figure 5 illustrates such a chain of signatures. Router AS1 sends a signed update to Router AS2. Router AS2 adds a hop in the AS path and signs everything. When Router AS3 receives the announcement, it receives it with the two previous signatures.

What are the challenges of deploying RPKI?

As we have seen, RPKI addresses the main security concerns with existing internet routing. To gauge the potential for success, it is important to understand the challenges pertaining to acceptance and adoption of RPKI. The challenges are not only technical and operational, but also political. One straightforward problem is data consistency of internet number resources. This is a major

```
...
as_id: 37668

prefixes:
  41.222.48.0/20

signing certificate:
  serial:    124 (0x7C)
  not before:      2013-07-17T11:21:40
  not after:       2037-07-17T11:36:40
  ...
```

Figure 6: Simplified view of an ROA’s content

impediment and is believed to be one barrier to successful adoption. Resource certificates form the basis of secure interdomain routing and it is important that accurate information is fed into the routing system. However, RPKI has no mechanism to prevent misconfigurations or unintentional mistakes. It is up to the service provider, for instance an ISP, to put in the necessary checks to ensure correctness of data to prevent IP reachability downtime.

Hardware and software compliance of routers to RPKI is certainly an issue that has the potential to hamper rapid deployment of RPKI. However, there is some indication that operators would be willing to upgrade routing equipment in the event of RPKI deployment within their organisation.

Fundamentally, resource holders need to understand the importance of having their resources certified for proof of ownership. This is especially important at a time when illegal trading of the IP address space is becoming more and more common, resulting in the depletion of IPv4 addresses. General training and awareness is also required because RPKI is quite a complex system; network operators will need training to prevent RPKI misconfigurations, which would almost certainly lead to network reachability outages.

Finally, on a more political note, studies have shown that some people are nervous of potential "bad intentions" behind the motivation of RPKI. The main driver behind interdomain security is the US Department of Homeland Security. The RPKI "root of trust" is currently intended to be hosted by US incorporated organisations (IANA or ICANN) and this subject has long caused political turmoil. People are increasingly wary about governments using RPKI as a censorship tool to bring down networks or to divert and eavesdrop internet traffic. Continuing publicity about the NSA's mass surveillance programme has certainly not alleviated these concerns.

Resource holders need to understand the importance of having resources certified for proof of ownership

Conclusion

RPKI is a new protocol that is still in an experimental phase. The main players, especially the RIRs, have already adopted RPKI and are offering this service on a pilot basis. It is important to make the transition as seamless as possible so that it does not disrupt the current *modus operandi* of internet routing. Having a fully secure internet routing system end-to-end is practically impossible, but RPKI seems to be a framework that will allow us to move in a positive direction towards that ideal. ■

About the authors

Amreesh Phokeer has an MEng in computer science from Telecom Nancy, University of Lorraine, France and an MSc in information security from Royal Holloway University of London. He currently works as an applications unit manager at AFRINIC, the African Regional Internet Registry, where he heads a software development team that supports the organisation's services as an RIR. Before this, he worked as a guest researcher at NIST, where he took part in research into computational biology. His areas of focus are network security, security technologies, secure software development and PKI. He is also passionate about research in information and communication technologies for developing and less-developed countries.

Professor Keith Martin is director of the Information Security Group at Royal Holloway University of London. He received his BSc (Hons) in mathematics from the University of Glasgow in 1988 and a PhD from Royal Holloway in 1991. After a period with the COSIC research group of the Katholieke Universiteit Leuven in Belgium, he returned to Royal Holloway in 2000 and was given a chair in 2007. His current research interests include key management, cryptographic applications and securing lightweight networks. He is the author of *Everyday Cryptography*, recently published by Oxford University Press. As well as conventional teaching, Martin is a designer and module leader on Royal Holloway's distance learning MSc information security programme, and regularly presents to industrial audiences and schools.