Study for Ofcom

Review of Resource Public Key Infrastructure (RPKI) to verify ownership and authenticity of telephone caller ID over Voice over Internet Protocol

Authors:
J Scott Marcus
Richard Shockey

WIK-Consult GmbH
Rhöndorfer Str. 68
53604 Bad Honnef
Germany

The opinions expressed are solely those of the authors.
They do not necessarily reflect the views of Ofcom.

Bad Honnef, 10 June 2015

wik CONSULT

Contents

Figures

Food for Thought

Executive summary

The public places great confidence in the modern global telephone network, and rightly so. The growing problem of falsification of the caller's number, however, risks eroding public confidence in the UK *Public Switched Telephone Network (PSTN)* (and in those of the US and Canada), as well as in the public safety networks that depend on the UK PSTN.

The core problem of Caller ID spoofing is the ability of a (possibly malicious) caller to either impersonate someone else, or to anonymise a call in such a way that fraud and abuse can occur. The network operator, the regulator, and law enforcement may find it difficult to track and trace the source of the abusive call.

This is essentially a problem of *authentication* of the identity of the caller. *Is the caller really who he claims to be?*

For the somewhat related problem of route hijacking, the implementation *of Resource Public Key Infrastructure (RPKI)* as the first phase of realisation of *Secure Inter-Domain Routing (SIDR)* is widely viewed as the most promising technical means of mitigating the problem.

In this study, we have been asked to consider whether the RPKI/SIDR tools that have been developed to ensure the integrity of the Internet routing system, and to authenticate that parties announcing routes for a range of IP addresses actually have rights (or *holdership*)[1] to those addresses, could perhaps be re-purposed to ensure that parties claiming to be placing a call from a given telephone number truly have rights to use that telephone number. To what extent are these two problems the same? To what extent are they different?

Effectiveness of RPKI-based solutions

The European implementation of RPKI/SIDR by RIPE NCC can be deemed to be a success. Their emphasis on ease of use was well placed. Some 20% of their members routinely download new copies of their data cache. Implementation costs were modest.

At the same time, it is important to bear in mind what RPKI/SIDR does *not* (yet) do. RPKI/SIDR authenticates holdership, and helps to protect the Internet routing system from inadvertent configuration errors; however, it does not protect against malicious attacks.

---

1 The responsible administrators are generally careful to avoid creating a precedent that parties who have usage rights to IP addresses actually own them, as if they were tangible property. Analogous issues relate to telephone numbers and also to spectrum. Property-like rights may be involved, but whether they should be legally treated as property is a complex question, the answer to which could vary from country to country. For that reason, RPKI documents refer to *holdership* rather than *ownership*.

These things take time. Understanding the problem of securing *Internet Numbering Resources (INR)* and the *Border Gateway Protocol (BGP)* using SIDR and RPKI took a great deal of time. The standards process alone took almost eight years.

Beyond that, industry adoption has been slow. Implementation and use is totally voluntary on the part of the ISPs. Complexity in the IPv4 address blocks has hindered deployment.

Applicability of RPKI/SIDR-based techniques to the Caller ID spoofing problem

These limitations notwithstanding, we believe that the RIPE NCC experience provides a solid proof of concept, and the relevance of the experience is clear. The application of RPKI/SIDR-based mitigation techniques to Caller ID spoofing for the UK phone system is inevitable in our view.

There is little doubt that an RPKI-based solution for validation of Caller ID, and perhaps CNAM information, is feasible. Modern database technology is more than adequate to support an RPKI-based data repository for every active phone number in the UK Numbering Plan.

It is important to bear in mind, however, that validating holdership of a telephone number is not in and of itself a complete solution to the problem of Caller ID spoofing.

Any concerted policy approach would necessarily consider measures to mitigate Caller ID Spoofing together with two interrelated topics: (1) Local Number Portability (LNP), and (2) IP-based Network-to-Network Interconnection. LNP complicates the problem space, and may necessitate a larger, more volatile, and more complex RPKI certificate repository than would otherwise be needed. A holistic view of the space is needed.

Ofcom's course going forward

The Caller ID spoofing problem is a complex and multifaceted problem. Mitigating the problem will surely be just as multifaceted, if not more so. No single "silver bullet" is likely to magically solve the problem.

We are strongly of the view that RPKI/SIDR is almost certain to play a key role over time in any comprehensive solution to the VoIP Caller ID spoofing problem. At the same time, key limitations must not be forgotten:

- The RPKI/SIDR capabilities that are deployed today authenticate holdership of IP addresses and Autonomous System numbers, but fall well short of an fully automated system, and do not (yet) protect against a serious, malicious attack.

- The pace of standards development, software implementation, and network deployment is such that even capabilities comparable those of current RPKI/SIDR would likely require at least five to seven years to deploy widely. The time to achieve a truly effective system might perhaps be considerably longer.

Following the typical Internet-based pattern of voluntary standards adoption for RPKI for UK phone numbers may not be sufficient. Protection against Caller ID Spoofing is of limited value to consumers until it is widely, if not universally, deployed. As we have already seen, RPKI-based authentication is for analogous reasons, of little utility to network operators until a large enough number of network operators deploy. For services such as these, *network effects* are crucial.[2]

Our strong belief is that *this problem needs to be addressed at national level in each country, i.e. at a level corresponding to that of ITU country codes, because that is the level to which responsibility for the national numbering plan has been delegated*.

Other countries – who may be less impressed with the immediacy of the problem than are the UK, Canada, and the UK – are more likely to take voluntary action if there is some demonstrably workable framework that they can join.

2 See J. Scott Marcus (2004), "Evolving Core Capabilities of the Internet", *Journal on Telecommunications and High Technology Law*. See also Jeffrey H. Rohlffs (2001), *Bandwagon Effects in High-Technology Industries*.

Glossary

Name	Acronym	Definition
Autonomous System	AS	An Autonomous System (AS) is a group of IP networks that use a single and clearly defined routing policy.
Autonomous System Number	ASN	Every Autonomous System must have an Autonomous System Number. ASNs are globally unique numbers used to identify these groups of networks. ASNs allow an autonomous system to exchange routing information with neighbouring autonomous systems. (Source: NRO)
Border Gateway Protocol	BGP	The primary inter-domain routing protocol used by the Internet.
Business Support Systems	BSS	Business Support Systems are operational support systems (OSS) that are used by telecommunications providers to support the management of their business processes.
Caller ID		Caller ID (caller identification, CID), also called calling line identification (CLID), is a telephone service, available in analogue and digital phone systems and most voice over Internet Protocol (VoIP) applications, that transmits a caller's number to the called party's telephone equipment during the ringing signal, or when the call is being set up but before the call is answered. (Source: Wikipedia)
Calling Name Delivery	CNAM	The service whereby caller ID provides a name associated with the calling telephone number is called CNAM. The information made available to the called party may for instance be displayed on a telephone's display or on a separately attached device. (Source: Wikipedia)
Holdership		Denotes that a party has been assigned rights of use to a telephone number or an IP address, but not necessarily ownership.

IP Multimedia System or Integrated Multimedia System	IMS	A standards-based platform based on IP and SIP protocols that seeks to employ common, reusable modules for commonly used functions.
International Telecommunications Union	ITU	The ITU is the United Nations specialised agency for information and communication technologies.
Internet Engineering Task Force	IETF	The Internet Engineering Task Force (IETF) develops and promotes Internet standards, dealing in particular with standards of the Internet protocol suite. It is an open standards organization, with no formal membership or membership requirements. It started out as a US federal government organization, and today it operates as a non-commercial not-for-profit non-governmental organization. (Source: Wikipedia)
Internet Number Resources	INR	Internet number resources include Internet Protocol (IP) address space and Autonomous System Numbers (ASNs). (Source: NRO)
Internet Protocol	IP	The Internet Protocol is a data communications standard that allows computers to communicate with one another over digital networks. Together with the TCP protocol, IP forms the basis of the Internet
IP Address		An IP address is a numeric identifier that includes information about how to reach a network location through the Internet routing system. Every device directly connected to the Internet must have an IP address. Every IP address must be unique for devices to connect to the Internet and to each other. (Source: NRO)
Internet Service Provider	ISP	An ISP is a firm that enables other organizations to connect to the global Internet.
Operational Support Systems	OSS	A system to support network operations or management.

Public Key Infrastructure	PKI	A public key infrastructure (PKI) is a set of hardware, software, people, policies, and procedures needed to create, manage, distribute, use, store, and revoke digital certificates. (Source: IEEE)
Public Land Mobile Network	PLMN	The circuit switched mobile network.
Public Safety Access Point	PSAP	A public-safety answering point (PSAP), sometimes called "public-safety access point", is a call centre responsible for answering calls to an emergency telephone number for police, firefighting, and ambulance services. Trained telephone operators are also usually responsible for dispatching these emergency services. (Source: Wikipedia)
Public Switched Telephone Network	PSTN	The circuit switched fixed network. Sometimes the term meant to encompass the mobile network as well.
Request for Comments	RFC	A Request for Comments (RFC) is a publication of the Internet Engineering Task Force (IETF) and the Internet Society, the principal technical development and standards-setting bodies for the Internet. All IETF standards are RFCs, but not all RFCs are standards.
Resource Public Key Infrastructure	RPKI	Resource Public Key Infrastructure (RPKI) is a specialized public key infrastructure (PKI) framework designed to secure the Internet's routing infrastructure. RPKI provides a way to connect Internet number resource information (such as Autonomous System numbers and IP Addresses) to a trust anchor. (Source: Wikipedia)
Secure Inter-Domain Routing	SIDR	SIDR is an IETF WG whose purpose is to reduce vulnerabilities in the inter-domain routing system. SIDR builds on existing infrastructure, including RPKI, and will also specify security enhance-ments for inter-domain routing protocols.

Session Border Controller	SBC	A session border controller (SBC) is a device regularly deployed in Voice over Internet Protocol (VoIP) networks to exert control over the signalling and usually also the media streams involved in setting up, conducting, and tearing down telephone calls or other interactive media communications. (Source: Wikipedia)
Session Initiation Protocol	SIP	An application-layer data communi-cations control protocol for creating, modifying, and terminating sessions with one or more participants. It can be used to create two-party, multiparty, or multicast sessions that include Internet telephone calls, multimedia distribution, and multimedia conferences. SIP is designed to be independent of the underlying transport layer; it can run on TCP, UDP, or SCTP. It is widely used as a signalling protocol for Voice over IP, along with H.323 and others.
Signalling System 7	SS7	Signalling System No. 7 (SS7) is a set of telephony signalling protocols which are used to set up most of the world's public switched telephone network (PSTN) telephone calls. The main purpose is to set up and tear down telephone calls. (Source: Wikipedia)
	STIR	The IETF STIR working group seeks to specify Internet-based mechanisms that allow verification of the calling party's authorisation to use a particular telephone number for an incoming call. (Source: STIR WG Chater)
Time Division Multiplexing	TDM	Time-division multiplexing (TDM) is a method of transmitting and receiving independent signals over a common signal path by means of synchronized switches at each end of the transmission line so that each signal appears on the line only a fraction of time in an alternating pattern. (Source: Wikipedia)

Transmission Control Protocol	TCP	A data communications protocol used to assure reliable delivery of data in an IP network.
Voice over IP	VoIP	A set of data communications protocols and technologies to enable voice to be sent over individual IP-based networks or over the Internet.
Working Group	WG	The IETF is organised into a large number of working groups and informal discussion groups, each dealing with a specific topic. Each group is intended to complete work on that topic and then disband. Each working group has a charter that describes its focus, and what and when it is expected to produce. (Source: Wikipedia)

1 Introduction

This is the Final Report of a study conducted by WIK-Consult GmbH (Germany) and Richard Shockey (US).

The topic is a complex one: Could RPKI/SIDR tools that have been developed to ensure the integrity of the Internet routing system, and to authenticate that parties announcing routes for a range of IP addresses actually have rights (or "holdership")[3] to those addresses, perhaps be re-purposed to ensure that parties claiming to be placing a call from a given telephone number truly have rights to use that telephone number? To what extent are these two problems the same? To what extent are they different?

1.1 The challenge of VoIP call spoofing

The public places great confidence in the modern global telephone network, and rightly so. The growing problem of falsification of the caller's number, however, risks eroding public confidence in the UK *Public Switched Telephone Network (PSTN)*[4] (and in those of the US and Canada), as well as in the public safety networks that depend on the UK PSTN.

The core problem of spoofing is the ability of a (possibly malicious) caller to either impersonate someone else, or to anonymise a call in such a way that fraud and abuse can occur. The network operator, the regulator, and law enforcement typically are unable to track and trace the source of the abusive call.

This is essentially a problem of *authentication* of the identity of the caller. *Is the caller really who he claims to be?*

It is not the purpose of this report to document the various types of fraud and abuse that have emerged; rather, we seek to investigate whether there are technical solutions available that might mitigate the problem, and if so in what time frame and at what cost they might be deployable in the UK. Within that space, the solutions of primary interest are based on the *Resource Public Key Infrastructure (RPKI)* solutions that are beginning to be used to authenticate *holdership* of IP addresses.

---

3 The responsible administrators are generally careful to avoid creating a precedent that parties who have usage rights to IP addresses actually own them, as if they were tangible property. Analogous issues relate to telephone numbers and also to spectrum. Property-like rights may be involved, but whether they should be legally treated as property is a complex question, the answer to which could vary from country to country. For that reason, RPKI documents refer to *holdership* rather than *ownership*.

4 When we refer to the Public Switched Telephone Network, we mean to refer both to the fixed switched network and to the mobile switched network unless we explicitly state otherwise. Some authors refer to only the fixed switched network as the PSTN, while the mobile switched network is the *Public Land Mobile Network (PLMN)*.

1.2 Our methodology

Our methodology for this report is largely the classic mix of desk research and expert interviews, together with a considerable base of knowledge and experience with the problem that we had at the outset.

We are greatly indebted to a number of first tier experts who gave generously of their time. Among them are:

- **Russ Housley**: IETF STIR co-chair, Chair of the Internet Architecture Board, former IETF Chair, former IETF Security Area Director. Author of numerous IETF security-related RFCs, including key documents related to X.509 certificate profiles.
- **Dr Steven Kent**: Fellow, BBN (USA). Dr. Kent wrote the original, key white papers on the use of RPKI for BGP security beginning in 1995.
- **Dr Geoff Huston**: Chief Scientist, APNIC (Australia).
- **Andrew de la Haye** and **Alex Band**: COO and Product Manager for RPKI, respectively, for RIPE NCC.
- **Hadriel Kaplan:** VP Technology, Oracle/Acme Packet (USA), a leading vendor of the Session Border Controllers (SBCs) used extensively in VoIP networks.
- **Jon Peterson**: VP Technology, NeuStar (USA). Mr Peterson is co-author of several of the principal problem statements in the IETF STIR working group, and co-author of many SIP RFCs. NeuStar is the contractor for both the North American Numbering Plan Administration and the US Number Portability Administration Centre.
- **John Curran:** CEO, ARIN (USA).
- Various senior voice network engineers in AT&T, Verizon and Comcast who wish to remain anonymous.

1.3 Structure of this report

The report is comprised of an extensive presentation of the background to the problem, including an introduction to the Resource Public Key Infrastructure (RPKI) (in Chapter 2), followed by our assessment of the feasibility of using RPKI-like solutions to verify the authenticity of phone numbers asserted as the caller ID for Voice over Internet Protocol (VoIP) calls (in Chapter 3). We then conclude with conclusions and recommendations (in Chapter 4).

2 Caller ID spoofing and possible means of dealing with it

Key Findings

- Falsification of the Caller ID is a serious growing problem for *Voice over IP (VoIP)* calls in the UK and in a number of other countries.
- The Caller ID spoofing problem has arisen due to the massive technical transition from traditional *Time Division Multiplexing (TDM)* and *Signalling System 7 (SS7)* to all-IP-based technologies such as the *Session Initiation Protocol (SIP)* and the *IP Multimedia System (IMS)*. In addition, various regulatory actions have appropriately encouraged new entrants into the voice marketplace. This evolution benefitted consumers, but also undermined the "Circle of Trust" among providers of traditional voice communications.
- For the somewhat related problem of Internet route hijacking, the implementation of *Resource Public Key Infrastructure (RPKI)* as the first phase of *Secure Inter-Domain Routing (SIDR)* is widely viewed as the most promising technical means of mitigating the problem.
- The European implementation of RPKI/SIDR by RIPE NCC can be deemed to be a success. Their emphasis on ease of use was well placed. Some 20% of their members routinely download new copies of their data cache. Implementation costs were modest.
- RPKI/SIDR authenticates "holdership", and helps to protect the Internet routing system against inadvertent configuration errors; however, it does not protect against malicious attacks.
- Understanding the problem of securing *Internet Numbering Resources (INR)* and the *Border Gateway Protocol (BGP)* using SIDR took a great deal of time. The standards process alone took almost eight years.
- Industry adoption has been slow for a number of reasons. Implementation and use is totally voluntary on the part of the ISPs. Complexity in the IPv4 address blocks has hindered deployment.
- The use of mitigation techniques to Caller ID spoofing that incorporate RPKI for the UK phone system is inevitable in our view. Likely deployment time lines for RPKI-based authentication, if mandated today, would be at least five to eight years.
- Any concerted policy approach would necessarily consider measures to mitigate Caller ID Spoofing together with two interrelated topics: (1) *Local Number Portability (LNP)*, and (2) IP-based *Network-to-Network Interconnection (NNI)*. LNP complicates the problem space, and may necessitate a larger, more volatile, and more complex RPKI certificate repository than would otherwise be needed. A holistic view of the problem space is needed.

In this section, we consider the root causes of the VoIP call spoofing problem, and then review the key technical elements of PKI, RPKI/SIDR, and future IETF STIR-based solutions that we and most experts consider to be among the most promising avenues for addressing the challenge of call spoofing.

2.1 Root causes of the VoIP spoofing problem

Just a few years ago, one could typically rely on the correctness of the phone number displayed when one received a call. Today, the migration to IP-based voice services has made a wealth of new services available to consumers, but at the same time and as an unexpected form of "collateral damage" has made it much easier for malefactors to falsify the caller ID. This problem could potentially threaten consumer trust in the UK Public Switched Telephone Network (PSTN).

Food for Thought 1: Falsification of Caller ID with VoIP is not a new problem in the UK

Falsification of Caller ID under VoIP has been an issue in the UK for a long time. The following is excerpted from a 2008 article in *The Guardian*.

"Do you find the Caller ID display on your home telephone useful? It's a great way to see who's ringing before answering. But thanks to a cold-calling American holiday company, some people now realise they cannot rely on Caller ID to tell the truth. Marketers can manipulate the telephone system to appear whomever they want to be - including non-existent numbers in Stratford-upon-Avon. Jack Wraith, the chief executive officer of the Telecommunications UK Fraud Forum, says: 'For some time now the UK as well as the rest of Europe has been a target for these sort of calls.' …

The calls from Tropical Grand Vacations of Florida go like this. A recorded voice says: 'Congratulations. You have won an all-inclusive cruise to the Caribbean.' It then asks you to press nine. Glancing at the telephone's caller display, you notice 01789 0000000 … The dialling code, if not the suffix, is the code for Stratford-upon-Avon. … For those using Caller ID to screen calls, it's a worrying development: can it still be trusted?

Presentation numbers (the number you see) are sometimes changed to 0800 or 0845 numbers. "However, they must conform to Ofcom's rules to ensure that they are not used to facilitate scams or malicious calls," says Ofcom. But if the caller is outside the UK, no action may be possible.

'Often this type of marketing call is sent via VoIP [Voice over IP] and VoIP networks can be less than first class. It would be simple for this marketer to ask his VoIP provider to mask or insert any chosen CLI,' says BT. …

So could anyone stop the calls? 'In theory. In practice, going after them could be difficult and challenging," says Wraith. "No matter what processes or procedures we put into place, whether they're voluntary, industry-led or whether they're legislated for by government, people will find ways round them.' …

> Since July, the Information Commissioner's Office has had 600 complaints where it's been unable to identify such callers and 1,200 calls about automated calls in general. 'We're looking into these calls and we are investigating the complaints we have received,' says the ICO which may contact US authorities. …"[5]

It is worth noting how the problem began in the first place.

For several decades, national regulators, including Ofcom, have attempted to inject competition into the telephone services marketplace by encouraging new competitive firms to enter the market, and by introducing procompetitive policies such as *Local Number Portability (LNP)*.

Where there was once a single national telephone monopoly, now there is a more mixed environment with new companies providing classic telephone service over a wide variety of access platforms (fixed, cable, and mobile) and new technologies such as *Voice Over IP (VoIP)* based on the *Session Imitation Protocol (SIP)* originally developed by the *Internet Engineering Task Force (IETF)*. These procompetitive regulatory efforts have succeeded in reducing costs of voice telephony for both UK consumers and for the network operators themselves; however, an unintended consequence of this new environment is that the traditional *circle of trust* among network operators has become frayed at the edges. The migration to IP-based voice has permitted a new class of traffic to inject itself into the UK PSTN. In order to increase their wholesale termination revenue from other network operators, some network operators have opened up their networks to less than altogether savoury operators, especially from overseas.

In other words, "No good deed goes unpunished". The same procompetitive initiatives that introduced competition and yielded consumer benefits also opened the door to certain forms of abuse.

Traditional telephony technologies such as *Time Division Multiplexing (TDM)* and classic telephony signalling such as *Signalling System 7 (SS7)* have proven ill equipped to firmly maintain trust and identity relationships among service providers as new call origination and termination strategies have been deployed.

The expansion of the Internet has also reduced the cost of transporting voice to near zero. Studies from Cisco Systems and Georgetown University have noted that voice communications represents just 1% of all Internet traffic, but nonetheless represents up to one third of the aggregate revenue in the system.[6]

---

5 Michael Pollitt (2008), "Who's really on the phone? With marketers able to manipulate what appears on your Caller ID, can the technology be trusted?", *The Guardian*, 11 December 2008, at http://www.theguardian.com/technology/2008/dec/11/caller-id-faking-fraud.

6 Anna-Maria Kovacs (2013), Telecommunications competition: the infrastructure-investment race, http://internetinnovation.org/images/misc_content/study-telecommunications-competition-09072013.pdf.

2.1.1 The transition of the PSTN to an IP-based Next Generation Network (NGN)

Complicating the problem is the worldwide Transition of the PSTN to all-IP technologies such as the Session Initiation Protocol (as specified in RFC 3261) and the *IP Multimedia Subsystem (IMS)* developed by the *3rd Generation Partnership Project (3GPP)*. IMS is a superset of SIP that was developed to address the specific problems in all-IP mobile networks, and which has also been adapted to apply to fixed network voice services (and potentially other services as well).

Recent data compiled by the US FCC indicates that nearly 30% of all voice calls in the US originate on and/or interconnect with all-IP networks.[7] US cable operators are nearly 100% SIP/IMS-based.

In particular, the US mobile industry is ramping up investment to deliver *Voice over LTE (VoLTE)*, which uses IP-based packet technology based on SIP to deliver high definition voice over the mobile networks.[8] A similar evolution can be expected in due time in Europe. This will mean that mobile calls in the UK and internationally will have audio quality roughly equal if not surpass services such as Skype with suitable Quality of Service (QoS) built in. Major enterprises have also adopted SIP as the primary technology for PBX systems within their organizations and service providers have responded by deploying *SIP trunking* (a technique that uses SIP technology to connect enterprise voice systems directly to service provider networks).

A number of interrelated factors are pushing network operators to migrate voice services rapidly to an IP base:

- Lower unit costs for IP-based networks.
- The aging condition of the existing Class 5 Voice switches and SS7 infrastructure. Some of this network equipment is now nearly 25 years old.
- Difficulties in maintaining (to say nothing of enhancing) the PSTN. The companies that manufactured these products are either no longer in business or have radically restructured. Nortel no longer exists. Siemens has exited the voice communications business. Alcatel-Lucent continues to struggle financially.
- Parts are in short supply. ATT admitted in recent FCC filings that it has been forced to buy line cards for its core 5ESS phone switches from eBay.
- The personnel that designed or maintain this TDM and SS7 equipment are retiring at an accelerating rate.

---

7 FCC (2013), "FCC Releases New Data on Local Telephone Competition", http://www.fcc.gov/document/fcc-releases-new-data-local-telephone-competition-4.

8 Carl Weinschenk (2014), "VoLTE approaches the starting line", http://www.itbusinessedge.com/blogs/data-and-telecom/volte-approaches-the-starting-line.html.

- Needlessly high operational costs to maintain TDM and IP-based network concurrently in parallel.

Many of these same considerations could make it difficult to introduce new technology into the existing PSTN to deal with Caller ID spoofing.

2.1.2 New network, new threats, new responses

The transition of voice telephony from the TDM world to world of IP has given criminal elements a vast new tool box with which to prey on the innocent, and has also enabled a new level of abuse such as *Caller ID Spoofing* and telephony *Denial of Service (DoS)* attacks against specific individuals and businesses. Modern VoIP protocols such as the IETF-developed *Session Initiation Protocol (SIP)*[9] and its superset the 3GPP-developed *IP Multimedia Subsystem (IMS)*[10] hold great promise to deliver higher quality, together with new forms of service delivery, that potentially represent a major advance for UK consumers and businesses. Unfortunately, the same transition from traditional TDM and SS7 phone networks to all IP systems has created ample attack vectors not only in Caller ID Identity spoofing, but also for large scale fraud that first attacks social media to gain Personal Identifying Information, and then uses that information to bypass traditional security measures in Call Centre environments.[11]

Food for Thought 2: Caller ID falsification can be used to exploit other security flaws

A recent article in *The Register* documents a reporter who successfully hacked into voice mail accounts at two major UK network operators. In this case, inadequate security arrangements at the network operators in question was arguably the root problem. "Voicemail inboxes on two UK mobile networks are wide open to being hacked. An investigation by *The Register* has found that even after Lord Leveson's press ethics inquiry, which delved into the practice of phone hacking, some telcos are not implementing even the most basic level of security. …

[The] 'calling line identification' (CLI) shown at the receiving end … [is] a bit of a misnomer because it can be changed as required. I'd long suspected that miscreants were hacking voicemail by spoofing their CLIs to fool the phone system into thinking it was the handset collecting the messages – but surely that's too easy? It is trivial to set an arbitrary CLI when making a call. I had to find out if voicemail systems were vulnerable to spoofing. …

The special sauce here is how does the mobile phone network know which phone you are calling from? The easy way is to look at the CLI sent when establishing a call. Unfortunately, as our reader found out, this caller identification isn't at all secure and can be spoofed …

9 http://www.ietf.org/rfc/rfc3261.txt.
10 ttp://en.wikipedia.org/wiki/IP_Multimedia_Subsystem.
11 http://www.cio.com/article/726310/How_Emerging_Technology_Fights_Fraud_in_the_Call_Center.

It's not like the networks have not been warned. … [T]he mobile networks' own industry body, the GSMA, … warned of the danger in its voicemail security guidelines published in February 2012. In this document the GSMA talks about fraud as well as security. It points to the danger that a crook could register a premium-rate number and then use that number to leave a message on the mark's voicemail. By spoofing the CLI, the miscreants can then pick up the message and return the call, raking in the profits from the premium-rate call."[12]

Of special concern are attacks against public safety Institutions such as Emergency Dispatch Centres and *Public Safety Access Points (PSAPs)*, hospitals, schools, and other public institutions. These are now well documented in North America.[13] A recent study on the Future of Public Safety networks by the CRTC in Ottawa highlighted many of these issues.[14]

Ofcom understands the background in which Caller ID spoofing is occurring and the growing problem it presents; moreover, Ofcom recognises that the problem is accelerating in other jurisdictions. Indeed, Ofcom is cooperating with the United States Federal Communications Commission (FCC) and the Canadian Radio-television and Telecommunications Commission (CRTC) in this matter.[15] [16]

The United States FCC is about to undertake a substantial regulatory proceeding leading to the virtual elimination of TDM and SS7 from the US Public Switched Telephone Network and replacing it with an all IP-based SIP network within five to seven years.

Food for Thought 3: Impersonation of tax collection agents in the United States

An egregious case of impersonation was recently reported in the press in the United States. "Treasury officials have warned of a widespread and sophisticated phone scam involving callers who impersonate Internal Revenue Service representatives and demand immediate payments with pre-paid debit cards and wire transfers. IRS inspector general Russell George described the ruse as the 'largest ever' of its kind last month, noting that thousands of victims had already lost more than $1 million through the ploy. ... [T]he callers hide their location by using Voice Over Internet Protocol (VOIP) services, which basically allow phone conversations to take place

12 Simon Rockman (2014), "Reg probe bombshell: How we HACKED mobile voicemail without a PIN: Months after Leveson inquiry, your messages are still not secure", *The Register*, 24 Apr 2014.

13 Chris Nussman (2013), "DHS Bulletin on Denial of Service (TDOS) Attacks on PSAPS", NENA, https://www.nena.org/news/119592/DHS-Bulletin-on-Denial-of-Service-TDoS-Attacks-on-PSAPs.htm.

14 Timothy Denton (2013), "A Report on Matters Related to Emergency 9-1-1", CRTC, http://www.crtc.gc.ca/eng/publications/reports/rp130705.htm.

15 http://media.ofcom.org.uk/2013/10/21/ofcom-joins-international-taskforce-to-tackle-number-%E2%80%98spoofing%E2%80%99/.

16 Ofcom (2013), "Ofcom joins international taskforce to tackle number 'spoofing'", http://www.crtc.gc.ca/eng/com100/2013/r131021.htm.

over the internet. The scammers use this to pretend they are calling locally or from an IRS's assistance number: 1-800-829-1040."[17]

2.2 Candidate solutions based on Resource Public Key Infrastructure (RPKI)

How could technology restore the ability of one carrier to trust the origin of a voice call, irrespective of whether the call transits multiple intermediary networks?

Any solution must rest crucially on two pillars:

- Robust verification that a call truly originates from the network from which it purports to have been originated; and
- Robust verification that the originating network truly holds the telephone number that is asserted to be the Caller ID.

In the world of the PSTN, this verification was implicitly linked to the physical topology of the switched network; in the world of IP, however, network traffic flows are any-to-any by design. Moreover, the Internet Protocol (IP) includes the source of each packet (datagram), but *provides no inherent validation of the source address of any IP packet*. As a practical matter, the physical topology of the network no longer serves to assure the called network of the origin of the call.[18]

In the Internet, the normal preference is to prefer end-to-end solutions over point-to-point solutions, especially where the services in question are themselves inherently end-to-end.[19] Voice calls are clearly end-to-end between caller and called party, and also between calling network and called network (although the "ends" are not the same in the two cases). It is therefore natural to turn to end-to-end solutions for the required forms of verification.

On an end-to-end basis, *authentication* (providing assurance that the source of a communication is as claimed) and the closely related problem of *authorisation* (providing assurance that the party seeking services is entitled to them) are routinely verified using *public key* cryptographic techniques. It is therefore natural that the Internet community has turned to public key technology for potential solutions to ensuring the needed (but technically challenging) verifications for VoIP voice calls.

---

17 Josh Hicks (2014), "Listen to the 'largest ever' phone scam involving IRS impersonators", *Washington Post*, 16 April 2014, at http://www.washingtonpost.com/blogs/federal-eye/wp/2014/04/16/listen-the-largest-ever-phone-scam-involving-irs-impersonators/?hpid=z4.

18 Some network operators choose to implement direct paths (usually virtually circuits) to one another in order to maintain a trustworthy connection. This cannot be relied on in the general case.

19 J.H. Saltzer, D.P. Reed and D.D. Clark (1981), End-to-End Arguments in System Design, at http://web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf.

Public Key Infrastructure (PKI) is everywhere. It is an integral part of Electronic Commerce worldwide. PKI is used to cryptographically secure and authenticate online commerce (SSL Certificates), for credit cards (chip and PIN), and for digital content.[20] Every television set top box, Blu-ray player, and satellite receiver system has an encryption system associated with it. Even the electric utilities are looking at PKI to solve problems with the power grid, securing electric meter data and other vital components.

PKI has been shown to work, so it is perfectly natural that the technical community would look to PKI as a promising possible solution to the spoofing problem. It is proven, generally reliable (with proper policy controls), and has proven to be highly scalable. If one assumes that each telephone number allocated within the United Kingdom and currently in service (probably about 150 million, if one follows the rule of thumb that the number of active telephone numbers is about 2.5 times as great as population) must be secured, the magnitude of data is not unreasonable for a PKI system.

As we explain in Section 2.22.3, solutions based on PKI are beginning to be used today, in the form of *Resource Public Key Infrastructure (RPKI)*, to achieve robust verification that a given network (as identified by its *Autonomous System (AS) number (ASN)*) truly holds the IP address blocks that appear in its *Border Gateway Protocol (BGP)* routing announcements. This is quite similar to the problem of establishing *holdership* for a telephone number, or block of numbers. Authentication that the source of the announcement is also as claimed is a related but separate problem that is currently being worked on and standardised by the IETF (see again Section 2.22.3).

Beyond that, the IETF has also developed other related PKI systems to solve other problems. For instance, *Domain Keys Identified Mail (DKIM)*[21] helps to avoid e-mail spoofing. The problem in e-mail SPAM is very similar to that in Caller ID spoofing, where there is a need to validate and authenticate identification information associated with the origin of a message, including the author's name and address.

There are a variety of other IETF PKI based protocols in wide deployment (including DNSSEC and DANE among many others), and the basic technologies and concepts for X. 509 certificate revocation and notification have been developed in the IETF and are well understood.

All things considered, we are strongly of the view that the use of PKI-based technology for secure telephone number authentication and validation is suitable for purpose.

---

20 Wikipedia contributors, "High-bandwidth Digital Content Protection," Wikipedia, The Free Encyclopedia, http://en.wikipedia.org/w/index.php?title=High-bandwidth_Digital_Content_Protection&oldid=613566350 (accessed June 20, 2014).

21 IETF (2011), “Domain Keys Identified Mail (dkim) … Charter for Working Group”, http://datatracker.ietf.org/wg/dkim/charter/.

With that said, the basic design decisions to be made in any PKI system include:

- The design of X.509 certificates to meet the application requirements.
- The structure of the data object that is to be signed – in this case, the SIP INVITE Message.
- A repository / distribution system for the certificates.
- Security for the repository / distribution system.
- The design of appropriate policies for *Certificate Revocation Lists (CRLs)*.
- The choice of encryption material.

2.3 Securing the Internet routing system: SIDR and RPKI

In this section, we explain the promising work of the IETF *Secure Inter-Domain Routing (SIDR) Working Group (WG)*, and explain the roots of the work in PKI and RPKI/SIDR infrastructure to secure *Internet Numbering Resources (INR)*.

The problem of securing numbering resources in an Internet-centric environment is associated with a long history. As early as 1997, it was already clear that the Internet routing system was under a spoofing threat from entities that wanted to manipulate the *Border Gateway Protocol (BGP)* routing system to redirect IP traffic away from its intended source.[22] Various approaches to hardening the Internet routing system were attempted, notably including Secure BGP (S-BGP); however, none achieved widespread acceptance or deployment.[23] Eventually, the IETF formed the *Secure Internet Domain Routing (SIDR)* working group in 2006.[24]

Routing in the Internet, as with interconnection in the classic PSTN model, is largely based on trust. The Internet relies on an ISP being able to assert and validate that it is authoritative for a particular *Autonomous System (AS)* number and for one or more blocks of IP addresses. The Internet routers then propagate that information throughout the Internet using BGP.

The work undertaken by the SIDR WG is very clearly elucidated by Huston and Bush (2011).[25] As they explain, "the approach the SIDR Working Group has taken … was undertaken in three stages: the first concentrated on the mechanisms to support

---

22 Geoff Huston and Randy Bush (2011), "Securing BGP with BGPsec", *The Internet Protocol Journal*, Volume 14, Number 2, http://www.internetsociety.org/articles/securing-bgp-and-sidr.

23 Stephen Kent, Charles Lynn, Joanne Mikkelson, and Karen Seo (2000), "Secure Border Gateway Protocol (S-BGP)—Real World Performance and Deployment Issues".

24 See IETF (2014), "Secure Inter-Domain Routing (sidr)", http://datatracker.ietf.org/wg/sidr/.

25 Geoff Huston and Randy Bush (2011), "Securing BGP with BGPsec", *The Internet Protocol Journal*, Volume 14, Number 2, op. cit.

attestations relating to addresses and their use; the second looked at how to secure origination of routing announcements; and the third looked at how to secure the transitive part of Border Gateway Protocol (BGP) route propagation."

It is crucial to bear in mind that *only the first of these three stages is standardised and deployed today*. Certificates can now be used to validate attestations relating to Autonomous System Numbers, IP address blocks and their use (see Section 2.3.1); however, mechanisms to validate the origin of BGP announcements are not yet deployed, and work to secure the AS path is at a still earlier stage of development. *The existing mechanisms can be used today by human experts to help identify inadvertent routing configuration errors; however, deployed solutions are by no means far enough along to deter malicious hijacking of the Internet routing system.*

Having said this, it is interesting to note that the hierarchy of authority for INR roughly is somewhat analogous to the hierarchy of authority for telephone numbers. *ICANN IANA* is the central root for all INR. It in turn allocates numbering resources to *Regional Internet Registries (RIRs)* who then allocate resources to *Internet Service Providers (ISPs)* who then ultimately provide those resources to consumers and enterprises. It is a simple and very neat tree of authority. SIDR-based RPKI is applied at the RIR layer, and each of the five global RIRs is now a full *Certificate Authority (CA)* that manages the RPKI infrastructure on behalf of its regional members.

In addition, it is worth noting that the RPKI system could potentially facilitate the creation of a trading platform that could enable ISPs to trade IPv4 address allocations among themselves.

2.3.1 Experience with RPKI to date

The SIDR system is now in full deployment among all five RIRs, with however very uneven results in terms of adoption.

Each RIR maintains its own repository of RPKI certificates. The size of this repository, then, represents an important measure of the adoption of RPKI. A crucial element of the RPKI infrastructure is the *Route Origination Authorization (ROA)*. An ROA is "an authority created by a prefix holder that authorizes an AS to originate one or more specific route advertisements into the interdomain routing system." It thus links an AS number to an IP address block.

That the five RIRs have achieved very different levels of take-up by their respective members becomes obvious when you consider the statistics, as shown in Figure 1. RIPE NCC has been particularly active in making RPKI easy to use; in consequence, their members have registered a substantial fraction of their total *IP version 4 (IPv4)* address assignments (roughly equivalent to six /8 address blocks, each of which contains 2^{24} IPv4 addresses). The other RIRs have not achieved anywhere near the

same level of adoption by their members – the number of IPv4 addresses covered is an order of magnitude less.

Figure 1: IPv4 address space covered by ROAs in each of the RIRs

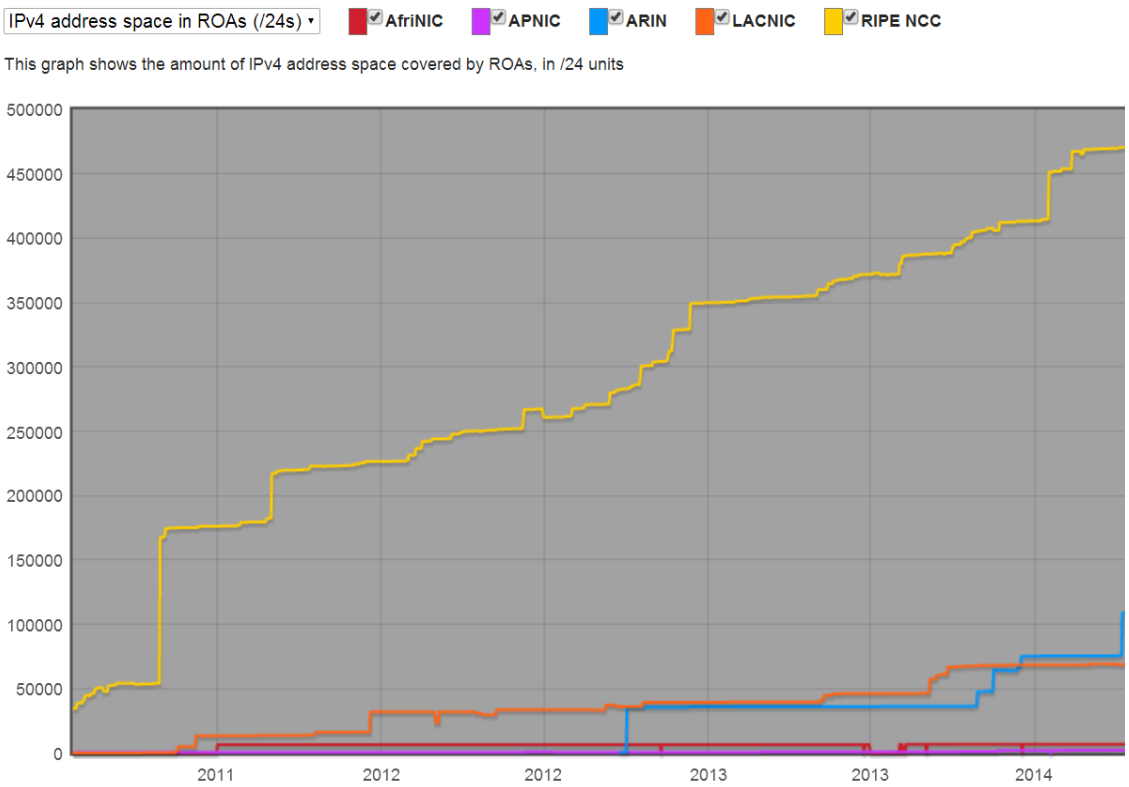

Source: RIPE NCC[26]

The adoption of protocols in the Internet is purely voluntary. The introduction of RPKI and other SIDR mechanisms requires the ISP to implement new tools, and to re-think various operational procedures. It often takes considerable time to achieve sufficient adoption of a new capability in order to benefit from the *network effects* that flow from widespread take-up.[27]

Indeed, it is for this reason that RIPE's success in getting their members to create ROAs, and also to download the collection (cache) or ROAs as it is updated, is crucial to the ultimate success of RPKI. Our sense is that they are well on their way to achieving critical mass; moreover, the expectation of continued success and growth in

26 See http://certification-stats.ripe.net/.
27 J. Scott Marcus (2004), "Evolving Core Capabilities of the Internet", *Journal on Telecommunications and High Technology Law*.

adoption by RIPE NCC members probably encourages more members and more equipment and software vendors to take up RPKI.

2.3.1.1 RIPE NCC

Among the five RIRs, the RIPE NCC (the European RIR, based in Amsterdam) is widely viewed as the thought leader. Based on our discussion with them, we share the assessment.

RIPE NCC recognised early on that achieving substantial adoption was crucial, and that *perceived ease of use* would be in turn be crucial in achieving that adoption. In consequence, their efforts focused on several key areas:

- Providing a user-friendly web-based interface for administrators.
- Out-sourcing all of the PKI infrastructure management issues, so as to relieve their members of the burden (for example, of dealing with key roll-over).
- Avoiding needless legal and administrative complexity.
- Training and awareness-raising activities on behalf of their members.

Their user interface is simple, and is based in terms and vocabulary with which an IP address administrator would be fully familiar. The underlying complexity, including in particular the cryptographic detail, is hidden from the IP address administrator.

For analogous reasons, RIPE NCC decided at the outset to offer all administration themselves. Their members can alternatively, if they so choose, implement their own delegated RPKI infrastructure, but only one of their members has chosen to (partially) do so. There is simply no incentive for most of their member organisations to take that level of complexity on themselves, nor to keep the specialised staff with specialised skills on board (and on call). RIPE NCC is trusted by its members, who understandably prefer to have RIPE NCC take on the complexity and potential headaches of RPKI administration.

Finally, RIPE NCC has taken a simple approach to legal liability. When a member first downloads the software, they acknowledge in doing so their acceptance of RIPE NCC's terms and conditions. These terms and conditions notably provide a disclaimer of those legal liabilities that can properly be disclaimed.[28]

About 20% of RIPE NCC's members download new versions of RIPE NCC's RPKI cache. This represents a very high level of adoption for a new and nascent service. At the same time, it must be understood that RPKI and SIDR *do not mandate any*

28 In general, one cannot disclaim liability for gross negligence or wilful misconduct.

particular use of the certificates, and the general sense is that automated enforcement would be premature; thus, they are used to flag apparent problem cases to human administrators, not to automatically enforce policy.

It must be recalled that, as long as RPKI is only partially implemented, all routes will be in one of three states, not two: "If a given route matches exactly the information contained in an ROA whose EE certificate can be validated in the RPKI (a 'valid' ROA), then the route can be regarded as a 'valid' origination. Where the address prefix matches that in a valid ROA but the origination AS does not match the AS number in the ROA, and there are no other valid ROAs that explicitly validate the announcing AS, then the route can be considered to be 'invalid.' Also, where the address prefix is more specific than that of a valid ROA, and there are no other valid ROAs that match the prefix, then the route can also be considered 'invalid.' Where the prefix in a route is not described in any ROA and is not a more specific prefix of any ROA, the route has an 'unknown' validation outcome."[29]

This three state model is one of several factors that makes fully automated verification of BGP advertisements challenging if not impossible. It is often said (for example, by those who develop *decision support systems*) that problems can be view as being unstructured, and thus requiring human analysis; fully structured, and thus suitable for solution entirely by computer (for example, inventory management); or semi-structured, which is the realm where automated tools can help to support a human analyst. The RPKI arguably converts the verification of routing announcements from an unstructured problem to a semi-structured problem, but not yet to a fully structured problem.

The RIPE NCC experience makes clear additional challenges to deployment, some of which could also be relevant to the application of similar technology to the VoIP spoofing problem.

As a notable example, many of the address allocations are in effect fragmented. In the nineties, IPv4 address assignments were flexibly aggregated into the largest feasible blocks in order to make BGP routing more efficient, and to avoid premature exhaustion of the most popular address (class B). These large blocks are gradually breaking down over time. Some large ISPs may no longer even have good records to indicate to whom the addresses were assigned. Maintaining certificates for large address blocks would be very efficient; however, as the assignments begin to look more like Swiss cheese (i.e. full of holes), more records are needed to cover a given number of addresses.

This fragmentation issue has an analogy to the world of telephone numbers, where Local Number Portability (LNP) potentially requires a far larger and more complicated authentication data repository than would have been required if telephone numbers were still maintained in the pristine blocks of 1,000 or 10,000 numbers in which they were originally assigned to network operators.

29 Huston and Bush (2011), op. cit.

2.3.1.2 APNIC

APNIC Is the Regional Internet Registry for the Asia Pacific Area. APNIC has a RPKI system in place based on the SIDR architecture for over four years now. Geoff Huston is APNIC's Chief Scientist, and is an expert on SIDR.[30] [31] [32]

APNIC began its work on RPKI from a slightly different perspective than that of the security of the BGP. The original intention was to create a secure and authoritative WHOIS for IPv4 numbers, which are nearly exhausted. The idea was that secure WHOIS would enable a stable trading market between Asia-Pac ISPs, thus conserving IPv4 number resources until IPv6 gained adoption. The work within APNIC began in 1998, went into production in late 2010, and is still under active development.

According to Huston, the major issues were reaching consensus on the actual X.509 certificate profile and the Certificate Revocation Lists (CRLs). Theoretical discussions began on this in the early 2000's, even before the IETF SIDR Working Group was formed in 2006. Consensus in the IETF on just the X.509 certificate profile took nearly two years, and CRL Policy took even longer.

The arguments centred on keeping as much extraneous information out of the X.509 profile as possible, specifically financial data. What should the maintenance time frame be for the certificates within the CRL? How should synchronisation be maintained for the repositories?

APNIC's potential community of RPKI users is the 500 or so members. As of today, less than 2% of the IPv4 address space has been signed. A key reason for lack of traction to date is the fractured nature of the allocations by the largest network operators including Telstra, NTT, Singtel and of course China Mobile and China Telecom. Some network operators do not have reliable records as to what they have actually deployed, and in the case of China they are nervous about having critical security infrastructure in the hands of third parties.

The actual cost of APNIC's SIDR infrastructure was several million AUDs, with three or four full time equivalent (FTE) employees currently implementing the program.

As regards the potential applicability of RPKI technology to VoIP, Huston emphasises that that the IP address space is different from that of telephone numbers; nonetheless, it is perfectly reasonable to make the comparison for the purposes of RPKI/STIR technology selection. The structure of telephone numbering is actually clearer and easier to define than the IP address space, making the Chain of Authority easier to manage. Ultimately, each individual telephone number could be signed.

---

30 http://www.apnic.net/publications/research-and-insights/geoff-huston.
31 http://www.potaroo.net/.
32 http://www.potaroo.net/ispcol/2014-04/rpkiv.html.

When asked whether RPKI/SIDR could be made to work for STIR, Huston's answer was simple and direct. "Of course. It will probably be a lot easier for you guys since the cert repositories for TN's will be so much easier to define. … Databases are cheap these days. … You'll end up learning from our mistakes."

Huston felt that some five to seven years might be needed to achieve RPKI standardisation for VoIP.

2.3.2 How long did the SIDR process take?

Based on our interviews with key principals in the process, from the development and recognition of a clear problem statement, though the technical standardisation, development and testing of software, to actual deployment in the field took some eight years.

It is important to note, however, that the process is far from over. Even as regards the use of RPKI, there is still ongoing technical work based on operational experience that is refining the processes that the RIRs are using. There is still a strong need to educate ISPs on the benefits of using RPKI/SIDR.

More to the point, however, is that *RPKI/SIDR in its present form does not prevent route hijacking*. RPKI/SIDR enables statements about ASN and address ownership to be cryptographically authenticated, which can help to prevent many forms of inadvertent misconfiguration; however, RPKI/SIDR origin validation "…does not provide cryptographic assurance that the origin AS in a received BGP route was indeed the originating AS of this route. A malicious BGP speaker can synthesize a route as if it came from the authorized AS. Thus, [origin validation] is very useful in detecting accidental misannouncements, but origination validation does little to prevent malicious routing attacks from a determined attacker."

Recall that RPKI/SIDR addresses only the first of the three planned phases of SIDR. It would still be necessary to secure the origination of routing announcements, as well as the transitive part of Border Gateway Protocol (BGP) route propagation.

2.4 Securing Voice over IP (VoIP): STIR and RPKI

What would it take to apply SIDR principles to phone numbers? This has been the focus of the IETF's STIR Working Group.

2.4.1 IETF's STIR Working Group

Current technical efforts to address Caller ID spoofing have centred on the STIR Working Group of the *Internet Engineering Task Force (IETF)*.[33] STIR's work should be viewed as a 'Work in Progress'" at a very, very early stage. We would conservatively estimate that it will be 18 to 24 months before the most basic technical outline within SIP/IMS is defined to determine what signalling mechanism goes "on the wire". The first order of business will be to define what fields in the SIP INVITE are signed by whom, and how that signed data is ultimately validated.

STIR benefits from the work already done on SIDR. There is a rough analogy between the hierarchical structure of telephone numbers and that of AS numbers and IPv4 address allocations. The hierarchical structure of telephone numbers is convenient in terms of legal and regulatory authority. From the global root of ITU E.164, the UK has plenary authority over its portions of the global numbering plan (associated with country code 44) in accordance with UK national law. The same is true in the United States, in Canada, and in every other country that we are aware of.[34]

The problem of origin identification in SIP and other forms of real-time communications, such as XMPP text messaging, has been well understood for years. The earliest attempt was the development of the "P-Asserted Identity" (PAI) headers in the SIP signalling message defined in RFC 3325.[35] This technique has been implemented in most carrier based SIP/IMS systems; however, it has a fatal flaw. SIP signalling is ASCII text and as such is potentially subject to modification in transit by various network elements. The ASCII text could easily be stripped out or modified. *The Internet Architecture Board (IAB)* recognized this limitation in 2012:[36]

> Even in a SIP-only environment, the choice of syntax, made separately by different implementers and users, impacts the security mechanisms that can be used for attesting to the authenticity of the identifier. Without any form of cryptographic identity assertion, the 'From' header can be easily forged, and headers are often stripped or modified by intermediaries in transit. Attempts at enhancing the integrity protection of SIP identity have not seen wide deployment.

The STIR Working Group has initially focused on two general problems. The first is the issue of authentication and validation of the SIP signalling between two *SIP Service*

---

33 https://datatracker.ietf.org/wg/stir/charter/.

34 Note that not every country code corresponds to a single country. The United States share country code 1 with eighteen non-US entities. For background on legal authority, see: http://www.gpo.gov/fdsys/pkg/USCODE-2011-title47/html/USCODE-2011-title47-chap5-subchapII-partII-sec251.htm; http://laws-lois.justice.gc.ca/eng/acts/T-3.4/page-16.html#s-46.1; and http://www.legislation.gov.uk/ukpga/2003/21/part/2/chapter/1/crossheading/general-conditions-telephone-numbers.

35 http://www.ietf.org/rfc/rfc3325.txt.

36 http://tools.ietf.org/html/draft-cooper-iab-secure-origin-00.

Providers (SSP), where either the originating or terminating network are principally using SIP to interconnect calls. This is known as the In Band solution. The second is a reflection as to whether or how RPKI could be used where one or more legs of the call setup are still using legacy TDM infrastructure. This is known as the Out of Band solution.

Central to the In Band solution is an understanding of how the SIP message originating a call (known as an INVITE) would actually be signed using a private key by the originating service provider, how the data would be carried in the SIP messaging, and how the terminating provider would process the message retrieved by the public key to authenticate and validate the transaction and connect the call. This is trying to re-establish the Circle of Trust between two operators that says in essence, "I am Vodafone and this number has been assigned to my customer and I vouch for that". Because the message has been signed by the originator, it is irrelevant how many other service providers the INVITE has to transit to reach the terminating operator. The INVITE has been signed and the terminating operator would be able to detect any tampering with the message.

The early work is centring on a major rewrite of what is known as "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474.[37]

No work has currently started on what an X.509 profile for telephone numbers would look like, or how a Certificate Revocation List Policy would be implemented. It is the opinion of the experts we have interviewed that these issues will take considerable time.

What would the Telephone Number Certificate Repositories be, and who would manage them, are even larger questions. These issues would presumably involve regulatory policy at the Ofcom level.

2.4.1.1 The STIR communication protocol

The following diagram represents an extremely simplified view of how the STIR system would work.

37 http://tools.ietf.org/rfc/rfc4474.txt.

Figure 2: A simplified view of the STIR protocol in action

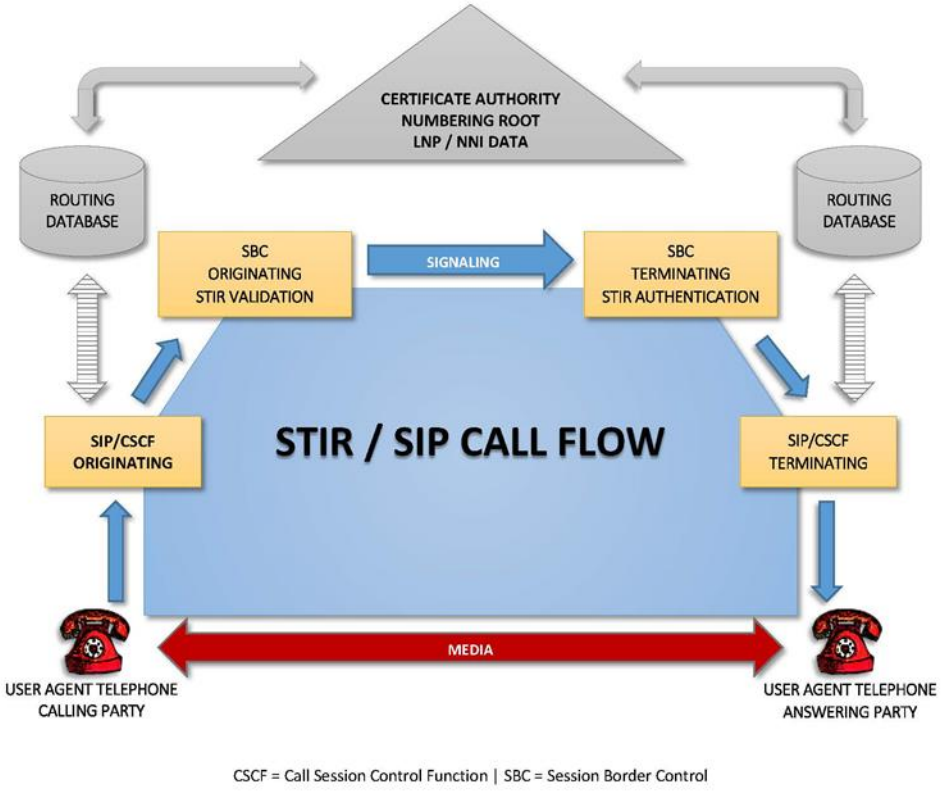

The steps in the process are as follows:

- When a session is originated on the SIP network, it first goes to what is known as a *SIP Proxy* (or a *Call Session Control Function (CSCF)* in IMS terms).
- The CSCF looks into its various databases (there are many of them) and ultimately makes a decision about how to route the call based on a variety of factors including policy and billing.
- In the case where the call is to be terminated outside its internal network, the session signalling then moves to the edge *Session Border Controller (SBC)*. SBCs provide an endless variety of functions in a modern SIP/IMS environment: Security, Quality of Service prioritisation, and media transcoding, among others. We believe that every single SIP/VoIP service provider on Earth uses SBCs.

  The emerging STIR architecture assumes that the Originating SBC will be the network element that validates or signs the phone number in the From: originating SIP INVITE message using its Private PKI Key issued by the designated Certificate Authority for the Numbering Plan Administration. Other

headers may be signed as well what headers need to be signed has not been formally standardized.

- The INVITE is then sent to the Terminating SBC where the signed signalling is inspected. After retrieving the Public Key from the designated Certificate Repository in the RPKI system, the SBC the Authenticates the From: phone number or other messages in the signalling. At that point, the terminating network operator is cryptographically assured as the identity of the sender of the INVITE, and that they have been authenticated to use the telephone number.

  In other words: “Ah ha … this message really did come from (e.g.) Vodafone and I trust Vodafone so I (e.g. BT) will complete the call.”

- From there, the normal process of media session establishment between the endpoints begins, and the phone rings.

Figure 3: Format of a SIP INVITE using a telephone number

```
INVITE sip:+19725552222@ss1.a.example.com;user=phone SIP/2.0
 Via: SIP/2.0/TCP client.a.example.com:5060;branch=z9hG4bK74bf9
 Max-Forwards: 70
 From: Alice <sip:+13145551111@ss1.a.example.com;user=phone>;tag=9fxced76sl
 To: Bob <sip:+19725552222@ss1.a.example.com;user=phone>
 Call-ID: 2xTb9vxSit55XU7p8@a.example.com
 CSeq: 1 INVITE
 Contact: <sip:alice@client.a.example.com;transport=tcp>
 Proxy-Authorization: Digest username="alice", realm="a.example.com",
  nonce="dc3a5ab25302aa931904ba7d88fa1cf5", opaque="",
  uri="sip:+19725552222@ss1.a.example.com;user=phone",
  response="ccdca50cb091d587421457305d097458c"
 Content-Type: application/sdp
 Content-Length: 154

 v=0
 o=alice 2890844526 2890844526 IN IP4 client.a.example.com
 s=-
 c=IN IP4 client.a.example.com
 t=0 0
 m=audio 49172 RTP/AVP 0
 a=rtpmap:0 PCMU/8000
```


2.4.2 Current status and unresolved questions

STIR Working Group technical activity in the Internet Engineering Task Force is now under way. Considerable progress has been made as of this date, but there are still any number of unanswered questions.

Standards development organizations, much like regulators, tend to work at their own pace as they seek a multitude of opinions, weigh options, and ultimately seek consensus on a sound course of action.

Determining the structure of the Resource Public Key Infrastructure to support STIR is a non-trivial exercise.

2.4.2.1 Time to implement RPKI/STIR

These things take time. Key elements of the process would likely include:

- **SIP protocol enhancements.** In our judgment, the IETF will probably take not less than 18 months to reach consensus on a technical specification as to how STIR protocols would work "on the wire" within SIP itself. This would be based on a revision of RFC 4474.
- **X.509 certificate profile.** In addition, there is the process of designing the X.509 certificate profile for phone numbers. Based on our interviews, this will possibly require two years.
- **Certificate Revocation List (CRL).** After that, there would be the matter of establishing policy for the Certificate Revocation List. Again, this might take two years.
- **Selection of Cryptographic Material.** Probably six months, but in parallel with the other standards work. Early thinking is focusing on Elliptical Curve Cryptography (ECC).[38] ECC has enormous advantages including a smaller key size, reducing storage and transmission requirements and reduced computational requirements at validation.
- **Regulatory consultations on Certificate Repositories.** Who issues private keys? Where are the public keys stored, in numbering databases?
- **Actual implementation.** Since the *Session Border Controller (SBC)* is the key to carrier implementation the following is a short version of how an SBC vendor might view the process.
    - Q. "How long would it take to implement RPKI in the SBC?"
    - A: "Just give me the darn key." "If I can see the IETF RFC? And I get a requirement from a carrier and they have a budget?" " 12-18 months to get something ready for a General Availability (GA) release" "Then you know what happens. You have to get into the carriers testing cycle window." "Then they beat it up for a year or so. Then maybe after a year they decide to

38 http://arstechnica.com/security/2013/10/a-relatively-easy-to-understand-primer-on-elliptic-curve-cryptography/.

> actually put it in the network." "That presumes the keys are ready at the same time "Carriers are very fussy about what they put in their networks. … That is why the phone system works."

It is reasonable to assume that, at a bare minimum, the process could take between five and seven years.

2.4.2.2 Cost to implement STIR solutions

Costs for an RPKI system of this scale are dependent on two factors that cannot yet be fully evaluated. Does the system sign only blocks of numbers assigned to UK network operators, or is the system going to be designed to sign each and every one of the roughly 150 million telephone numbers in the UK Numbering Plan? There are rational arguments for each option. There are a limited number of network operators in the UK, so signing 10,000 number blocks would be a fairly straightforward proposition with minimal direct RPKI structural costs. Signing each number within each block would increase the relative complexity of the system; however, it may be necessary in order to rigorously ensure that the terminating network knows unambiguously which network is permitted to originate a call for each number, and it might well be necessary in any case if Ofcom reconsiders its strategy for Local Number Portability.

It is our view that each and every operational number in a numbering plan should be signed, rather than attempting to sign blocks of numbers only. There are many reasons for this. First, it facilitates simple, straightforward and unambiguous validation of holdership of the number.[39] Second, signing each and every number in a nationally deployed RPKI system might stimulate innovation in the form of new products and services that have yet to be imagined. There could for instance be linkages to mobile payments, and to the Internet of Things (IoT).

Based on our interviews with the RIRs, we believe the core RPKI/STIR infrastructure costs, such as the root Certificate Authority and the signing and validation mechanisms within electronic communication networks, represent a rather small portion of the overall cost of implementation from the perspective of network operators. The majority of RPKI/STIR implementation costs are likely to rest with the integration with the *Operational Support Systems* and *Business Support Systems (OSS/BSS)* that are central to the provisioning of telephone service. These are extremely expensive systems that activate numbers, perform service activation, implement changes (adds, moves and drops), deal with mobile LNP billing and other critical elements that network operators rely on to maintain the UK PSTN. These undergo constant modification and enhancements as new requirements emerge. At this time, it is unclear what the costs to

39 Signing at the number block level does not absolutely preclude validation, but it leads to a complicated and messy system.

upgrade those systems to accommodate STIR requirements would be, and will remain unclear until the fundamental design architecture is understood.

2.4.2.3 RPKI/STIR management issues

At some point, it will be necessary to consider various ancillary issues to the management of any RPKI system for telephone numbers. Among them are key management, and anticipated transactions in the Certificate Repository. Based on our interviews, the RIRs have generally considered a one year time frame for a full roll over of the keys within their system. In our opinion, this seems reasonable. Transaction volumes are likely to roughly reflect the level of Local Number Portability transactions in the system. Experience in the US suggests that these volumes are relatively low; however, given the size of the database, a high degree of automation will be required.

2.4.2.4 Integration with existing UK numbering databases and resources

It is our strong opinion that the implementation of RPKI for the UK numbering plan may require a significant rethinking of how the UK relies on centralised numbering databases for a variety of issues.

There is no question that the deployment of a RPKI solution for the United Kingdom will be intimately linked with the structure and usage of the UK Numbering Plan and with future UK requirements for *Local Number Portability (LNP)* databases for the fixed and mobile network, as well as with IP-based *Network-to-Network Interfaces*.

Any transaction with any of these databases has ripple effects. With the allocation of new numbers or number blocks, public keys need to be associated with them.

Even now, in the United States and Canada, there is a general assumption that the RPKI infrastructure could easily slip into the existing real time Local Number Portability Databases (the *NPAC*) or the fixed database (the *LERG*). Existing data provisioning and distribution systems for these databases could perhaps be quickly leveraged for this new application.[40]

As we noted at the outset, the Caller ID spoofing problem is intimately associated with the PSTN transition to all-IP technologies. Consequently, some restructuring of the UK National Telephone Numbering Plan and its associated databases may be required.

Network-to-Network Interfaces are under technical consideration in several jurisdictions, and central to that is the need for more specific IP routing data associated with a specific telephone number. The design and deployment of RPKI should be coordinated with that of IP routing data for the PSTN transition.

40 http://www.atis.org/PRESS/pressreleases2014/010814.asp.

2.4.3 Consumer access to RPKI validation data

Can consumers themselves use the data derived from the RPKI/STIR system to make informed decisions on whether a call comes from a trusted source?

We consider it unlikely that the problem of Caller ID Spoofing can be solved solely within the network itself. It is likely that it will ultimately need the active participation of consumers. One concept under consideration is a modification of the current technical standards to provide an enhanced *Calling Name Delivery (CNAM)* capability that would give the Called Party additional validated and authenticated information on who is calling, and would potentially display network based reputation data on consumer mobile handsets, desktop handsets, or even television screens.[41]

CNAM is an SS7 service that displays a 15 character ASCII string on a phone display.[42] It is an extremely popular service. Consumers can simply look at their telephones and decide whether they wish to answer the call or not based on who is calling.

Today, many network operators display UNKNOWN, for instance, when they do not have accurate data on either the Calling Party ANI or CNAM.

Food for Thought 4: Impersonation over the phone contributes to the growing problem of voice phishing

> "The BBC has obtained exclusive figures from the financial ombudsman that show there have been nearly 100 complaints about 'vishing' (voice phishing) and courier fraud in the past three months. …
>
> Vishing and courier fraud target some of the most vulnerable people in society by duping them into transferring money directly into criminals' accounts, or handing over bank cards and personal identification numbers (PINs) to couriers. …
>
> Gangs posing as police or bank employees ring people at home telling them there has been a fraud and to ring their bank. But the criminal does not hang up, so when the victim tries to ring out they are still connected to the fraudster.
>
> Although the scams have been around for a while, there has been a huge surge in reported cases in recent months. The Metropolitan Police alone has had 2,200 cases this year and in 2012 victims lost £3.5m in its area. …"[43]

In the SS7 world, the CNAM service is a *terminating carrier* service, meaning that the terminating network operator must perform the lookup before the call is placed.

41 See for instance XFINITY (2014), "Caller ID from XFINITY® Voice", http://xfinity.comcast.net/callerid/.
42 Cisco (2007), http://www.cisco.com/c/en/us/td/docs/voice_ip_comm/pgw/9/feature/module/9-7_3_/cnam.html.
43 BBC (2013), "'Vishing' and courier scam complaints increase", 14 December 2013, at http://www.bbc.com/news/uk-25365698.

Typically, this is a *TCAP* query to the originating carriers *LIDB* service; however, several third party vendors have emerged that create their own CNAM databases that competitive network operators can then use to look up the data. In SIP/IMS, this is reversed, and the verbose CNAM data can be delivered in the originating SIP INVITE message by any of several means.

As of now, CNAM is nothing more than 15 characters of ASCII. In the modern age, this seems absurd, and logically should change. Would it be possible, for instance, to add validation data to CNAM so that the network operator could display “TRUSTED” or “UNTRUSTED”?

The IETF has begun some discussion on how this might be enhanced in the future. The idea would be to use some new data object that could inserted and ultimately signed in the SIP signalling INVITE, using STIR, that would allow for more enhanced Caller Name Delivery.

3 Suitability of RPKI-like solutions for validation of VoIP caller ID

Key Findings

- There is little doubt that an RPKI-based solution for validation of Caller ID, and perhaps CNAM information, is feasible.
- Modern database technology is more than adequate to support an RPKI-based data repository for every active phone number in the UK Numbering Plan.
- We are strongly of the view that RPKI/SIDR is almost certain to play a key role over time in any comprehensive solution to the VoIP Caller ID Spoofing problem.
- At the same time, key limitations must not be forgotten. First, the RPKI/SIDR capabilities that are deployed today authenticate holdership of IP addresses and Autonomous System numbers, but fall well short of a fully automated system, and do not (yet) protect against a serious, malicious attack. Second, the pace of standards development, software implementation, and network deployment is such that even capabilities comparable those of current RPKI/SIDR would likely require at least five to seven years to deploy widely. The time to achieve a truly effective system might perhaps be considerably longer. Third, Ofcom needs to take realistic account of what is technologically achievable; however, this is not a cause for complacency. Ofcom's decisions (together with those of the FCC and the CRTC) also influence the pace of change.
- *Our strong belief is that this problem needs to be addressed at national level in each country, i.e. at a level corresponding to that of ITU country codes, because that is the level to which responsibility for the national numbering plan has been delegated.*
- Following the typical Internet-based pattern of voluntary standards adoption for RPKI for UK phone numbers may not be sufficient. Protection against Caller ID Spoofing is of limited value to consumers until it is widely, if not universally, deployed. As we have already seen, RPKI-based authentication is for analogous reasons, of little utility to network operators until a large enough number of network operators deploy. For services such as these, network effects are crucial.

Every knowledgeable stakeholder with whom we discussed the issue was strongly of the view that RPKI-based solutions were the only practical way forward; at the same time, we found widespread recognition that no single "silver bullet" is likely to solve the problem of VoIP call spoofing.

3.1 Likely deployment scenarios: Voluntary, or Mandated?

Should Ofcom decide that some form of RPKI is necessary to secure the UK PSTN, there are two likely scenarios. As we have already noted, the timing of standards specification and of implementation into network equipment poses limitations on what can be done, and when it can be done; at the same time, Ofcom's decisions (together with those of the US FCC, the Canadian CRTC, and other NRAs) will also influence the speed with which equipment manufacturers implement the capabilities.

In our judgement, following the typical Internet-based pattern of voluntary standards adoption for RPKI for UK phone numbers may not be sufficient. There is every reason to believe UK carriers will resist implementing RPKI for Caller ID Spoofing remediation the way they fiercely resisted LNP.

Protection against Caller ID Spoofing is of limited value to consumers until it is widely, if not universally, deployed. As we have already seen, RPKI-based authentication is for analogous reasons, of little utility until a large enough number of network operators deploy. For services such as these, *network effects* are crucial.[44]

3.2 Linkages with Local Number Portability

It is useful to compare the issue of RPKI and Trust in the UK PSTN with the historical issue of Local Number Portability.

The Regulatory Framework for Electronic Communications that was enacted in 2002-2003 required LNP, and every National Regulatory Authority (NRA) transposed this obligation into national law; nonetheless, the road to implementation was rough.

Ofcom began its journey into LNP in 2000.[45] In the US, it took an Act of Congress and a major FCC Report and Order to implement.[46] Australia and Canada began their LNP process in 1997.[47]

A nationwide Gallup survey of LNP on behalf of MCI conducted in anticipation of the US Telecommunications Act of 1996 found that "… 83 percent of business customers and 80 percent of residential customers would be unlikely to change local service providers if they had to change their telephone numbers."[48]

LNP is thus a valuable capability; however, it would not have deployed without an active public policy intervention on the part of national regulatory authorities. Analogous arguments would appear to apply to measures to mitigate Caller ID Spoofing.

Any implementation of techniques to mitigate Caller ID Spoofing should strive either to improve the ease with which LNP (which is closely related to the VoIP spoofing problem) is implemented, or at least, following the oath of Hippocrates, to "do no harm".

---

44 See J. Scott Marcus (2004), "Evolving Core Capabilities of the Internet", *Journal on Telecommunications and High Technology Law*. See also Jeffrey H. Rohlffs (2001), *Bandwagon Effects in High-Technology Industries*.

45 Ofcom (2000), Numbering Directive: Number Portability Requirements, http://www.ofcom.org.uk/static/archive/oftel/publications/numbering/port0100.htm.

46 FCC (1996), In the Matter of Telephone Number Portability, http://transition.fcc.gov/Bureaus/Common_Carrier/Orders/1996/fcc96286.txt.

47 Telecom Decision CRTC 97-8 (1997), Local Competition.

48 FCC (1996), op. cit., paragraph 29.

3.3 Likely international coordination requirements

The same considerations of *network effects* that appeared in Section 3.1 argue equally strongly for an internationally coordinated approach. Protection against VoIP spoofing would be most effective if they were global; conversely, as long as some countries fail to implement protective measures, malefactors might be motivated to operate from those countries.

A particular concern is that any successful efforts to close the holes in the UK (and other supportive countries) might well serve to encourage malefactors to simply move their operations off-shore. This naturally inspires a few key questions:

- To what extent could a UK-only solution, or a solution in a few (sufficiently large) countries, be effective in mitigating the problem?
- What kind of international cooperation is needed?
- How likely is it that the necessary cooperation would be forthcoming in time?

Our strong belief is that *this problem needs to be addressed at national level in each country, i.e. at a level corresponding to that of ITU country codes, because that is the level to which responsibility for the national numbering plan has been delegated*.

Other countries – who may be less impressed with the immediacy of the problem than are the UK, Canada, and the UK – are more likely to take voluntary action if there is some demonstrably workable framework that they can join.

We think that a UK-only solution might have substantial effect if properly implemented, and if supported not only be networks but also by consumer education and perhaps by intelligent handset software. As noted in Section 2.4.3, network operators already are able to display something like “UNKNOWN” when they do not trust the ASCII character string that purports to be the Caller ID. One can easily imagine a slightly more sophisticated scheme where the displayed Caller ID might start with, for instance, with

- a check mark (√) where Caller ID is at least plausible;
- a question mark (?) where Caller ID cannot be verified;
- two question marks (??) where Caller ID is probably false.

If Ofcom were successful in getting most or all network operators, including VoIP operators, on board with a validation scheme, it might be possible to validate at least UK numbers (i.e. country code +44) with moderate confidence.[49] Off-shore operators could still falsify the Caller ID, but might have difficulty falsifying a UK Caller ID without detection. For consumers, this would at least bound the problem.

49 It might still be necessary to provide some form of safe harbour to protect network operators from legal liability for errors made in good faith, and in the absence of gross negligence.

The UK has already found a cooperative attitude in the US and Canada. The question is, to what extent will other countries be concerned?

As it happens, many developing countries are motivated to address Caller ID spoofing. Unreliable Caller ID undermines their ability to benefit from high international settlement rates. For this reason, there was strong interest in anti-spoofing measures at the 2012 WCIT of the *International Telecommunication Union (ITU)* in Dubai, which sought to update the *International Telecommunications Regulations (ITRs)*. Unfortunately, we are quite far from achieving a global consensus as to what the ITRs should like look going forward (or even as to whether their continued existence is necessary); thus, it seems unlikely at this point that Caller ID spoofing can be somehow solved in the near term through ITU action. In the medium term, however, the ITU may well be the most appropriate forum.

The UK's immediate interests likely rest more in the nearer term with the UK's major communications "trading partners" among the developed countries, especially those within the European Union. The EU, the CEPT and ETSI all represent potentially interesting avenues for European coordination. Our sense is that the issue has not yet become prominent at European level, but it is a safe bet that its relevance and visibility will steadily increase over the next few years.[50] For Europe, this is perhaps an issue whose time has not yet come.

3.4 Implications for Public Safety

Securing the UK PSTN and the UK numbering plan will ultimately impact public safety and law enforcement as well. The UK has not yet experienced this kind of attack, so far as we can determine, but it is already an issue in North America.[51] [52] Authentication and validation of the caller's telephone number will be needed at the Public Safety Answering Point (PSAP) in support of first responders, including the National Health Service (NHS).

The PSAP needs to be able to validate the 999/112 call from the originating network, since the signalling transaction would be signed by the originating network to the PSAP.

The application of tools to mitigate Caller ID spoofing has many implications for public safety, and raises many questions.

- How feasible is it to reliably display accurate network call validation data at the PSAP terminal?

50 The German BNetzA, for instance, is well aware of the long term importance of the issue.

51 Chris Nussman (2013), "DHS Bulletin on Denial of Service (TDOS) Attacks on PSAPS", NENA, https://www.nena.org/news/119592/DHS-Bulletin-on-Denial-of-Service-TDoS-Attacks-on-PSAPs.htm.

52 Timothy Denton (2013), "A Report on Matters Related to Emergency 9-1-1", CRTC, http://www.crtc.gc.ca/eng/publications/reports/rp130705.htm.

- How would call validation data be delivered in the 999/112 signalling path?
- In what time frame is the UK PSAP itself likely to evolve to an all-IP structure?
- Can the RPKI/STIR system be used to validate public alerts? How should a consumer know that an alert is really coming from the Metropolitan Police?

3.5 Implications for deployment by Ofcom

Of great importance is the potentially time consuming work to determine exactly how an RPKI infrastructure would deploy within the United Kingdom. Key questions include:

- Who are the Certificate Authorities?
- Do these Certificate Authorities need to be certified by Ofcom? What requirements must they pass?
- How would the private/public keys be provisioned and distributed?
- What security mechanisms would need to be associated with these systems?
- How would the system recover its costs, and who would bear those costs?

These are not technical questions, but rather policy questions.

The experiences at RIPE NCC and at APNIC are consistent with implementation costs in single digit number of millions of euro, and a maintenance staff of perhaps four FTEs for a system with that number of entries and that level of activity. At that level, Ofcom might perhaps choose to operate the system itself. If however a voice-oriented RPKI infrastructure required entries down to the level of individual phone numbers, and if database changes occurred with a frequency reflecting some 150 million entries, that might imply a much larger undertaking, which in turn might argue that the function needs to be out-sourced to a more specialised organisation.

Network operators will likely resist mandatory imposition of measures to mitigate Caller ID spoofing, but the cost of RPKI databases will probably not be their biggest worry. It is likely that their concerns will have far more to do with the costs they incur to upgrade their internal *Operational Support Systems (OSS)* and *Business Support Systems (BSS)*.

4 Conclusions and recommendations

The Caller ID spoofing problem is a complex and multifaceted problem. Mitigating the problem will surely be just as multifaceted, if not more so. No single "silver bullet" is likely to magically solve the problem.

We are strongly of the view that RPKI/SIDR is almost certain to play a key role over time in any comprehensive solution to the VoIP Caller ID Spoofing problem. At the same time, key limitations must not be forgotten:

- The RPKI/SIDR capabilities that are deployed today authenticate holdership of IP addresses and Autonomous System numbers, but fall well short of a fully automated system, and do not (yet) protect against a serious, malicious attack.

- The pace of standards development, software implementation, and network deployment is such that even capabilities comparable those of current RPKI/SIDR would likely require at least five to seven years to deploy widely. The time to achieve a truly effective system might perhaps be considerably longer.

Following the typical Internet-based pattern of voluntary standards adoption for RPKI for UK phone numbers may not be sufficient. Protection against Caller ID Spoofing is of limited value to consumers until it is widely, if not universally, deployed. As we have already seen, RPKI-based authentication is for analogous reasons, of little utility to network operators until a large enough number of network operators deploy. For services such as these, *network effects* are crucial.[53]

Any concerted policy approach would necessarily consider measures to mitigate Caller ID Spoofing together with two interrelated topics: (1) Local Number Portability (LNP), and (2) IP-based Network-to-Network Interconnection. LNP complicates the problem space, and may necessitate a larger, more volatile, and more complex RPKI certificate repository than would otherwise be needed. The migration to a pure IP-based Network-to-Network Interconnection that is already ongoing in many European countries implies a further opening up of voice telephony infrastructure to potential mischief. A holistic view of the space is needed.

Our strong belief is that *this problem needs to be addressed at national level in each country, i.e. at a level corresponding to that of ITU country codes, because that is the level to which responsibility for the national numbering plan has been delegated*.

Other countries – who may be less impressed with the immediacy of the problem than are the UK, Canada, and the UK – are more likely to take voluntary action if there is some demonstrably workable framework that they can join.

53 See J. Scott Marcus (2004), "Evolving Core Capabilities of the Internet", *Journal on Telecommunications and High Technology Law*. See also Jeffrey H. Rohlffs (2001), *Bandwagon Effects in High-Technology Industries*.