RTFM LLP

Study into UK IPv4 and IPv6 allocations

TENDER NO: ITT 12/2014

Submitted by:

Reid Technical Facilities Management LLP (RTFM)

St. Andrews House

385 Hillington Road

Glasgow

G52 4BL

Document Version: Ofcom/140701-00

TABLE OF CONTENTS

GLOSSARY

6to4	A mechanism for carrying IPv6 traffic over an IPv4 network
AFRINIC	African Network Information Centre (RIR)
API	Application Programming Interface
APNIC	Asia Pacific Network Information Centre (RIR)
ARIN	American Registry for Internet Numbers (RIR)
ARPA	American Research Projects Agency
ARPAnet	Forerunner of the Internet created by ARPA
BGP	Border Gateway Protocol (core Internet routing protocol)
GCN	Carrier Grade NAT
CIDR	Classless Inter-Domain Routing
DE File	Delegated-Extended-*YYYYMMDD* file published daily by each RIR
DNS	Domain Name System
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
IETF	Internet Engineering Task Force
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISP	Internet Service Provider
IXP	Internet Exchange Point
JANET	Joint Academic Network (UK education/research network)
JSON	JavaScript Object Notation
LACNIC	Latin America/Caribbean Network Information Centre (RIR)
LIR	Local Internet Registry
LRH	Legacy Resource Holder
NAT	Network Address Translation
RARE	Réseaux Associés pour la Recherche Européenne
RDBMS	Relational Database Management System
RFC	Request For Comments (Internet Standards Document)
RIPE	Réseaux IP Européens (European IP networking forum)
RIPE NCC	RIPE Network Co-ordination Centre (RIR)
RIPEstat	Facility for accessing statistics/data provided by RIPE NCC
RIR	Regional Internet Registry
RIS	Routing Information Service
SQL	Structured Query Language
Teredo	An IPv6 tunnelling/transition mechanism
TERENA	Trans-European Research and Education Networking Association

EXECUTIVE SUMMARY

PROJECT OVERVIEW

Ofcom issued a Tender to gather information on the use of IPv4 and IPv6 address space by UK-based organisations. The key motivations for this Tender are the run-out of the remaining IPv4 address space, how that is influencing the uptake of IPv6 and how these factors are affecting UK IP address holders. Ofcom wished to understand the current status of IPv6 and IPv4 address blocks allocated to UK address holders and their respective visibility within the global Internet routing tables. The scope of this study included UK-held IP address resources obtained through the RIR system and those which were distributed to legacy resource holders prior to the introduction of the RIR system. This report and analysis is the outcome of that study.

The nature of the Internet makes it difficult to precisely define its address space and usage by geographic and territorial boundaries. UK-based organisations will be using some of their IP numbering resources outside the UK: for instance UK multinationals with a corporate network, part of which is located overseas. Similarly, overseas organisations will be using their IP addresses to support their business activities in the UK. An added complication is the dynamic nature of the Internet itself. IP addresses can and do move around the world to wherever an organisation decides is the most appropriate use of those resources. Those criteria can change rapidly too.

The study gathered data from RIPE NCC, the Regional Internet Registry for Europe and the Middle East. In addition to its role in allocating IP addresses and other numbering resources, RIPE NCC is a neutral and trusted source of tools and statistics on various aspects of Internet addressing and routing. Samples from findings were cross-checked against secondary data sources. This was used for quality assurance and to verify that independently gathered information was in line with the findings produced from the RIPE NCC's data sources.

RIR-ALLOCATED UK IPv4 ADDRESS SPACE

The study established that UK-based organisations hold 70.3 million IPv4 addresses that have been allocated by RIPE NCC. Of that 70.3 million, some 53.8 million (77% of them) are estimated to be “in use” and 16.5 million are considered “free” or available for assignment. This estimate is based on RIPE NCC's criteria for IP address assignment and allocation. Other estimates could have been used — for example by scanning large ranges of addresses — but these were rejected for technical and practical reasons. Ofcom wanted to use a stable baseline so that further data gathering exercises were not complicated by the introduction of new variables or changes to the metrics used on previous ones.

The number of IPv4 addresses issued to UK organisations has generally increased by around 4 million year-on-year. Though there have been exceptions where the annual increase was as low as 2 million or as much as 10 million. The increase in allocated address space stopped in 2013 after a new address policy had been adopted. This was a consequence of the IPv4 run-out. The old policy which issued addresses on the basis of need was replaced with one which restricts organisations to a single allocation of what is known as a /22: 1,024 IPv4 addresses. This new policy took effect on September 14th 2012.

At present, UK address holders are estimated to have in total 16.5 million IPv4 addresses available for assignment, excluding any legacy resources. There are just too many unknowns and potential variables to make predictions on actual run-out rates with any reasonable degree of certainty. The IPv4 assignment rate in the UK has averaged 3.4 million addresses per year since 2000. If this rate continues and no other factors change, that suggests the UK could in theory be able to continue assigning IPv4 addresses for about 5 years or so from the time of writing. The actual crunch points will vary. Some organisations will exhaust their remaining IPv4 sooner than then, others may be able to hold out for longer.

The actual end-point for IPv4 could be much further away. Allocation and assignment rates have reduced dramatically since the new address policy was adopted: only 500,000 addresses were assigned in 2013 and under 200,000 for 2014 to date. If these rates continue, all else being equal, current UK IPv4 address space might even last for 25-30 years.

Both these scenarios exclude the possibility that the UK’s legacy resources of IPv4 could be made available. It would be unwise to make long-term plans or assumptions based on these estimates because address usage rates will vary from time to time and shortages of IPv4 address space for key ISPs could have significant market impacts.

Just under 80% of the IPv4 address space issued by RIPE NCC to UK organisations, is routed on the public Internet. This may overestimate the amount of address space that is actually active because ISPs generally announce routes for an entire allocation of addresses from RIPE NCC even if some have still to be assigned to business units or end users. This percentage has remained reasonably static in recent years. In 2003, 66% of UK address space was being routed on the Internet. This gradually increased by one or two percentage points until it reached 77% in 2011 where it has remained ever since.

LEGACY UK IPv4 ADDRESS SPACE

In addition to the address space allocated by RIPE NCC, a further 53.1 million IPv4 addresses were issued to legacy holders. These are generally early adopters of IP-based networking who acquired IP addresses prior to the introduction of the RIR system at the end of 1997. Legacy resources are not subject to RIR address policy and are not obliged to use assignments in the way that RIR-allocated space does. 63% of these UK legacy addresses are held by two government departments, 16.7 million each by the Ministry of Defence and the Department of Work and Pensions. These two address blocks have never been routed on the Internet.

A further 6.3 million legacy IPv4 addresses were allocated to UK academia. Almost all of these are routed on the public Internet. This accounts for 12% of the 20% of the UK legacy address space that is routed on the Internet. The remaining 8% of routed legacy space mostly comes from corporate networks.

Apart from the MoD, DWP and academia space mentioned above, 20 million IPv4 legacy addresses have been allocated to UK-based organisations. Most of that space is not routed on the public Internet and is presumably used on private, internal networks. The proportion of routed legacy space has remained essentially static since 2004 at 20% of the total UK legacy address space: i.e. around 10 million of the UK's 53 million legacy IP addresses are routed on the Internet.

OVERALL UK UPv4 ADDRESSING

UK-based organisations hold a total of 123.4 million IPv4 addresses, 70.3 million in the RIR system and 53.1 million as legacy space. It is difficult to assess how much legacy space is used. Only 20% of that is routed on the Internet. Assuming the two blocks of 16.7 million held by the Ministry of Defence and the Department of Work and Pensions are just 1% utilised then amount of unused legacy space in the UK is therefore likely to be at least 33 million addresses. An estimated 16.4 million IPv4 addresses from RIR-allocated space are thought to be available for assignment. This means there should be a minimum of 50 million unused IPv4 addresses held by UK organisations. However it does not follow that all of these addresses can be made available. It should also be noted that intermixing RIR-issued and legacy address resources is unwise. These are discrete commodities subject to different policies and business considerations.

ADDRESS TRANSFERS

A more liberal address transfer policy has been in effect because of the IPv4 run-out to help redistribute address space. To date, this has had little impact on UK address holders. Total transfer activity is less than 1% of overall UK RIR-issued address space and most transfers have been for blocks of 4,096 addresses or fewer. This broadly matches the pattern of behaviour across the RIPE service region.

RIPE NCC offers a transfer listing service for its members. It shows demand for address space is greater than supply.

IPv6 ADDRESSING IN THE UK

Uptake of IPv6 in the UK remains lower than in many similar economies. Only 74% of the UK membership of RIPE NCC have IPv6 allocations and less than half of these are announcing routes for their IPv6 address space. Overall, only 28% of UK LIRs are offering direct IPv6 connectivity to the Internet. Although there has been a growth in the amount of IPv6 address space allocated to UK organisations, only 0.4% of that IPv6 space appears to be "in use".

RIPE NCC's RIPENESS measurements show that the UK lags behind the rest of the RIPE service region. Perhaps the most thorough measurement of IPv6 usage is performed by APNIC Labs, part of the RIR for the Asia-Pacific region. Their analysis shows the UK at 40th place in a list of 175 countries for which they have statistically significant IPv6 usage data.

Routing visibility of UK IPv6 space has diminished from 94% in 2010 to 71% this year. The most likely explanation will be that more UK organisations have recently acquired IPv6 space because of the IPv4 run-out. These organisations have yet to make use of their recent IPv6 allocations. It should also be noted that reachability as measured by routing announcements does not reflect actual usage. Route announcements are generally made for an entire allocation even when very little of that address space is active.

The recent ICC study for Ofcom suggests that large UK ISPs and CSPs are considering the deployment of Carrier Grade NAT (CGN) as an intermediate step to IPv6 deployment and that this may be one of the reasons why IPv6 utilisation is currently relatively low in the UK.

INTRODUCTION

The Internet is facing a major challenge. It is running out of IPv4 address space. Every device on the Internet needs to have at least one unique IP address. Although users identify resources on the Internet by a name (e.g. `www.ofcom.org.uk`), the underlying network relies on unique IP addresses (194.33.160.25 in the case of Ofcom's web site) to route data between devices that are connected to the Internet.

The current addressing scheme in use today (Internet Protocol version 4 or IPv4) was invented and deployed over 30 years ago. This uses 32-bit addresses, providing approximately 4 billion IP addresses. At the time IPv4 was introduced, it was thought the address space would be far, far larger than could ever be utilised. However the Internet has vastly outstripped those early expectations and become a ubiquitous, indispensable resource for substantial numbers of people, businesses and applications. As a result, almost all of this address space has been allocated and the remaining capacity is not enough to allow for future growth of the Internet. Put simply 4 billion addresses is far too small for a global population of over 7 billion and growing. Furthermore consumer demand to connect ever-increasing numbers of devices — smart phones, tablet computers, games consoles, home computers, entertainment systems, etc. — to the Internet coupled with increasing amounts of content and new services being delivered over the Internet will exacerbate the demand for IP addresses.

The current shortage of IPv4 addresses and the unavoidable truth that the IPv4 address space is too small are likely to present barriers for the deployment and availability of new services and technologies on the Internet.
Clearly, the long-term solution to this problem will be a global migration to a new version of the core Internet protocol, IPv6: Internet Protocol version 6. [IPv5 was a short-lived and largely forgotten experiment 25 years ago.] IPv6 uses 128-bit addresses, creating an unimaginably large address space. This allows for approximately 340 trillion, trillion, trillion (3.4×10^{38}) unique addresses. Even if the Internet continues its explosive growth of recent years, it is highly unlikely the IPv6 address space could be exhausted in the foreseeable future. IPv6 will also allow Internet capabilities to be embraced by new technologies such as the use of RFID tags, contactless payment systems, smart metering, intelligent utility grids, sensor networks and the Internet of Things.

However, the deployment and use of IPv6 has been slower in the UK when compared with similar economies. Uptake of IPv6 in the UK is particularly slow because the UK's biggest Communications and Internet Service Providers seem to favour deployment of Carrier Grade NAT (CGN) as an intermediate step ahead of a deployment of IPv6.[1] Mobile operators and some fixed-line ISPs are conducting trials of GCN. One of the largest DSL providers in the UK is planning a limited deployment of CGN.
Availability of IPv4 address space has become increasingly scarce in recent years. IANA, the master address registry, exhausted its supply of IPv4 addresses in February 2011 when it distributed its last remaining address blocks (known as /8s) to the 5 Regional Internet Registries (RIRs). On 14 September 2012 RIPE NCC, the RIR for Europe and the

[1] http://stakeholders.ofcom.org.uk/binaries/research/technology- research/2013/cgnat.pdf

Middle East, began to allocate IPv4 address space from the final /8 of IPv4 address space — approximately 16 million addresses — it obtained from IANA.
Address allocation policies have tightened at each RIR as IPv4 address space began to run out. Historically, these policies distributed addresses to organisations on the basis of need. An organisation would supply the RIR with an addressing plan justifying a certain amount of address space and provided this looked reasonable, the RIR would make a suitably sized allocation. Since RIPE NCC began allocating space from its last /8, a much stricter regime has applied. RIPE NCC members can now only request a single allocation of a /22 of IPv4 space (1,024 addresses) and will only receive that allocation if they already have an IPv6 allocation.[2]

A market in IPv4 address trading has emerged and RIR policies on address transfers have loosened. RIPE NCC reports that 631 IPv4 address blocks have been transferred since 2012. These transfers appear to be the result of actual sales of address space rather than routine business activity such as reorganisations, mergers, acquisitions and spin-offs which happen from time to time. It is also unclear if the size of these transferred blocks is significant. One aspect of this Ofcom study is to identify the nature of these transfers and their potential to affect the overall distribution and usage of IPv4 addresses. The shortage of IPv4 addresses might be mitigated by trading or transfer mechanisms.

In addition, there are blocks of IPv4 addresses which were allocated to organisations before the RIR system was introduced at the end of 1997. These are known as legacy address resources and legacy resource holders respectively. In general, legacy resource holders were very early adopters of IP-based networking and are typically drawn from the academic or research community. These address resources are not subject to RIR policies.

Some of these legacy IPv4 address allocations are very significant. For instance two UK government departments each have a /8 (16.7 million addresses) of IPv4 space. Although it is improbable these legacy /8s are fully used it is unclear if some or all of that space could be made available for public use. It is not widely understood how well or poorly legacy allocations are utilised or even if they are reachable on the public Internet. It is also unclear just how much legacy space is held by UK-based organisations, who holds that space or how much they hold.

With this background, Ofcom issued a Tender to gather information on the use of IPv4 and IPv6 address space by UK-based organisations. Ofcom wished to understand the current status of IPv6 and IPv4 address blocks allocated to UK address holders and their respective visibility within the global Internet routing tables. The scope of this study included UK-held IP address resources obtained through the RIR system and those which were distributed to legacy resource holders prior to the introduction of the RIR system. This report and analysis is the outcome of that study.

[2] http://www.ripe.net/internet-coordination/ipv4-exhaustion

OVERVIEW OF IP ADDRESSING

INTERNET ADDRESS FUNDAMENTALS

Every device connected to the Internet has a unique number, its IP address. This is used to identify the source and destination of data crossing the network. Any activity on the Internet —say for a web session or video streaming — involves splitting the data into packets which are then routed between two IP addresses: for instance from a web browser so some web server and back again. It therefore follows that the end-points for this activity must have unique IP addresses to identify them to ensure that traffic gets directed to the correct location. The Internet's Domain Name System (DNS) maps the memorable strings that people use to identify Internet resources (e.g. `www.ofcom.org.uk`) into the unique IP address for that resource so the underlying network can route data to and from that location.

The vast majority of devices on the Internet today use IPv4, Internet Protocol version 4, which was standardised in 1981[3]. It has a 32-bit address space, allowing for $2{\times}2^{32}$ unique IPv4 addresses: just over 4 billion. At the time IPv4 was invented, the IPv4 address space was expected to be far, far bigger than could ever be used. In hindsight, that assumption turned out to be wrong. As uptake of the Internet increased, it was clear a new protocol with a larger address space would be needed.

That protocol was developed in the early 1990s: Internet Protocol version 6 (IPv6). The first specification for IPv6 was RFC1883[4] which was published in December 1995. IPv6 uses 128-bit IP addresses. $2{\times}2^{128}$ is a huge number. It represents an unimaginably large address space: approximately 340 trillion, trillion, trillion ($3.4{\times}10^{38}$) unique addresses. Even if the Internet continues its explosive growth of recent years, it is highly unlikely the IPv6 address space could be exhausted in the foreseeable future. Although IPv6 is mature and widely supported, deployment has been slow. This is expected to change now that IPv4 resources approach exhaustion.

INTERNET ADDRESS NOTATION

Every packet on the network includes a header which contains the source and destination addresses as 32- or 128-bit values. This makes it easy for computers and routers to handle them. However these are not easily understood by anything else.

IPv4 addresses are usually written in dotted-decimal notation: four numbers separated by dots. The first of those numbers represents the first 8 bits of the address, the second represents the next 8 bits of the address and so on: e.g. 194.33.160.25. Each of those numbers is written in decimal and can only represent values between 0 (8 bits of all zeros) and 255 (8 bits of all ones).

---

[3] http://tools.ietf.org/html/rfc791
[4] http://tools.ietf.org/html/rfc1883

A colon notation is generally used for presenting IPv6 addresses. These take the form *H*:*H*:*H*:*H*:*H*:*H*:*H*:*H*, where each *H* represents the hexadecimal value of the eight successive 16-bit pieces of the address: for example 2001:7FE:0000:0000:0000:0000:0000:53. The first 16 bits of the IPv6 address are 2001 and the next 16 are 07FE. Then there are 80 bits which are zero and the final 16 bits of the address are 0053. Whenever there are groups of 16-bit values of zeros in an IPv6 address, these can be represented by the use of a double colon: "::". A double colon can only appear once in the string representing a given IPv6 address. This would allow the IPv6 address above to be written as 2001:7FE::53.

PREFIX NOTATION

A standardised notation for representing address prefixes first appeared in RFC1519[5] and later revised in RFC4632[6] which was published in 2006.
The prefix part of the address is written in the usual dotted-decimal or colon notation, followed by a slash character and then a number representing the number of bits in the top-most part of the address. This is known as a prefix. For example, the historic Class A network 10.0.0.0 could be represented as 10.0.0.0/8 or perhaps just 10/8 to indicate that the top-most 8 bits of the address share the same prefix, the value 10 expressed in decimal (00001010 in binary). In the case of IPv6, 2001:DB8::/32 and 2001:DB8/32 represent a prefix where the top-most 32 bits have the value 2001:0DB8 in hexadecimal. In the context of IP allocation and allocation polices or route filtering, it is customary to use terms like a /*N* to denote a specific prefix length — and by implication the size of some address block — when the value of the actual prefix itself is not relevant. Thus, the master IP registry IANA issued /8s of IPv4 space to each of the RIRs. The RIRs then divided these /8s into longer prefixes (say /20s) and allocated these to their members. For IPv6, IANA issues a /12 of IPv6 space to each RIR and they in turn allocate a minimum of a /32 from that /12 to each of their members.

VARIABLE LENGTH PREFIXES AND SUBNET MASKS

The concepts of variable length prefixes and subnet masks were first described in RFC1519 which was published in 1993. This standard was superseded by RFC4632. This technique is sometimes known as CIDR — Classless Inter-Domain Routing — though the term is not widely used nowadays. There were two key reasons for using variable-length subnet masks: concerns about the growth of the global routing table and the prospect of the exhaustion of the IPv4 address space. It is remarkable that these two concerns are still significant over 20 years later!

When IP was first introduced, the divisions between the network and host part of an address were aligned on 8-bit boundaries. Class A networks had an 8-bit prefix and a 24-bit host part. Class B networks used a 16-bit prefix and 16-bit host part. Class C networks had 24-bit prefixes and 8-bit host parts. The distinction between Class A, B and C addresses was hard-coded in software. When the top-most bit of an IPv4 address was zero, it meant the address was for a Class A network and the network prefix was a /8. When the top two bits were 10 in binary, that meant a Class B network with 16-bit network and host parts of the address. For a Class C network, the top 3 bits of the address were

[5] http://tools.ietf.org/html/rfc1519
[6] http://tools.ietf.org/html/rfc4632

110 and the network prefix was a /24. This arrangement was somewhat clumsy. It was also wasteful.

Few organisations were likely to ever have a network with millions of hosts requiring an allocation of a Class A network. Even so, earliest adopters of IP-based networking were allocated these networks: IBM, Xerox, HP, Apple, Ford Motor Company, AT&T and some US universities. MIT is still using the Class A network 18/8 today. A university campus that might hope to have a few thousand hosts in the early 1990s would get allocated a Class B network big enough for 65,536 addresses even if it had no realistic expectation at that time of using more than 10% of the allocated space. Similarly, a small business with 2 or 3 computers would have been allocated a Class C network of 256 addresses.

Once variable-length subnet masks were introduced, these constraints were removed. The prefix — i.e. the boundary between the host and network part of the address — could be placed anywhere, not necessarily on an 8-bit boundary. That meant less wastage of IP address space. For instance, a small office with 10 hosts could be given a /28 as its network prefix which is enough for up to 16 addresses instead of a Class C network that had space for 256. Likewise, a campus-sized network of 10,000 hosts might get a /18 (up to 16,000 addresses) instead of a Class B allocation that used 65,536 addresses.

The second aspect of CIDR was route aggregation. It was no longer necessary to announce routes using prefixes aligned on 8-bit boundaries in an IP address. When many networks had a common prefix of whatever length, a single routing announcement could be made for that common prefix instead of discrete ones for each Class A, B or C network. That reduced the size of the global routing table which was a far more serious problem in the 1990s when computer memory cost several hundred pounds per megabyte and the most powerful microprocessors might only be rated at 5-10 million instructions per second.

Suppose an organisation had been allocated 8 contiguous /8s: 80/8, 81/8… and 87/8. Prior to the introduction of CIDR, there would have to be 8 entries for these in the global routing table, one for each /8. However with CIDR, there would be just one route announcement and router table entry for 80/5 because the top 5 bits of these 8 networks have the same value. This aggregation was particularly valuable to the first ISPs who were issuing small blocks of IP addresses to their customers. They had neither the equipment nor inclination to make discrete routing announcements for every prefix issued to each customer. The ISP could just make a single route announcement covering all of its address space, further reducing the pressure on the global routing table.
This practice continues today. An ISP will almost certainly announce a routing prefix for each allocation they get from an RIR even when not all of that space is active.

INTERNET ADDRESS ROUTING AND PREFIXES

IP addresses are notionally split into two, a network (or subnet) part and a host part. In simple terms, the top-most bits of the address are a prefix which generally identifies a physical network: a section of Ethernet perhaps or a distinct wireless network. The low-end parts of the address identify a specific host on that physical network. When all the bits of the host part of the address are zero, that represents the address of the network itself. When all the bits of the host part are one, that is the broadcast address for the network. All the devices connected to that network will receive traffic for this broadcast address as well as that for the actual address of the device itself.
Address prefixes are sometimes called subnet masks because they enable routers to perform bit-string masking of an IP address to split it into a host part and a network part. In

simple terms, routers use the network part of an address to decide where to send each packet. The final router in the chain then uses the host part of the IP address to direct the packet to the correct destination on that network.

Routers are configured to announce the prefixes for the networks they are connected to. They exchange this routing information with adjacent routers which in turn pass on that information to the routers they are adjacent to and so on. Routing protocols such as BGP, the Border Gateway Protocol, use these announcements to compute the optimal path to use when sending traffic for each prefix. This information changes frequently based on a large number of continually changing metrics: the underlying topologies of the physical networks and links; the addition/removal or reconfiguration of routers; operational problems such as physical faults or power failures; changes to routing policies by Internet Service Providers (ISPs) and/or at Internet Exchange Points (IXPs); congestion; traffic shaping and engineering; availability of bandwidth; contractual/legal considerations and so on. While this continually and rapidly moving target makes Internet routing very complex, it is also one of the Internet's greatest strengths. Routers are always using BGP to make adjustments to how traffic flows across the network based on how the underlying network of networks are behaving and what routing announcements are being made. This makes the Internet remarkably resilient because it automatically adapts to changed circumstances such as cable breakages or congested data links.

On August 12th 2014, the number of global routing table for IPv4 exceeded 512,000 for the first time: i.e. there were over that number of discrete routes (prefixes) being announced. While this event did not present a significant problem overall, it did create difficulties for some ISPs who had old equipment and software which could not handle so many routes. It is not necessary for every router to maintain a complete copy of the global routing table. However core routers must do this. By definition, these have to know how to route traffic for all known prefixes that are presently in use on the Internet. Few, if any, of those core routers were exposed to this problem since this 512,000 limit had been known about for some time and plans were made to upgrade them long before this threshold was reached.

HISTORY AND OVERVIEW OF IP ADDRESS POLICIES

LEGACY USAGE

When IP networking was invented, its earliest adopters were generally researchers, academics and contractors working on the ARPAnet. One of those contractors oversaw the management of the network and co-ordinated the distribution of IP addresses to the organisations which participated.

The concept of Class A, B and C addresses was used. If an organisation expected to need fewer than 256 addresses, it would probably be allocated a Class C address block, a /24 in today's jargon. An organisation expecting to have more than 256 but less than 65,536 addresses, would usually be assigned a Class B block. Organisations needing more than 65,536 addresses got a /8: a Class A address block.
This was thought to be "good enough" at the time though it was somewhat wasteful of what was then considered to be an abundant address space. Several organisations were allocated Class A addresses even though they had no expectation of ever using the majority of the 16 million addresses in that prefix. With the benefit of hindsight today, this initial scheme was unfortunate. Though in fairness to the engineers of the time, nobody really envisaged that a fledgling IP network would become the Internet as it is known and experienced today. It simply didn't matter 30 years ago or more if a major university like MIT was allocated a Class A address block. The then network consisted of a small number of organisations and there were a variety of competing network protocols and systems in use around the world. IP based networks were just another component in that protocol Tower of Babel along with SNA, X.25, DECnet and many others.

INITIAL MANAGEMENT OF INTERNET NUMBERING RESOURCES

The administrative task of allocating addresses at the outset of what became the Internet was handled by one individual, Jon Postel, on a volunteer basis. When that task grew too large, IANA, the Internet Assigned Names and Numbers Authority, was established to manage Internet resource registration and various co-ordination activities. Part of IANA's responsibilities included an Internet Registry function to oversee registration of IP address space. This Internet Registry function was first handled by the Information Sciences Institute at the University of Southern California (USC/ISI) and was finally managed by Network Solutions. That arrangement was terminated with the introduction of the RIR system. The IANA function was ultimately taken over by ICANN where it continues to be operated under a zero-dollar contract with the US Department of Commerce. IANA became the "master" registry for IP addressing when the RIR system was introduced.
As the network continued to grow, the model of central co-ordination by a contractor funded by the US government became unsustainable. Organisations were using IP-based networking even if they were not directly connected to the ARPAnet. They needed to get globally unique IP addresses. The nature of the ARPAnet was also changing as it was no longer limited to organisations working on ARPA-funded contracts. The US National Science Foundation set up a national IP-based backbone network, NSFnet, so that its grant-holders could be interconnected to supercomputer centres, universities and various national/regional academic/research networks, including ARPAnet. That resulting network of networks was the beginning of today's Internet.

INTRODUCTION OF THE RIR SYSTEM

The Regional Internet Registry (RIR) system was devised to take over the management of numbering resources for the emerging Internet. This was also the beginning of the governance model that is generally found in Internet organisations: self-regulating, multi-stakeholder participation, open bottom-up policy making with minimal barriers to participation.

ARIN, the American Registry for Internet Numbers, was founded in 1997 and started operations after taking over responsibility for the Internet Registry function previously handled by Network Solutions. Initially ARIN had global scope. Some of its functions were later devolved to other RIRs: first to RIPE NCC mainly serving Europe and the Middle East; then APNIC for Asia-Pacific; then LACNIC for Latin America and finally AFRINIC for Africa. Each of these RIRs is a non-profit membership association whose members develop IP addressing polices specific to that RIR. These policies are generally aligned or very similar across the RIRs.

Although RIPE NCC began operations in 1992 it did not formally exist until it was incorporated under Dutch law and became an RIR in 1997[7]. It started out as a RARE project and co-ordinated IP-based network activity in Europe. (RARE was an association of European national research and education networks, now known as TERENA[8].) In the early 1990s, organisations could either obtain IP addresses from IANA directly or get an allocation from one of the /8s that IANA had issued to the RIPE NCC. Space allocated via that latter route is considered “RIR space” and subject to RIR policies even though it would have been allocated before the introduction of the RIR system.

THE RIR-LIR RELATIONSHIP

A member of an RIR is known as a Local Internet Registry, LIR. It enters into a contractual arrangement with the RIR. The LIR pays fees in exchange for services from the RIR which include the allocation of IP numbering resources. LIRs are expected to become members of the RIR for their service region. Some organisations are members of more than one RIR — for example when they have a global presence and need numbering resources from each of the service regions where they operate.
In the case of RIPE NCC, LIRs elect the organisation’s board and vote each year on the annual activity plan. They also decide the fee structure.
Policy-making on Internet addressing matters is not controlled by the RIPE NCC membership however. This is determined by RIPE, a forum that is open to anyone. It is not necessary to be a member of the RIPE NCC or indeed any RIR to take part in RIPE policy making, attend its meetings or subscribe to RIPE’s (electronic) mailing lists. That said, the bulk of the participation at RIPE comes from individuals representing LIRs who are members of the RIPE NCC or other RIRs. Although many people use the terms RIPE and RIPE NCC interchangeably, these are in fact different things.

The introduction of the RIR system coincided with the deployment of variable-length subnet masks and CIDR. It was no longer necessary to allocate IP space aligned on 8-bit boundaries. Allocations could be issued that made more efficient use of the address space. This is one of the key reasons why IPv4 addressing has been able to keep pace with the rapid and exponential uptake of the Internet for so long.

---

[7] ftp://ftp.ripe.net/ripe/docs/ripe-164.ps
[8] http://www.terena.org

Addressing policies followed the same original principle: needs-based. Whenever an LIR presented a justification for needing *X* IP addresses, the RIR would allocate them an address block which would have an address space of at least *X*. This would usually be rounded up to the next convenient power of 2 to simplify both the management of IP addresses and the announcement of routing prefixes. In concept an organisation that needed 2,500 addresses might get allocated a /20 by the RIR, allowing the organisation to use up to 4,096 addresses. That allocation would provide enough “headroom” for the LIR to grow into the remaining 1,500 or so addresses in that /20.

RIR IP ADDRESS POLICY IN OUTLINE

The key concepts for IP addressing policy are allocation and assignment. In outline, an RIR issues an allocation to a member, a Local Internet Registry (LIR). The LIR then makes assignments from that allocation to its customers or for its internal use. When most of the address space in an allocation has been assigned, an LIR can apply for another allocation. Small LIRs are usually expected to get approval from the RIR before making assignments. Larger LIRs tend to be able to make assignments without first getting approval from the RIR.

When an LIR needed IPv4 space, it submitted a template containing an addressing plan. This outlined how the addresses were expected to be used and how that was projected to change over time. The LIR would state how many addresses were needed immediately and how many were expected to be used 1 and 2 years later. This information would be checked by an IP resource analyst at the RIR and then used to determine the size of the allocation and, if appropriate, decide an initial assignment. These address templates are confidential because they usually contain business-sensitive information.
If the LIR already had an allocation or allocations, they would be required to show that their existing space had been substantially used before the RIR issued them with another allocation. This reduced the potential for address hoarding and encouraged LIRs to make efficient use of the space they were allocated.

For IPv4, RIR-allocated addresses fit into two categories. The first of these is Provider Aggregatable space (PA space). Most IP addresses are PA space. The LIR is usually able to announce a single routing prefix for each block of PA space. How that space is divided up internally does not matter to those outside that network even if it is shared out amongst many customers or business units. The second is Provider Independent space (PI space). The main rationale for that was to allow organisations to change ISP without having to renumber their network. They could take their PI space to another ISP and have them announce a route for that space. There is effectively no meaningful difference between PI and PA space for the purpose of this study.

Nowadays there is less pressure for route aggregation than there was when PI space was devised. Most ISPs are prepared to accept or announce routes for /24 prefixes. Organisations are able to take small blocks of PA space and have that routed independently from whatever the original LIR announces as a prefix for that (aggregated) allocation.

A third address category — LEGACY — is used by the RIRs to identify IPv4 resources which were issued before the introduction of the RIR system. The RIRs maintain database entries for these so that other parts of the Internet’s addressing infrastructure are able to deal with these resources. The DNS and whois servers that provide public lookup services for all IP resources are populated with information from the RIR database.

IPv4 ADDRESS POLICY IN THE RIPE SERVICE REGION

The principle of needs-based address allocation for IPv4 addresses no longer applies in the RIPE service region. It had to be abandoned as a result of the IPv4 run-out. Since September 14th 2012 RIPE NCC has been allocating space from the last /8 it received from IANA[9] and a new address distribution policy is in force. This is explained in Section 5.1 of "*IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region*"[10], the current policy that has been adopted by RIPE. The other RIRs have similar policies. However the point at which these get put into effect changes because the trigger is when each RIR begins to allocate space from its last /8 from IANA.

Today, an LIR in the RIPE service region can receive only one /22 allocation of IPv4 space (1,024 IPv4 addresses), even if they can demonstrate a need which would otherwise justify a larger allocation. This /22 allocation will only be made to the LIR if it has already received an IPv6 allocation from an upstream LIR or from the RIPE NCC. In other words, an LIR can only get its final /22 of IPv4 if it has an IPv6 allocation. The main objective of this policy is to conserve the remaining IPv4 space for a long time so that new entrants will still be able to get some IPv4 space when connecting to what should then be a mostly IPv6-based Internet. The new policy also reduces the likelihood of existing LIRs "panic buying" the final stock of IPv4 address space at the RIPE NCC. Having IPv6 as a prerequisite for obtaining this final allocation of a /22 of IPv4 is intended to encourage uptake and deployment of IPv6.

It is anticipated that this current address policy should ensure small amounts of IPv4 space remain available to new entrants for several years. RIPE NCC has around 10,000 members and if each one obtains a final /22, that would account for approximately 10 million of the 16 million IPv4 addresses in that last /8. RIPE NCC report approximately 40% of its membership have already obtained their final /22 allocation of IPv4.[11]

---

[9] http://www.ripe.net/internet-coordination/news/announcements/ripe-ncc-begins-to-allocate-ipv4-address-space-from-the-last-8

[10] http://www.ripe.net/ripe/docs/ripe-622

[11] https://www.ripe.net/ripe/meetings/roundtable/october-2014/presentations/iana-stewardship-and-the-ripe-community/view

IPv6 ADDRESS POLICY IN THE RIPE SERVICE REGION

The current IPv6 policy for the RIPE service region is described in RIPE document 589[12]: "*IPv6 Address Allocation and Assignment Policy*". Although the principle of needs-based allocation still applies for IPv6 numbering resources, the address space is so vast that the concept of need is essentially meaningless. In the current IPv6 policy, the minimum IPv6 allocation to an LIR is a /32: 8×10^{27} addresses. This is 18 billion, billion times the size of the entire IPv4 space. Some LIRs in the RIPE NCC service region have received far larger IPv6 allocations than that. BT received a /22 in 2007 and in 2005 France Telecom and Deutsche Telekom both obtained a /19. These companies demonstrated valid use cases to justify getting so much extra space.

Current thinking is each LIR customer would be assigned a minimum of a /48, yet another unimaginably large amount of address space — 120 trillion, trillion addresses — that could never be fully utilised. This is likely be the amount of IPv6 space each domestic DSL or cable customer can expect to receive once their ISP deploys IPv6.

When an organisation holding an IPv6 address allocation makes IPv6 address assignments, it must register these assignments in the NCC database. These registrations can either be made as individual assignments or by inserting an object with a status value of 'AGGREGATED-BY-LIR' where the assignment-size attribute contains the size of the individual assignments made to end users. The AGGREGATED-BY-LIR status is for administrative convenience. The LIR and RIR do not need to manage discrete database entries for every /48 that is issued to each customer or end-user. When more than a /48 is assigned to an organisation, it must be registered in the database as a separate object with status 'ASSIGNED'.

[12] http://www.ripe.net/ripe/docs/ripe-589

DATA GATHERING & PROCESSING

DEFINING AND IDENTIFYING UK ADDRESS SPACE

One of the first challenges for this project was defining and identifying UK IP addresses. This is not as straightforward as may be first thought. The RIRs publish data which shows the “nationality” of address allocations. In principle this should provide data about IP addresses allocated to UK address holders. However it does not follow that some of all of those IP addresses are actually used in the UK. Or if they were used in the UK today, they would still be used there tomorrow. Similarly, there will be non-UK organisations who have a network footprint in the UK: global connectivity and content providers for example. In addition, some UK organisations use non-UK LIRs. BT for instance uses a Reg-ID `eu.bt` **f**or part of its operations. Geographic considerations are further muddied by multinationals. IP address space allocated to British Airways (say) might notionally be viewed as UK even though parts of that space are used for the company’s corporate network in North America.

The dynamic nature of the Internet means that is effectively impossible to get definitive statistics on IP address usage or reachability in the UK. Routing announcements and IP address distribution are in a continual state of flux. Routing policies and network topologies change frequently. IP addresses move around the world to meet the immediate (internal) needs of address holders. So by the time an exhaustive survey of UK IP space completed — assuming this was feasible — the results would inevitably be out of date. Therefore any data gathering and analysis can only be a reasonable snapshot and approximation of what is actually happening. In short, the UK’s address space is a continually changing and rapidly moving target.

DATA SOURCES

The RIPE NCC publishes data which help identify UK-held IP resources. One data source is the `delegated-ripencc-extended-`*`YYYYMMDD`* file (DE File) that gets published every day.[13] The format of these plain-text files is fairly obvious: the name of the RIR, the ISO-3166 two-letter country code for the address holder (“GB” for UK organisations), the address prefix and number of addresses, date the allocation was last updated and whether the space has been allocated or is available for allocation. Figure 1 shows a sample entry from this file.

ripencc|GB|ipv4|2.24.0.0|524288|20100921|allocated

Figure 1 - extract from delegated-ripencc-extended-YYYYMMDD file

A second data source is a URL which shows countries where RIPE NCC members offer their services[14] and the UK in particular.[15] This second page contains the name of each

---

[13] ftp://ftp.ripe.net/pub/stats/ripencc
[14] https://www.ripe.net/membership/indices/
[15] https://www.ripe.net/membership/indices/GB.html

LIR active in the country and a hyperlink to a page containing its contact details. It also contains the LIR's Reg-ID, a useful identifier for lookups in the RIPE NCC database. These pages are best harvested with a web scraper. Unfortunately this second source excludes EU-based LIRs and it is unclear why RIPE NCC does not provide this facility for them.

A further source of data is the `alloclist.txt` file[16] which is updated and published on a daily basis by the RIPE NCC. This plain text file lists the names and Reg-IDs of each LIR along with some details of each LIR's allocations. However the Reg-IDs use "UK" as their prefix even though their IP space in the DE File is tagged as "GB". This inconsistency is an irritant. The file only holds data about address space allocated by the NCC. It does not have details of legacy space. An excerpt from this file is shown in Figure 2.

```
uk.100percentit
        100 Percent IT

        20001013    217.25.0.0/20   ALLOCATED PA
        20051011    87.254.0.0/19   ALLOCATED PA
        20120501    2a00:eac0::/32
```

Figure 2 - extract from RIPE NCC `alloclist.txt` file

The above example shows that the Reg-ID `uk.100percentit` was issued to a UK-based organisation called 100 Percent IT. The LIR has two IPv4 allocations, 217.25.0.0/20 issued on October 13th 2000 and 87.254.0.0/19 which was allocated on 11th October 2005. It received an allocation of IPv6, 2a00:eac0::/32, on May 1st 2012.

After discussion with Ofcom, it was agreed the entries in the daily DE File that were listed as "GB" resources would be a satisfactory approximation to the overall address space used in the UK. Any space held by UK organisations but used elsewhere would probably be balanced out by address space held by non-UK LIRs and used in the UK. This pragmatic approach was a reasonable compromise given that it is unrealistic to get a definitive, precise and timely assessment of IP address usage in the UK.

In addition to the GB entries in the DE file, Ofcom asked for other potentially relevant LIRs to be included. This meant manually checking the above data sources for any EU Reg-IDs belonging to the main UK Communication Service Providers. Two were identified, `eu.bt` for British Telecom and `eu.en` for Cable & Wireless. The address resources for these two LIRs were included in this analysis.

Identifying potentially relevant Org-IDs and Reg-IDs to use for database lookups was challenging. Establishing which UK-based organisations use which Org-IDs and Reg-IDs is not straightforward. Ofcom asked for special attention on the resources held by the main UK Communication Service Providers: BT, BSkyB, TalkTalk, Virgin Media, KCom, O2, EE, Vodafone, C&W and Three. Finding all the numbering resources for these organisations by automated means is impractical because the relevant database objects for these organisations are somewhat unstructured. For instance EE uses Reg-IDs and LIRs issued to T-Mobile and Orange even though those companies merged their UK business some years ago. Talk Talk uses Reg-IDs, Org-IDs and LIRs that were issued to businesses such as Tiscali, Opal Solutions and Opal Telecom which have been acquired by Talk Talk. British Telecom and Cable & Wireless have `eu.*` Reg-IDs in addition to their `uk.*` ones.

---

[16] ftp://ftp.ripe.net/ripe/stats/membership/alloclist.txt

A list of the relevant LIRs and their allocations had to be generated by manually checking the RIPE NCC's `db.organisation` file[17] and cross-checking LIR contact details. Developing scripts or tools to find these would have taken longer than manual checking.

RIPE DATABASE

The RIPE Database contains registration information for IP addresses and networks in the RIPE NCC service region and related contact details. It is actually a set of lookup tables rather than a conventional relational database (RDBMS) such as Oracle or MySQL. It is therefore not possible to use some form of Structured Query Language (SQL) to interrogate the database or perform conventional database operations such as a join or merge. Instead, lookups are done to retrieve objects from the database keyed on specific attribute values. For example a lookup for an organisation object for a given Org-ID will return some information about the LIR which had been assigned that Org-ID.
This means that bulk retrieval of data needs to be done using scripts to make lookups on the RIR database. Although RIPE NCC provides a web-based GUI for making database lookups,[18] it is a little clumsy and impractical for sustained data-gathering exercises which require large numbers of queries.

RIPEstat

RIPEstat[19] is a powerful and versatile tool for interrogating the RIPE database and other sources of information maintained by the RIPE NCC. It presents registration and routing data, DNS data, geographical information, abuse contacts and related information from the RIPE NCC's internal data sets as well as from external sources such as other RIRs and IANA. RIPEstat's main web-based interface presents this information in the form of widgets that can be embedded on any webpage. It also provides a JSON API (JavaScript Object Notation Application Programming Interface) to access the raw data for use in advanced applications. Possible input for RIPEstat queries include domain names, IP addresses, Autonomous System Numbers (ASNs), IP address prefixes and ISO 3166 country codes.

Initial evaluation of RIPEstat suggested it would be the key resource for this project. Scripts and tools were then developed to gather the bulk of the data needed for this report and produce the spreadsheets to analyse the results. These tools would also be part of the project deliverables.

Following consultation with Ofcom, it was agreed that RIPEstat should be the foundation of this study. First, the RIPE NCC is a trusted, impartial source of data on IP addressing and networking statistics. Second, RIPEstat was expected to provide most if not all of the IP address information that Ofcom was likely to need for this study. Third, RIPEstat is formally supported and has stable service and APIs. This means it provides a reliable baseline for any future data gathering and analysis that Ofcom chooses to do. Ofcom can be reasonably confident that the service will be available in the long term. Other possible data sources may not be as dependable because some depend on volunteer, best efforts commitments or lack the support of a well-resourced organisation which has a mission to publish freely accessible information on the Internet resources it oversees. Finally, the metrics used by RIPEstat are stable and unlikely to change. This should mean there are

---

[17] ftp://ftp.ripe.net/ripe/dbase/split/ripe.db.organisation.gz
[18] https://apps.db.ripe.net/search/query.html
[19] https://stat.ripe.net

fewer variables or parameters which could be different when Ofcom repeats this data gathering exercise in the future.
Discussions then took place with RIPE NCC staff on how to make best use of RIPEstat and the potential impact of making a large number of API calls. It was important that the project would not overwhelm RIPEstat or for its expected data mining to be subject to rate-limiting. Once these matters were resolved, scripts were developed and tested to exploit RIPEstat's JSON API and extract most of the data needed for analysis.

EXTRACTING AND ANALYSING IPv4 & IPv6 DATA

Almost all the LIRs analysed in this study have Reg-IDs of the form `uk.`*`something`*. The `alloclist.txt` file was checked by hand and identified two additional LIRs of interest: `eu.bt` and `eu.en`. and these were added to the list of UK-based LIRs for analysis. This yielded a total of 1,186 unique LIRs.
These two `eu.*` Reg-IDs refer to LIRs belonging to British Telecom and Cable & Wireless respectively. Some organisations in the RIPE NCC service region choose to have "European" Reg-IDs `eu.`*`something`* to reflect operations that are not necessarily linked to just one country or to indicate that the LIR has international scope.[20] Further manual checks of the RIPE NCC's `alloclist.txt` and `db.organisation` files did not identify any other major UK-based ISP with non-UK Reg-Ids.

A script was used to examine each of the UK LIR allocations to determine their address space usage and reachability using data obtained using the RIPEstat API. The script read the RIPE NCC `alloclist.txt` file and generated a text file with one line for each allocation. Some information was extracted directly from the `alloclist.txt` file. Other elements from this file were used by the script to make RIPEstat API calls to gather additional data. The file produced by this script has the following fields:
Information derived from the `alloclist.txt` file:
LIRcode — the Reg-ID, e.g. `uk.example`
IPv4/IPv6 — either "IPv4" or "IPv6" used to select and filter data in later analyses
Prefix — the prefix for the allocation (e.g., a.b.c.d/e or a:b:c::/d)
IP addresses — the number of addresses in the allocation (IPv4) or the number of /48 subnets (IPv6)
Date allocated — the date the allocation was made
Information returned from the `address-space-usage` RIPEstat API call:
`ip_stats` values — number of addresses (IPv4) or number of /48 subnets (IPv6)
`status` — ASSIGNED PA/PI, FREE (IPv4) or ASSIGNED, AGGREGATED BY LIR, FREE (IPv6)
Information derived from the `routing-status` RIPEstat API call:
`first_seen` — when the prefix was first visible in RIPE NCC's RIS database (blank if never seen)
`last_seen` — when the prefix was last visible in RIPE NCC's RIS database (blank if never seen)
Amounts for IPv6 resources — number of IP addresses, numbers of addresses in each IPv6 status (FREE, ASSIGNED, AGGREGATED BY LIR) — are scaled in units of /48 so that the resulting numbers can be understood.
The "free" space is the amount the RIPE NCC has calculated as the difference between the amount of space that was allocated to an LIR and how much the LIR has assigned

[20] BT and C&W are presumably using these LIRs to provide globally managed network services.

from that allocation. This is a reasonable estimate of the amount of address space that remains in an allocation for assignments to customers or use internally. However the data quality for an LIR's assignments can vary. Even when assignment data are correct, it does not necessarily mean that all the addresses in a given assignment are used.
This simplifying assumption is a trade-off between the need for accurate information and consistent results. Precise measurements of "free" address space — assuming these are viable — will introduce extra variables which may change for future data gathering exercises and analyses.
The resulting text file was further processed by spreadsheet software to produce the summary counts, tables, and charts shown in this report. That processing also generated some of the `UKLIRs` spreadsheet that is part of the deliverables.

TRANSFER DATA

RIPE NCC publishes a web page which is updated weekly. It contains details of address transfers between LIRs: what prefixes were transferred and when and who the donors and recipients were.[21] The web page provides a link to download that information as a JSON file. However this file is somewhat clumsy. It is not straightforward to process the download with common JSON tools.

Helpful information is not supplied — for instance the Reg-ID or Org-ID of the LIRs. That makes it awkward to determine the actual LIRs who were involved in the transfer or their nationality. Sometimes the names shown on the list of transfers does not quite match the organisation name of the LIR. This means subjective judgements are needed to identify any UK LIRs who have transferred IPv4 resources using the new transfer policy.
The JSON file was edited by hand to make it more amenable to scripting and editing. A further edit deleted the entries which clearly did not involve UK-based organisations. The prefix data for the remaining entries were checked against whois to confirm that the names were for UK-based organisations or not. Finally, the details for UK-based transactions were given to a script which converted them into CSV files, one for overall transfer activity and one for UK-specific transfers, `Transfers.csv` and `UKTransfers.csv` respectively.

IDENTIFYING UK LEGACY ADDRESS SPACE

Legacy address space broadly refers to IP addresses issued before the introduction of the RIR system in 1997. However this is not always the case. Legacy allocations may have been updated and changed their status since then. Some may have been merged into the address resources held by an LIR.
Address space was also issued by RIPE NCC before the introduction of the RIR system and in general these allocations are not considered to be legacy space. RIPE NCC began operations in 1992. Until RIPE NCC's incorporation in 1997 and the start of the RIR system, organisations could either obtain address space from IANA directly or get it from one of the /8 allocations IANA had made to the RIPE NCC in the early 1990s. Space allocated via that latter route is deemed to be "RIR space" and subject to RIR policies even though it was allocated before the introduction of the RIR system.

Therefore it is not possible to only depend on the timestamps in the `alloclist.txt` or daily DE file to identify legacy space: i.e. select all entries where the date precedes 22 December 1997, the date ARIN began co-ordination of the oversight of IP address resources and the RIR system was introduced. These complications mean the only

---

[21] http://www.ripe.net/lir-services/resource-management/ipv4-transfers/table-of-transfers

definitive way of identifying legacy address space is to check each allocation's status shown in the output from whois. A `status:` line saying LEGACY indicates legacy address space.

PROCESSING UK LEGACY ADDRESS SPACE

GB-tagged allocations from the `delegated-ripencc-extended-20140910` file were extracted by a script and whois lookups were performed for these allocations. The results were stored for further analysis, one file for each allocation. The output from these whois lookups were checked to find allocations which had the status LEGACY. This resulted in a list of 731 UK legacy allocations. A script was written to generate `LegacyHolders.csv` containing the name of the resource holder, the prefix, its size and the date of allocation for each of these allocations.

Another script used the RIPEstat `routing-status` API call to harvest routing data for these legacy address holders. The outputs from all these scripts were merged into another spreadsheet, `LegacyReachability`. A summary and analysis of the data in these files is presented in Section 8: Legacy IP Resources.

For each legacy allocation, the `routing-status` RIPEstat API call was used with the base address of the allocation to determine if there was any first_seen and last_seen dates returned. If so, that represented any first and last date when that base address was seen in a BGP routing announcement. If no dates were found and if the legacy allocation could be expressed in CIDR form, the RIPEstat `routing-status` call was invoked with the routing prefix and the dates of `first_seen` and `last_seen` to obtain any relevant routing data.

This pragmatic approach was taken after discussion with Ofcom. It was impractical to make RIPEstat `routing-status` API calls for every potential prefix that a legacy holder might have used to route part or parts of their address space. If this had been done, it is doubtful if the results would make much difference to the overall picture.

The result is a text file comprising one line for each legacy allocation with the following columns:

Address Holder — from whois
Base Address — from delegated-ripencc-extended-20140910 file
Addresses — from delegated-ripencc-extended-20140910 file
Date Allocated — from delegated-ripencc-extended-20140910 file
First Seen — from RIPEstat `routing-status` call
Last Seen — from RIPEstat `routing-status` call
Whois Routes — relevant whois route object (if any)

A spreadsheet tool further analysed this text file to generate a table of allocated and reachable addresses by year and a corresponding chart. An allocation is considered Allocated if the corresponding "Date Allocated" indicates it was allocated for the whole year. For example, a legacy allocation with a Date Allocated of June 1997 would be counted as allocated every year from 1998 (the first full year) onwards. An allocation is considered reachable in a year if the whole year is covered between the First Seen and Last Seen dates. For example, if First Seen is June 1997 and Last Seen is April 2003, the allocation would be counted as reachable for the years 1998–2002.

REACHABILITY FOR UK ADDRESS SPACE

RIPE NCC's Routing Information Service (RIS) provides the underlying data to RIPEstat about which address prefixes are routed on the public Internet. The RIS consists of a number of collector nodes at key Internet Exchange Points (IXPs) which monitor the BGP routing announcements made by the ISPs and other actors using the exchanges. Each collector node stores a copy of its view of the global routing table every eight hours. It also continuously logs the BGP traffic (routing announcements) that have been seen in the last 5 minutes.[22] This provides an excellent, fine-grained source of data about routing activity. RIPEstat holds data for IPv4 reachability since 2004. It has no routing data before then. Data on IPv6 routing are available from RIPEstat from 2008 onwards. Other initiatives such as the RouteViews[23] project hold earlier routing data.

A script has been supplied (`RoutingHistory.php`) so that Ofcom can inspect the routing history for any address prefix or LIR over the lifetime of RIPE NCC's Routing Information Service (RIS). This script retrieves additional data for an individual subnet or prefix by retrieving data returned from the RIPEstat `routing-history` API. The results from this script generated a large amount of data for some subnets/prefixes, some of which introduced further levels of inconsistency and complexity.

It was impractical to gather and present routing history for every UK allocation and assignment. RIPEstat's `routing-history` API offers just too much data to present here. However this was carried out for some address blocks that were of particular interest to Ofcom, namely allocations to leading UK communications providers and for the most significant elements of legacy space. First, the RIPEstat `routing-status` API for was used for each prefix or probable subnet to determine any first-/last-seen dates. RIPEstat's `routing-history` API was then used determine which portions of each of these prefixes or probable subnet had routing announcements. This yielded the ASN (Autonomous System Number) — a routing identifier — generated these announcements. The earliest start date and latest end date from all of the entries returned for a given subnet by a `routing-history` API call was then calculated. The name of the holder of the ASN was also looked up to identify the ISP making the route announcement.

Data from each of these larger legacy allocations was combined into a single spreadsheet. A similar spreadsheet was prepared for the leading UK communications providers. Visual inspection of the resulting information yielded some results. Although many of the larger allocations had no first-seen or last-seen dates, smaller subnets within those allocations did appear in the results of `routing-history` API calls. In some cases organisations that were unlikely to be associated with the legacy holder were announcing routes for portions of that legacy address space for short periods. For other resource holders, the routing history appeared to be consistent with activity by the address holder or a related organisation.

[22] http://www.ripe.net/data-tools/stats/ris/ris-raw-data
[23] http://routeviews.org

ANALYSIS OF IPv4 ADDRESSING

UK IPv4 ALLOCATIONS

The IPv4 addresses issued by RIPE NCC to UK organisations are analysed in this section. These are the UK-held IPv4 addresses which are subject to RIPE address policies. Although some of these address allocations pre-date the introduction of the RIR system, RIPE address policies apply to these resources. This is because these addresses were allocated by RIPE NCC before it formally became an RIR.
Address space acquired directly from IANA prior to 1997 is generally considered to be legacy space. These resources are not considered here and are analysed separately in Section 8: Legacy IPv4 Resources.

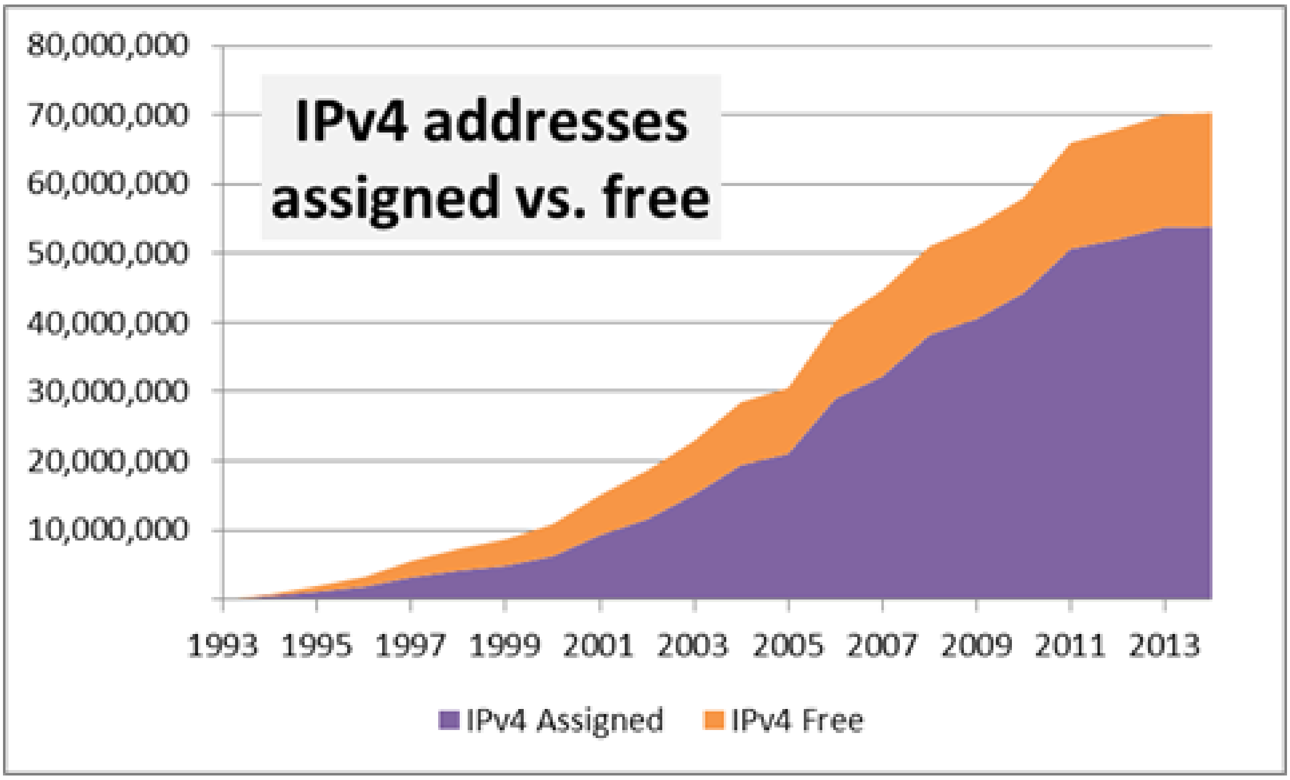

Figure 3 - Allocated/Assigned UK-held IPv4 Address space

The spreadsheet `UKLIRs` contains the data on IPv4 allocations, assignments and route visibility for the 1,186 UK LIRs which are summarised in this section of the report.

Table 1 shows the extent of IPv4 address space issued by RIPE NCC to UK-based organisations over the last 20 years. It explains the amount of IPv4 address space allocated to UK LIRs, how much of that space has been assigned, how much is notionally “free” and the ratio of assigned (i.e. notionally used) to allocated address space. In this context, “free” is a reasonable assessment of how much address space remains available for assignment by an LIR. This is explained in more detail below.

Year	Number of Addresses **(in thousands)**			
	Allocated	**Assigned**	**Available**	**% Assigned**
1994	750	526	223	70%
1995	1,937	1,134	804	59%
1996	3,236	1,836	1,399	57%
1997	5,562	3,187	2,376	57%
1998	7,303	4,147	3,157	57%
1999	8,667	4,853	3,814	56%
2000	10,869	6,240	4,629	57%
2001	15,047	9,267	5,779	62%
2002	18,610	11,573	7,037	62%
2003	22,903	15,142	7,761	66%
2004	28,517	19,426	9,092	68%
2005	30,559	21,035	9,524	69%
2006	40,211	29,015	11,197	72%
2007	44,691	32,234	12,458	72%
2008	51,028	38,240	12,788	75%
2009	53,959	40,552	13,406	75%
2010	58,049	44,314	13,735	76%
2011	65,960	50,672	15,287	77%
2012	67,898	52,052	15,847	77%
2013	70,092	53,782	16,309	77%
2014	70,327	53,857	16,470	77%

Table 1 - IPv4 Allocation, Assignment & Availability in the UK

Growth in the number of IPv4 addresses allocated to UK-based organisations has been fairly steady, usually increasing by approximately 4 million year on year. However that trend stopped for 2013-14 where the increase has been just 200,000 IPv4 addresses. This is a consequence of the introduction of the restrictive address policy that applies now that the RIPE NCC is allocating from its last /8.

The amount of assigned address space is a reasonable estimate of the amount of address space that is "in use". This has also generally shown steady growth of around 4 million addresses year on year. That is to be expected. The old needs-based address policy required LIRs to demonstrate that they had a valid use case for address resources. An LIR had to show it had used most of its previous allocation before they could get another. If an LIR has been allocated a /12 (say), it would not get a further allocation until it had assigned most of the space in that /12. Therefore assignments by LIRs would broadly track allocations to LIRs by the RIPE NCC.

The amount of assigned address space has hardly changed for 2013-14, an increase of just 70,000 addresses. This too will also be a reflection of the current address policy. LIRs will be adopting a much more conservative approach to managing their remaining IPv4 resources because they cannot get any more IPv4 allocations from the RIPE NCC. Workaround solutions such as Network Address Translation and/or Application-Level Gateways (ALGs) may well be getting deployed to eke out those remaining reserves. Checks were made of the corresponding DE files maintained by the other RIRs and these showed very little amounts of UK-held address space. ARIN allocated a /16 for Royal Bank of Scotland and around 20 /24s to other UK organisations, all in the mid-1990s. None of the space issued by the other RIRs was allocated after 2005. To date there does not appear to have been any forum shopping by UK organisations at the other RIRs. That could change if the demand for more IPv4 address space becomes a serious business risk.

The "free" space is the amount the RIPE NCC has calculated as the difference between the amount of space that was allocated to an LIR and how much the LIR has assigned from that allocation. It is a measurement of how much space in the allocation that remains for assignment to an LIR's customers or business units. It is a reasonable measure of the amount of address space that remains in an allocation for assignments to customers or use internally. This estimate is not foolproof however because the data quality for an LIR's assignments can vary. Some LIRs will ensure these are reasonably accurate and up to date. Others won't bother. Even when assignment data are correct, it does not necessarily mean that all the addresses in a given assignment are used. For instance an LIR might assign a /27 (32 IPv4 addresses) to a customer who only uses 20 of those addresses. Despite these limitations, relying on the RIPE NCC's definition of "free" address space provides a stable basis for any future data gathering and analysis because it avoids depending on variables and assessment criteria which may change later.

It is also possible that some LIRs will be using more of their allocated address space than they have actually assigned. Whenever assignments have been made, it does not necessarily follow that all of the address space in each of these assignments was fully utilised. A reasonable working assumption is that these two scenarios broadly cancel each other out and the counts of "free" space reported by the RIPE NCC are a good approximation to the amount of allocated address space that is still to be activated. Other metrics for assessing the amount of free space could have been chosen. However these may not provide more accurate or consistently reproducible results. For instance, although it is possible to perform address scans using facilities like the Internet-Wide Scan Data Repository,[24] it does not necessarily follow that these accurately measure active address space. Firewalls and access routers may well block these probe packets or return false information, reporting addresses as being in use when they are not or *vice versa*. High volumes of probe packets might be viewed as a Distributed Denial of Service attack, attracting adverse publicity or complaints. Techniques like this may also be subject to change, say by introducing different types of probe packets to get round the latest firewall defences. If so, what gets measured in future could well be different from what can be measured today.

There may be some variation in those usage thresholds. Relative to the size of their allocations, some LIRs will have little space remaining for assignment, others may have a lot. However the ratio of unassigned to allocated space has remained fairly stable over the last 8 years, starting at 72% in 2006 and gradually rising to 77% by 2011 where has remained to date. Details of allocations, assignments and reachability of the address space held by the main UK providers are supplied as discrete spreadsheets.

Year	Allocated IPv4 addresses **(thousands)**	Reachable IPv4 addresses **(thousands)**	**% Reachable**
2004	28,517	22,025	77%
2005	30,559	24,198	79%
2006	40,211	32,685	81%
2007	44,691	35,834	80%
2008	51,028	40,864	80%
2009	53,959	43,526	81%
2010	58,049	46,307	80%
2011	65,960	50,225	76%
2012	67,898	56,786	84%
2013	70,092	56,187	80%
2014	70,327	56617	81%

Table 2 - Reachability of UK-held IPv4 Address Space

Note that the entries for 2014 in Tables 1 and 2 refer to the data gathered at the time of writing, September 2014, rather than the year-end that was used for previous years. Table 2 presents the year-by year breakdown of route visibility for UK address space issued by the RIPE NCC as determined by the RIPEstat `routing-status` API, looking at the First Seen and Last Seen dates for each prefix. This shows that approximately 80%

[24] https://scans.io

of those IPv4 resources have been in prefixes that have had routing announcements which were detected by the RIPE NCC's Routing Information Service[25], RIS, and logged in its database. This service began in 2004, so it holds no data for earlier routing activity. These statistics indicate that approximately 80% of the IPv4 address space issued by RIPE NCC to UK LIRs is routed on the public Internet. There are some important qualifiers that apply to this statement.
First, it does not mean that the remaining 20% of that space is not being used. That space is simply not being announced or routed on the Internet. The address space could well be used for internal purposes, say for a corporate Intranet which is not directly reachable from the public Internet. There would be little if any justification to announce routes for these publicly unreachable prefixes.

Second, it does not mean that the 80% of the IPv4 address space issued by RIPE NCC to UK LIRs and routed on the public Internet is actually in use. It is common practice for LIRs to announce routes for the entire prefix for each allocation, even if only some of that allocated space has been assigned and is active. When LIRs announce a route for an entire prefix, they can avoid reconfiguring their routers and/or having further negotiations with upstream providers whenever they assign more space from that allocation/prefix.

The consequence of that common practice is the amount of address space which is being routed is likely to be larger than the amount of space that is currently in use. LIRs will usually accept incoming traffic for any unused assignments and just discard it, perhaps returning a "destination unreachable" error to the originating source address.

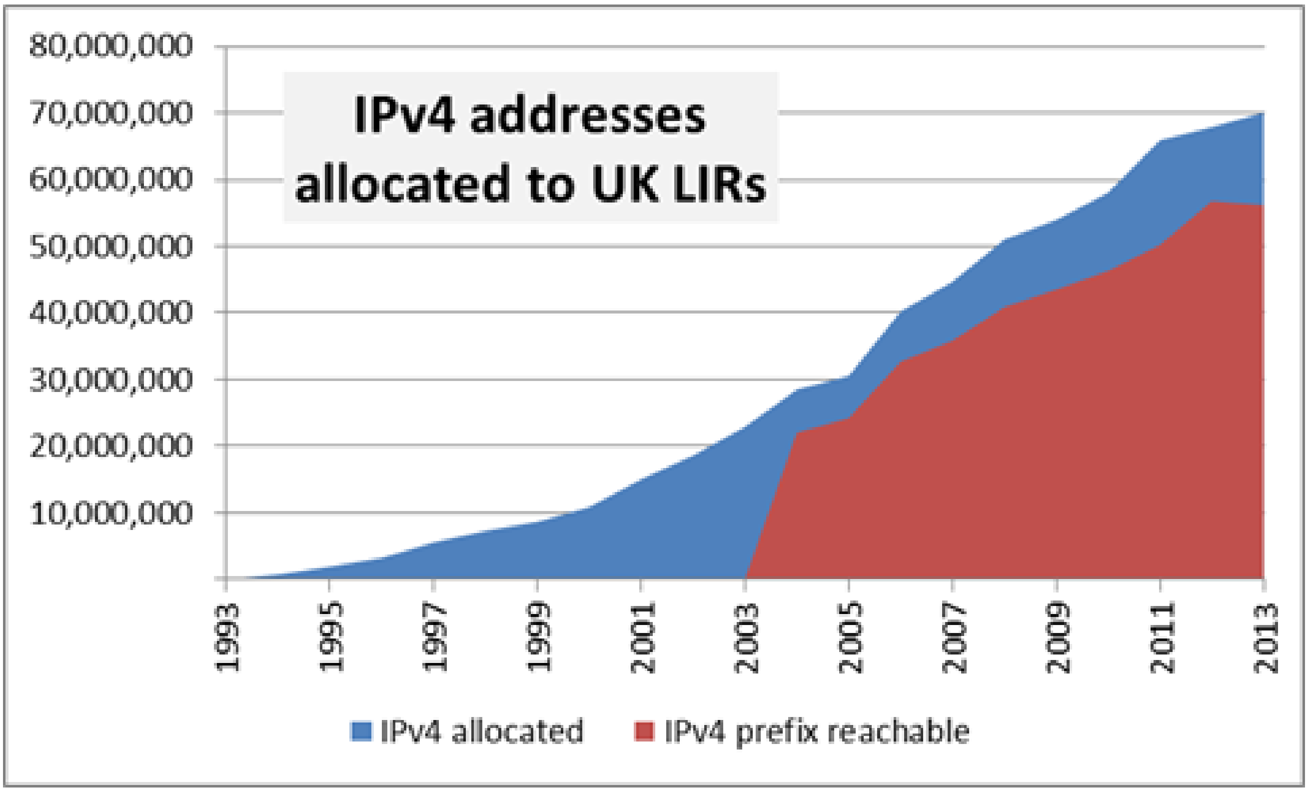

Figure 4 - Allocations and reachability of UK-held IPv4 address space

[25] http://www.ripe.net/data-tools/stats/ris

IPv4 TRANSFERS

CURRENT RIR ADDRESS TRANSFER POLICY

RIR policies on IP address transfers have loosened as a result of the run-out of IPv4 space. Until recently, address resources could only change hands whenever some sort of business reorganisation such as a merger or acquisition took place. These restrictions have been removed to allow for address distribution to be rebalanced by LIRs who wish to do so. LIRs who have extra IPv4 addresses can make these available to LIRs who need more addresses than they can obtain from their RIR. The gaining LIR does not have to represent need or provide an addressing plan as they would be expected to do when requesting space from an RIR. Any transaction that might take place between the donor and recipient LIRs is invisible to the RIR. The RIR simply keeps track of who holds transferred address blocks, ensuring the continued integrity of its database of IP addressing resources.

Details of RIPE's current transfer policy and its conditions are given in Section 5.5 *Transfers of Allocations* of RIPE Document 612, *"IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region"* [26]. This updated policy was adopted in August 2104. In outline any LIR can re-allocate complete or partial blocks of IPv4 address space that were previously allocated to them by the RIPE NCC or otherwise through the RIR system. Address space may only be re-allocated to another LIR that is also a member of the RIPE NCC. These re-allocations must be reflected in the RIPE Database.

The RIPE NCC provides a web page where its members can exchange information about IPv4 address space which could be re-allocated under this new transfer policy.[27] At the end of August 2014, 800,000 IPv4 addresses were available for transfer outwith the established mechanisms for routine business mergers, acquisitions and spin-offs. Most of these are small blocks of addresses. The largest block on offer is a /16: 65,536 addresses. LIRs had lodged requests for over 13 million IPv4 addresses. The vast majority of these are for 4,096 addresses or less (a /20) though the largest is for a /11: around 2 million addresses. Information on who is advertising or requesting space is limited to RIPE NCC members and only available via its LIR portal.

The RIPE NCC maintains a public list with details of these transfers: donor, recipient, the address block or blocks in question and the date when they were transferred.[28] This list is currently updated on a weekly basis. It provides information on all LIR-LIR transfers in the RIPE service region since the restrictions on address transfers were first relaxed at the end of 2012. This information does not include transfers that were due to a change in company structure such as a merger or acquisition.

Table 3 is derived from that list and summarises the transfer activities involving UK LIRs. It shows that 71 address blocks — and by implication a similarly small number of UK LIRs — have been involved in this new transfer regime. Overall, the amount of address space that has been transferred is low. Transfers involving UK LIRs account for a total of 527,360

[26] http://www.ripe.net/ripe/docs/ripe-622
[27] http://www.ripe.net/lir-services/resource-management/listing
[28] http://www.ripe.net/lir-services/resource-management/ipv4-transfers/table-of-transfers

IPv4 addresses. This corresponds to less than 1% of the UK IPv4 address space allocated by RIPE NCC. Only four of those 71 transfers involved more than 5,000 addresses. The vast majority of transfers involving UK address holders were for 4,096 addresses or fewer. Almost all of these activities involved UK LIRs as the recipients of transferred address space rather than providers of that space.

Size of Transferred Address Block	Number of Transfers
131,072	2
65,536	1
32,768	1
4,096	24
2,048	24
1,024	19

Table 3 - UK Address Transfers

There are some significant exceptions. BSkyB acquired approximately 100,000 addresses from Easynet and 200,000 from O2/Telefonica. An industry contact confirmed the /15 from Easynet was sold to BSkyB. In 2003 BSkyB bought the consumer broadband and fixed-line business from O2 and BE which were owned by Telefónica UK at that time. Transfers of those /16 and /15 blocks may be related to that acquisition. Further details are in the `UKTransfers.csv` spreadsheet.

Across the RIPE service region 731 address blocks have been transferred under this new policy, some 6.6 million IPv4 addresses in total. Approximately 500 of these transfers were for 4,096 addresses or fewer. However 23 large blocks have been transferred, accounting for 3.5 million (53%) of the transferred IPv4 addresses. Four /14's (262,144 addresses) and 19 /15's (131,072) have changed hands. The spreadsheet `Transfers.csv` provides full details. A summary of transfer activity across the region is shown in Table 4.

The format of the `UKTransfers.csv` and `Transfers.csv` spreadsheets are identical. They contain the address block that was transferred, its size, the name of the donor and recipient LIRs and the date of transfer. In this context donor just means the organisation which released the address space. It does not imply the space was donated. The LIR might well have sold it to the recipient and details of any payment would be a private matter between both parties.

To date, UK LIRs have not made much use of this new transfer policy. The evidence shows that few UK LIRs have offered to transfer or acquire addresses through this mechanism. Likely explanations could include: (a) UK LIRs feel they have few "spare" addresses for transfer; (b) address holders are waiting to see how the new policy works out and what market opportunities, if any, emerge; (c) UK LIRs do not have suitably sized blocks of addresses which match the needs of the LIRs who are requesting IPv4 space; and (d) address holders are unsure if the costs in network restructuring to make addresses available for transfer are worth the rewards. Similar considerations apply on the demand side. UK LIRs might feel they have enough IPv4 or can work around any shortfall by deploying NAT or IPv6. They may be adopting a “wait and see” approach. Or they may be hoping to obtain address blocks of a size that are not currently on offer. They might be obtaining address space from an organisation that has plenty of unused addresses: for instance by coming to an arrangement with a legacy resource holder or by acquiring an existing LIR.

Checks were made of the corresponding DE files maintained by the other RIRs and these showed very little amounts of UK-held address space. None of this space was allocated after 2005. To date there does not appear to have been any forum shopping by UK organisations at the other RIRs. That might change if the demand for more IPv4 address space becomes a serious business risk.

The transfers shown on the above RIPE NCC web page are for address space in the RIR system and managed by RIPE NCC. None of the address blocks listed there or in the supplied spreadsheets are legacy space.

It is possible uptake of this transfer mechanism will increase as the demand for IPv4 addresses continues and the effects of the IPv4 run-out become more acute. On the other hand, supply might not be able to keep pace with demand. If so, there could be fewer transfers as address holders hoard their remaining reserves of IPv4 space in expectation of realising a higher price.

Size of Transferred Address Block	Number of Transfers
262,144	4
131,072	9
65,536	24
32,768	25
16,384	39
8,192	39
4,096	138
2,048	150
1,024	203

Table 4 - Transfer activity across the RIPE service region.

Ofcom may want to monitor the use of the new transfer policy. First, transfer activity will show how much address space is involved and who the actors are. That would allow Ofcom to identify the active UK LIRs. Second, it may be prudent to monitor this emerging market. Finally, the transfer policy itself may be changed in light of operational practice. That might not necessarily be in the best interests of UK address holders as a whole or some segment of the UK LIR community.

The NCC transfer listing web page also shows transfers of address space to Vodafone subsidiaries in Malta, Portugal, The Netherlands and Germany. Its German subsidiary Vodafone GmbH received 750,000 addresses, 2 /14s and 2 /15s. Vodafone’s Portuguese subsidiary/joint venture obtained 375,000 addresses, a /14 and a /15. While these transfers are outside the scope of this study, it does help to illustrate one of the problems of defining UK IP address space. Vodafone may or may not be using these transferred addresses for its UK customers or for running parts of its corporate network which may or may not be in the UK. Such information is not in the public domain.

LEGACY IPv4 RESOURCES

LEGACY ADDRESS RESOURCES AND THE RIR SYSTEM

Legacy IP space was issued before the RIR system was introduced at the end of 1997. Those addresses and their holders are not subject to RIR policies. Legacy address holders would have acquired space from IANA which co-ordinated the early days of what became the Internet. However not all pre-RIR allocations are considered legacy space. Some have been merged into the address resources held by an LIR.

Address space was also issued by RIPE NCC before the introduction of the RIR system. These allocations are generally not considered to be legacy space. Some organisations obtained IP address space from RIPE NCC prior to its formal incorporation in 1997. During the early 1990s organisations could either obtain address space directly from IANA or via RIPE NCC which had been given /8 allocations by IANA. Space obtained via RIPE NCC before 1997 is deemed to be "RIR space" and subject to RIR policies even though it had been allocated before the RIR system was in place.

Legacy address holders may be members of RIPE NCC or another RIR. Some became LIRs because they later had a need for additional numbering resources. In these cases, the LIR's legacy resources usually remain outside the RIR system from a policy perspective. Those legacy addresses would not be taken into consideration by the RIR when considering the member's allocation or assignment requests. These constraints would only apply to the resources that the LIR had obtained through the RIR system. Legacy address space is not bound by RIR policies for transfers and the address holders are not obliged to return any of that unused space for redistribution. Even so, this has sometimes happened. In 2000 Stanford University renumbered its campus network and the Class A network 28/8 it had been allocated at the beginning of the ARPAnet was returned to ARIN for general allocation.

A further reason for keeping legacy space outside the RIR system was to avoid fees and extra paperwork. Until recently, the NCC's fee structure was related to the size of the organisation's number resources. Those with more numbering resources paid higher fees. There was therefore an incentive for holders of large legacy resources to keep these outside the RIR system. The NCC now uses a flat fee structure. When legacy address space gets merged into an LIR, usage of that space would be assessed whenever the LIR requested more address space from the RIR. That might prevent the LIR from getting another allocation until it had used most of its existing address space, including its legacy resources.

Another incentive for keeping legacy resources outside the RIR system is the discretion to redeploy them without interference from RIR policies. While it is hard to know for certain, it is reasonable to assume that some legacy holders may be taking a "wait and see" approach until the emerging market in IPv4 address trading has matured. In 2011 Microsoft acquired 666,000 IPv4 addresses — a /19 and a /18 — which had been held by bankrupt telecoms firm Nortel for $7.5M.[29] Legacy holders may well be hoping to realise

---

http://www.bbc.co.uk/news/technology-12859585

similar sums by selling off excess address space. This may be more likely in North America where some large US-based telcos and ISPs are believed to have substantial reserves of unused address space.
Some organisations have voluntarily merged their legacy space into their LIRs. Sometimes this has been done for administrative convenience so that a single point of contact can be used for managing the RIR database entries. In other cases, the organisations felt this was the right thing to do. Since the RIRs were providing service for that legacy space, it was only fair that the organisation paid some fee in return for those services. When these mergers take place, the LIR's legacy space become subject to RIR policy.

Table 5 shows the largest legacy address allocations to UK-based organisations, those larger than a /16.

Holder	Prefix/ **Address Block**	Number of **Addresses**	Allocation Date	RIR Database **Route Object**
Ministry of Defence	25/8	16,777,216	1/1/95	
Department for Work & Pensions	51/8	16,777,216	1/9/93	
BP	149.177.0.0	1,245,184	13/5/91	
British Telecommunications plc	147.147.0.0	393,216	1/1/84	147.147.0.0/16
Global Crossing	171.28.0.0	196,608	9/1/95	
DEFRA	148.252/15	131,072	1/9/93	
HM Customs and Excise	163.171/15	131,072	14/2/92	

Table 5 - Largest UK Legacy Address Allocations

There are a total of 53.2 million UK legacy IPv4 addresses. 63% of that is accounted for by the /8s allocated to DWP and the MoD, some 33.4 million addresses. There are 261 /16 (Class B) legacy allocations, a further 17.1 million addresses — 32% of the overall total.

96 of those /16 allocations were made to UK universities. The largest allocations shown in Table 5 apart from the 2 /8s account for another 2.1 million addresses: 4% of the overall UK legacy space. Small allocations, mostly groups of contiguous /24s, account for the remaining 5% of legacy space, some 2.6 million IPv4 addresses.

Some of the allocations shown in Table 5 can't be represented in standard slash notation because they are not neatly aligned on *N*-bit boundaries: some power of 2 multiple of a /24 with 256 addresses or a /16 with 65536 addresses. Organisations could be allocated several contiguous /16s or /24s. BP's legacy allocation consists of 19 contiguous /16s, Global Crossing has 3 contiguous /16s and BT has 6 /16s. Two contiguous /16s were allocated to both DEFRA and HMRC.

The presence of a route object or objects in the RIR database is a good indication that the address holder intends (or intended) to have some or all of the address space routed on the public Internet. These objects can be looked up using RIPEstat or other tools and services provided by RIPE NCC. They can also be found in the output from a whois lookup. Many ISPs use the presence of a route object in the RIR database to decide if they will accept a routing announcement or route traffic for that prefix on behalf of a customer. It provides a simple check that an impostor is not trying to fake a route or redirect traffic for some prefix. Although anyone is in principle able to announce a route for anything, only the address holder would be able to insert or update a route object for their address space

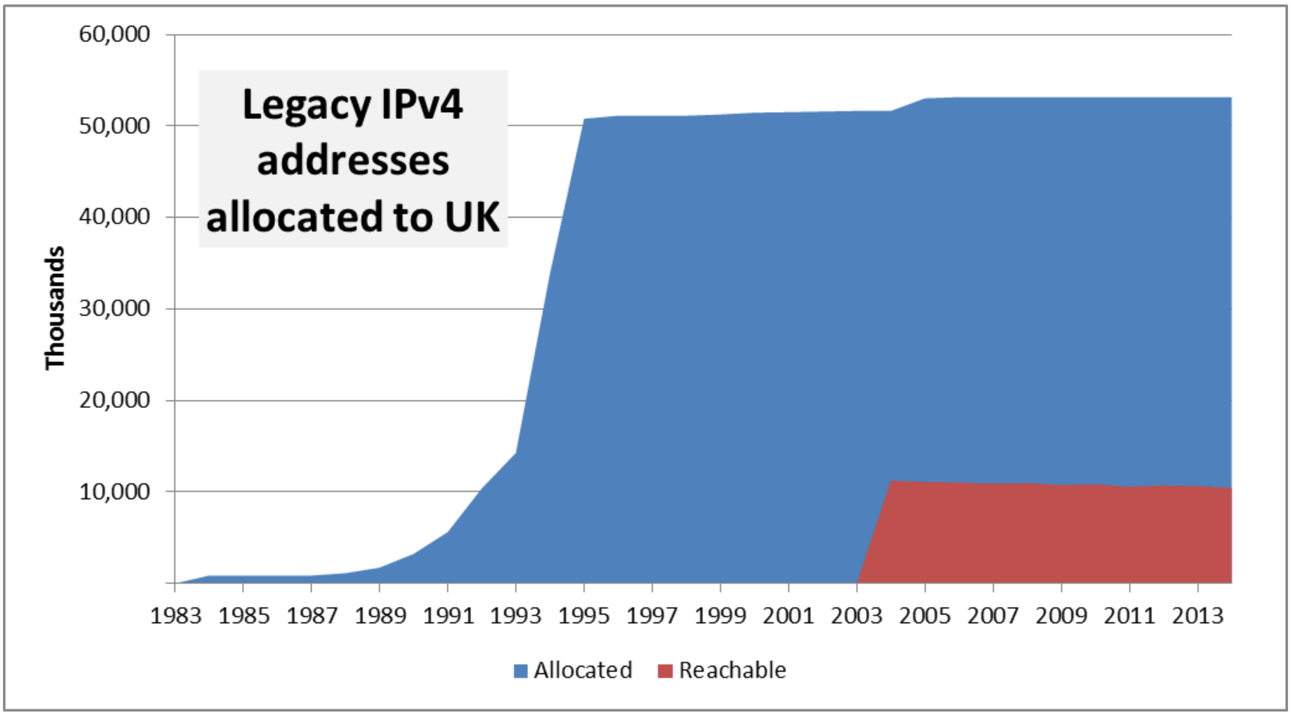

Figure 5. Legacy UK IPv4 address allocations

in the RIR database. Route announcements that are not backed by a corresponding route object in the RIR database will usually be treated with scepticism.

The `LegacyReachability` spreadsheet gives a complete breakdown of UK Legacy Space. It contains one line for each legacy allocation with the following columns: name of the address holder, the prefix or base address, the number of IP addresses, date of allocation, first and last seen dates in the RIS database and any whois route object for the prefix. A more detailed report of IP address usage and reachability for the largest UK legacy resource holders is in the `LargestLegacyAllocationAnalysis` spreadsheet.

ADDRESS UTILISATION BY LARGEST UK HOLDERS OF LEGACY ADDRESSES

In a UK context, two legacy address allocations are by far the most significant. These accounts for 63% of the total UK legacy space: 33.4 million out of 53.1 million IPv4 addresses. The original Class A network 25/8 was issued to the Ministry of Defence in January 1995. The DHSS, now the Department of Work and Pensions, obtained the Class A network 51/8 in August 1994. Both of these allocations contain 224 addresses: just over 16 million. Neither is reachable from the public Internet or has even been seen to be reachable as there have been no observed valid routing announcements for these prefixes.

It is improbable these address blocks are being used efficiently. A reasonable metric for gauging the size of a network is one address per employee. According to their respective web sites, the MoD has just over 61,000 civilian employees and DWP has a staff of approximately 160,000. Applying this metric would suggest that both networks are likely to be using no more than perhaps 1% of their available address space. However it does not automatically follow that the remaining 99% could be available for redistribution or that the DWP and MoD would be prepared to make that spare address space available.

There are some smaller but still substantial legacy allocations. BP plc has 149.177/12: just over 1 million addresses. British Telecom plc has 6 contiguous /16s: approximately 400,000 IPv4 addresses. Two government departments each have 2 contiguous /16s: HM Customs and Excise and The Department for Environment, Food and Rural Affairs. DEFRA employs 10,000 staff so perhaps only 10% or so of its legacy IPv4 allocation is in use. HMRC has 67,000 employees. A reasonable estimate for IPv4 usage by HMRC would be around 50% of its allocated legacy address space. There are 261 Class B allocations, totalling 17.1 million IP addresses. Royal Mail holds 4 legacy Class B allocations: 144.87/16, 147.119/16, 147.77/16 and 163.169/16. A further Class B allocation, 163.36/16 is held by the Department of Business, Enterprise and Regulatory Reform. It is very unlikely these addresses could be made available for redistribution for the reasons outlined above. Assuming half the UK's legacy allocations of Class B space could be redistributed — which seems very optimistic — it would amount to 8.6 million addresses.

Of the remaining Class B legacy allocations, 96 /16s were allocated to UK academia. This accounts for 6.3 million UK legacy addresses, 12% of the total legacy space. Most of the address space held by academia is in use and routed on the public Internet via the UK academic/research network, JANET.

CONSIDERATIONS FOR THE RELEASE OF LEGACY ADDRESS SPACE

There are many challenges and risks for holders of legacy address space who wish to release that space or parts of it. They are likely to incur significant costs and there may well be disruption to their networks and the systems using them.

Audits would be necessary to get an accurate and timely picture of how the existing space is being used in those networks. These are likely to be expensive and time-consuming because large networks tend to be complicated and difficult to manage. Networks may need to undergo restructuring or a redesign if internal address assignment policies are sub-optimal. Parts of the address holder's network might have to be renumbered. Addressing plans for the networks would probably require major overhaul and/or review. Complex migration strategies may be necessary: i.e. temporary deployment of NAT or proxy solutions to help move systems and/or end users away from address space that is to be released.

Significant security considerations can be expected too. For instance if parts of the MoD's 25/8 network was released, systems and firewalls on that network would need to be checked and probably get reconfigured. They would no longer be able to consider the whole 25/8 network as "trusted" because some of that address space would be getting transferred away from the MoD to a third party. Any released space would therefore be outside rather than inside the MoD's network.

Addressing these issues will be costly and troublesome. Put simply, redistributing some of that space might be too risky. If the freed up space was then sold, that might not raise enough money to cover the cost of doing so.

It is also unclear what terms and conditions would apply if legacy space was to be released. IP addresses are not considered to be property from a legal perspective and it is uncertain if suitable legal documents exist to underpin the provenance or “ownership” of IP address space. Lawyers may decide there can be no basis for selling legacy address space because the address holders do not legally own it. In 2004, the then DTI’s lawyers reached this conclusion when assessing the status and ownership of the UK’s ENUM delegation (4.4.e164.arpa) even though this was notionally under Departmental control. Similar legal considerations may well apply to the legacy space held by government departments.

The case for redistribution of legacy Class B and C space or parts of that space is even weaker than it is for a Class A network. The same concerns over costs and disruption that are outlined above still apply. However the amount of space that could be freed up and/or sold would be proportionately much smaller. That suggests there is unlikely to be much enthusiasm for redistributing parts of legacy Class B and C space. The costs would be far more likely to outweigh the potential benefits. In addition, the holders of those legacy resources may well want to hold on to any unused space within those blocks. That would permit the legacy holders to continue growing their IPv4 networks with the minimum of fuss when there were so many uncertainties about the dwindling reserves of IPv4 space at the RIRs and elsewhere.

REACHABILITY OF UK LEGACY ADDRESS SPACE

Table 6 gives details of the total amount of UK legacy address space and how much of that is routed on the Internet. This has been remarkably stable over the last 10 years with about 18% of the legacy space being routed. Much of that is accounted for by the hundred or so Class B networks — some 6.3 million addresses — that were allocated to UK academia. The two government /8s have not been routed and these alone comprise 63% of the UK legacy space.

It is conceivable that legacy holders could be making that space available to third parties. However since this amounts to a tiny percentage of the overall UK legacy space, any transactions that might have taken place to date seem likely to be insignificant from a market perspective. This might well change in the future. Similarly, UK organisations may be obtaining legacy resources from holders outside the UK. They may be using these resources internally or on the public Internet. It is not known if this is occurring and it would be hard to find out. In any case, this is out of scope for this project.
The `RoutingHistory.php` script was used to gather data about the largest legacy allocations. These results introduced further levels of inconsistency and complexity. The `LargestLegacyAllocationAnalysis` spreadsheet provides further details of these allocations and their reachability.

Year	Legacy IPv4 Addresses (thousands)		
	Allocated	**Reachable**	**% Reachable**
2004	51,656	11,222	22%
2005	53,033	11,091	21%
2006	53,164	11,026	21%
2007	53,164	10,896	20%
2008	53,164	10,965	21%
2009	53,164	10,768	20%
2010	53,165	10,834	20%
2011	53,165	10,572	20%
2012	53,165	10,704	20%
2013	53,165	10,637	20%
2014	53,165	10,459	20%

Table 6 – Reachability of Legacy Address Allocations

Some salient points from that data gathering are worth reporting however. Small prefixes, typically a /24, have been announced as routes for parts of the DWP's legacy /8. These have tended to be short-lived, perhaps lasting for a few weeks. That suggests there was either a short-term requirement to route some of that space on the public Internet or else there was some administrative error that caused these prefixes to temporarily "leak" into the core Internet routing tables. Parts of the DWP's legacy /8 were found to have been announced as routes by a US-based LLC for 11 days at the end of December 2012, and by a Bulgarian company for 23 days between May and June 2014. It is highly doubtful either of these announcements were authorised or requested by DWP. These are almost certainly bogus and most likely have been caused by mistakes in router configuration at these companies that went unresolved for several days. These errors probably had no impact on the DWP's network because that appears to be isolated from the public Internet Examples of obviously bogus routing announcements were also found. Prefixes larger than a legacy allocation were reported as reachable for some periods of time — for example an announcement of a route for a /1 or a /7. They were ignored for this analysis. These were probably caused by human error.

Almost all of the /16s for UK academia are found in RIPE NCC's RIS database. This is to be expected because universities will be connected to the Internet via the national academic network, JANET. Some universities hold more than a single /16 and a small number of these have not been visible in the routing table. According to the RIPEstat `routing-history` API 18% of UK legacy space is routed on the Internet and this proportion has remained essentially static for the last 10 years. Since 12% of routed legacy space can be accounted by academia, the remaining 6% or so will be routes that are announced by other legacy holders.

Route visibility data for corporate legacy allocations are interesting. For instance, some of BP’s legacy space is visible on the public Internet. This does not mean that the rest is unused. BP may well be using that on their internal network. Many large organisations use globally unique IP addresses even though they are not directly connected to the Internet. This helps to avoid renumbering problems when acquiring or divesting business units or departments.

IPv6 UPTAKE

IPv6 ALLOCATIONS AND ROUTE ANNOUNCEMENTS

The spreadsheet `UKLIRs` contains the data on IPv6 allocations, assignments and route visibility for the 1,186 UK LIRs which are summarised in this section of the report. Uptake of IPv6 by UK LIRs is somewhat disappointing. The extent of IPv6 allocation by UK LIRs is presented in Table 7. Only 74% of UK LIRs have an IPv6 Allocation. A quarter of UK LIRs do not have IPv6 space at all. Interestingly, 20 UK LIRs ONLY have IPv6.

Total number of UK LIRs	1186
LIRs with IPv6 allocations	876
LIRs with no IPv6 allocations	310
LIRs with only IPv6 allocations (no IPv4)	20
LIRs with reachable IPv6	332
LIRs with unreachable IPv6	544

Table 7 - Extent of IPv6 deployment.

It should be noted that allocated IPv6 address space is not required to be advertised in the routing system. Some public IPv6 addresses may be used in private networks and therefore are not visible in public routing tables. Furthermore the routing tables only indicate the capability to support IPv6 in routing, rather than actual use of IPv6 for services or traffic.

62% of UK LIRs who have IPv6 space are not announcing routes for their IPv6 allocations. In other words, they are not using IPv6 natively on the Internet. The absence of an appropriate routing table entry does not necessarily mean an LIR has not deployed IPv6. The addresses might be used internally. While the routing table may provide a good indication about the deployment of "native" IPv6 addresses, it does not consider “special case” IPv6 addresses used in transition/migration mechanisms such as 6to4 and Teredo. It is possible that the LIRs who do not have IPv6 allocations or announce IPv6 routes may be using some form of tunnelling or migration technology to facilitate IPv6 connectivity. Tunnel brokering services such as SixXS[30] or Hurricane Electric[31] might be being used. These encapsulate IPv6 traffic inside IPv4 packets which then traverse a mostly IPv4 Internet. Some organisations may be using these tunnel broker services or migration technologies like 6to4 and Teredo to gain operational experience with IPv6 or carry out testing and training before fully deploying IPv6.

Of the UK LIRs who have IPv6 space, less than half of these are announcing routes for their space. The majority of UK-held IPv6 allocations are unreachable from the Internet

---

[30] https://www.sixxs.net
[31] https://ipv6.he.net

unless some sort of tunnelling is used. Overall, only 28% of UK LIRs are able to offer direct IPv6 connectivity to the Internet.
This compares unfavourably to most other countries in the RIPE NCC service region. The RIPE NCC maintains and publishes a rating system which awards stars to its members depending on key indicators of IPv6 preparedness.[32] Stars are awarded to LIRs for:

- Having an IPv6 address space allocation or assignment from the RIPE NCC
- Visibility in the Routing Information Service (RIS)
- Having a route6 object in the RIPE Database
- Having reverse DNS delegation set up for their IPv6 address space

Table 8 presents a breakdown of the percentage of LIRs for each of the 5 star ratings. A zero star rating means no IPv6 and a 4 star rating indicates IPv6 is fully deployed. The percentages are shown for all LIRs and then for those based in the UK. Overall, the UK's IPv6 rating is below the average across the RIPE NCC service region.

STAR RATING	PERCENTAGE OF LIRS ACROSS NCC SERVICE REGION	PERCENTAGE OF LIRS IN UNITED KINGDOM
0	29	31
1	25	32
2	8	7
3	14	12
4	21	17

Table 8 - RIPENESS metrics

APNIC Labs operate a long-running experiment to track IPv6 uptake. It uses carefully designed adverts to gather data on whether web browsers across the Internet are using IPv4 or IPv6 when making a total of 400-500,000 connections per day. The results are published as aggregated information that shows the uptake of IPv6 on a global basis, and on a country-by-country basis, over time across as much of the Internet as practically possible.[33] This gives perhaps the most comprehensive picture of overall IPv6 activity. At present, the UK is ranked at 40th out of 175 countries by this APNIC Labs experiment. [There are insufficient data for the other 50 or so ISO 3166 country codes to rank them.] UK-based IPv6 prefixes account for 0.18% of the detected IPv6 traffic and an average of 49,000 web hits per month over the last three months. By contrast, Belgian IPv6 prefixes account for 31.5% of the detected traffic. China and Turkey each generate an average of over 400,000 web hits per month over the last three months, almost 10 times as much as the UK.

The Department of Business, Innovation and Skills set up 6::UK, an organisation to raise awareness in the UK about the IPv4 run-out and need for IPv6 deployment in 2010. It

[32] http://ipv6ripeness.ripe.net
[33] http://labs.apnic.net/ipv6-measurement/Economies/

failed to get support from UK business or the public sector and was wound up at the end of 2012.

A 2010 OECD report, *"Internet Addressing: Measuring Deployment of IPv6"* provides a good explanation of the challenges for IPv6 deployment and how it can be measured.[34]

IPv6 ALLOCATIONS AND ROUTE ANNOUNCEMENTS

Table 9 gives an overview of how UK LIRs have acquired and allocated IPv6 space since 2000. The values for thousands of /48s have been rounded up or down for presentation in this report.

For 2000 to 2003, this means the number of free, assigned and aggregated /48s do not equal the count of allocated /48s. The actual results are provided in the `UKLIRs` spreadsheet. The entries for 2014 refer to the data gathered at the time of writing, September 2014, rather than the year-end that was used for previous years.

Since the IPv6 address space is huge, it is difficult to present these numbers in a way that is straightforward to understand. It simply isn't meaningful to discuss absolute numbers of IPv6 addresses or utilisation rates. The minimum IPv6 assignment an end user is probably a /48: 1.2×10^{24} (1,208,925,819,614,629,174,706,176) addresses. This is approximately 300 trillion times the size of the entire IPv4 address space. It is improbable an end user such as a DSL customer or 3/4G subscriber would be able to use more than a few hundred of the IPv6 addresses in that assignment. After discussion with Ofcom, it was decided to use the /48 as the granularity of IPv6 analysis in this report since this was likely to be the unit of assignment to an end user or customer.

The results in Table 9 indicate there was little interest in IPv6 from UK LIRs before 2005, which was when the prospect of IPv4 exhaustion began to get attention. Since then, there has been a steady growth of IPv6 allocations. However the number of assignments has remained largely static and very low, generally less than 1% of the allocated IPv6 space. Assignment rates are actually decreasing because more IPv6 space has been allocated and very little of that extra space is being assigned. Aggregated IPv6 allocations have not changed much since 2005 and suggest only around 1,100 end users have been assigned IPv6 space in the last 10 years. This may well be explained by the lack of support for IPv6 in the consumer market. One encouraging sign however is a small increase in the number of direct IPv6 assignments larger than a /48. This indicates there is some interest from the non-consumer market. LIRs are assigning small amounts of IPv6 to customers or business units.
It should also be noted that even though only 0.5% of the /48s in UK IPv6 space are "in use", it does not necessarily mean these addresses are reachable on the public Internet. Some may be used internally or for interconnecting equipment at Internet exchanges. Others might be used for testing and evaluation as an LIR prepares for IPv6 adoption.

[34] http://www.oecd.org/internet/ieconomy/44953210.pdf

When IPv6 space has been assigned, it cannot be assumed that all those who have been issued with IPv6 space by an LIR are actively using it by enabling IPv6 on their equipment. Sometimes this might happen by accident. Vendors are now supplying systems that automatically use or try to use IPv6 by default. So IPv6 can sometimes be in use with minimal deliberate effort by the local network administrator to enable it. For instance most Apple products have a default setting to try Stateless Auto-Configuration (SLAAC) to

YEAR	Addresses **(in thousands of /48)**		% Visibility
	Allocated	**Reachable Prefixes**	
2008	214,041	139,198	65%
2009	216,334	140,771	65%
2010	222,429	208,798	94%
2011	229,966	211,157	92%
2012	258,277	216,400	84%
2013	280,822	223,216	79%
2014	330,826	235,799	71%

Table 9 - IPv6 Allocations & Availability

obtain an IPv6 address and usually have to be configured to disable that behaviour. If one of these devices finds a valid IPv6 prefix on the local network, they will automatically use that to configure a global IPv6 address for themselves that should in theory “just work”.

IPv6 REACHABILITY & ROUTE VISIBILITY

Reachability of the UK’s IPv6 address space as measured by the RIPE NCC’s Routing Information Service is shown in Table 10. It is encouraging that routing announcements for well over 80% of the address space have been made in recent years. Of course it does not necessarily mean all that space is actually reachable. Since very little IPv6 space has been assigned, there will be hardly any activated IPv6 space that’s being used by end users and organisations.

Year	Addresses (in thousands of /48s)					% In Use
	Allocated	Free	Assigned	Aggregated	Assigned + Aggregated	
2000	66	65				0.3%
2001	66	65				0.30%
2002	197	194	2		3	1.30%
2003	983	973	9	1	10	1.00%
2004	2,294	2,283	9	1	10	0.40%
2005	141,427	140,302	9	1,115	1,124	0.80%
2006	142,672	141,546	11	1,115	1,126	0.80%
2007	143,524	142,397	11	1,115	1,127	0.80%
2008	214,041	212,880	12	1,148	1,160	0.50%
2009	216,334	215,174	12	1,148	1,161	0.50%
2010	222,429	221,268	12	1,149	1,161	0.50%
2011	229,966	228,798	13	1,155	1,168	0.50%
2012	258,277	256,990	21	1,266	1,287	0.50%
2013	280,822	279,403	152	1,266	1,418	0.50%
2014	330,826	329,399	152	1,274	1,427	0.40%

Table 10 - IPv6 Reachability & Route Visibility

The level of route visibility has tailed off in 2013 and 2014. This is best explained by the recent increase in the number of IPv6 allocations. In part that may be due to the new IP address policy which requires LIRs to have an IPv6 allocation in order to get a final allocation of IPv4 space. The IPv4 run-out suggests LIRs are becoming more concerned about IPv6 and are applying for IPv6 allocations so that they can prepare for IPv6 adoption. They may not be ready yet to carry traffic for recently obtained IPv6 allocations even they are announcing routes for that space.
As before, Tables 9 and 10 report data for 2014 up to the time of writing (September 2014) rather than the end of the year as used for earlier entries.

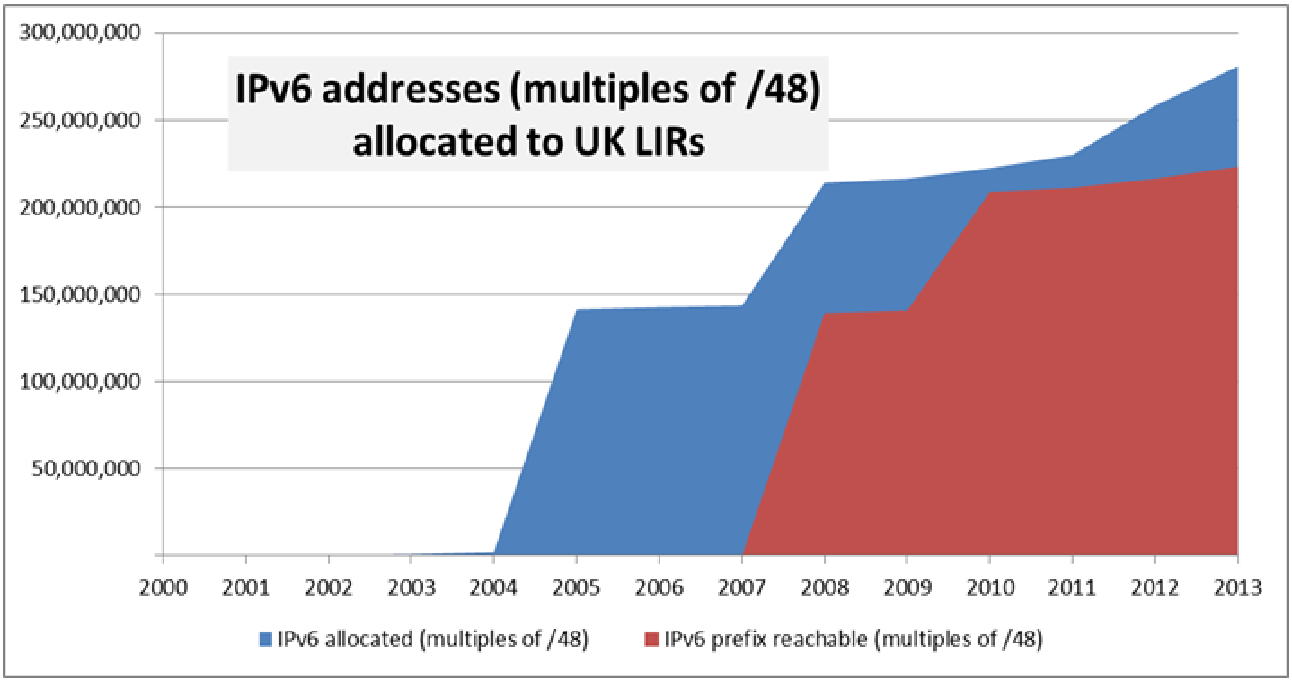

Figure 6 - IPv6 Allocations & Reachability