September 2024
Geoff Huston

Looking for 240/4 Addresses

If you look through the IANA's IPv4 address registry you will find a set of reservations which collectively are encompasses by the address prefix 240/4, and are annotated in the registry for "Future Use." These entries reference RFC 1112 section 4, which states: "Class E IP addresses, i.e., those with "1111" as their high-order four bits, are reserved for future addressing modes." This address prefix encompasses some 268,435,455 IPv4 addresses. From time it has prompted the obvious question: "If we have run out of available IPv4 addresses, then why are some quarter of a billion IPv4 addresses still sitting idle in an IANA registry waiting for an undefined Future Use?" Surely, if there was to be some "future addressing mode" to be defined in IPv4, then we would've done it by now. Why can we just add this pool of IP addresses into the all-but fully depleted pool of available IPv4 unicast addresses and relieve, to some small extent, some of the pressures that we have been experiencing with IPv4 depletion over the past decade?

The major points of discussion on this topic were recorded in a couple of Internet drafts from 2008. One of these, draft-wilson-class-e, advocated the redesignation of this address block for private use, extending the set of such local addresses to "assist in the IPv6 transition of larger networks who are using IPv4 in the context of a dual stack deployment." In such contexts it was reported that the reuse of network 10/8 was not an option because of existing use and potential address clashes [1918bis]. The use of 240/4 offered a more conventional method to connect Consumer Premises Equipment (CPE) Network Address Translators (NATs) to the network's border Carrier NATs without having to use more involved solutions such as Dual-Stack Lite (RFC 6333), NAT464 or 464XLAT (RFC 6877). Another reason why a private use context was advocated for this address prefix was that it was believed that many IP implementations had implemented this reservation of the 240/4 address block within the IP code itself within end hosts, discarding the processing of any IP packet that had an address from this prefix as either the source or destination address. The address prefix was unsuitable for general use while significant populations of host protocol stacks contained this discard code. The other draft, fuller-240space, advocated the reclassification of this address block as conventional unicast address space, noting that "given the current consumption rate, it is clear that the block should not be left unused."

Given that by 2008, when these drafts were submitted, the prospect of IPv4 address depletion was estimated to occur between 2010 and 2012, the discussion in the IETF turned to what would be the most productive use of the available time before the pools of available IPv4 addresses ran out. Consumption of IPv4 addresses was rising dramatically with the transition of mobile voice networks into mobile IP networks, and these networks were slow to adopt a dual stack mode of operation. By 2009 the annual IPv4 address consumption rate was rising to some 190M addresses per year (see Addressing 2009), so an additional pool of 268M addresses would apparently defer the inevitable IPv4 address exhaustion by only some 16 months or so. There was a prevalent view that the cumulative effort to update the hundreds of millions of IP hosts to accept IP addresses drawn from the address prefix 240/4 would probably take a comparable period, if not significantly longer, and such as effort would form a distraction to the overall objective to rapidly transition all hosts and networks to support IPv6 before we experienced acute IPv4 address exhaustion pressures.

The mobile industry had gathered an unprecedented level of momentum at this stage, so our collective attention then turned to dual stack transition mechanisms (at one point around 2010 there were more than 30 such IPv6 dual stack transition mechanisms being proposed!) and the associated effort to manage the remaining pools of available IPv4 addresses in a responsible manner took up far more attention than the plight of this little corner of address space. These proposals to revive 240/4, either as private use space or as general unicast space, quietly languished.

Languished, yes, but the topic still resurfaces from time to time.

Measuring Usability of 240/4

There have been number of exercises in recent years to see to what extent this address prefix is usable. A 2022 measurement exercise was reported in the RIPE Labs blog. An analysis of the traceroute data collected by the Atlas project indicated that at the time Amazon AWS (AS14618 and AS6509) were using this address block internally to address customer assets. Other instances of internal private use of this address prefix were also apparent at the time, as evidenced by the traceroute reports of router interface addresses reporting the use of this address prefix in these traceroutes. It appears that the proposal described in draft-wilson-class-e for use in private contexts had apparently not fallen on totally deaf ears!

The second part of this measurement experiment involved setting up a server using an IP address drawn from the 240/4 address and directing some 7,600 Atlas probes to perform a traceroute to this destination. They reported that some 34 probes were able to reach this server, and all of these probes were hosted in the AS701 (Verizon Business) network.

As a follow-up for this article, I have conducted a similar, but smaller scale experiment using Atlas in July 2024. Just 1 of the 190 tested probes reached a host server (hosted in AS29208, Quantcom, Czech Republic). This network peers with DE-CIX, and the one successful traceroute appeared to transit DE-CIX to reach the server. A larger re-run of this experiment, using 1,000 randomly selected probes from the Atlas collection did not fare any better. Of some 967 probes that responded, all of them reported failure in reaching this server.

In measurement terms, the Atlas network is of medium size with some 12,500 probes, but it is extremely flexible in terms of what the probes can be programmed to measure. At APNIC Labs we use a somewhat different measurement approach, based on enrolling users to perform a simple web object retrieval. This approach is less flexible, but by using an online ad network (Google Ads) to pass these retrieval tasks to end users, we are able to undertake measurements at a significantly larger scale both in volume and in coverage. The ad program is currently running with a total of some 25M ad impressions per day, which is significantly larger than the Atlas probe set. With this measurement system we can place a simple web object on a server (a 1x1 pixel image using a *png* image format, or a "blot"), and direct users to attempt to retrieve this object within the ad's script. For this measurement the server is addressed with a host address drawn from the 240/4 prefix and the server's network service provider advertises reachability to this server with a BGP announcement to its routing peers.

Before looking at the results of this measurement it may be useful understand the problem space in reaching a destination that has an address drawn from this 240/4 prefix. There are a number of potential issues in trying to use such an address, including:

- Routers may reject packets with a destination address from 240/4.
- The configuration of a network's routing environment may not accept a route advertisement for this prefix, or for more specifics of this prefix.
- CPE equipment that performs a NAT function may reject packets with source or destination addresses drawn from this prefix.
- For mobile networks other forms of network middleware, such as carrier grade NATs, may reject packets with source or destination addresses drawn from this prefix.
- End systems may reject packets with source or destination addresses drawn from this prefix.

Of the 788 total BGP peers for the Route Views and RIS systems, just 3 peers have propagated a route to a prefix drawn from 240/4, namely 242.242.0.0/16, originated by AS 8747 and propagated via AS 29208 (both networks are operated by Quantcom, in the Czech Republic). Another IPv4 network originated by the same AS, 109.235.180.0/24, is seen by a total of some 702 separate peers of Route Views and RIS, so it would be reasonable to infer that there is widespread BGP filtering taking place for the 242.242.0.0/16 route.

As a quick illustration of the issues in network reachability , lets compare two traceroutes from a Akamai Linode server located in Frankfurt, Germany. The first is to a conventional IPv4 host address:

```
$ traceroute 109.235.180.1
traceroute to 109.235.180.1 (109.235.180.1), 30 hops max, 60 byte packets
 1                    (10.210.2.210)              0.102 ms  0.025 ms  0.063 ms
 2                    (10.210.35.30)              0.215 ms  *         *
 3                    (10.210.32.1)               0.221 ms  0.423 ms  0.272 ms
 4  lo0-0.gw2.fra1.de.linode.com (139.162.129.102) 0.476 ms  0.326 ms  0.268 ms
 5  decix.quantcom.cz (80.81.192.217)             1.078 ms  0.995 ms  1.094 ms
 6  cz-prg-p1sit-be3.quantcom.cz (82.119.246.102)  7.832 ms  7.848 ms  7.815 ms
 7  * * *
```

The first three hops appear to pass through the Linode infrastructure, which uses network 10 to number its internal routers. The packets then pass out through a Linode egress router to the Frankfurt DE-CIX switching infrastructure, and is next seen at Quantcom's ingress, and from there into a Quantcom router in Prague.

```
$ traceroute 242.242.100.1
traceroute to 242.242.100.1 (242.242.100.1), 30 hops max, 60 byte packets
 1                (10.210.2.210)                  0.100 ms  0.050 ms  0.036 ms
 2                (10.210.35.30)                  0.669 ms  0.475 ms  0.290 ms
 3                (10.210.32.1)                   0.203 ms  0.200 ms  0.180 ms
 4  lo0-0.gw2.fra1.de.linode.com (139.162.129.102)  0.500 ms  0.398 ms  0.402 ms
 5  ae18.r02.fra03.ien.netarch.akamai.com (23.210.54.18) 0.587 ms  0.611 ms  0.477 ms
 6  * * *
```

The difference here is in hop 5, where the outbound packet is passed into the Akamai infrastructure rather than to DE-CIX. It is likely that these initial hops are following the internal *default* route. As it appears that Linode is not accepting a route to any prefix in 240/4 from DE-CIX, then the default outbound path points to an Akamai router, which discards the packet as there is no explicit route and no further *default* routes to follow.

Measurement Results for Testing Unicast Reachability for 240/4

Now let's look at the results of this experiment, shown at a country level in Table 1.

CC	Tests	Hits	Rate	CC Name
RO	1,313,452	42,814	3.2597%	Romania
CZ	985,470	15,162	1.5386%	Czech Republic
SK	198,416	532	0.2681%	Slovakia
RU	48,574	270	0.5559%	Russian Federation
AE	137,308	92	0.0670%	United Arab Emirates
US	4,539,376	34	0.0007%	United States of America
BH	32,302	8	0.0248%	Bahrain
GR	61,106	2	0.0033%	Greece
Total:	**130,298,477**	**58,914**	**0.0452%**	

Table 1 – 240/4 Accessibility by Country

Over a 14-day period from late August through early September we presented some 130 million unique clients with tests to users drawn from across the entire Internet. The test included the direction to fetch a blot from this server addressed within the block 242.242.0.0/16.

Our expectations were understandably low, as we have already noted that the interdomain routing space has not propagated the route for 242.242.0.0/16 very far, and just 3 RIS and RouteViews peers report visibility of this route prefix, compared to some 702 peers for other prefixes advertised from the same origin AS. On the other hand, this network directly peers with 913 other networks (https://stat.ripe.net/ui2013/AS29208#tabId=at-a-glance), so even if this route is not extensively propagated over transit networks so that it is seen at various route collectors, there is still a relatively rich domain of local propagation. Table 1 shows that access to an endpoint addresses in the 240/4 prefix is largely limited to hosts located in Romania and the Czech Republic where the host's network either peers directly with AS 29208 or is very closely connected.

There are a small number of more anomalous entries in this table. What is going on where we see just a handful of connections from networks that are geolocated to the United States, Bahrain and Greece? The most likely explanation is a geolocation failure, where the user in question is indeed located within the realm if visibility of 240/4 in Eastern Europe. It is entirely possible that there is some form of infrastructure route tunnelling or web proxy activity where the route is leaking via a route tunnel, but the very small hit counts would tend to suggest some form of individual customised network or application configuration as distinct from a whole-of-network tunnelling configuration.

We can increase the level of detail to look at the extent of propagation of access of this service by access network rather than by geolocated country of origin. The results of this accessibility measurement are shown in Table 2.

AS	CC	Tests	Hits	Rate	AS Name
9050	RO	40,952	34,990	85.4415%	RTD Bucharest
29208	CZ	6,898	5,236	75.9061%	QUANTCOM-AS
48161	RO	15,374	4,682	30.4540%	NG-AS Sos. Bucharest
28725	CZ	8,406	2,768	32.9289%	CETIN-AS
39668	RO	2,652	2,290	86.3499%	AS-INTERSAT_CT
25424	CZ	2,778	2,202	79.2657%	INEXT-CZ
6740	CZ	1,648	1,500	91.0194%	INEXT-CZ-ADA
209947	CZ	976	936	95.9016%	MWIFI
205619	CZ	816	790	96.8137%	ASVESNET
48926	CZ	1,196	616	51.5050%	PE3NY-AS
60895	SK	638	384	60.1881%	LEKOS
196952	CZ	342	328	95.9064%	ASBEZVANET
35725	RO	37,528	202	0.5383%	TELEKOM
57825	CZ	174	164	94.2529%	MORAVANYNET-AS
52029	CZ	320	150	46.8750%	ASNOVOSEDLY
34315	CZ	132	132	100.0000%	MAXNET-AS
20485	RU	110	108	98.1818%	TRANSTELECOM MOSCOW
214529	RO	106	106	100.0000%	DSNET
12905	SK	106	106	100.0000%	ACS-SK-AS,
205275	RO	298	106	35.5705%	ROMARG HOSTING
206382	RO	142	100	70.4225%	NEXTSTART
34560	RO	100	92	92.0000%	SOFTEX-AS
197083	CZ	86	86	100.0000%	K2ATMITEC
199405	CZ	98	84	85.7143%	OSLAVANY-NET-AS
44081	RO	148	84	56.7568%	YUL-PRO-INTERNET-RASNOV-AS
138915	AE	124	80	64.5161%	KAPOKU Cloud
211137	CZ	74	72	97.2973%	ISPSERVICES
206438	CZ	182	60	32.9670%	MXNET-AS

60533	SK	136	40	29.4118%	SATELIX
49107	RU	34	34	100.0000%	TELKO-AS
207913	RO	32	32	100.0000%	CLAR-TELEVISON-SRL
205400	CZ	186	32	17.2043%	VIVOCONNECTION
57180	RO	32	30	93.7500%	STAR-NET-ALBA-AS
203574	RO	28	28	100.0000%	CONECTX-AS
210713	RO	26	26	100.0000%	CORESI-NETLINK
35512	RO	26	26	100.0000%	TELEMEDIA-AS
6856	RU	24	24	100.0000%	IC-VORONEZH-AS
61403	RU	24	24	100.0000%	SEVER-TELECOM-CHER
138915	US	154	20	12.9870%	KAOPU CLOUD
60840	RU	20	20	100.0000%	TELECOMSERVICEVRN
57411	RU	16	16	100.0000%	NOVOTEHNIKS-AS
47165	RU	14	14	100.0000%	OMKC-AS
62642	US	312	14	4.4872%	BIGLEAF
210616	RU	30	12	40.0000%	SIBMEDVED-AS
5384	AE	17,218	12	0.0697%	EMIRATES-INTERNET
51102	RO	12	10	83.3333%	IMPATT-AS
41087	RO	10	10	100.0000%	ROMPRIX-AS
47236	RU	24	6	25.0000%	CITYLINK-AS
56791	RU	6	6	100.0000%	CT-AS
138915	BH	4	4	100.0000%	KAOPU
13150	CZ	6	4	66.6667%	CATON
5416	BH	9,642	4	0.0415%	Internet Service Provider
49055	RU	2	2	100.0000%	NEWIT-AS
38949	SK	6	2	33.3333%	TRESTEL
57294	CZ	204	2	0.9804%	INTERNET_EXPER
41719	RU	2	2	100.0000%	SKTVSPEKTR-AS
60042	RU	2	2	100.0000%	ONTELECOM-AS
13150	GR	2	2	100.0000%	CATON

Table 2 – 240/4 Accessibility by Network

The first sever networks in Table 2 appear to have accepted the route to 242.242.0.0/16 into their network, and also have both a sizeable user base and have a high proportion of host platforms that also support communications with server addresses drawn from 240/4.

However, within the networks AS 48161 and AS 28725 the lower 240/4 hit rate of some 30% of tests suggests that there is a further factor here. A possible explanation lies in variations in the CPE equipment used in the interface between the client network and the service provider, on in the gateway equipment used to connect the network to the Internet. The user may also be using a "clean filter" DNS resolver service where a name will not be resolved if the name is on some exception list, or, as may be the case here, the IP address is in a reserved address block.

There is also a "long tail" in this table of more remote networks where just one or two tested users were seen to perform a successful fetch of this test blot. A possible explanation of these isolated anomalies may lie in the use of web proxy agents in the client-side network, as the IP address used to access this server are using networks that apparently have no route to this 242.242.0.0/16 prefix.

Observations

Is the address prefix 240/4 usable in a global unicast sense in the same way as all other IPv4 global unicast addresses? With a measured reachability rate drawn from across much of the Internet at just 0.0452%, it's clear that it is not a generally reachable prefix, which implies that it is just not a useful address.

The intent of a unique unicast IP address in the Internet model is that any other host is capable of sending packets to it and receiving packets from it. It's clear from these measurements that this is just not happening for this test server.

In the most general terms, there are three causes of reachability failure:

- Network routing, where the routing system does not propagate a route for a prefix drawn from 240/4
- Host filtering, where the host performs some form of address check on outgoing and possibly incoming packets and will discard IP packets with destination and/or source addresses drawn from 240/4.
- Middleware filtering, where various forms of network middleware, generally NATs, will not process a packet being directed to a 240/4 destination address.

The very limited propagation of the route 242.242.0.0/16 indicates that there is widespread practice of route filtering. This may be due to a particular block for the 240/4 prefix, or a more general route "bogon filter" where routes drawn from IANA reserved address blocks are rejected.

The very high hit rates in some networks in Table 2 appear to indicate that host filtering is not a major block for using addresses from this prefix.

The lower hit rates or around 30% for some networks pose an interesting question. Is filtering of this prefix a property of some consumer premises equipment (CPE) used by clients? Is this a behaviour of some network-level Carrier Grade NATS (CGN) used by some networks? In the case of Germany, Austria, Hungary and Romania, all of which are "close" to the Czech Republic there is a reasonable level of IPv6 adoption. Are the dual stack transition mechanisms being used in some of these networks dropping IPv4 packets addressed to this 240/4 address?

Should we do anything about this? Or not?

The Internet is big enough that any form of a coordinated change, such as a "flag day," to enable access to 240/4 in networks, middleware and hosts is just not a realistic approach. For adoption in today's Internet any technology must be deployable in a piecemeal fashion. However, in general, piecemeal adoption is generally motivated by the perception of early adopter rewards. These rewards motivate additional adopters and momentum gathers, if all is going well. But in the case of addresses the extremely limited visibility of services that use this address, and the limited ability of other users to access services that are addressed from this prefix suggest that there is an early adopter penalty rather than any form of reward. That tends to make piecemeal adoption for the use of this prefix as a general-purpose unicast address a forbidding proposition.

The 2008 advice contained in the wilson-class-e draft (where I'll own up to being a co-author), which advocated the designation of this address space as "Private Use" seemed to me to be the most sensible approach then, and now, some 26 years later the approach still makes some sense to me. Such a use allows a network operator to use these addresses in a controlled environment where the operator can assure themselves that the addresses are fully functional in their desired limited context of use. But in many ways, there is little that prevents a network operator from using addresses drawn from 240/4 in their internal environment already, as has been reported in traceroute data by Amazon collected by RIPE Atlas. Such a private use will not clash with any existing unicast addresses. This form of entirely private use of this IANA-reserved address block is, pragmatically speaking, already an option today and doesn't require any particular IANA re-designation of the address block, so the obvious question is why bother with an IANA re-designation in any case.

It appears that the status quo is entirely adequate for the 240/4 address prefix!

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net