



What do I need for Peering

This section of the Toolbox describes what a network operator needs before embarking on their peering journey.

Recap

A network operator embarking on their peering journey now knows:

- What their goals as a service provider are
- How the Internet Ecosystem is structured
- What Peering is
- What Transit is
- How Peering helps with operating costs, reducing latency, increasing bandwidth, improving relationships
- What Private Peering is
- What Public Peering is
- What an Internet Exchange Point is

What needs to be done next to start the journey? This is covered in the following section.

Internet Resources

The Peering Toolbox mentioned [elsewhere](#) that Network Operators required Internet Resources.

When the industry talks about Internet Resources, we mean our own independent IPv4 address space, IPv6 address space, and Autonomous System Number (ASN). A Network Operator needs their own independent Internet Resources to be able to take part in the global peering ecosystem.

The sections below describe each:

- [IPv4 address space](#)
- [IPv6 address space](#)
- [Autonomous System Numbers](#)

IPv4 Address Space

To take part in the peering ecosystem, a Network Operator will quite likely require its own independent IPv4 address space (and definitely will if it does not make any use of IPv6). This address space needs to be globally routable (known as public address space)

Background

IPv4 addresses have been used since the early days of the global Internet as we know it today. IPv4 addresses are distributed by the five Regional Internet Registries ([AfriNIC](#), [APNIC](#), [ARIN](#), [LACNIC](#), [RIPE NCC](#)).

To all intents and purposes the global Internet has exhausted the IPv4 address supply, with only AfriNIC and APNIC having limited IPv4 resources available, and only for new members now.

The other three Regional Internet Registries have no IPv4 address space to distribute, although from time to time may have limited amount of address space available, reclaimed from network operators who no longer require it.

Obtaining IPv4 Address space

A tutorial on obtaining IPv4 address space is well beyond the scope of the Peering Toolbox. Each RIR website has lots of information on obtaining IPv4 address space, whether directly from the RIR itself or by their respective transfer policies (IPv4 addresses transferred from one RIR member to another), and the RIR websites should be consulted for further information and guidance.

Using IPv4 Address space

Network Operators who have been operational for many years will quite likely have sufficient IPv4 address space for all their requirements, whether they received this address space from InterNIC (prior to the existence of the RIRs) or from one of the RIRs directly.

Describing how to plan the usage of IPv4 address space within a Network Operator is beyond the Peering Toolbox scope, and would in any case be very specific to the nature of the business, the type of end customers, and so on. However, it is essential that every Network Operator reserves sufficient IPv4 address space for their own infrastructure needs, be it their routers, servers, Network Operations Centre, and so on. Private address space cannot be used for public interconnects such as for peering or transit, as globally only public address space is can be routed (and is guaranteed to be unique).

Private IPv4 Address space

Private IPv4 address space is designed to be used in private network infrastructure, and network infrastructure that is not globally routed. The well known private address space includes:

- **10.0.0.0/8**

- **172.16.0.0/12**
- **192.168.0.0/16**

which are typically used by end-sites for devices which do not need to access the global Internet. Private address space is commonly used by end-users and enterprises for their internal addressing needs when it is no longer possible to obtain public address space. So that these internal devices can get Internet access, an intermediate device uses a technique known as Network Address Translation (NAT) to translate these private addresses into a public address pool that has been made available to the end site.

The other private address space is **100.64.0.0/10**, which is known as the Carrier Grade NAT (CGNAT) address pool. This is used by network operators (eg mobile providers) to provide addressing to their infrastructure so as not to confuse it with the three private address blocks that would typically be used by their customers. This CGNAT pool is then also translated by the operators CGNAT devices to public IPv4 address space the operator has available to it.

Be aware that none of these 4 blocks can be used on or announced to the public Internet, and all IP packets sourced from these address blocks must either be translated into public IPv4 addresses or blocked from reaching the public Internet.

IPv6 Address Space

To take part in the peering ecosystem, a Network Operator will quite likely require its own independent IPv6 address space (and definitely will if it does not make any use of IPv4). This address space needs to be globally routable (known as public address space)

Background

IPv6 was developed in the mid 1990s to replace IPv4. IPv4 was not designed with the current global Internet infrastructure in mind, being restricted in address size to realistically only a few million hosts. Not nearly sufficient to cover a human population approaching 8 billion, or an Internet requiring dozens if not hundreds of IP addresses per human.

IPv6 addresses are distributed by the five Regional Internet Registries ([AfrinIC](#), [APNIC](#), [ARIN](#), [LACNIC](#), [RIPE NCC](#)).

While each RIR has developed its own policies for distributing IPv6 address space, the overall concept is straightforward:

- Network Operators receive a /32
- End-sites receive a /48

End-sites don't participate in the peering ecosystem as noted elsewhere. The expectation and advice is that Network Operators distributed /48 address blocks to their end-site customers out of their /32 - and when this /32 is used up, they return to their RIR for a larger allocation (size depends on the RIR policy in place at the time of request).

Using IPv6 address space

Given the vastness of the size of the IPv6 address pool, and the large initial allocation made to Network Operators, constructing an address plan using IPv6 is less constrained than it is for IPv4.

As with IPv4, advising on the design of an address plan is beyond the scope of the Peering Toolbox. However, common best practice by Network Operators generally results in reserving a /48 out of their /32 for their own network infrastructure, and assigning the remaining 65535 /48s to their customers in a structured and meaningful manner that makes traffic engineering easy to manage. Examples of IP address planning are noted in this [IPv6 Address Planning](#) presentation.

Private IPv6 address space

There is no direct equivalent of private address space in IPv6. However, the address block FC00::/7 is set aside for what are known as Unique Local Addressing, and the application is similar to what was originally intended for the IPv4 private addresses, namely used for devices that are on isolated networks or do not need to communicate with the global Internet.

Autonomous System Numbers

To take part in the peering ecosystem, a Network Operator will require its own Autonomous System Number (ASN). The ASN is a unique identifier for that network, used by the Border Gateway Protocol (BGP).

Background

The ASN is a fundamental requirement for BGP, and is the globally unique identifier for the network using it. Like IP address space, ASNs are distributed by the five Regional Internet Registries ([AfriNIC](#), [APNIC](#), [ARIN](#), [LACNIC](#), [RIPE NCC](#)), or may have been assigned by the InterNIC prior to the existence of the RIRs.

As with IP address space, each RIR has developed their own policies about the distribution of ASNs. But the overall concept is that if an entity needs to connect to more than one other independent network, an ASN is granted. Connecting to more than one other network means that a dynamic routing protocol is required (to make choices between the two available paths), and that protocol is BGP which requires an ASN to function.

AS numbers for use on the public Internet range from 1 through to 458751 (with exceptions for documentation, private use, and those reserved between 65552 and 131071). ASNs from 458752 up to 4199999999 are reserved, with the remainder of the range above that intended for private use as well.

Using ASNs

An autonomous system defines a network with a unique routing policy. So a network operator would

use one ASN for all the network infrastructure which has a single routing policy towards other autonomous systems.

It's quite common for Network Operators to use more than one ASN in their network though. Each ASN will have different policies, for example:

- an ASN for a transit network
- an ASN for an access network (different peering requirements from the transit network)
- an ASN for a datacentre (different peering requirements from other networks)
- an ASN for its mobile infrastructure

and so on. In the advanced section of the Peering Toolbox we'll examine scenarios where an operator may need multiple ASNs to achieve their most optimum peering and transit goals.

Private ASNs

There are, as mentioned earlier, private ASNs. These are for use internally on infrastructure, usually for operational requirements, or on infrastructure that may benefit from the use of BGP and have a different policy from the parent AS but not have any need for that policy to be globally visible.

Private ASNs range from 64512 to 65534 and 4200000000 to 4294967294.

The Internet Routing Registry

The Internet Routing Registry (IRR) is used to document policy of autonomous networks taking part in the global Internet.

There is no one system that is the IRR, but is made up of several components. Each of the 5 Regional Internet Registries ([Afrinic](#), [APNIC](#), [ARIN](#), [LACNIC](#), [RIPE NCC](#)) runs their own component of the IRR, and this is done as a service to their members.

For operators who are not members of any RIR, likely those who received their IPv4 address space from InterNIC (prior to the existence of the RIRs), the component of the IRR they use is the [RADB](#) (Routing Arbiter Database) operated by Merit Network. Access to the RADB is a subscription service, compared with the RIR component of the IRR which is available to members as part of their membership fee. Note that most of the information contained in the RADB is considered inaccurate or out of date, so Network Operators would be recommended to treat content there with due caution.

Note also that some major operators and Tier-1s also operate their own IRR - but they also refer to information stored in the RIRs' instances of the IRR as well.

Our advice is as follows:

- Network Operators holding IP address distributed by an RIR should only use their RIR's instance of the Internet Routing Registry
- Network Operators holding IP address distributed by InterNIC (pre-existing the RIRs) means the Network Operator has to use RADB unless their RIR has a policy permitting them to use the RIR's instance of the IRR.

It is beyond the scope of the Peering Toolbox to provide a detailed tutorial about the operation of the Internet Routing Registry. However, we have to highlight the three key objects that all network operators need to be aware of, and one that is more or less mandatory in today's Internet. The following sections describe:

- the Route Object
- the AS Object
- the AS Set

Route Object

The Route Object documents which Autonomous System is originating the route listed. It is required by many major transit providers because they build their customer and peer filter based on the route-objects listed in the IRR. Operators will refer to at least the 5 RIR routing registries and the RADB to check for route-objects. Those who run their own IRR instance will generally check there first before consulting with the IRR instances run elsewhere.

A typical IPv4 route object may look like this:

```
route:      202.144.128.0/20
descr:      DRUKNET-BLOCK-A1
country:    BT
notify:     ioc@bt.bt
mnt-by:     MAINT-BT-DRUKNET
origin:     AS18024
last-modified: 2018-09-18T09:37:40Z
source:     APNIC
```

with the IPv6 version looking like this:

```
route6:     2405:D000::/32
descr:      DRUKNET-IPV6-BLOCK
origin:     AS17660
notify:     netops@bt.bt
mnt-by:     MAINT-BT-DRUKNET
last-modified: 2010-07-21T03:46:02Z
source:     APNIC
```

The key ingredients of a route-object are:

- route/route6: identifying the IP address block
- descr: describing what the block is about (useful but not essential)
- country: which country it is used in (can help with geolocation)
- notify: who to notify if anything with the object changes
- maint-by: who the maintainer of the object is
- origin: the ASN which is originating this address block
- last-modified: when the object was last changed
- source: which instance of the IRR provided the data

Operators who build their BGP filters based on the contents of the IRR will search all route-objects for their peer ASNs, and only accept BGP announcements from peers (and customers) which have matching and correct route-objects. No route-object or an incorrect route-object, and the BGP announcement will not be accepted.

It is vitally important that all Network Operators taking part in peering ensure that their route-objects are always up to date. Every prefix that is announced **must** have a matching route-object.

AS Object

The AS Object documents a Network Operator's peering policy with other Autonomous Systems. The AS Object lists network information, contact information, routes announced to neighbouring autonomous systems, and routes accepted from neighbouring autonomous systems.

Some operators pay close attention to what is contained in the AS Object. Some operators will configure their border router BGP policy based on what is listed in the AS Object.

A typical AS object may look like this:

```
aut-num:      AS17660
as-name:      DRUKNET-AS
descr:        DrukNet ISP, Bhutan Telecom, Thimphu
country:      BT
org:          ORG-BTL2-AP
import:        from AS6461      action pref=100;      accept ANY
export:        to AS6461        announce AS-DRUKNET-TRANSIT
import:        from AS2914      action pref=150;      accept ANY
export:        to AS2914        announce AS-DRUKNET-TRANSIT
<snip>
import:        from AS135666    action pref=250;      accept AS135666
export:        to AS135666      announce {0.0.0.0/0} AS-DRUKNET-TRANSIT
admin-c:      DN01-AP
tech-c:       DN01-AP
notify:       netops@bt.bt
mnt-irt:      IRT-BTTELECOM-BT
mnt-by:       APNIC-HM
mnt-lower:    MAINT-BT-DRUKNET
mnt-routes:   MAINT-BT-DRUKNET
last-modified: 2019-06-09T22:40:10Z
source:       APNIC
```

The key components here are the import and export statements. These use Routing Policy Specification Language ([RPSL](#)) to describe the BGP policies used by the operator. They bear some resemblance to BGP's attributes but there isn't intended to be a direct 1:1 correspondence.

Our advice is to at least create a minimal AS Object that describes the EBGp relationship that any Network Operator has with other autonomous systems.

Creation of an AS Object can be done via the RIR's member portal - consult the relevant RIR for more information.

AS Set

Route Origin Authorisation

The Peering Database

References

This content is sourced from many contributors, including:

- [Value of Peering Presentation](#) - Philip Smith
- Network Startup Resource Center
- Input from Mark Tinka, Kurt Erik Lindqvist, etc

[Back to Home page](#)

From:

<https://bgp4all.com/pfs/> - Philip Smith's Internet Development Site

Permanent link:

<https://bgp4all.com/pfs/peering-toolbox/how-to-peer?rev=1651750502>

Last update: **2022/05/05 11:35**

