
S1E14 - the Border Gateway Protocol

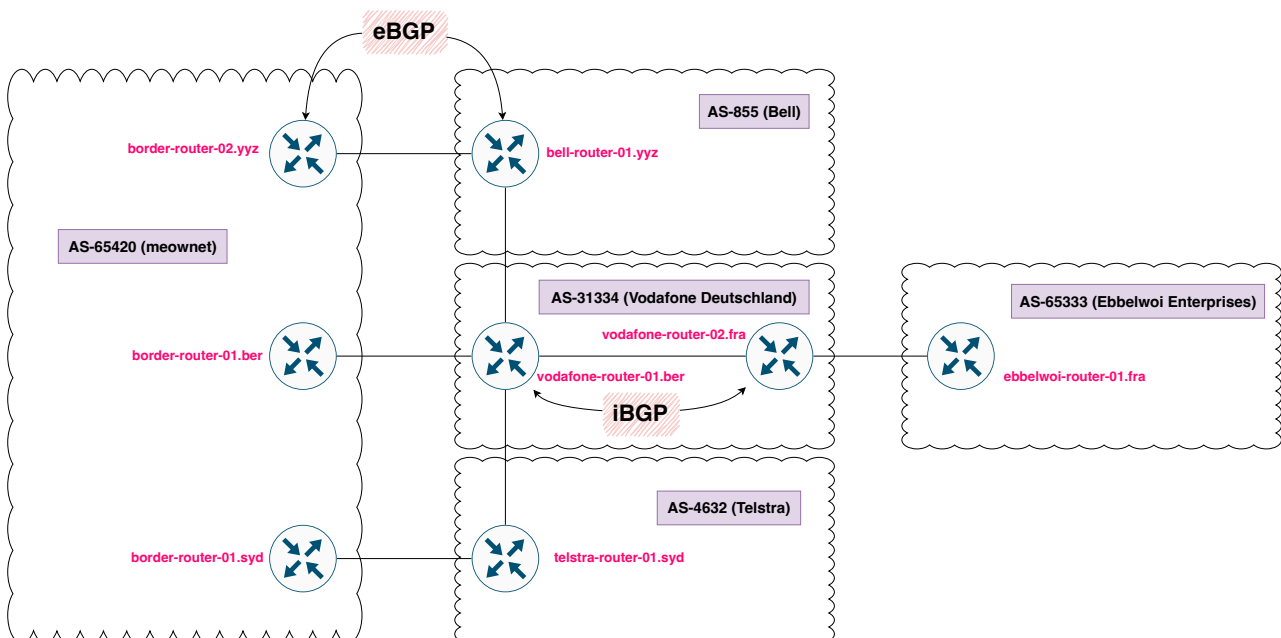
Nicholas Morrison nick@nanocat.net



The Border Gateway Protocol

- BGP is the Internet's routing protocol
- Autonomous System (AS): a network (or group of networks) managed and controlled by a single entity.
- Traditionally used for:
 - exchanging network information between different AS's => external BGP
 - distributing reachability information within an AS => internal BGP
- Lately also:
 - distributing reachability information within a datacentre (EVPN, MPLS)
 - high-speed failure detection and re-routing (with the help of BFD (Bi-directional Forwarding Detection))

Illustrated



- Ebbelwoi -> Meownet (Sydney) AS-PATH: 65333 31334 4632 65420
- Meownet (Toronto) -> Ebbelwoi AS-PATH: 65420 855 31334 65333
- Meownet (Berlin) -> Meownet (Sydney) AS-PATH: 65420 31334 4632 65420

IANA

- Internet Assigned Numbers Authority
- AS number allocations (and IP address allocations, and DNS TLDs) are governed by IANA
 - via the five Regional Internet Registries (RIRs)
 - * AFRINIC (Africa)
 - * ARIN (Antarctica, Canada, some of the Caribbean, US)
 - * APNIC (East Asia, Oceania, South Asia, Southeast Asia)

- * LACNIC (most of the Caribbean, Latin America)
- * RIPE (Europe, Central Asia, Russia, West Asia)
- via Local Internet Registries (LIRs)
 - * of which there are billions

Who uses BGP?

- ISPs
- IXs
- Organisations, who have their own public and portable IPv4 and IPv6 allocations
- potentially any company running a datacentre
 - modern layer 3 leaf-spine networks can optionally use BGP as an IGP

A brief history of BGP

- inspired by the Exterior Gateway Protocol (EGP), originally defined in RFC 827 (October 1982)
- "It is proposed to establish a standard for Gateway to Gateway procedures that allow the Gateways to be mutually suspicious."
- "1 INTRODUCTION
The DARPA Catenet is expected to be a continuously expanding system, with more and more hosts on more and more networks participating in it. Of course, this will require more and more gateways. In the past, such expansion has taken place in a relatively unstructured manner. New gateways, often containing radically different software than the existing gateways, would be added and would immediately begin participating in the common routing algorithm via the GGP protocol. However, as the internet grows larger and larger, this simple method of expansion becomes less and less feasible."
- (did you know: the internet used to be called the Catenet? <https://www.rfc-editor.org/in-notes/ien/ien48.txt>)

A brief history of BGP (cont...)

- Original BGP version defined in RFC 1105 (released June 1989)
- Updates defined in RFC 1163 (released June 1990)
- Version 3 defined in RFC 1267 (released October 1991)
- Version 4 defined in RFC 4271 (released January 2006)
- Multiprotocol BGP extension (MP-BGP) defined in 1998
 - Allows BGP to distribute different types (address families) of data, not just IPv4 prefixes
 - * <https://www.iana.org/assignments/address-family-numbers/address-family-numbers.xhtml>
 - * <https://www.iana.org/assignments/safi-namespace/safi-namespace.xhtml>

What's special about BGP?

- it makes connections to peers over TCP/IP (tcp/179)
 - as opposed to OSPF and RIP, which use IP multicast
- it can't discover peers
 - you have to manually configure them
- it has very strong filtering capabilities
 - route-maps, prefix-lists and access-lists
- it has a "tagging" system called BGP Communities
 - eg the BGP Community #666 means "please blackhole this prefix"
- it considers AS Path Length (among other things) when selecting the best path

Basic BGP configuration

```
!  
router bgp 64512  
  router-id 10.0.0.1  
  neighbor 10.0.0.2 remote-as 64838  
  neighbor 10.0.0.2 description Amazing ISP GmbH  
  network 192.168.0.0/24  
!  
  
r1#show ip bgp summary  
r1#show ip route bgp  
r1#show ip bgp neighbor 10.0.0.2 advertised-prefixes
```