



# One administrative domain using BGP

idr WG Interim (Apr, 2023)

[draft-uttaro-idr-bgp-oad@ietf.org](mailto:draft-uttaro-idr-bgp-oad@ietf.org)



# Problem statement

- Networks under a common administrative entity can span multiple autonomous systems
  - Evolution
  - Mergers/acquisitions
  - Partnerships, business relationships
  - ...
- Desire for better policy control and influence on route selection by carrying otherwise-stripped attributes across EBGp boundary



## Basic approach

- Define a new BGP peering type EBGP-OAD in addition to EBGP and IBGP
- The peering type is configured explicitly on both sides of a session
  - I.e. no capability negotiation
- Follow EBGP processing rules for the most part
  - E.g. AS\_PATH procedures remain the same
- In addition, IBGP and non-transitive attributes can be optionally exchanged
- Wherever required, enumerate specific rules for attribute processing
  - E.g. no propagation of ORIGINATOR\_ID and CLUSTER\_LIST attributes across EBGP-OAD peerings



# Misconfiguration

- If the two sides of a session are not configured with the same peering type
  - Attributes get ignored
  - Standard EBGP processing applies
  - Session remains operational



# Todo

Attribute enumeration from a transitivity standpoint	Almost done
Consider a BGP capability to negotiate peering type	
Consider new role definition as per RFC 9234	
(Optional) configuration to define all ASes in an OAD set to make consistent decisions (MED e.g.)?	
Consider adding a section for IXP/RS operation	
Add text on interworking between OAD and draft-ietf-idr-bgp-attribute-announcement	
Add text on guidelines for new attributes	



## Other points

Using ATTR\_SET to carry attributes

Disjoint ASes are out of scope.



## Next steps

- Request review and feedback