

upna


Universidad Pública de Navarra  
Nafarroako Unibertsitate Publikoa

**Redes de Nueva Generación**  
*Área de Ingeniería Telemática*

# BGP-4

# Enrutamiento jerárquico

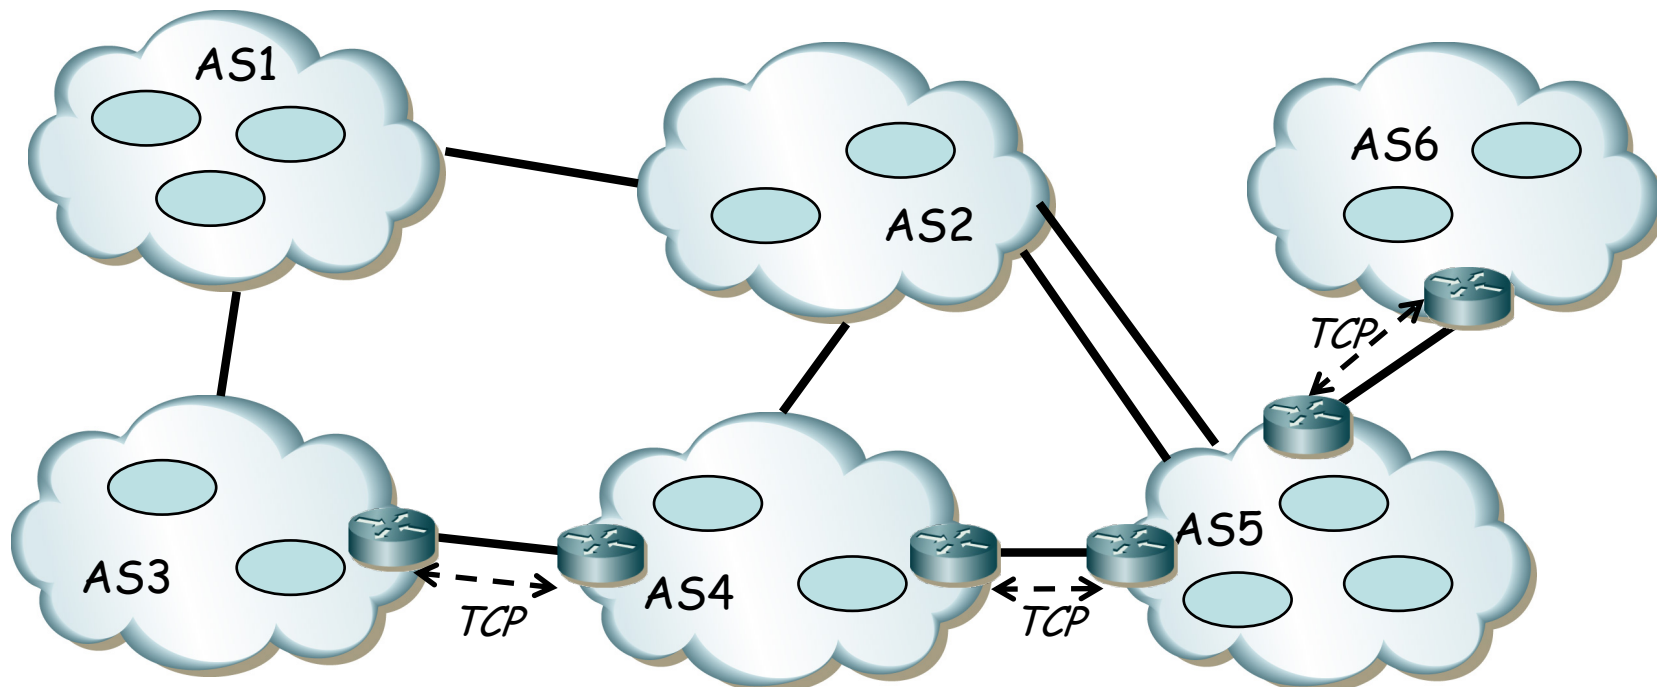
- ¿Un solo grafo para toda la Internet?
  - Problemas de escala
  - Problemas de coordinación (¿métrica?)
- Enrutamiento jerárquico
  - IGP: Interior Gateway Protocol
  - EGP: Exterior Gateway Protocol
  - Interior/exterior respecto a “sistemas autónomos” (*Autonomous Systems*)
  - *“An AS is a connected group of one or more IP prefixes run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy” (BCP 6)*



# BGP: Introducción

# BGP


- *Border Gateway Protocol*
- BGP-4, RFC 4271
- BGP-4 primera versión classless
- Protocolo Interdomain estándar *de facto*
- Comunicación fiable mediante conexión TCP entre routers adyacentes
- Puerto 179



# BGP

## Path Vector

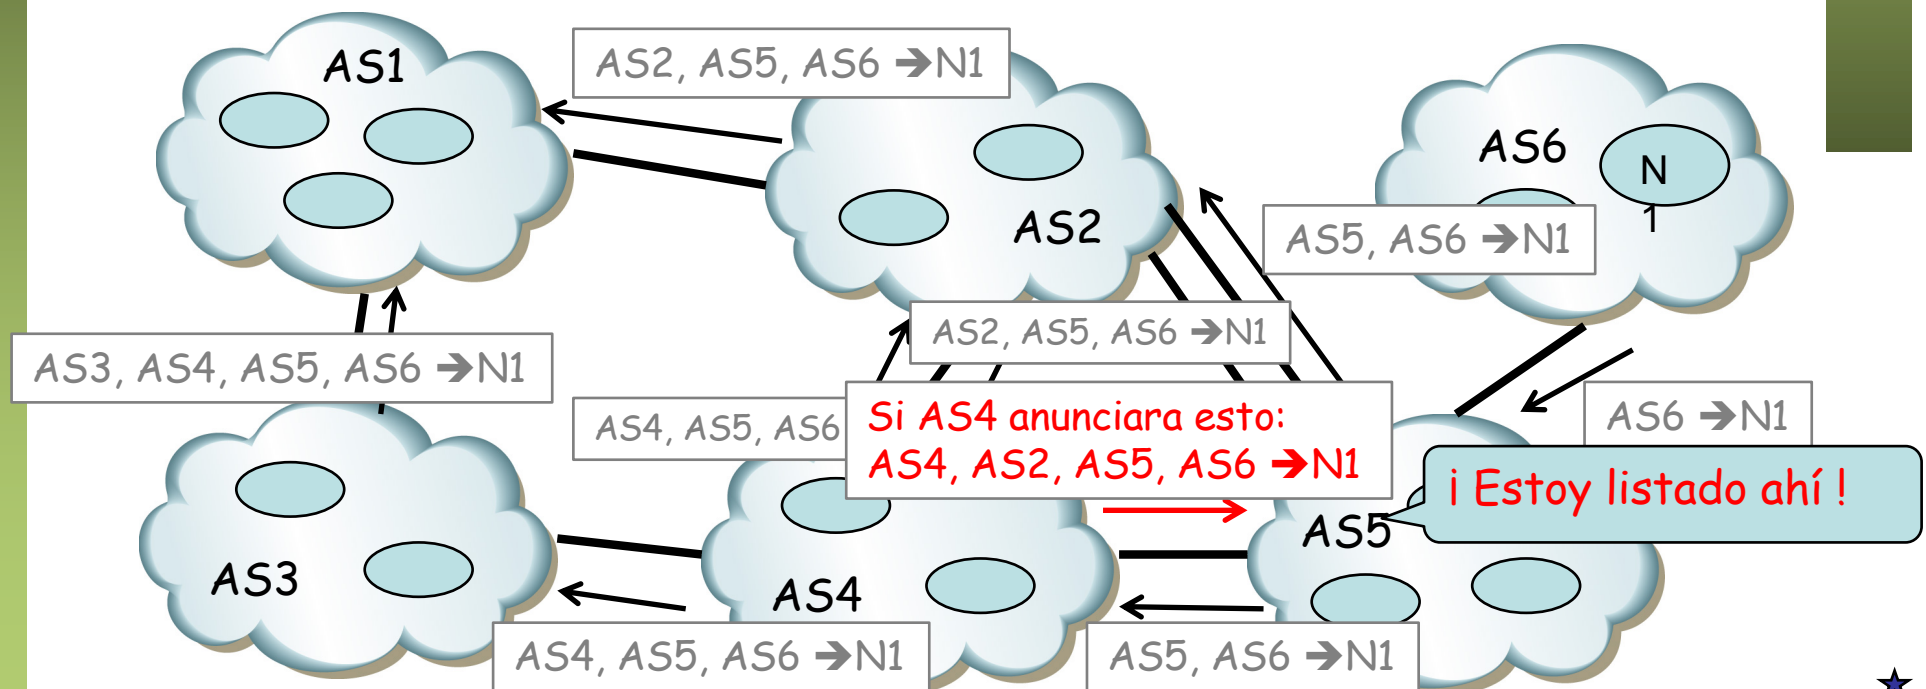
- Calcula caminos a prefijos
- Como DV recibe de vecinos, calcula sus rutas y envía a vecinos
- En vez de métrica anuncia la lista de AS en cada camino (. . .)
- Por defecto elige el camino que pasa por menor número de ASs



# BGP

## Path Vector


- Anunciar el camino permite evitar los ciclos
- El menor número de ASs no quiere decir que sea el menor número de saltos por routers



# eBGP vs iBGP

# Peering en BGP


- Los *peers* de un proceso BGP pueden estar:
  - En otro AS: *external peer*  $\Rightarrow$  **eBGP**
  - En el mismo AS: *internal peer*  $\Rightarrow$  **iBGP**
- (...)





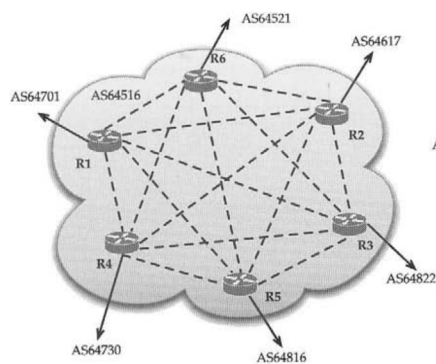
# Peering en BGP

- En el mismo AS el *peering* iBGP forma una malla porque...
- No se pasan por iBGP prefijos aprendidos por iBGP
- Reconoce si es del mismo AS porque en el OPEN anuncia el ASN
- No interesa difundir todas las rutas al IGP (escalabilidad)
- iBGP permite que otros ASBRs aprendan los prefijos a anunciar
- El ASN se añade a la ruta al hacer anuncio a otro *eBGP*

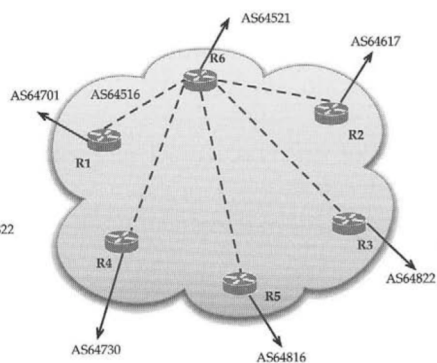


# Route Reflectors


- Problema de escalabilidad en iBGP debido al full-mesh
- RFC 4456 “BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)”
- En lugar de *full-mesh* conectan todos con el RR del *cluster*
- El RR sí reenvía rutas aprendidas por iBGP
- Un RR puede ser un cliente para otro RR



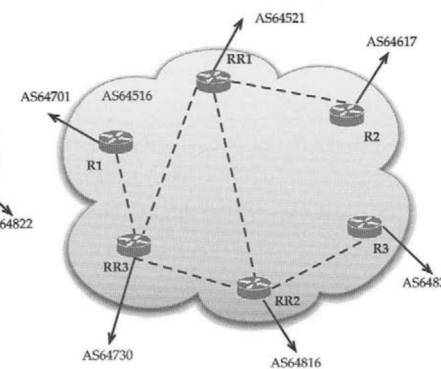
(a) Full mesh



(b) With one RR (R6)




(c) With two RRs



(d) With three RRs

# Confederations

- Otra solución al problema de escalabilidad de iBGP
- Internamente el AS se divide en sub-ASs, por ejemplo con ASNs privados
- Externamente se anuncia como un solo AS (el identificador de la Confederación)
- Internamente hay *full-mesh* en cada sub-AS pero no globalmente al AS
- La estructura interna no es visible externamente



# Atributos en BGP

# Path Attributes

- Son características de una ruta BGP, incluidos en el anuncio de la misma

## **Tipos según se soporten:**

- *Well-known: mandatory* (en update) o *discretionary*
- *Optional: transitive* o *nontransitive*

"*well-known*" : Debe soportarlo

"*Optional*" : No está obligado a soportarlo

"*mandatory*" : Debe aparecer en los mensajes

"*discretionary*" : Puede no aparecer en los mensajes

"*Transitive*" : Debe reenviarlo

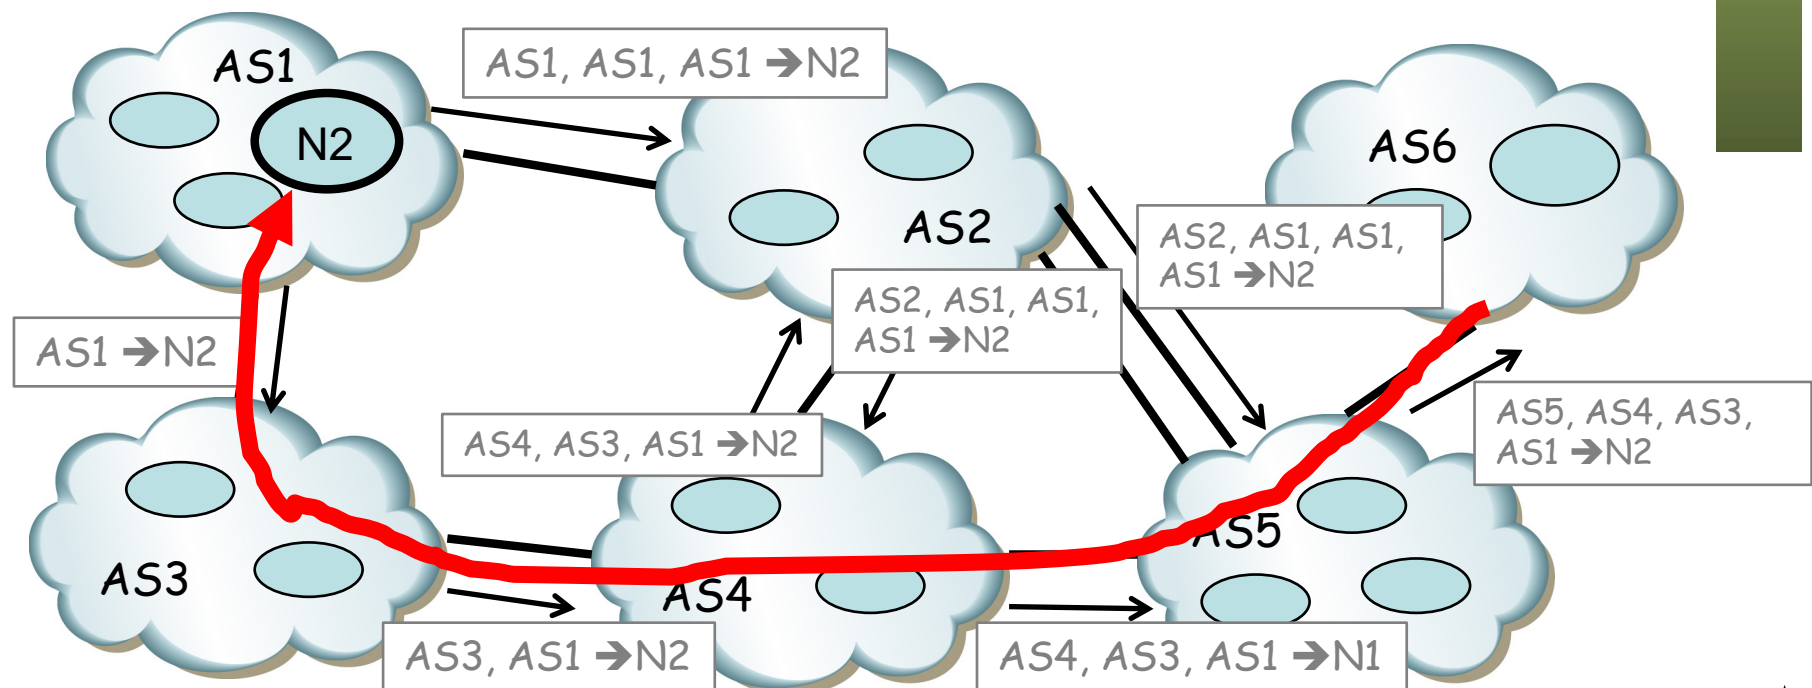
"*Nontransitive*" : No debe reenviarlo



# Path Attributes

## AS\_PATH (well-known mandatory)

- Secuencia de ASs hasta el destino
- Al mandar un *update* por eBGP se añade el ASN a la secuencia
- Si se manda por iBGP no se añade el ASN
- *AS path prepending*: añadir el ASN *más veces* para desalentar usar este camino (. . .)



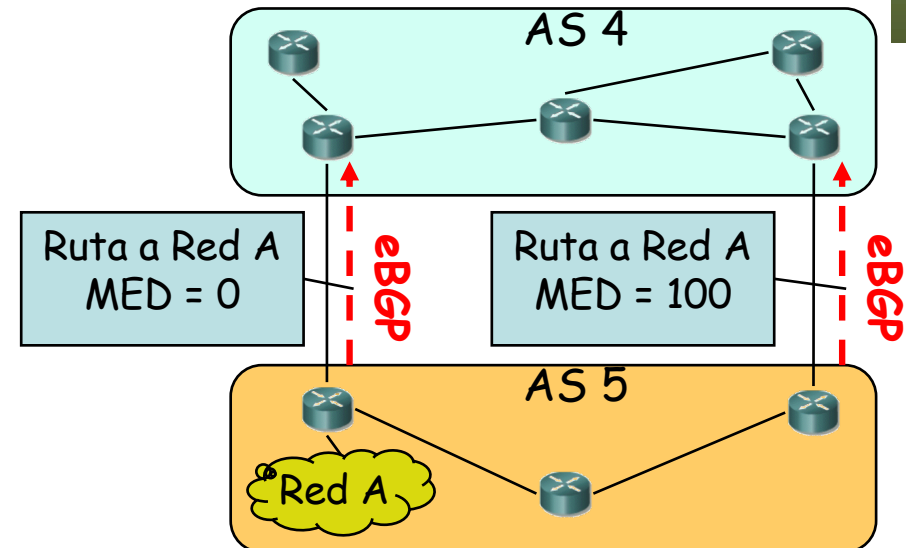
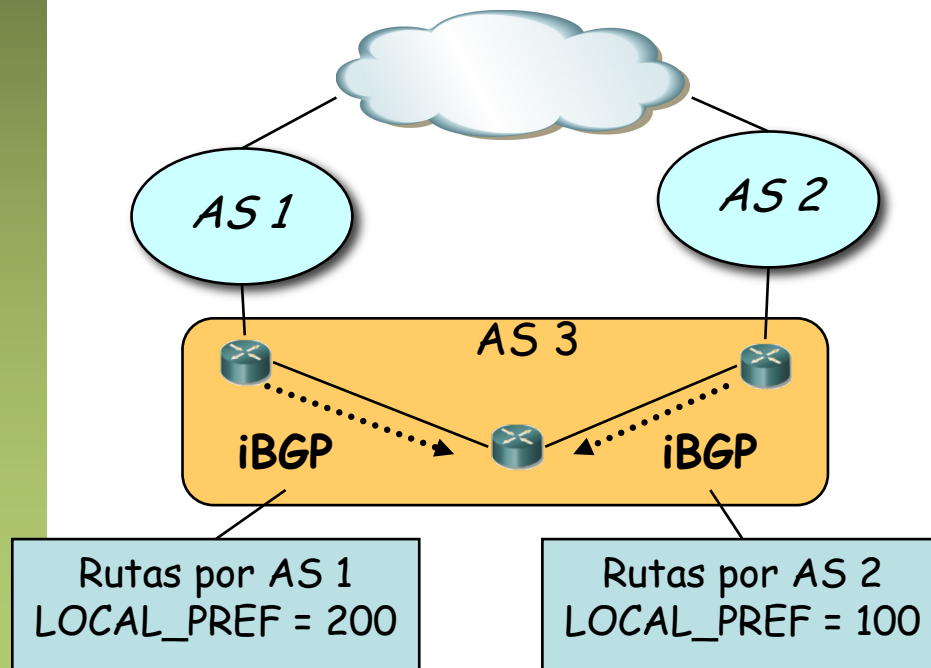
# Path Attributes

**LOCAL\_PREF** (well-known discretionary, nontransitive)

- Solo en iBGP
- Comunica el grado de preferencia por una ruta
- La ruta de mayor valor es seleccionada

**MED** (optional, nontransitive)

- Multi-Exit-Discriminator
- Cuando hay múltiples links a un AS
- Anuncia el *ingress point* preferido
- Es una métrica y se selecciona el de menor MED
- No se propaga a más ASs (debe borrarlo al pasar la ruta a otro AS)



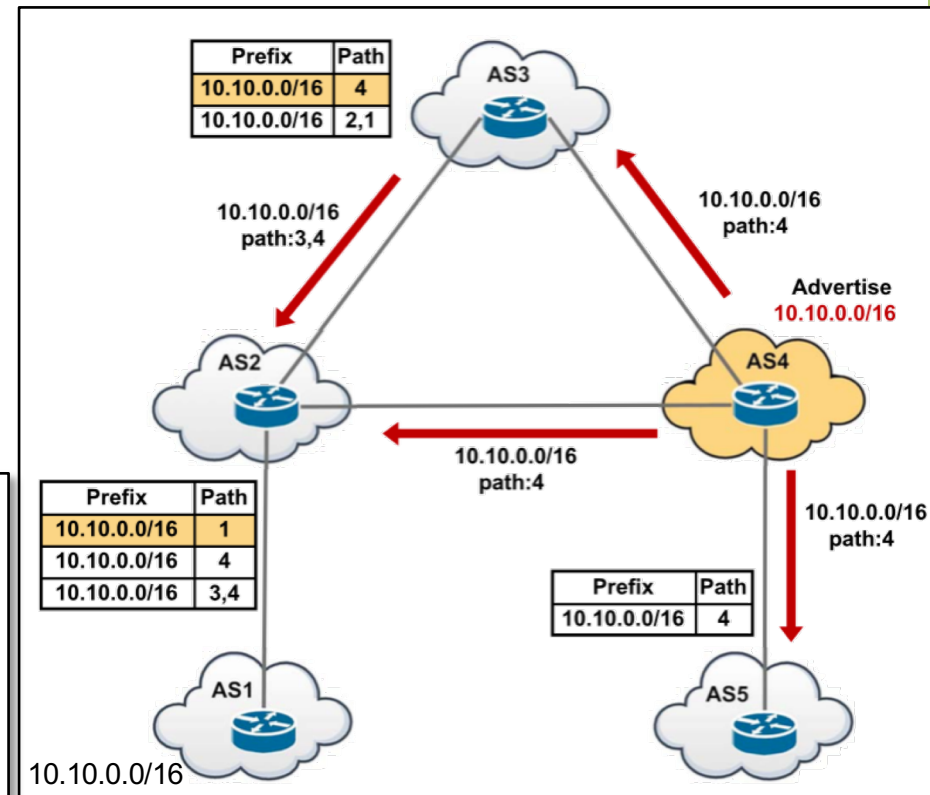
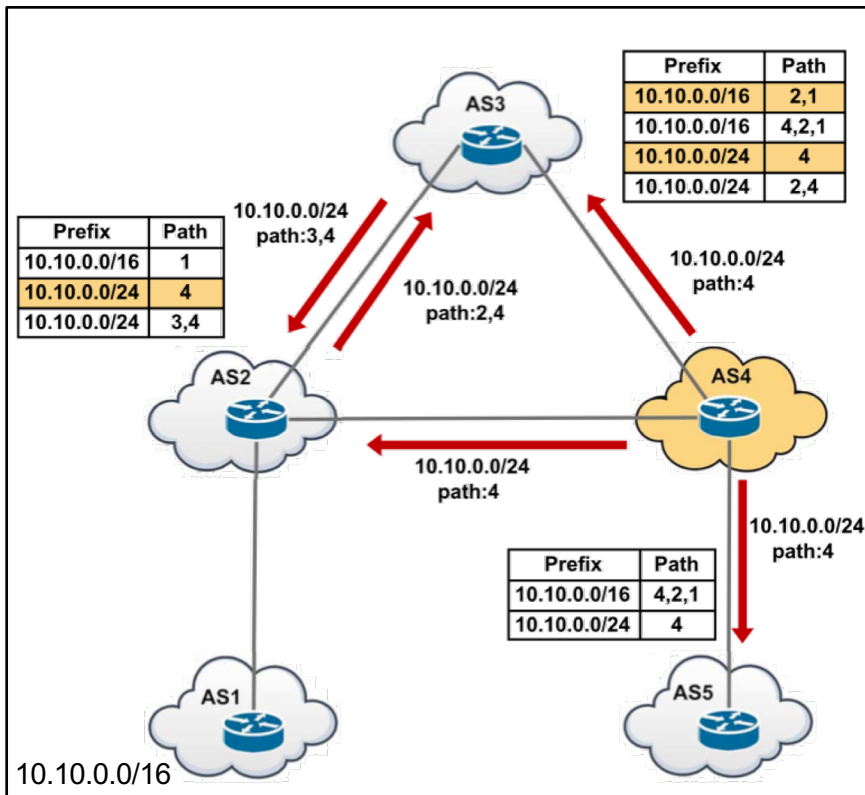
# BGP - Problemas



# BGP – (Sub-)Prefix hijacking

# Ejemplos de problemas

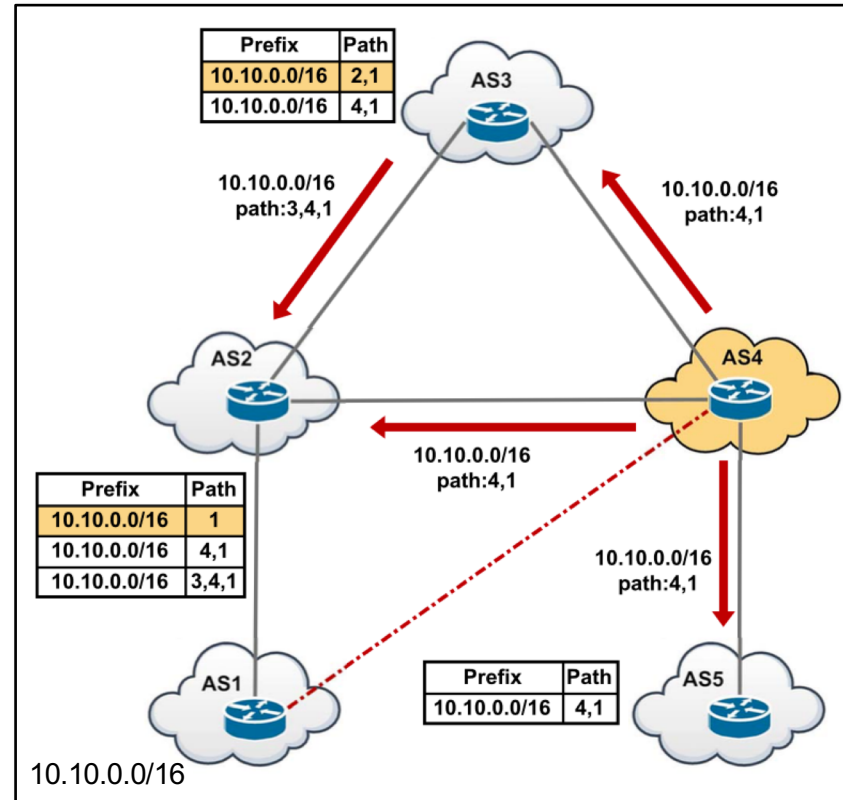
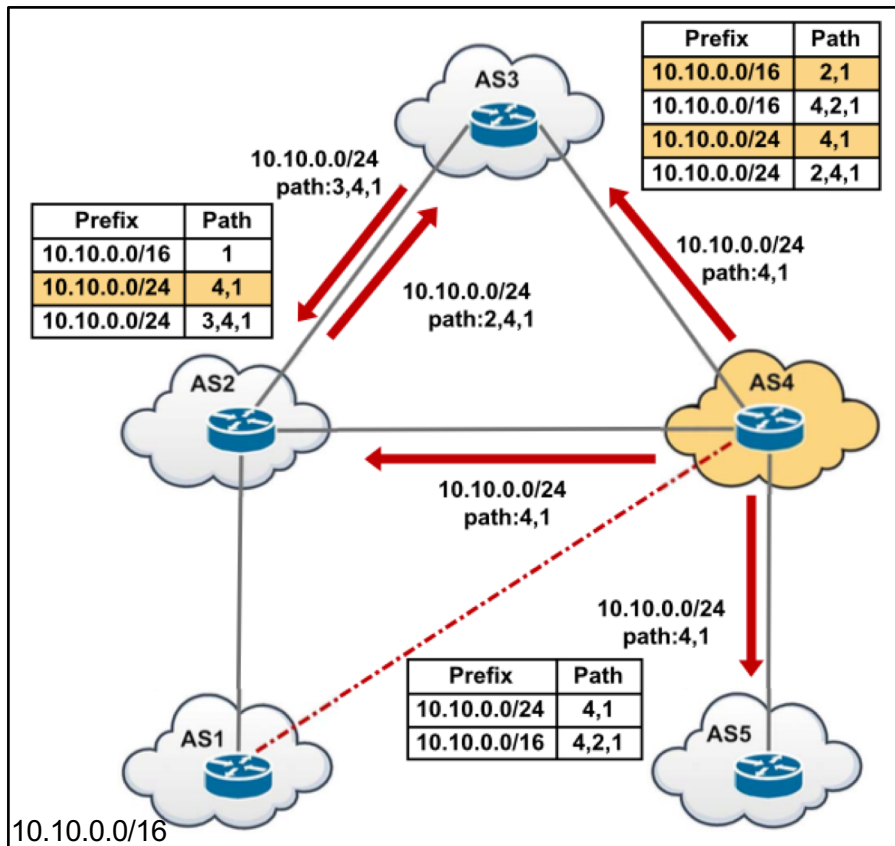
- Prefix hijacking
- Sub-prefix hijacking



# BGP – (Sub-)Prefix and its AS hijacking

# Ejemplos de problemas

- Prefix and its AS hijacking
- Sub-prefix and its AS hijacking



# Ejemplo de secuestro

- YouTube Hijacking: A RIPE NCC RIS case study
  - <https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study>
  - <https://youtu.be/lzLPKuAOe50>

The screenshot shows the RIPE NCC website interface. At the top left is the RIPE NCC logo (RIPE NETWORK COORDINATION CENTRE). To the right is a search bar for the RIPE Database (Whois) and a link to the Website. Below the search bar is a navigation menu with items: Manage IPs and ASNs, Analyse, Participate, Get Support, and Publications. A breadcrumb trail reads: You are here: Home > Publications > News > Industry Developments > YouTube Hijacking: A RIPE NCC RIS case study. On the left side, there is a sidebar menu with 'Publications' and 'News' sections. The main content area features the article title 'YouTube Hijacking: A RIPE NCC RIS case study' and a yellow warning box stating 'You're viewing an archived page. It is no longer being updated.' Below this is the 'Introduction' section, which describes an unauthorised announcement of the prefix 208.65.153.0/24 by Pakistan Telecom (AS17557) on Sunday, 24 February 2008. The text continues: 'In this report we show how the events were seen by RIPE NCC's Routing Information Service (RIS) and how, in general, one can use the RIS tools to obtain hard data on network events.' The 'Event Timeline' section follows, listing key events: 'Before, during and after Sunday, 24 February 2008: AS36561 (YouTube) announces 208.65.152.0/22. Note that AS36561 also announces other prefixes, but they are not involved in the event.' and 'Sunday, 24 February 2008, 18:47 (UTC): AS17557 (Pakistan Telecom) starts announcing 208.65.153.0/24. AS3491 (PCCW Global) propagates the announcement. Routers around the world receive the announcement, and YouTube traffic is redirected to Pakistan.'

# Más ejemplos

- <https://www.youtube.com/watch?v=HnCfQUzMzFM&feature=youtu.be>

**How Pakistan knocked YouTube offline (and how to make sure it never happens again)**  
YouTube becoming unreachable isn't the first time that Internet addresses were hijacked. But if it spurs interest in better security, it may be the last.  
BY DECLAN MCCULLOUGH · FEBRUARY 25, 2008 4:29 PM PST

**BORDER GATEWAY PROTOCOL — How 3ve's BGP hijackers eluded the Internet—and made \$29M**  
3ve used addresses of unsuspecting owners—like the US Air Force.  
DAN GOODIN · 12/21/2018, 12:30 PM

**BORDER GATEWAY PROTOCOL ATTACK — Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency**  
Almost 1,300 addresses for Amazon Route 53 rerouted for two hours.  
DAN GOODIN · 4/24/2018, 2:00 PM

**Criminals, Nation-States Keep Hijacking BGP and DNS**  
While Exploitable Protocols and Processes Persist, Adoption of Secure Fixes Lags  
Matthew J. Schwartz · euronews.com · February 18, 2019

**Why BGP Hijacking Remains a Security Scourge**  
Cyber criminals are stepping up their attacks against routing protocols, creating new problems for network operators.

**OmanTel hijacking of IP space**  
Jared Mauch [jared@puck.nether.net](mailto:jared@puck.nether.net)  
Wed Jan 11 15:50:49 UTC 2017

- Previous message (by thread): [Advice re network compromise and "law enforcement" \(PCI certification\)](#)
- Next message (by thread): [OmanTel hijacking of IP space](#)
- Messages sorted by: [\[date\]](#) | [\[thread\]](#) | [\[subject\]](#) | [\[author\]](#)

There is an ongoing pattern of OmanTel hijacking IP space and advertising it to many of their peers here. 42000. Please

**IPv4 and IPv6 hijacking by AS 6**  
Matt Harris [matt@netfire.net](mailto:matt@netfire.net)  
Thu Apr 12 16:34:31 UTC 2018

- Previous message (by thread): [\[date\]](#)
- Next message (by thread): [\[date\]](#)
- Messages sorted by: [\[date\]](#) | [\[thread\]](#) | [\[subject\]](#) | [\[author\]](#)

AS 6 is now announcing s like I'm not alone. Doe might be going on? The tremendous. The phone n non-functional. I've se (Mike Abbott and John Lu not optimistic.

**198.154.60.0/22 bogon/hijacked?**  
Jeremy Parsons [jeremyp@gmxx.us](mailto:jeremyp@gmxx.us)  
Mon Nov 14 00:49:29 UTC 2016

**AS3266: BitCanal hijack factory, courtesy of Cogent, GTT, and Level3**  
Ronald F. Guilmette [rfg@tristatelogic.com](mailto:rfg@tristatelogic.com)  
Tue Jun 26 04:49:15 UTC 2018

- Previous message (by thread): [Call for presentations RIPE 77](#)
- Next message (by thread): [AS3266: BitCanal hijack factory, courtesy of Cogent, GTT, and Level3](#)
- Messages sorted by: [\[date\]](#) | [\[thread\]](#) | [\[subject\]](#) | [\[author\]](#)

**AS9498 Bharti BGP hijacks**  
George William Herbert [george.herbert@gmail.com](mailto:george.herbert@gmail.com)  
Sat Apr 1 18:19:55 UTC 2017

- Next message (by thread): [AS9498 Bharti BGP hijacks](#)
- Messages sorted by: [\[date\]](#) | [\[thread\]](#) | [\[subject\]](#) | [\[author\]](#)

Hey, Bharti, knock that off.

**Prefix hijack by INDOSAT AS4795 / AS4761**  
Randy amps [amps@djab.com](mailto:amps@djab.com)  
Thu Mar 26 14:08:20 UTC 2015

- Previous message: [booster to gain distance above 60km](#)
- Next message: [Prefix hijack by INDOSAT AS4795 / AS4761](#)
- Messages sorted by: [\[date\]](#) | [\[thread\]](#) | [\[subject\]](#) | [\[author\]](#)

(going) we are seeing like seeing similar or  
1436 29889  
1436 29889