AGH University of Science and Technology

Inter-ALTO Communication Protocol

draft-dulinski-alto-inter-alto-protocol-00

Zbigniew Duliński
Rafał Stankiewicz
Piotr Chołda
Piotr Wydrych
Burkhard Stiller

IETF78 ALTO WG Session, 26 July 2010

Inter-ALTO Communication Protocol

Outline

Inter-ALTO Communication Protocol
Outline

Inter-ALTO Communication Protocol

Outline

Inter-ALTO Communication Protocol

Outline

Contents of the Draft

- Motivation.
- Definitions (parameters and communities).
- Request and response messages format.
- Considerations on:
  - ...
  - ...
  - ...
  - ...

Contents of the Draft

- Motivation.
- Definitions (parameters and communities).
- Request and response messages format.
- Considerations on:
  - ...
  - ...
  - ...
  - ...

Contents of the Draft

- Motivation.
- Definitions (parameters and communities).
- Request and response messages format.
- Considerations on:
  - Inter-ALTO server discovery,
  - [illegible],
  - [illegible],
  - [illegible].

Contents of the Draft

- Motivation.
- Definitions (parameters and communities).
- Request and response messages format.
- Considerations on:
  - Inter-ALTO server discovery,
  - reliability,
  - scalability,
  - IANA actions,
  - security.

Contents of the Draft

- Motivation.
- Definitions (parameters and communities).
- Request and response messages format.
- Considerations on:
  - Inter-ALTO server discovery,
  - reliability,
  - scalability,
  - IANA actions,
  - security.

Contents of the Draft

- Motivation.
- Definitions (parameters and communities).
- Request and response messages format.
- Considerations on:
  - Inter-ALTO server discovery,
  - reliability,
  - scalability,
  - IANA actions,
  - security.

Contents of the Draft

- Motivation.
- Definitions (parameters and communities).
- Request and response messages format.
- Considerations on:
  - Inter-ALTO server discovery,
  - reliability,
  - scalability,
  - IANA actions,
  - security.

AGH

Contents of the Draft

- Motivation.
- Definitions (parameters and communities).
- Request and response messages format.
- Considerations on:
  - Inter-ALTO server discovery,
  - reliability,
  - scalability,
  - IANA actions,
  - security.

Contents of the Draft

- Motivation.
- Definitions (parameters and communities).
- Request and response messages format.
- Considerations on:
  - Inter-ALTO server discovery,
  - reliability,
  - scalability,
  - IANA actions,
  - security.

Inter-ALTO Communication Protocol

Outline

1. Contents of the Draft

2. **Motivation**
   - Route Asymmetry
   - Remote ISPs' Preferences
   - Coordination of ISPs' Policies

3. Parameters and Communities

Motivation
Route Asymmetry

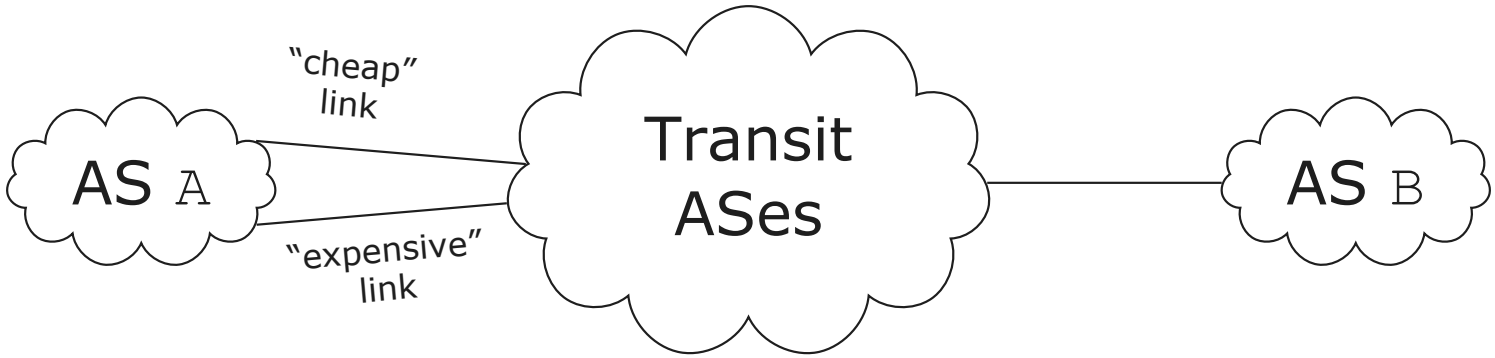

- AS A knows only the "up" route (from BGP tables).
- Assume cost-based ranking.
- Peers from AS B are preferred while they shouldn't be if the P2P application mainly downloads.
- We've done some simulations (topology is in the backup slides) and... it works!

Motivation
Route Asymmetry

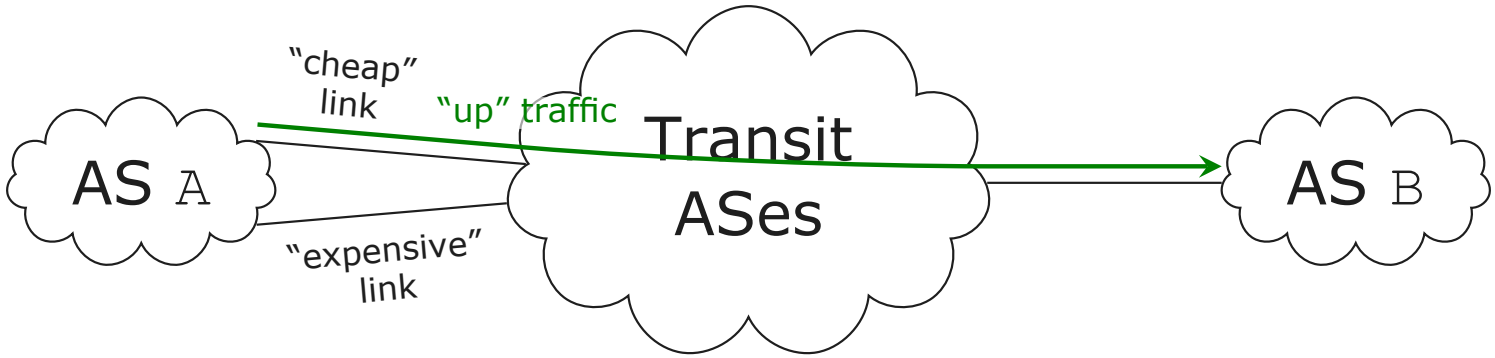

- AS A knows only the "up" route (from BGP tables).
- Assume cost-based ranking.
- Peers from AS B are preferred while they shouldn't be if the P2P application mainly downloads.
- We've done some simulations (topology is in the backup slides) and... it works!

Motivation
Route Asymmetry

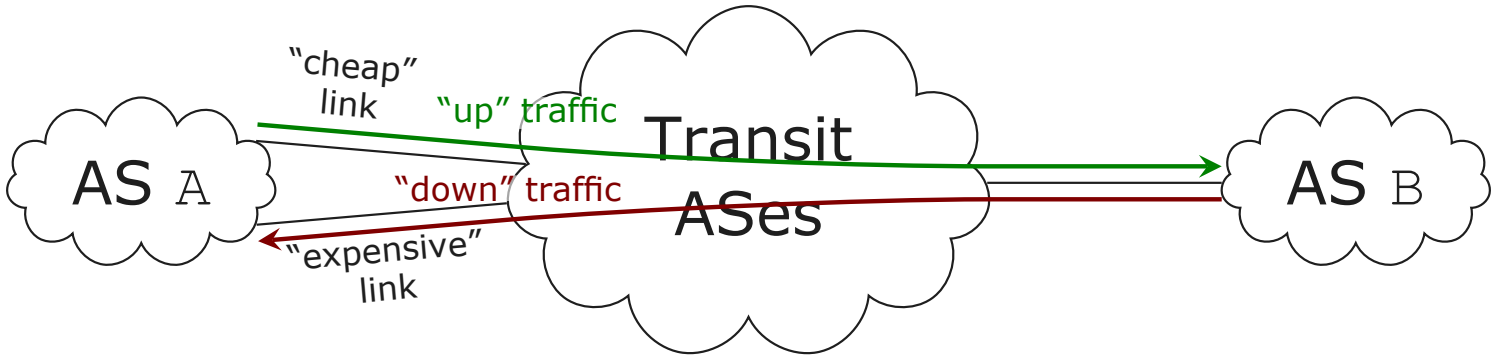

- AS A knows only the "up" route (from BGP tables).
- Assume cost-based ranking.
- Peers from AS B are preferred while they shouldn't be if the P2P application mainly downloads.
- We've done some simulations (topology is in the backup slides) and... it works!

Motivation
Route Asymmetry

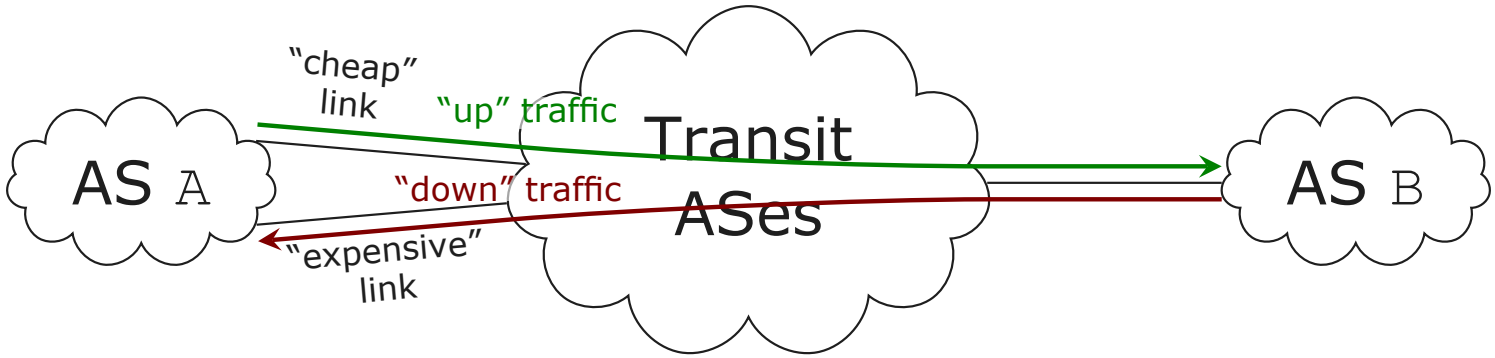

- AS A knows only the "up" route (from BGP tables).
- Assume cost-based ranking.
- Peers from AS B are preferred while they shouldn't be if the P2P application mainly downloads.
- We've done some simulations (topology is in the backup slides) and... it works!

Motivation
Route Asymmetry

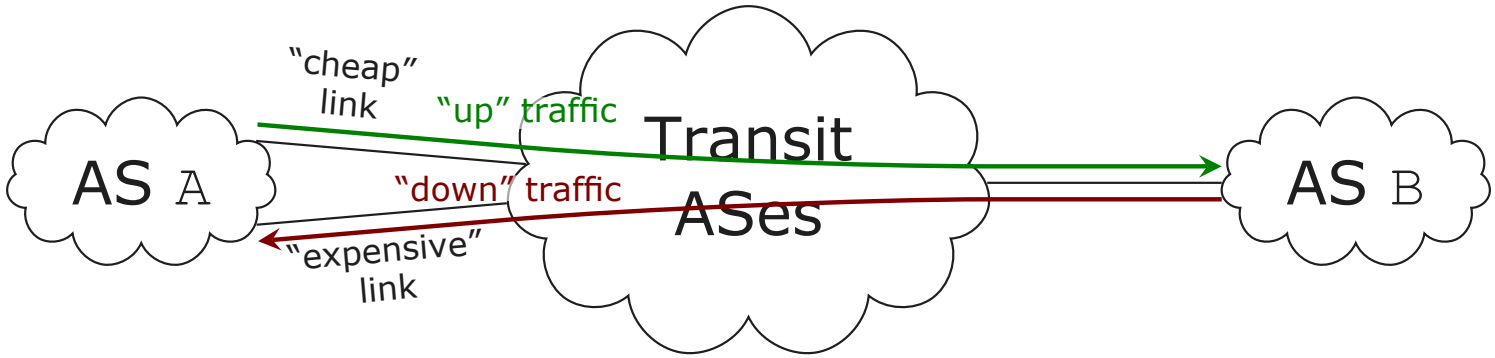

- AS A knows only the "up" route (from BGP tables).
- Assume cost-based ranking.
- Peers from AS B are preferred while they shouldn't be if the P2P application mainly downloads.
- We've done some simulations (topology is in the ▸ backup slides) and... it works!

Motivation
Route Asymmetry

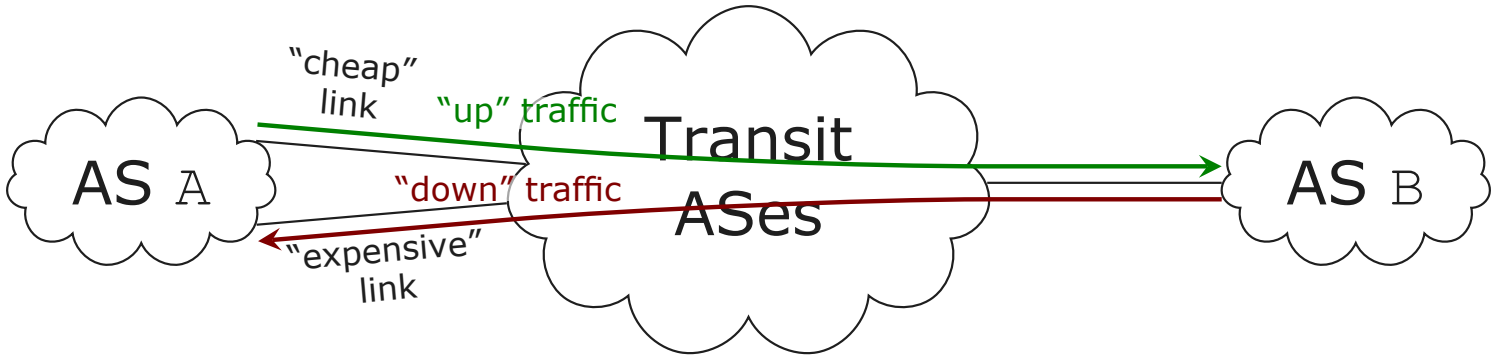

- AS A knows only the "up" route (from BGP tables).
- Assume cost-based ranking.
- Peers from AS B are preferred while they shouldn't be if the P2P application mainly downloads.
- We've done some simulations (topology is in the backup slides) and... it works!

Motivation
Route Asymmetry

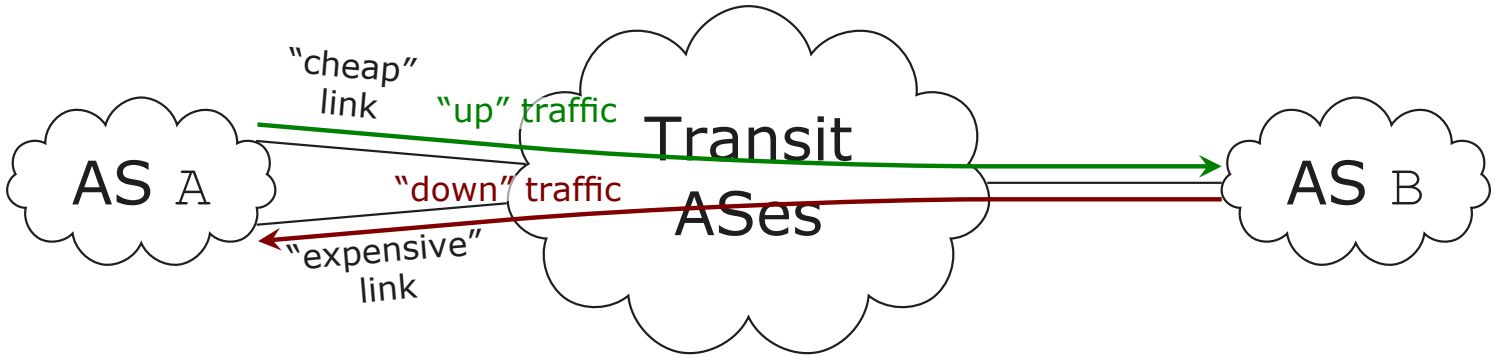

- AS A knows only the "up" route (from BGP tables).
- Assume cost-based ranking.
- Peers from AS B are preferred while they shouldn't be if the P2P application mainly downloads.
- We've done some simulations (topology is in the ▸ backup slides) and... it works!

Motivation
Remote ISPs' Preferences

- AS A can't distinguish between the peers from the α subnet and from the rest of AS B.

Motivation
Remote ISPs' Preferences

- AS A can't distinguish between the peers from the α subnet and from the rest of AS B.

Motivation
Coordination of ISPs' Policies

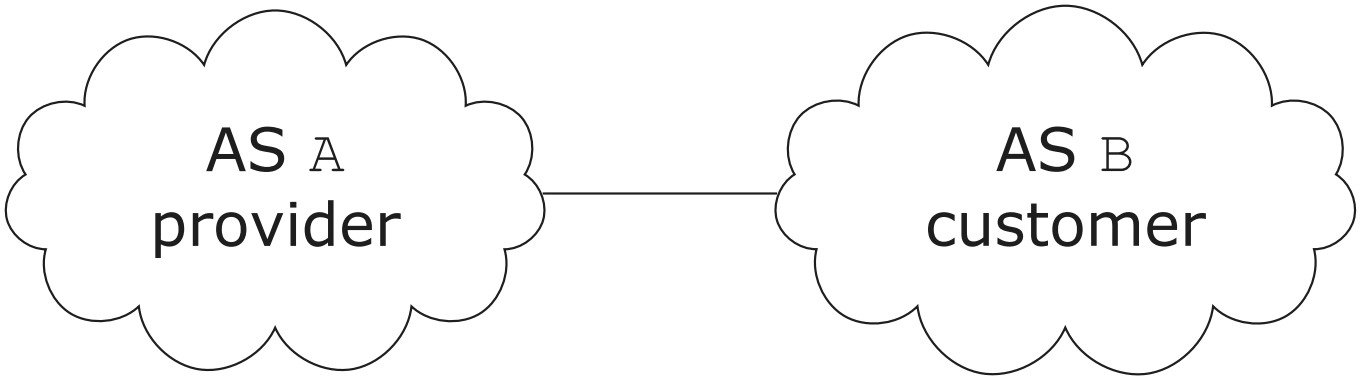

- AS B buys traffic from AS A.
- AS B does not prefer peers from AS A.
- AS A do prefer peers from AS B.

Motivation
Coordination of ISPs' Policies

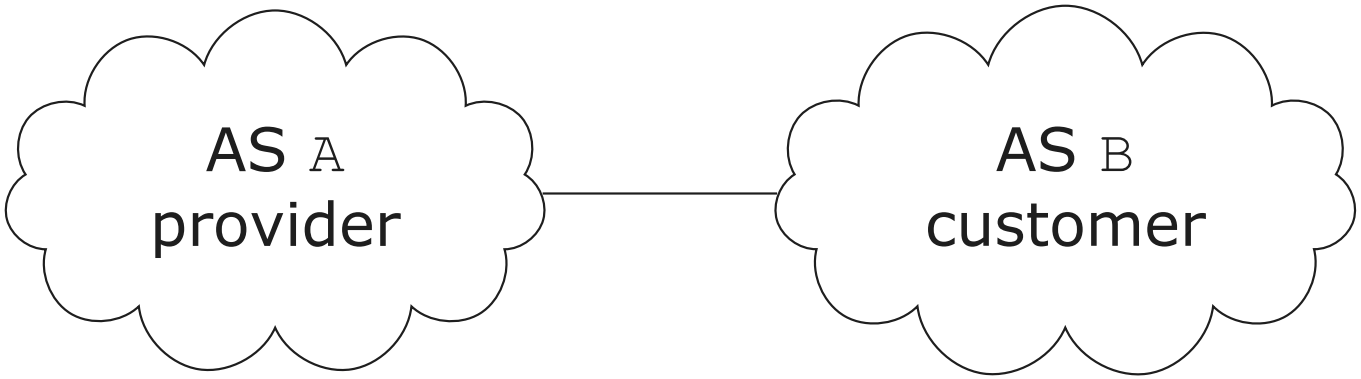

- AS B buys traffic from AS A.
- AS B does not prefer peers from AS A.
- AS A do prefer peers from AS B.

Motivation
Coordination of ISPs' Policies

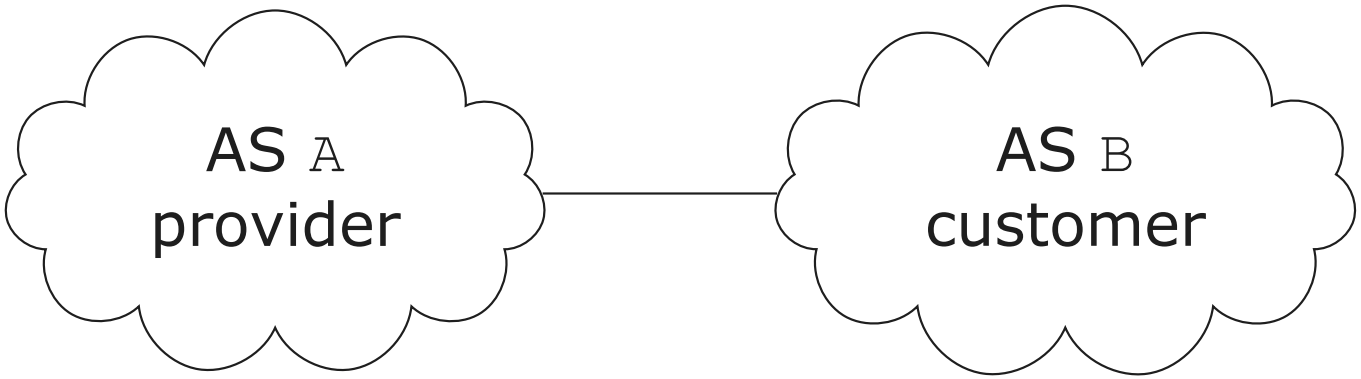

- AS B buys traffic from AS A.
- AS B does not prefer peers from AS A.
- AS A do prefer peers from AS B.

Motivation
Coordination of ISPs' Policies

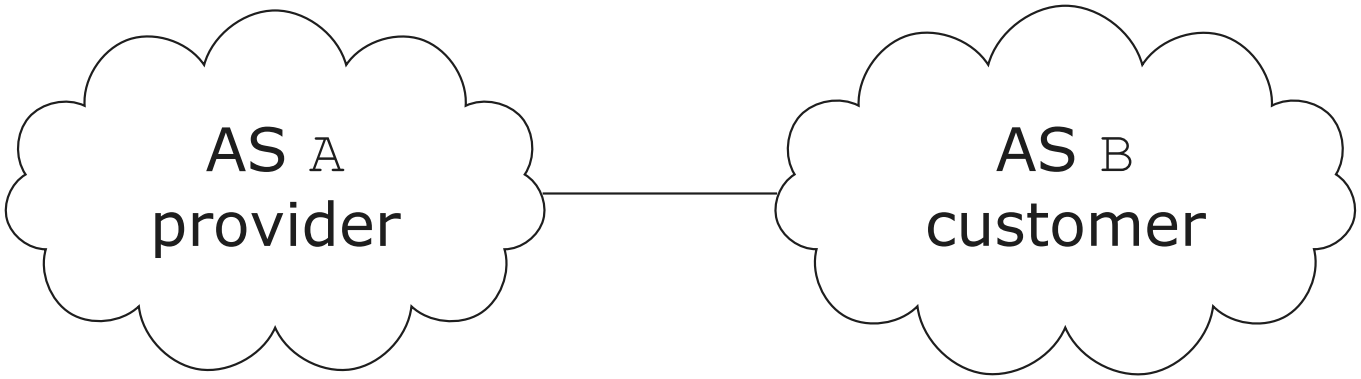

- AS B buys traffic from AS A.
- AS B does not prefer peers from AS A.
- AS A do prefer peers from AS B.

Inter-ALTO Communication Protocol

Outline

Parameters and Communities

Parameters

- Parameters are used to exchange the data.
- Each parameter has its name and meaning (info, asc, desc).
- Units are used to measure the values.

Parameters and Communities

Parameters

- Parameters are used to exchange the data.
- Each parameter has its name and meaning (info, asc, desc).
- Units are used to measure the values.

Parameters and Communities

Parameters

- Parameters are used to exchange the data.
- Each parameter has its name and meaning (info, asc, desc).
- Units are used to measure the values.

Parameters and Communities

Communities

- Communities:
  - organize parameters,
  - provide configurability and flexibility.
- Contain mandatory and optional parameters.
- Build a dependency tree.

Parameters and Communities

Communities

- Communities:
  - organize parameters,
  - provide configurability and flexibility.
- Contain mandatory and optional parameters.
- Build a dependency tree.

Parameters and Communities

Communities

- Communities:
  - organize parameters,
  - provide configurability and flexibility.
- Contain mandatory and optional parameters.
- Build a dependency tree.

Parameters and Communities

Communities

- Communities:
  - organize parameters,
  - provide configurability and flexibility.
- Contain mandatory and optional parameters.
- Build a dependency tree.
  - The root of the tree: GENERAL community.

Parameters and Communities

Communities

- Communities:
  - organize parameters,
  - provide configurability and flexibility.
- Contain mandatory and optional parameters.
- Build a dependency tree.
  - The root of the tree: GENERAL community.

Parameters and Communities

Communities

- Communities:
  - organize parameters,
  - provide configurability and flexibility.
- Contain mandatory and optional parameters.
- Build a dependency tree.
  - The root of the tree: GENERAL community.

Parameters and Communities

Communities

Sample community dependency tree:

Note: the step with an optional parameter is not required. Also a mandatory parameter can be added while creating a new community.

Parameters and Communities

Communities

Sample community dependency tree:

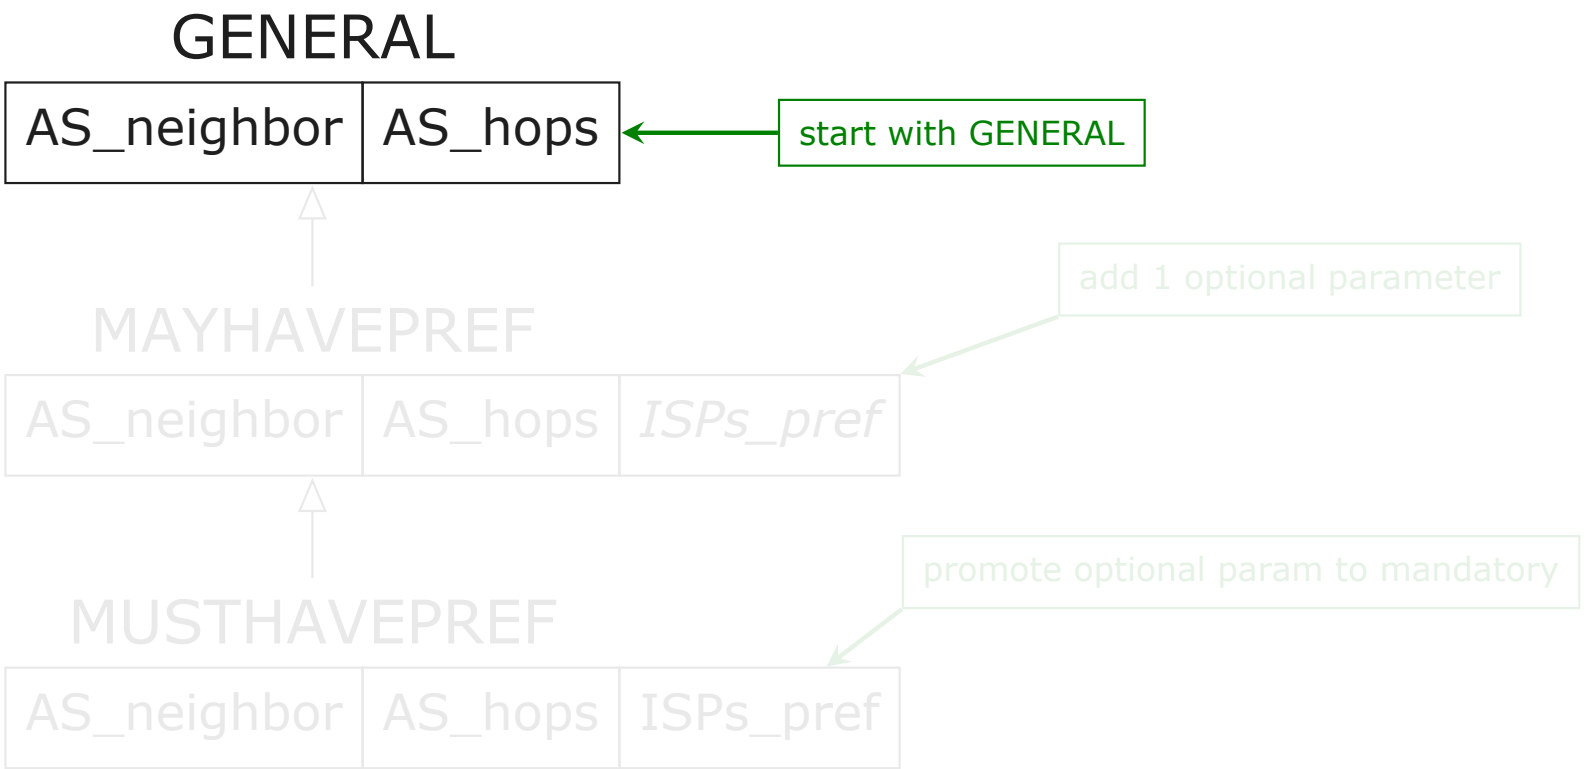

Note: the step with an optional parameter is not required. Also a mandatory parameter can be added while creating a new community.

Parameters and Communities

Communities

Sample community dependency tree:

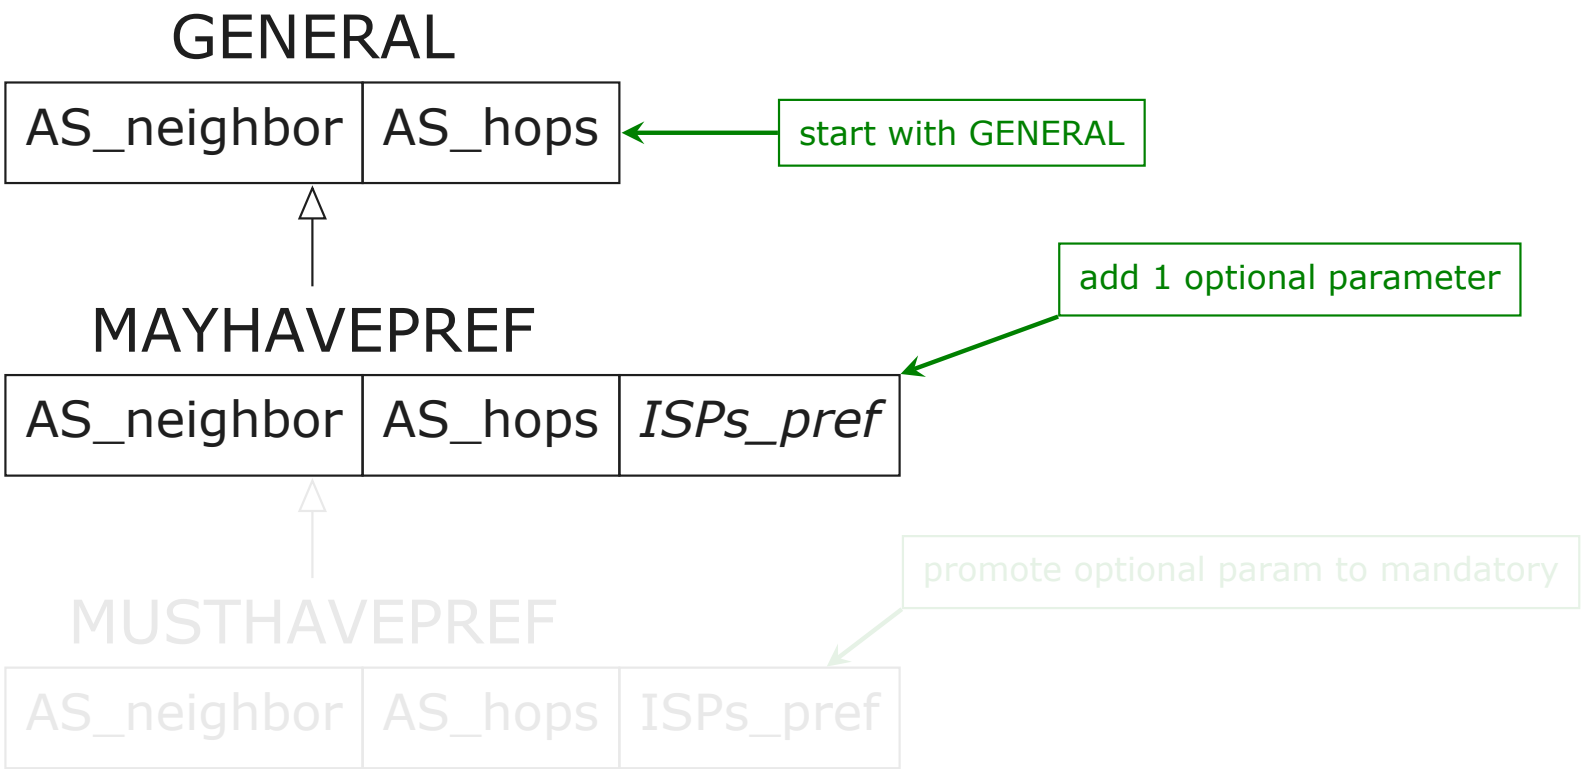

Note: the step with an optional parameter is not required. Also a mandatory parameter can be added while creating a new community.

Parameters and Communities

Communities

Sample community dependency tree:

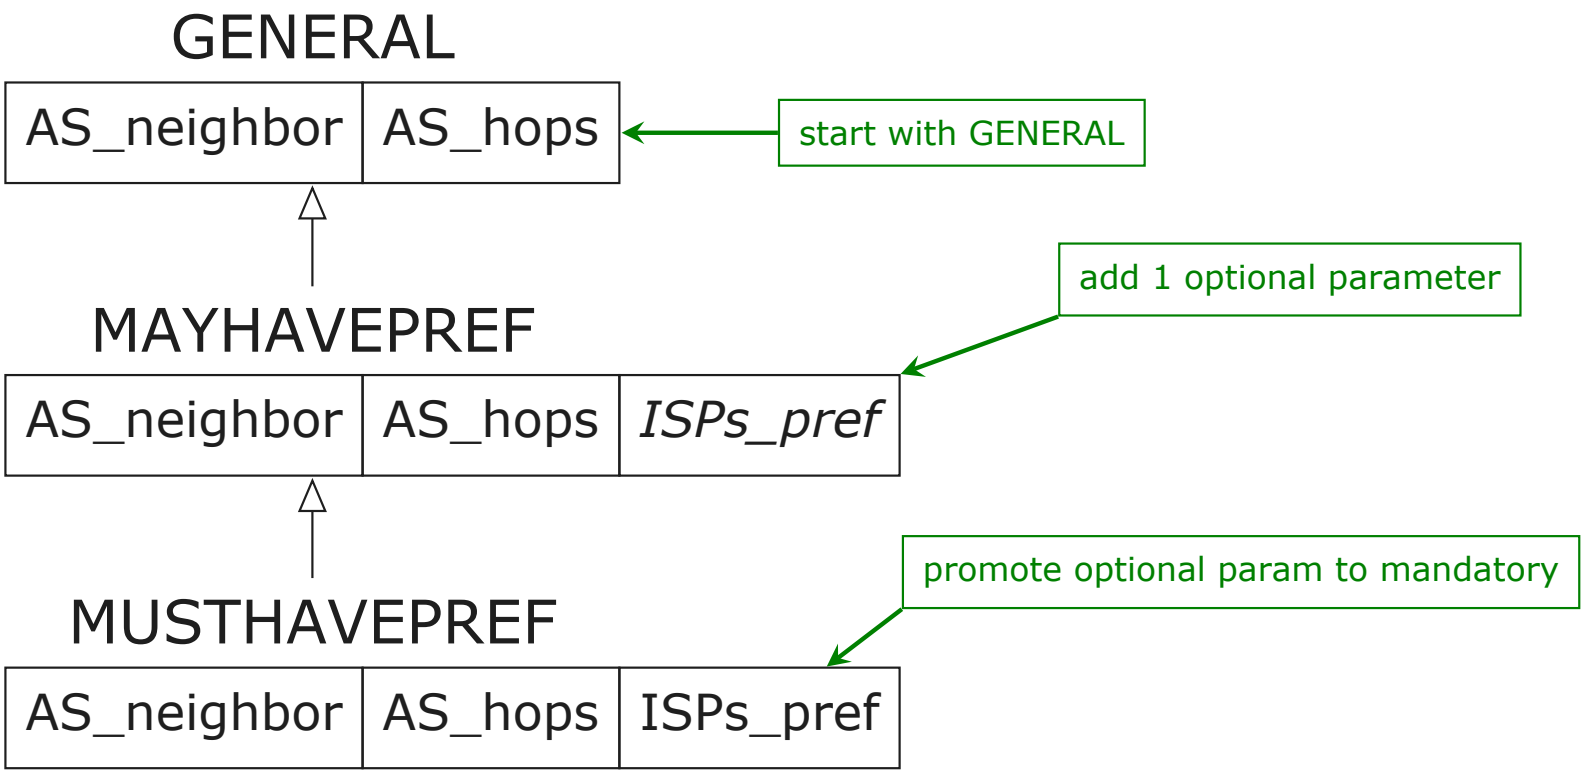

Note: the step with an optional parameter is not required. Also a mandatory parameter can be added while creating a new community.

Parameters and Communities

Communities

Sample community dependency tree:

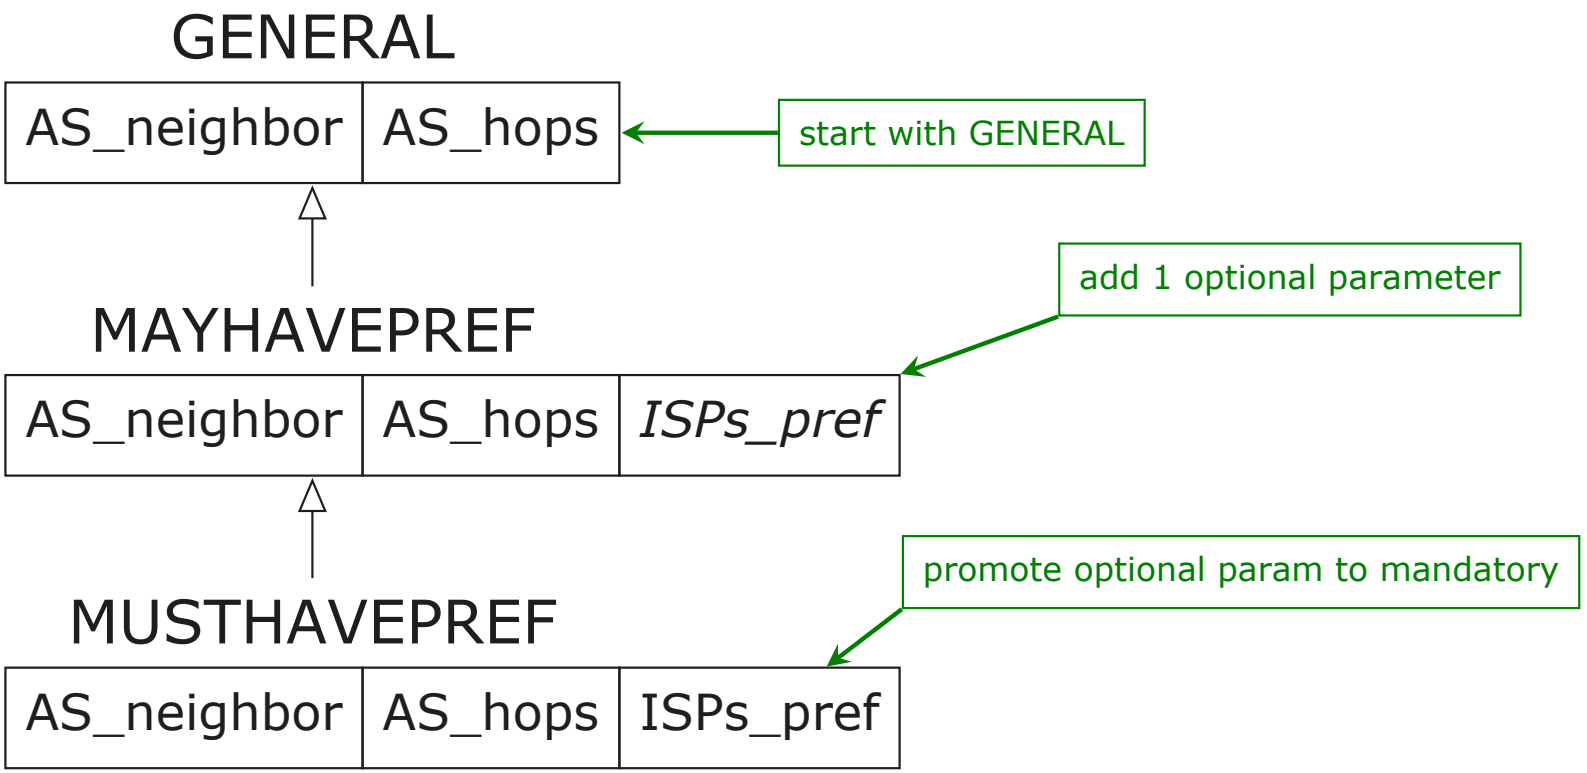

Note: the step with an optional parameter is not required. Also a mandatory parameter can be added while creating a new community.

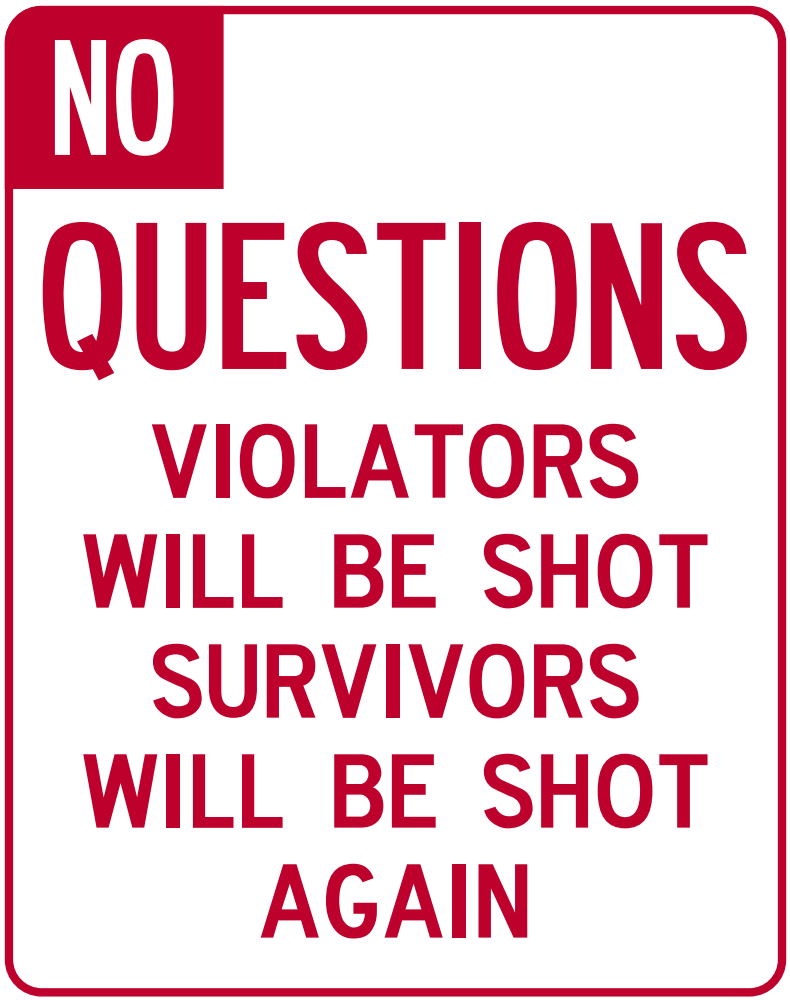
NO
QUESTIONS
VIOLATORS
WILL BE SHOT
SURVIVORS
WILL BE SHOT
AGAIN

AGH

Thank you!

Backup Slides
Outline

Backup Slides
Outline

4 **Requests and Responses**

5 **Simulated Topology**

Backup Slides
Outline

Requests and Responses

Request Types

- BASIC REQUEST
  - Any subset of parameters (defined for a specified community).
  - List of requested parameters for remote peers.
  - Designed for download-only P2P applications.
- EXTENDED REQUEST
  - Upload-only P2P applications do not need Inter-ALTO. We can rely on BGP info.

Requests and Responses

Request Types

- BASIC REQUEST
  - Any subset of parameters (defined for a specified community).
  - List of requested parameters for remote peers.
  - Designed for download-only P2P applications.
- EXTENDED REQUEST
  - Upload-only P2P applications do not need Inter-ALTO. We can rely on BGP info.

Requests and Responses

Request Types

- BASIC REQUEST
  - Any subset of parameters (defined for a specified community).
  - List of requested parameters for remote peers.
  - Designed for download-only P2P applications.
- EXTENDED REQUEST
  - Upload-only P2P applications do not need Inter-ALTO. We can rely on BGP info.

Requests and Responses

Request Types

- BASIC REQUEST
  - Any subset of parameters (defined for a specified community).
  - List of requested parameters for remote peers.
  - Designed for download-only P2P applications.
- EXTENDED REQUEST
  - All mandatory parameters + any subset of optional parameters.
- Upload-only P2P applications do not need Inter-ALTO. We can rely on BGP info.

Requests and Responses

Request Types

- BASIC REQUEST
  - Any subset of parameters (defined for a specified community).
  - List of requested parameters for remote peers.
  - Designed for download-only P2P applications.
- EXTENDED REQUEST
  - All mandatory parameters + any subset of optional parameters.
  - List of values of parameters for local peers + list of requested parameters for remote peers.
  - Designed for download-and-upload P2P applications.
- Upload-only P2P applications do not need Inter-ALTO. We can rely on BGP info.

Requests and Responses

Request Types

- BASIC REQUEST
  - Any subset of parameters (defined for a specified community).
  - List of requested parameters for remote peers.
  - Designed for download-only P2P applications.
- EXTENDED REQUEST
  - All mandatory parameters + any subset of optional parameters.
  - List of values of parameters for local peers + list of requested parameters for remote peers.
  - Designed for download-and-upload P2P applications.
- Upload-only P2P applications do not need Inter-ALTO. We can rely on BGP info.

Requests and Responses

Request Types

- BASIC REQUEST
  - Any subset of parameters (defined for a specified community).
  - List of requested parameters for remote peers.
  - Designed for download-only P2P applications.
- EXTENDED REQUEST
  - All mandatory parameters + any subset of optional parameters.
  - List of values of parameters for local peers + list of requested parameters for remote peers.
  - Designed for download-and-upload P2P applications.
- Upload-only P2P applications do not need Inter-ALTO. We can rely on BGP info.

Requests and Responses

Request Types

- BASIC REQUEST
  - Any subset of parameters (defined for a specified community).
  - List of requested parameters for remote peers.
  - Designed for download-only P2P applications.
- EXTENDED REQUEST
  - All mandatory parameters + any subset of optional parameters.
  - List of values of parameters for local peers + list of requested parameters for remote peers.
  - Designed for download-and-upload P2P applications.
- Upload-only P2P applications do not need Inter-ALTO. We can rely on BGP info.

Requests and Responses

Request Types

- BASIC REQUEST
  - Any subset of parameters (defined for a specified community).
  - List of requested parameters for remote peers.
  - Designed for download-only P2P applications.
- EXTENDED REQUEST
  - All mandatory parameters + any subset of optional parameters.
  - List of values of parameters for local peers + list of requested parameters for remote peers.
  - Designed for download-and-upload P2P applications.
- Upload-only P2P applications do not need Inter-ALTO. We can rely on BGP info.

Requests and Responses

Response Types

- NORMAL RESPONSE
  - Everything OK, like HTTP 200.
  - List of values of requested (+ non-requested) parameters.
  - Sending values of optional parameters is not a must.
- REFUSE RESPONSE
- ERROR RESPONSE

Requests and Responses

Response Types

- NORMAL RESPONSE
  - Everything OK, like HTTP 200.
  - List of values of requested (+ non-requested) parameters.
  - Sending values of optional parameters is not a must.
- REFUSE RESPONSE
- ERROR RESPONSE

Requests and Responses

Response Types

- NORMAL RESPONSE
  - Everything OK, like HTTP 200.
  - List of values of requested (+ non-requested) parameters.
  - Sending values of optional parameters is not a must.
- REFUSE RESPONSE
- ERROR RESPONSE

Requests and Responses

Response Types

- NORMAL RESPONSE
  - Everything OK, like HTTP 200.
  - List of values of requested (+ non-requested) parameters.
  - Sending values of optional parameters is not a must.
- REFUSE RESPONSE
  - Responding ALTO server won't communicate with the requesting ALTO server within the indicated community.
- ERROR RESPONSE

Requests and Responses

Response Types

- NORMAL RESPONSE
  - Everything OK, like HTTP 200.
  - List of values of requested (+ non-requested) parameters.
  - Sending values of optional parameters is not a must.
- REFUSE RESPONSE
  - Responding ALTO server won't communicate with the requesting ALTO server within the indicated community.
  - After receiving this response the requestor should "go up" in the community dependency tree.
- ERROR RESPONSE

Requests and Responses

Response Types

- NORMAL RESPONSE
  - Everything OK, like HTTP 200.
  - List of values of requested (+ non-requested) parameters.
  - Sending values of optional parameters is not a must.
- REFUSE RESPONSE
  - Responding ALTO server won't communicate with the requesting ALTO server within the indicated community.
  - After receiving this response the requestor should "go up" in the community dependency tree.
- ERROR RESPONSE

Requests and Responses

Response Types

- NORMAL RESPONSE
  - Everything OK, like HTTP 200.
  - List of values of requested (+ non-requested) parameters.
  - Sending values of optional parameters is not a must.
- REFUSE RESPONSE
  - Responding ALTO server won't communicate with the requesting ALTO server within the indicated community.
  - After receiving this response the requestor should "go up" in the community dependency tree.
- ERROR RESPONSE
  - An unrecognized parameter name has been found in a request.

Requests and Responses

Response Types

- NORMAL RESPONSE
  - Everything OK, like HTTP 200.
  - List of values of requested (+ non-requested) parameters.
  - Sending values of optional parameters is not a must.
- REFUSE RESPONSE
  - Responding ALTO server won't communicate with the requesting ALTO server within the indicated community.
  - After receiving this response the requestor should "go up" in the community dependency tree.
- ERROR RESPONSE
  - An unrecognized parameter name has been found in a request.

Requests and Responses

Response Types

- NORMAL RESPONSE
  - Everything OK, like HTTP 200.
  - List of values of requested (+ non-requested) parameters.
  - Sending values of optional parameters is not a must.
- REFUSE RESPONSE
  - Responding ALTO server won’t communicate with the requesting ALTO server within the indicated community.
  - After receiving this response the requestor should “go up” in the community dependency tree.
- ERROR RESPONSE
  - An unrecognized parameter name has been found in a request.

Backup Slides
Outline

Simulated Topology

A scrap of the real Internet topology (AS numbers were hidden):

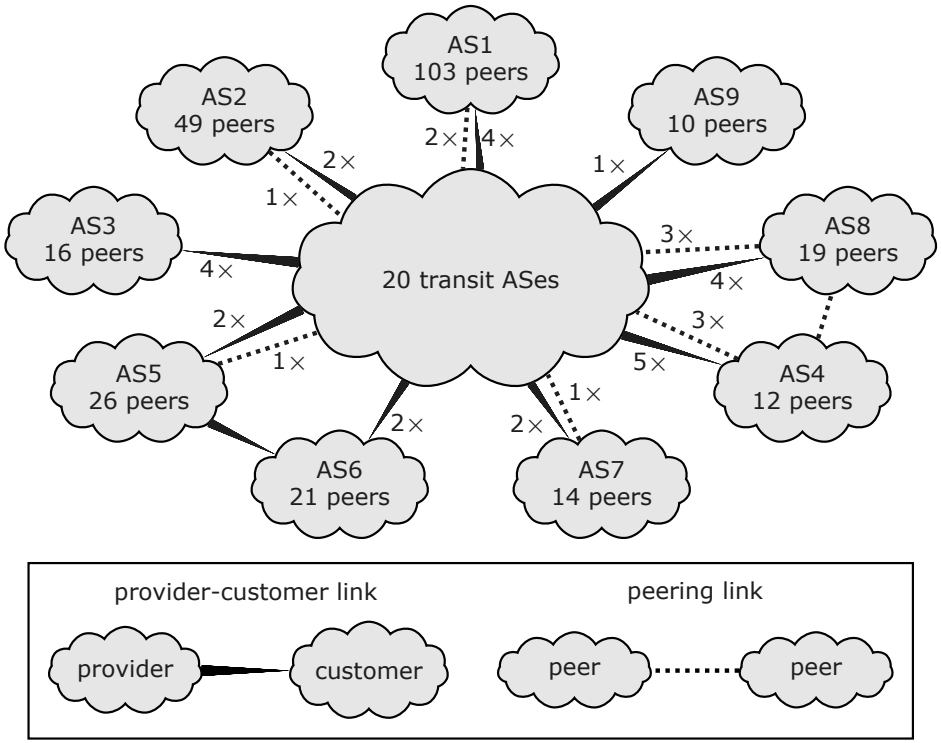

We managed to:

- Move a part of the traffic from provider-customer to peering links.
- Coordinate policies of AS5 and AS6.