

ISP Border Definition

Alexander Azimov

<aa@qrator.net>

ISP goal (one ASN)



BGP Decision Process

Step	Attribute	Transit
1	Highest Local Preference value	Non-transit
2	Lowest AS Path length	Transit
3	Lowest Origin type	Transit
4	Lowest MED	Non-transit
5	eBGP learned over iBGP learned	Automatic
6	Lowest IGP cost to border router	Local configuration
7	Lowest router Id	Local configuration

Send traffic with maximum benefit.

BGP Decision Process

Step	Attribute	Transit
1	Highest Local Preference value	Non-transit
2	Lowest AS Path length	Transit
3	Lowest Origin type	Transit
4	Lowest MED	Non-transit
5	eBGP learned over iBGP learned	Automatic
6	Lowest IGP cost to border router	Local configuration
7	Lowest router Id	Local configuration

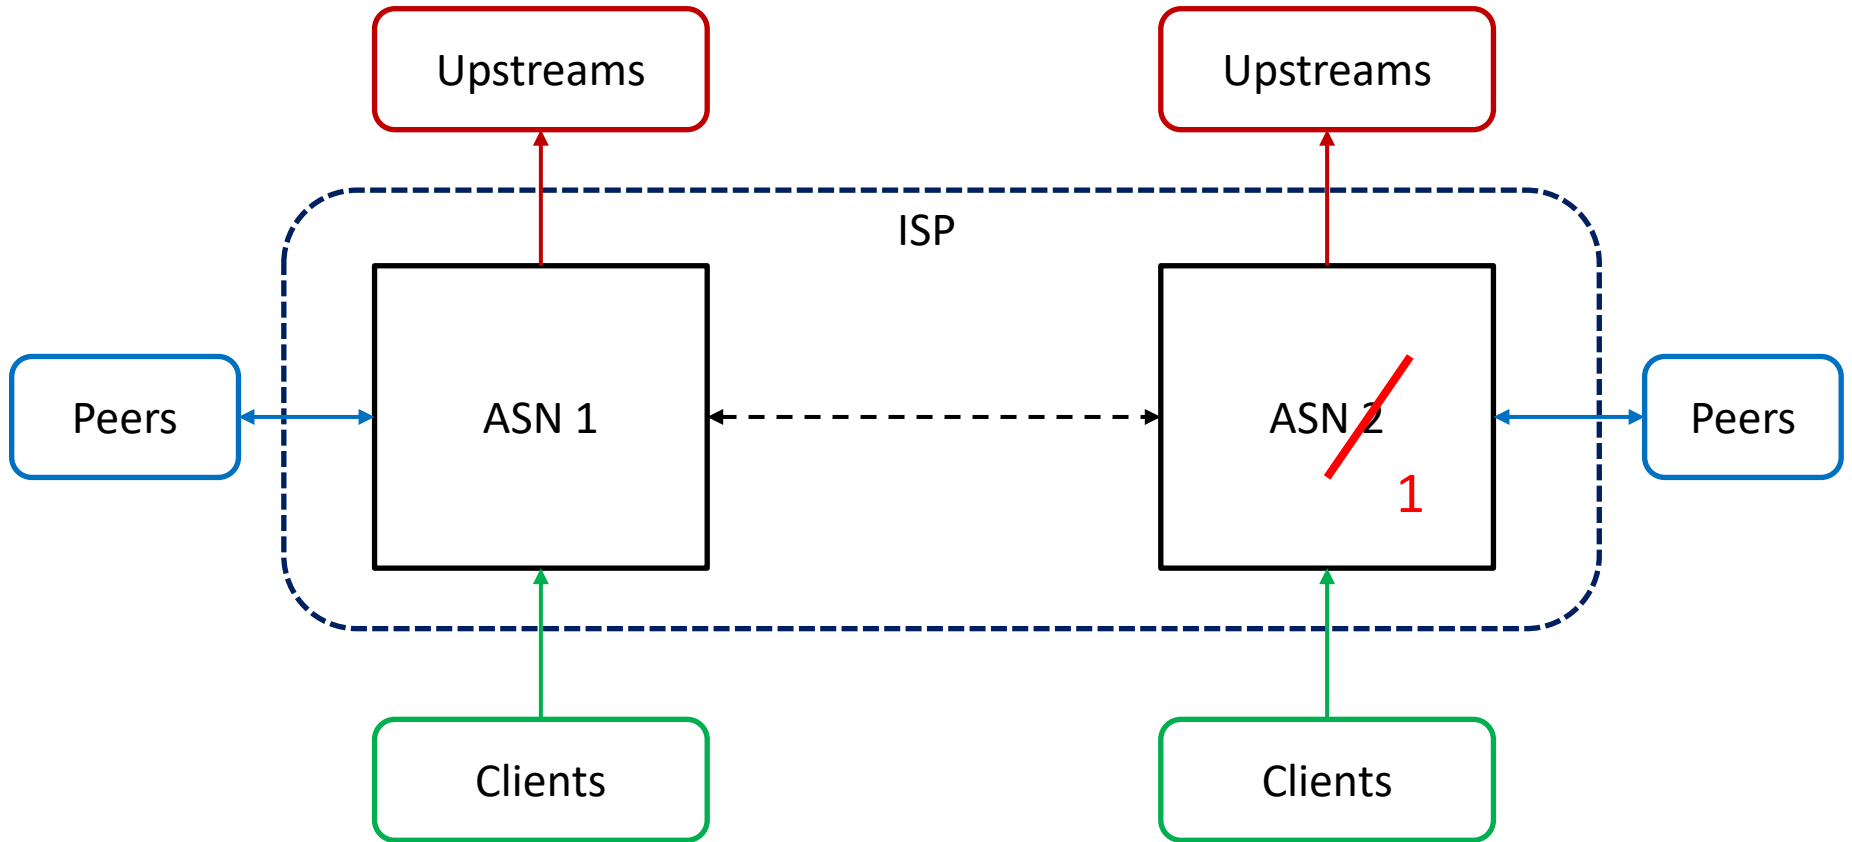
Get rid of traffic as soon as possible!

ISP goal (multiple ASNs)




$$Profit(ASN_1 + ASN_2) > Profit(ASN_1) + Profit(ASN_2)$$

Multiple ASNs – One ISP



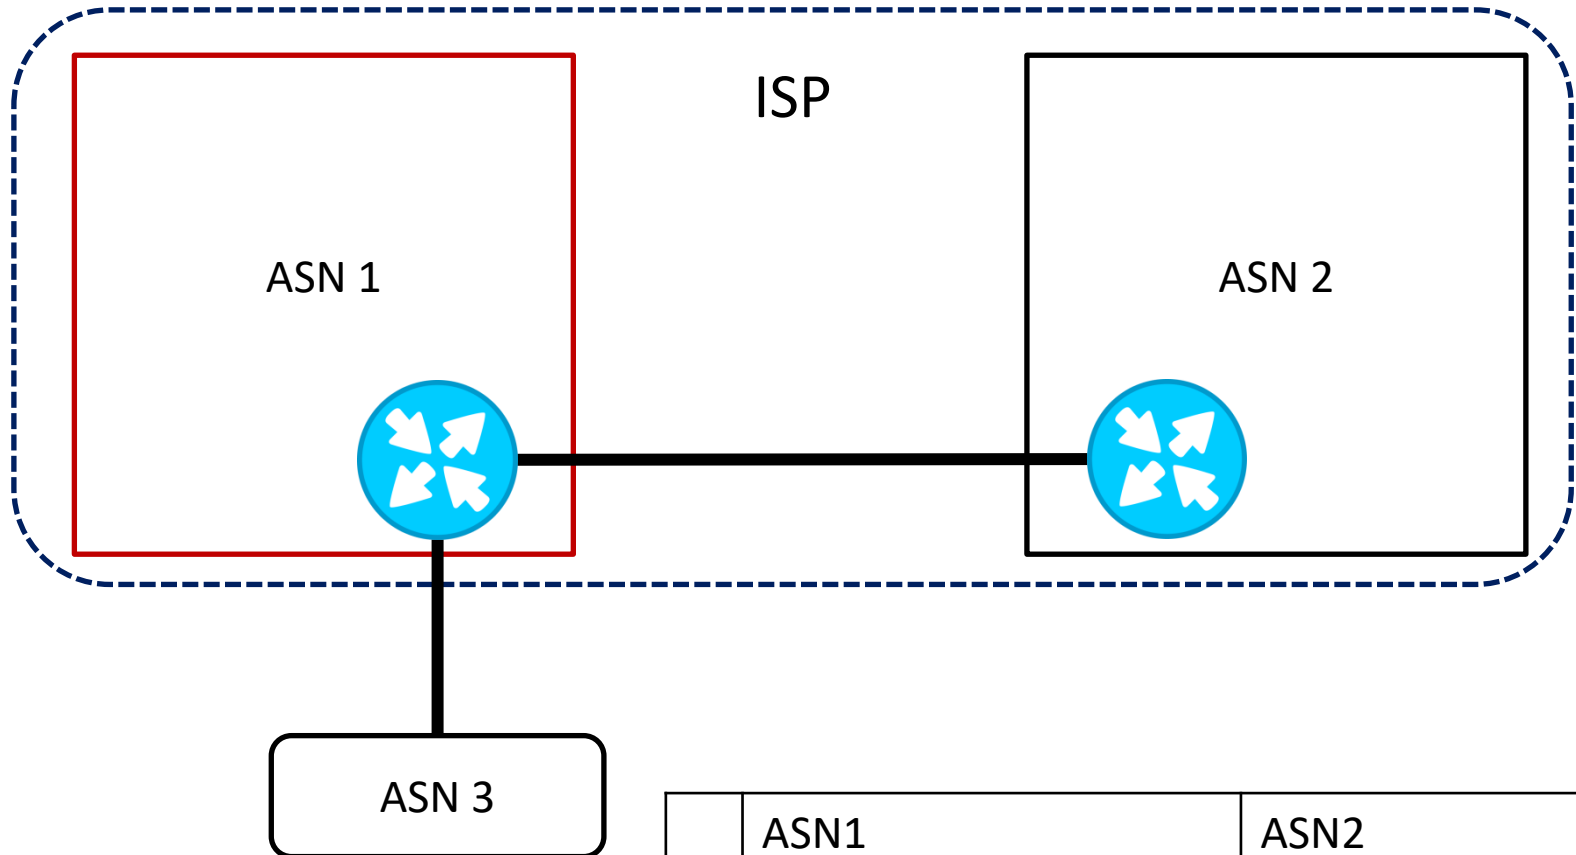
Not so fast...

Pref Transmission



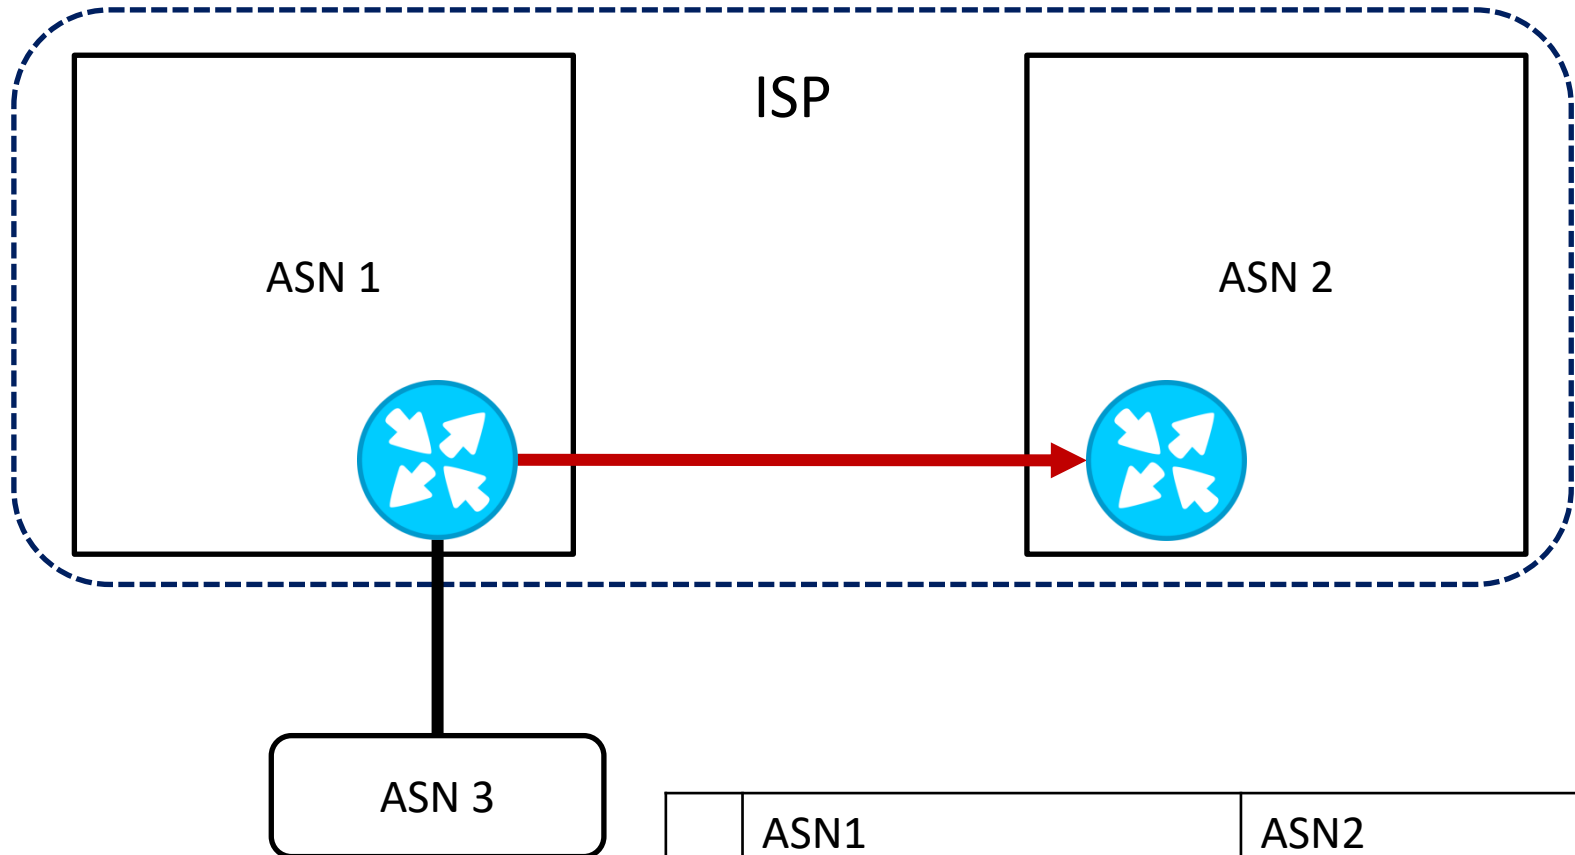
	ASN1	ASN2
a.	Assign pref	

Pref Transmission



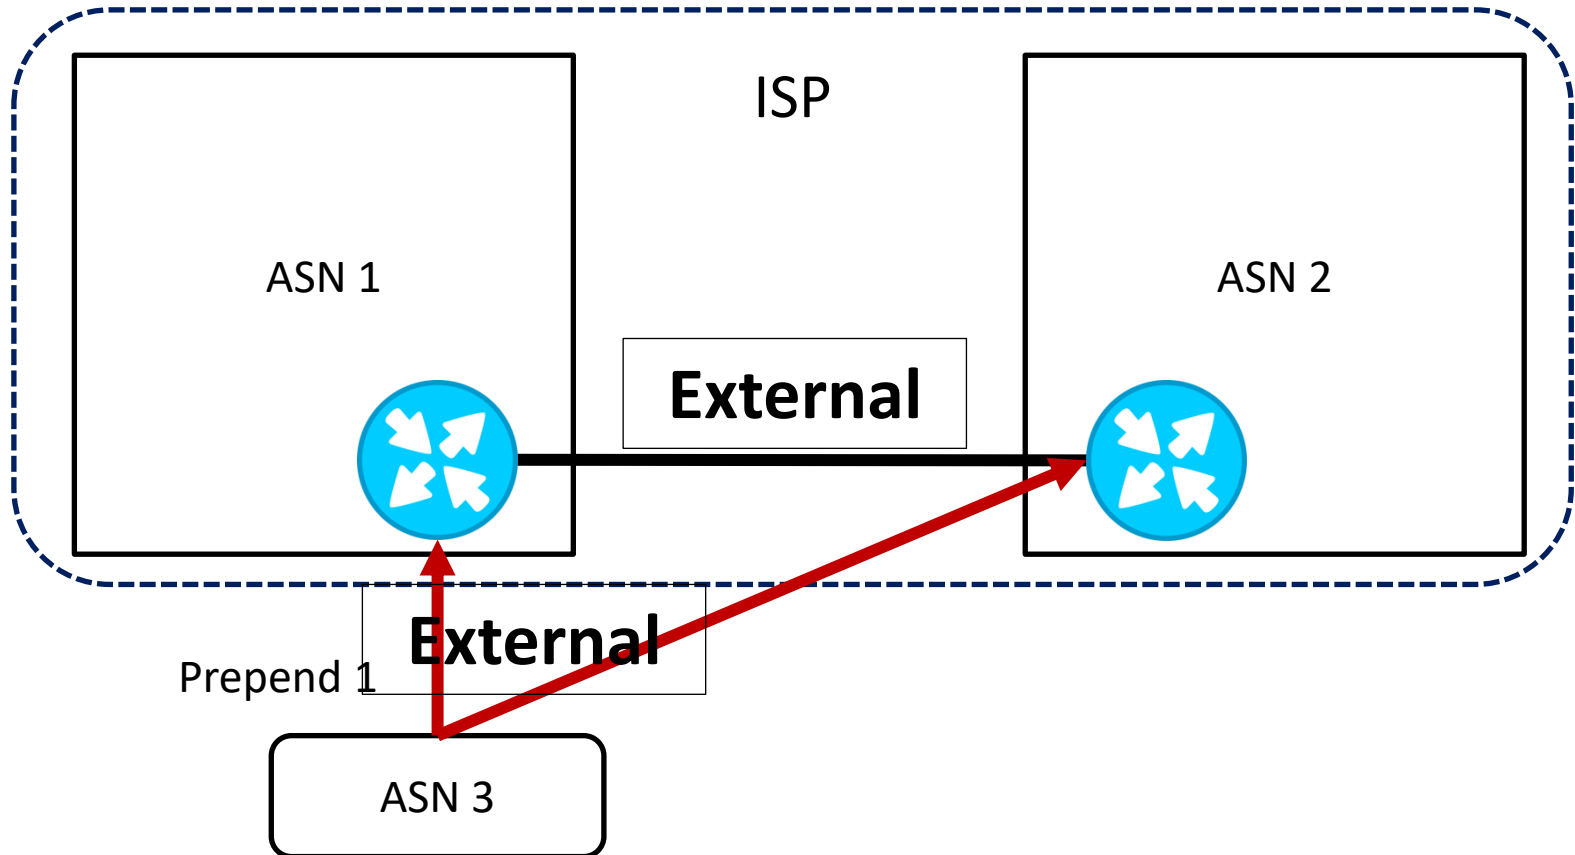
	ASN1	ASN2
a.	Assign pref	
b.	Community -> pref	

Pref Transmission




	ASN1	ASN2
a.	Assign pref	
b.	Community -> pref	
c.		Community -> Pref

'Hot potato' routing




	ASN 1	ASN 2
AS Path	3 3	3

'Hot potato' routing



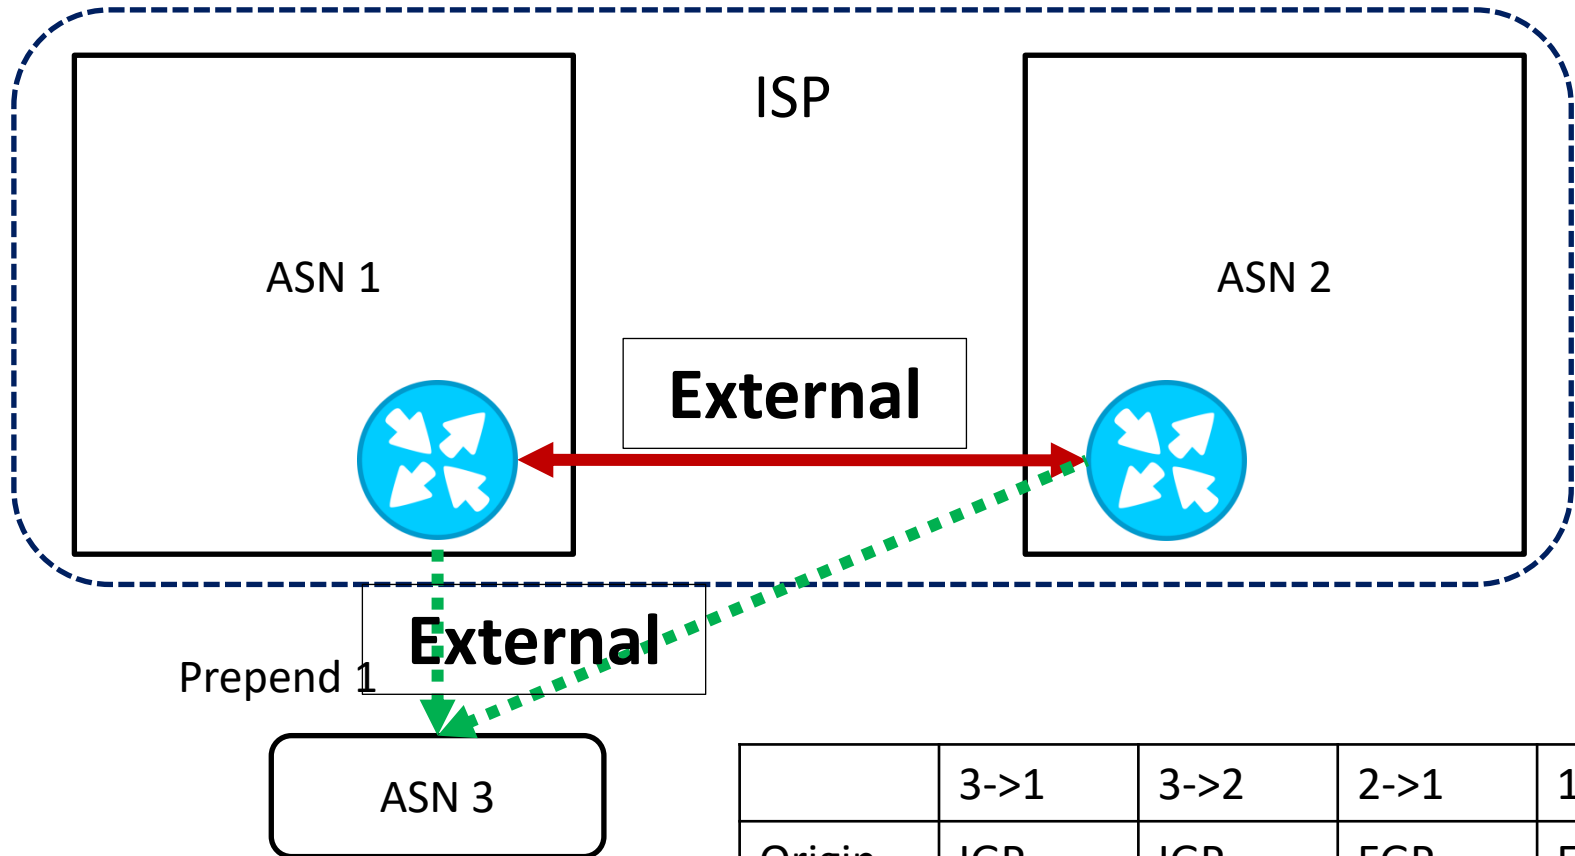
	ASN 1	ASN 2
AS Path	3 3	3

'Hot potato' routing



	ASN 1	ASN 2
AS Path	2 3	3

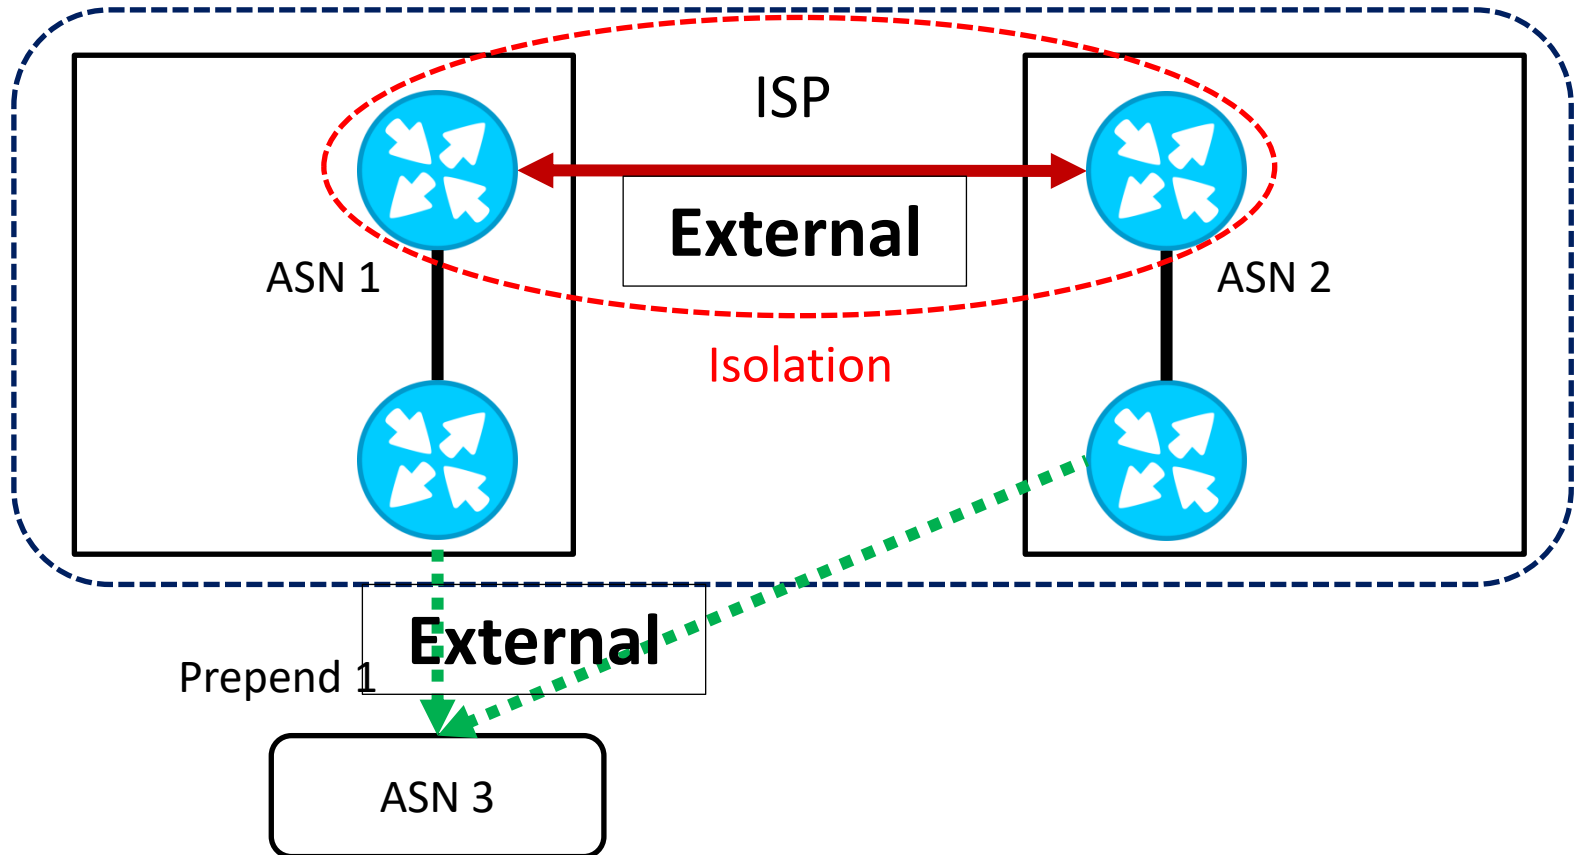
Securing 'Hot potato' routing



	3->1	3->2	2->1	1->2
Origin	IGP	IGP	EGP	EGP

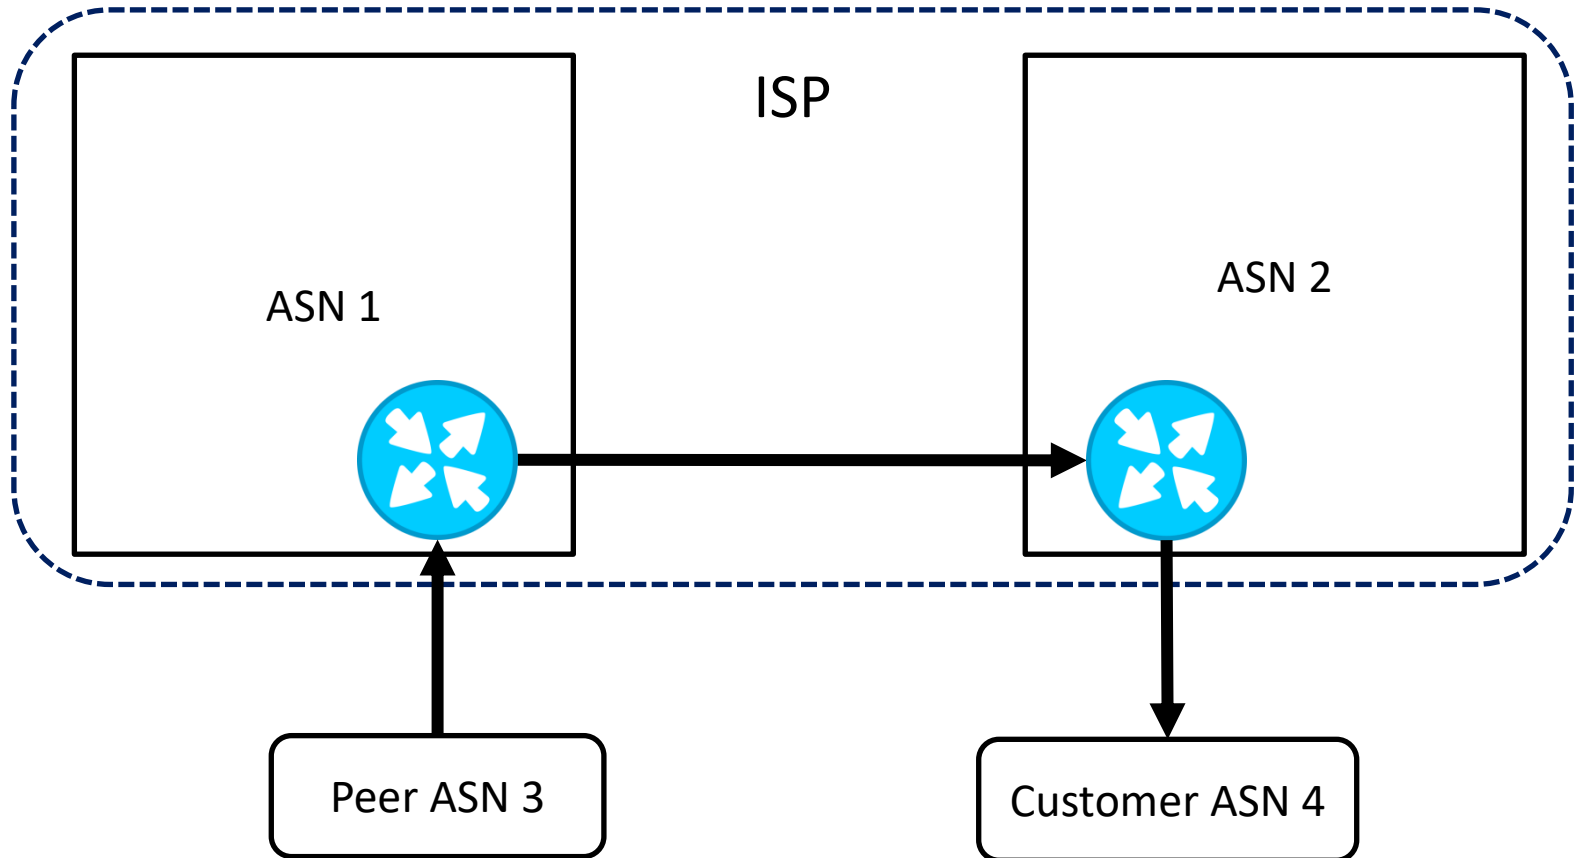
	ASN 1	ASN 2
AS Path	3 3	3

Securing 'Hot potato' routing



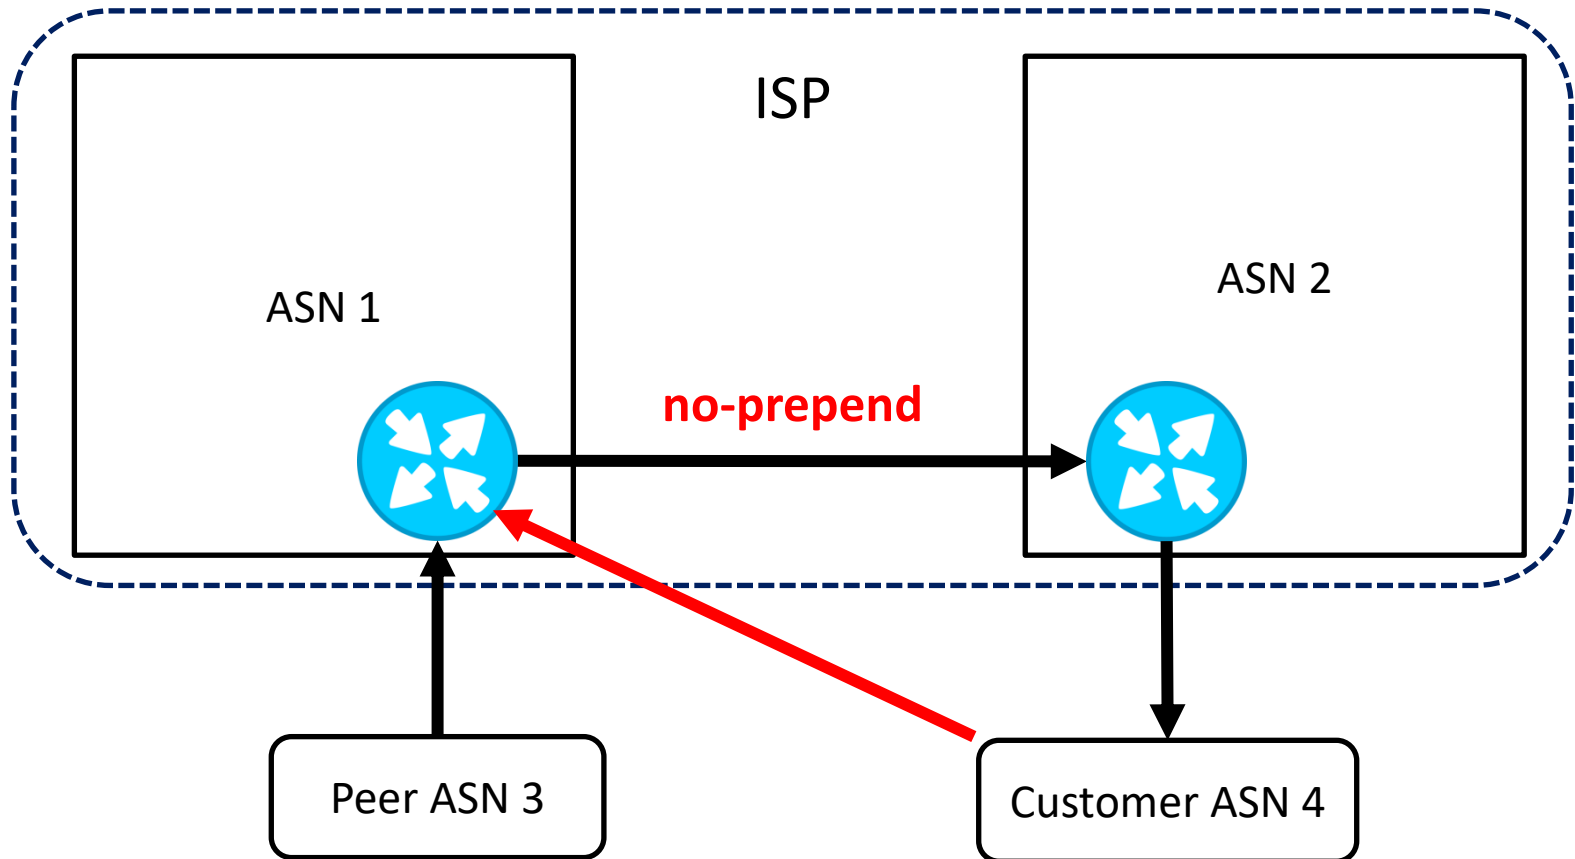
	ASN 1	ASN 2
AS Path	3 3	3

Path Length: Increased



	ASN 1	ASN 2	ASN 4
AS Path	3	1 3	2 1 3

Path Length: Route Loops



	ASN 1	ASN 2	ASN 4
AS Path	4 2 3	3	2 3

Multiple ASNs – One Network


Target	Plain BGP
Preference transmission	1. Synchronize Local Preference Values 2. Local-pref->Community->Local_pref
'Hot potato' routing	Origin modification / router isolation
Path length	Increased / Loop Risk
Route Leak prevention	Community synchronization

There is no simple way!

rfc7705: Autonomous System Migration Mechanisms

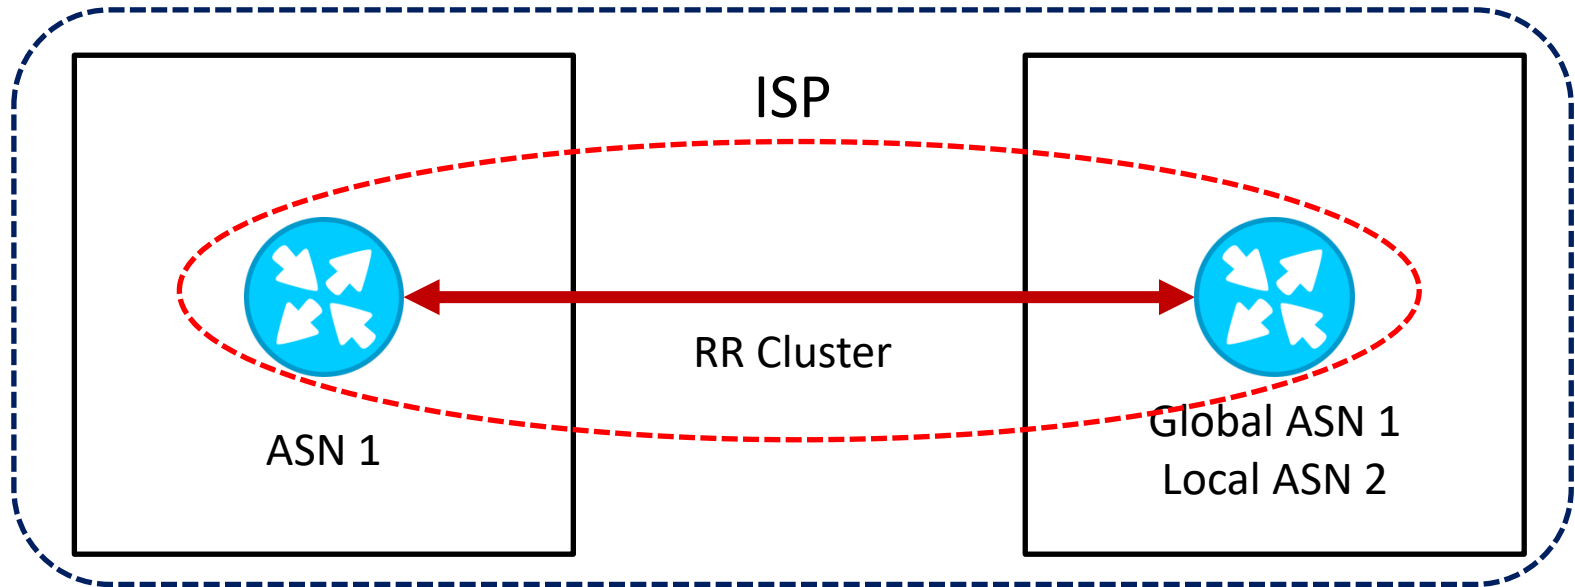


rfc7705: Local ASN capability



BGP session will be established if:
ASN in OPEN message = Local ASN **OR**
ASN in OPEN message = Global ASN

rfc7705: Migration ASN2 to ASN1



1. Add Local AS capability to RR
2. Unite **ALL** RR in one cluster
3. Slow migration to ASN 1

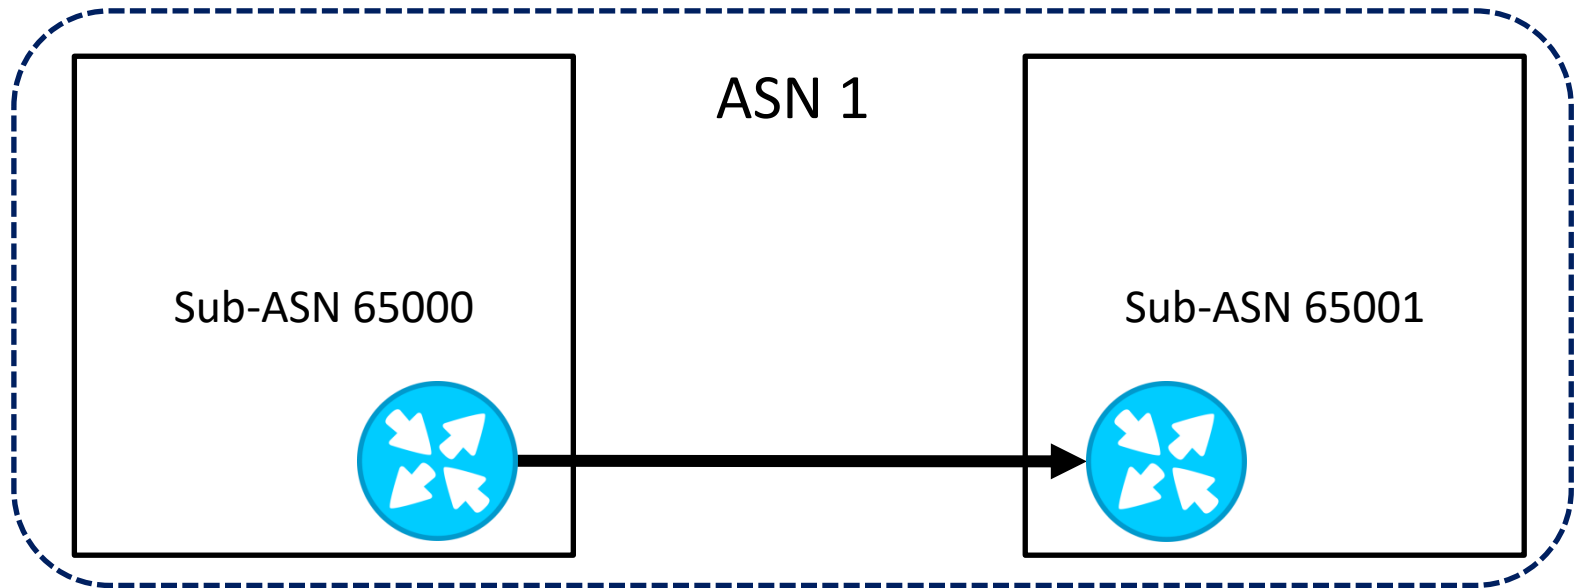
rfc7705: Multiple ASNs – One Network

Target	rfc 7705
Preference transmission	1. Synchronize Local Preference Values 2. Unite all RR in cluster
‘Hot potato’ routing	Built in
Path length	Partial increased / Loop Risk
Route Leak prevention	Community synchronization

What if you have single RR in legacy AS?

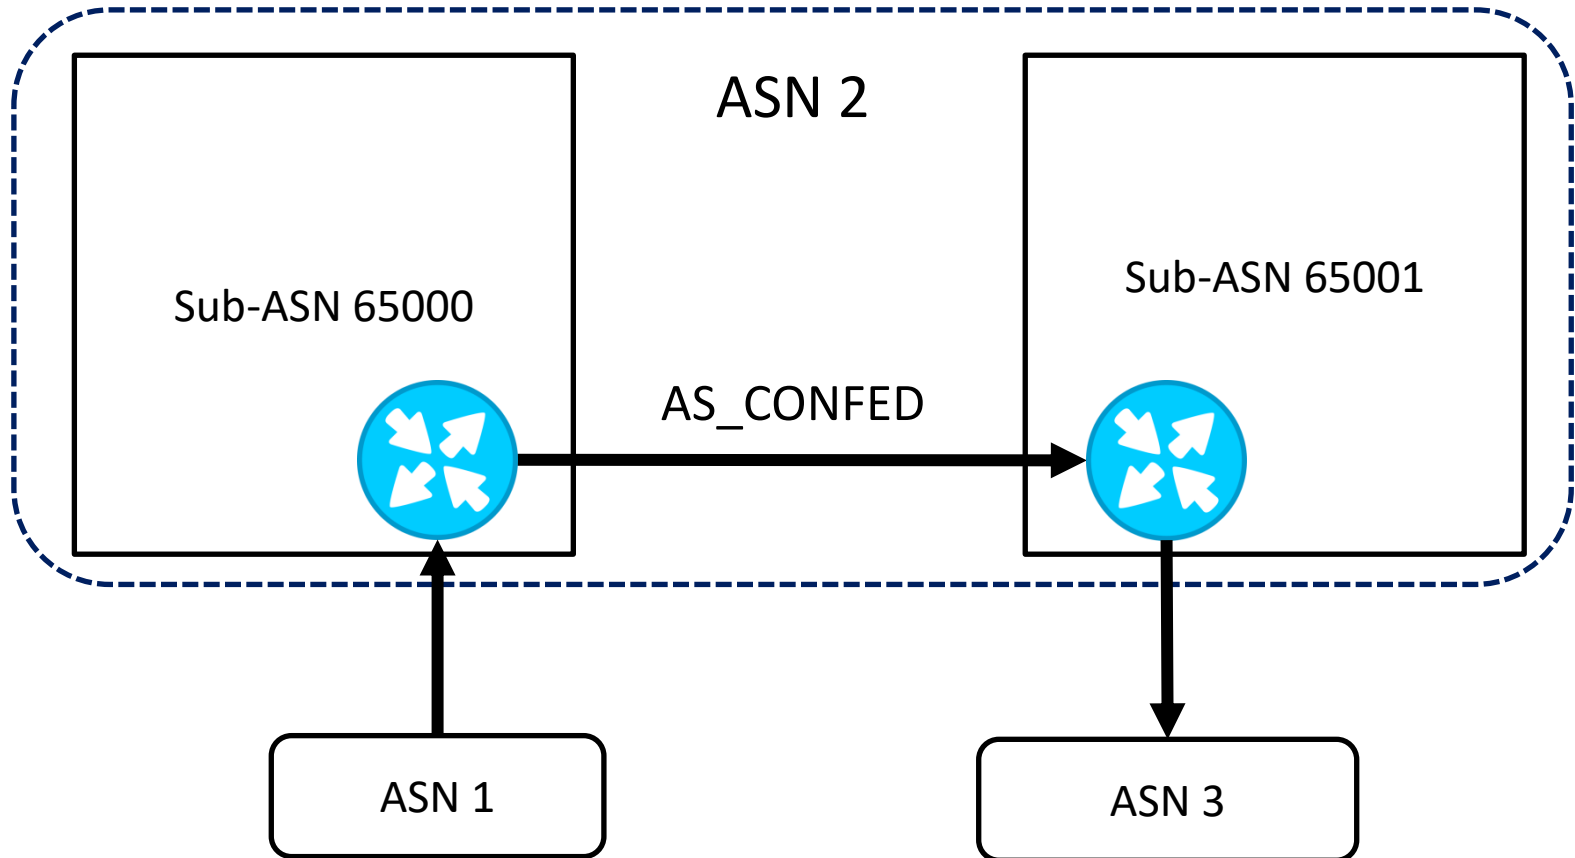
What if you have more than two networks to migrate?

rfc5065: ASN Confederations



A collection of autonomous systems represented and advertised as a single AS number to BGP speakers that are not members of the local BGP confederation.

rfc5065: Internal Routing



	ASN 1	ASN 65000	ASN 65001	ASN 3
AS Path	1	1	(65000) 1	2 1

rfc5065: Multiple sub-ASNs – One ASN

Target	rfc5065
Preference transmission	1. Synchronize Local Preference Values 2. Synchronize confederation list
‘Hot potato’ routing	Built in
Path length	Not increased
Route Leak prevention	Community synchronization

How to extend to multi-ASN level?


ISP Border Definition

BGP Session is Internal if:

Neighbor ASN = Local ASN	Plain BGP
Neighbor ASN = Global or Local ASN	rfc5065
Neighbor ASN in Confederations list	rfc7705

Next technical problem – next ad-hoc?

The Neighbor Roles




4 pairs of non-conflict roles:

1. Peer <---> Peer
2. Customer <---> Provider
3. Complex <---> Complex
4. **Internal <---> Internal**

Preventing Route Leaks

iOTC – internal Only To Customer attribute (zero size)




Internal Session
No iOTC change

Set iOTC if neighbor is
customer or peer

Filter routes if iOTC is set and
neighbor is customer or peer


If route was learned from a provider or peer it should not be
announced to another provider or peer

ISP Border Definition



Enable transmission of local attributes through
If both sides are **Internal**

Securing 'Hot potato' routing

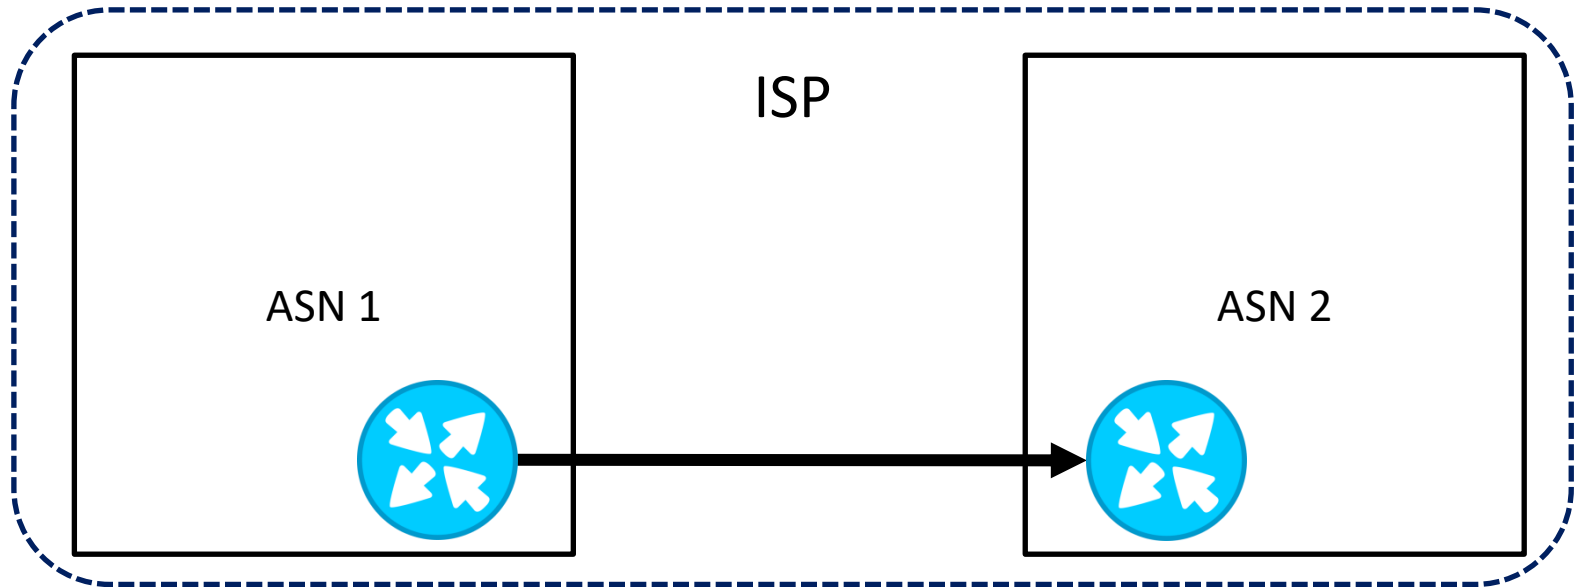


	ASN 1	ASN 2
AS Path	3 3	3

Multiple ASNs – One Network

Target	BGP + roles
Preference transmission	1. Synchronize Local Preference Values 2. Synchronize roles
'Hot potato' routing	Built in
Path length	Increased / Loop Risk
Route Leak prevention	Built in

ISP Confederations



A collection of autonomous systems ~~represented and~~ advertised as a single AS number BGP speakers that are not members of the local BGP confederation.

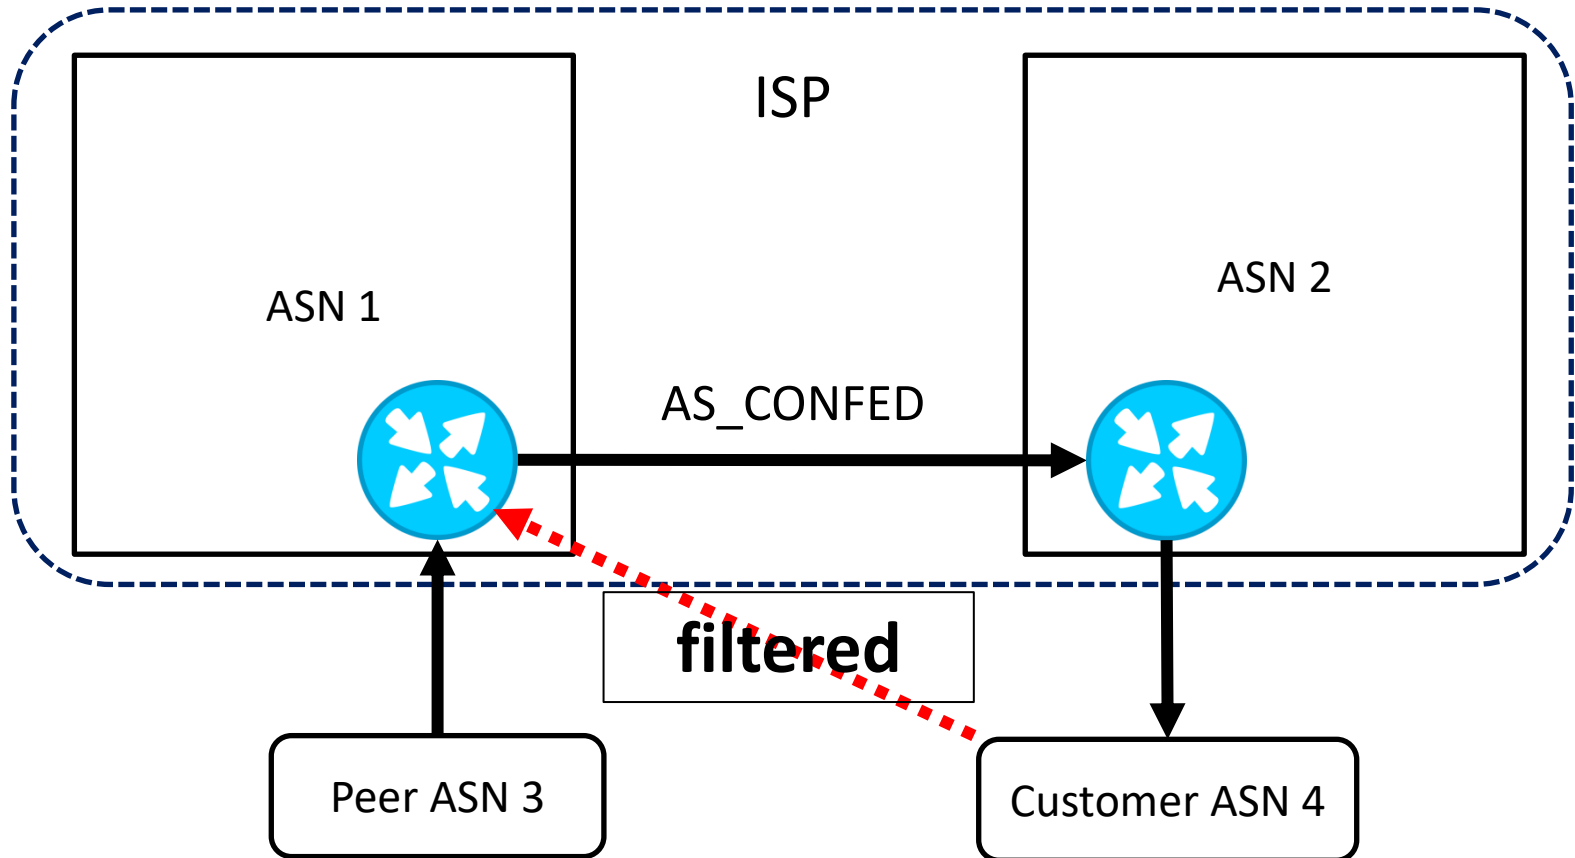
Loop Prevention

Internal Routing – likewise ASN confederation;

External Routing – filter routes with AS Path
segments from confederation list;

These two rules guarantees loop free scenario!

Path Length



	ASN 1	ASN 2	ASN 4
AS Path	3	3 (1)	2 3

A Simple Config

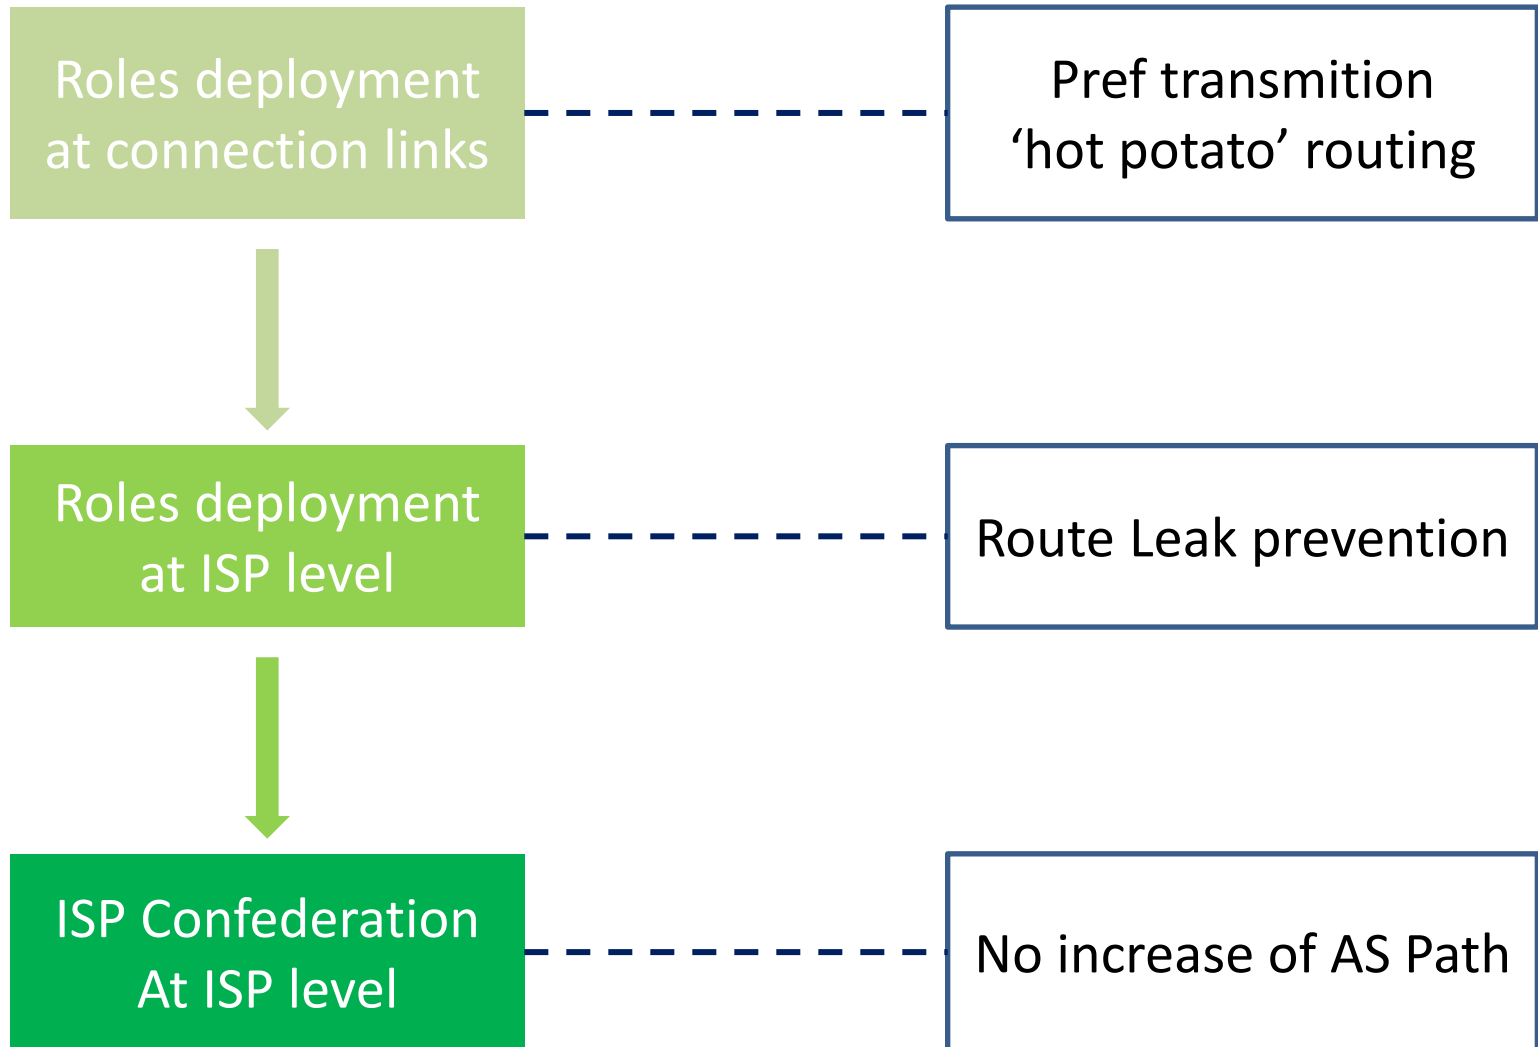
```
protocol bgp IAMOPERATOR_1 {  
    local as MY_AS1;  
    neighbor X.X.X.X as MY_AS2;  
    role internal;  
    confederation MY_AS1, MY_AS2;  
}
```

Multiple ASNs – One Network

Target	BGP + Roles + ISP Confederations
Preference transmission	<ol style="list-style-type: none">1. Synchronize Local Preference Values2. Synchronize roles3. Synchronize Confederation List
'Hot potato' routing	Built in
Path length	Not increased
Route Leak prevention	Built in

There was no simple way!

Migration Process



Summary

- ASN union vs ASN migration;
- Utilization with BGP confederations;
- Simplified configuration with no AS Path increase;
- Route Leak prevention;
- No affect on networks with single ASN.

Useful Links

- [Draft, description](#) and [implementation](#) of roles and route leak mitigation;
- [Description](#) and [implementation](#) of ASN union.