

# Enterprise Multihoming *back to the future ?*

Olivier Bonaventure

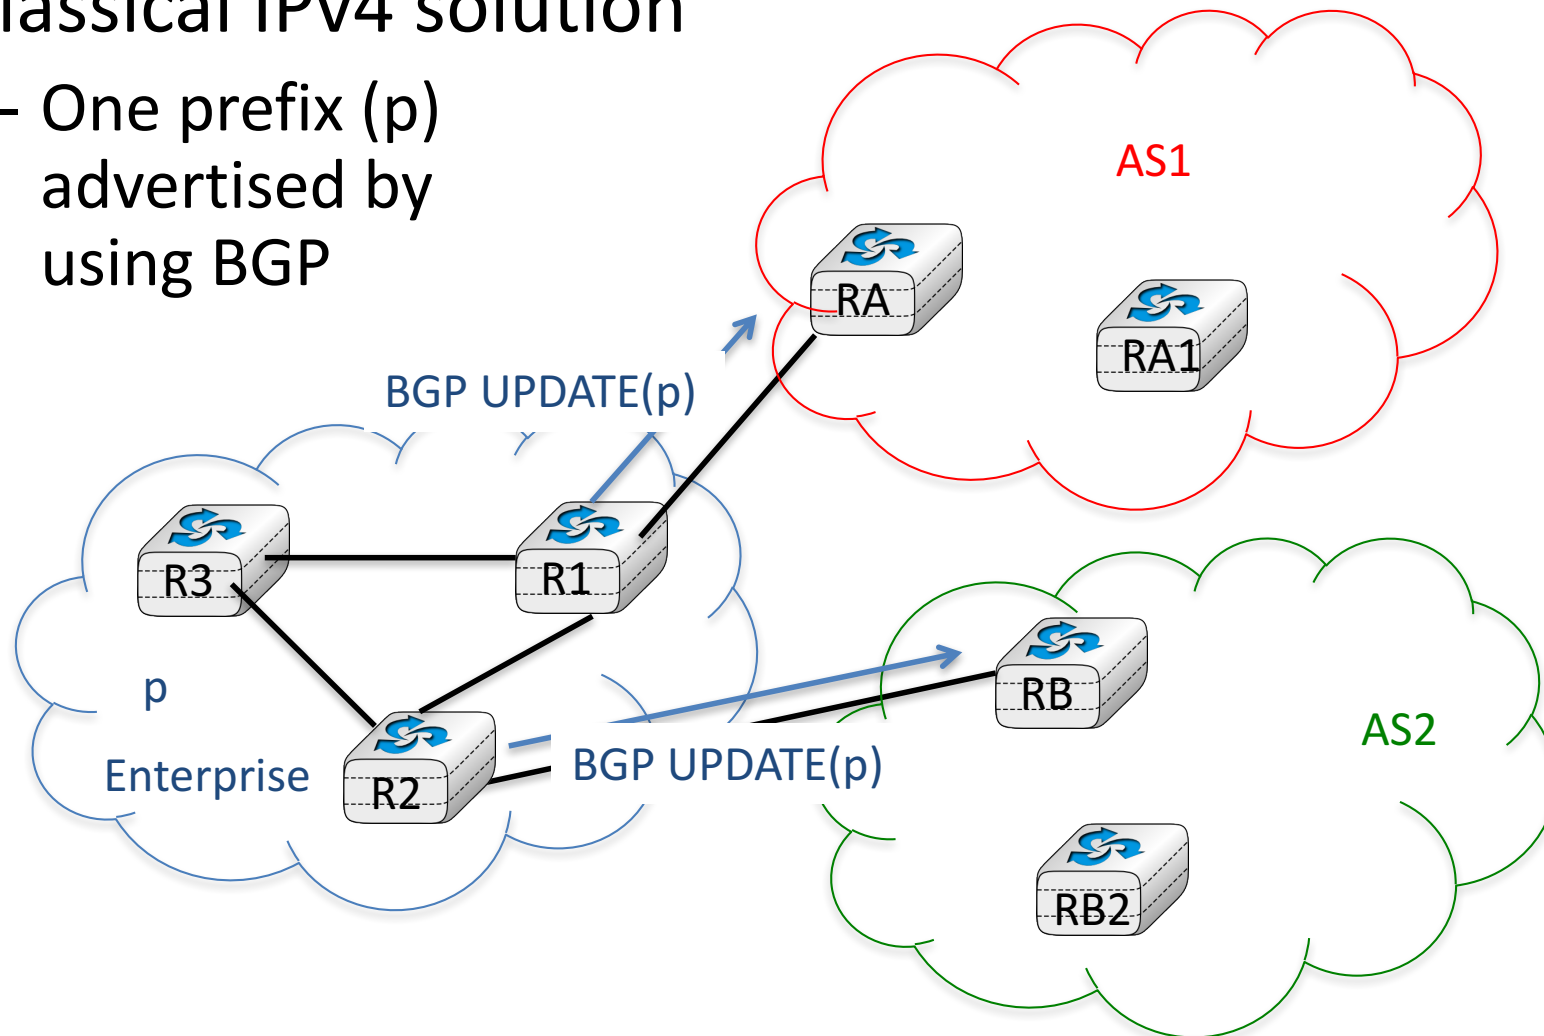
Mathieu Jadin

# Enterprise multihoming

- Why do enterprise want to be multihomed ?
  - Technical reasons
    - Redundancy against link/router failures
    - Performance
  - Economical reasons
    - Redundancy against providers becoming bankrupt
    - Cost

# Multihoming with PI prefixes

- Classical IPv4 solution
  - One prefix (p) advertised by using BGP



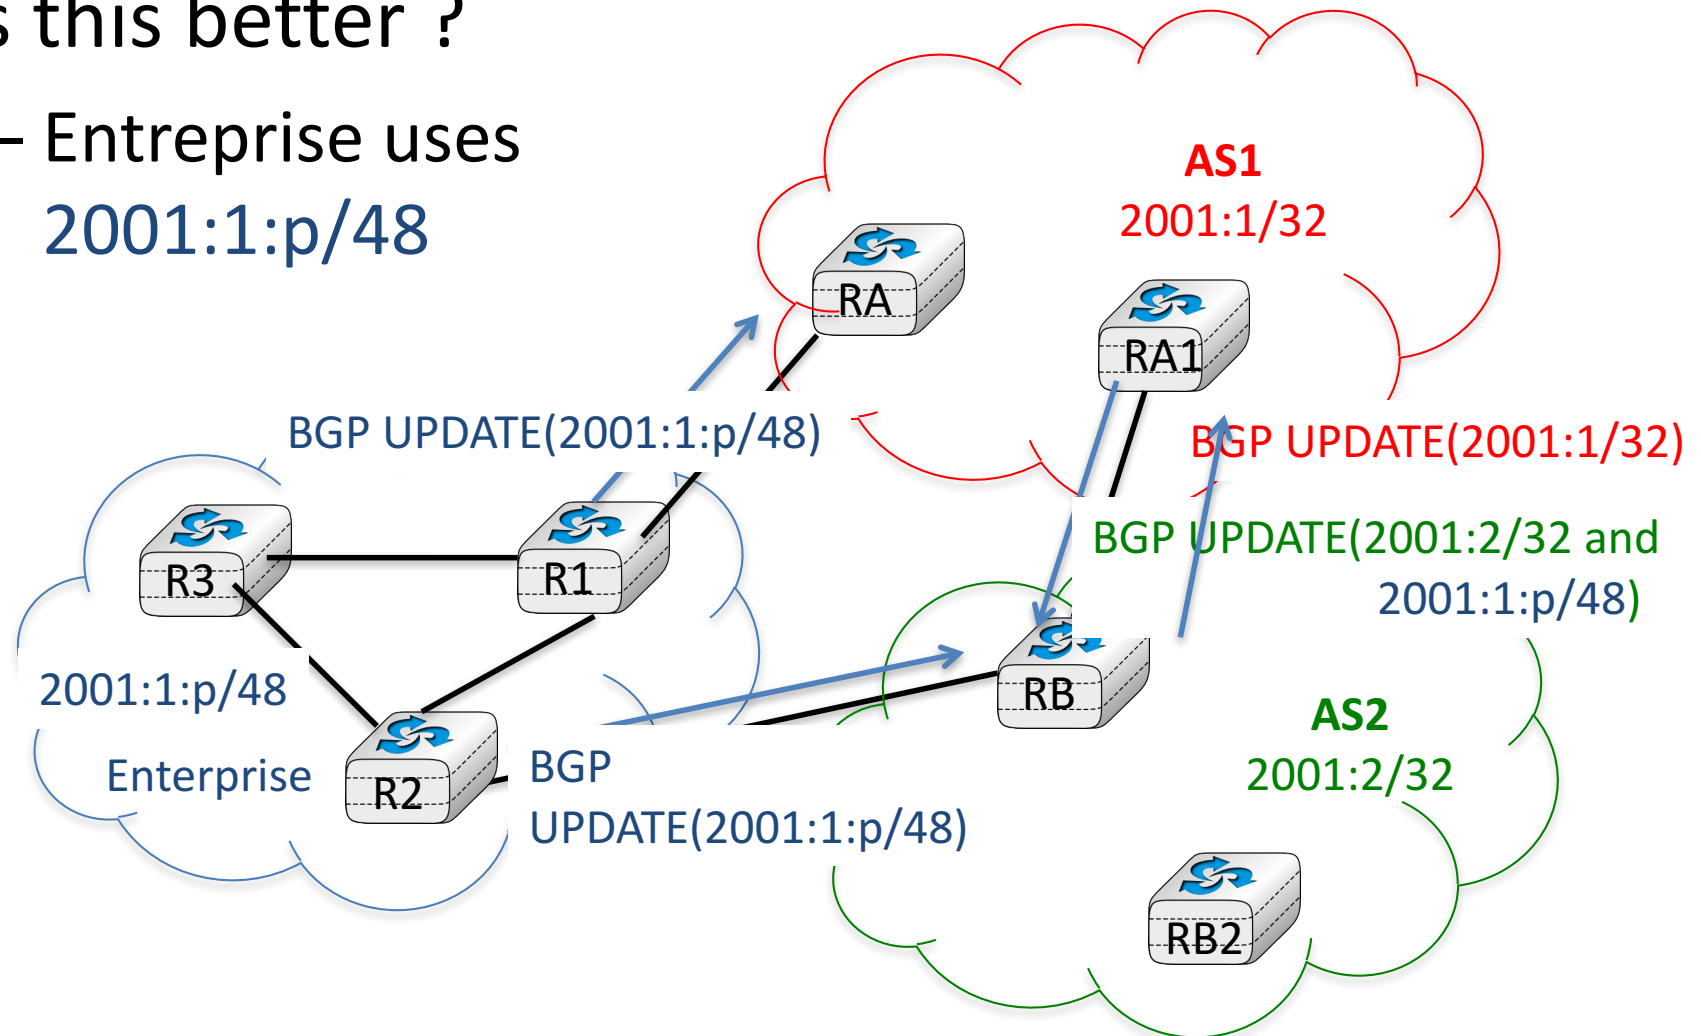
# Limitations of PI Multihoming

- Some enterprises may have difficulties in registering a PI prefix
- Traffic engineering
  - Control of the outgoing traffic is trivial
  - Control of the incoming traffic is much more difficult
    - AS Path prepending
    - More specific prefixes
    - BGP Communities
- Internet Routing scalability
  - All ASes need to carry your BGP routes
  - **Large ISPs bare the cost of growing BGP routing tables**



# Multihoming with one PA prefix

- Is this better ?
  - Enterprise uses 2001:1:p/48

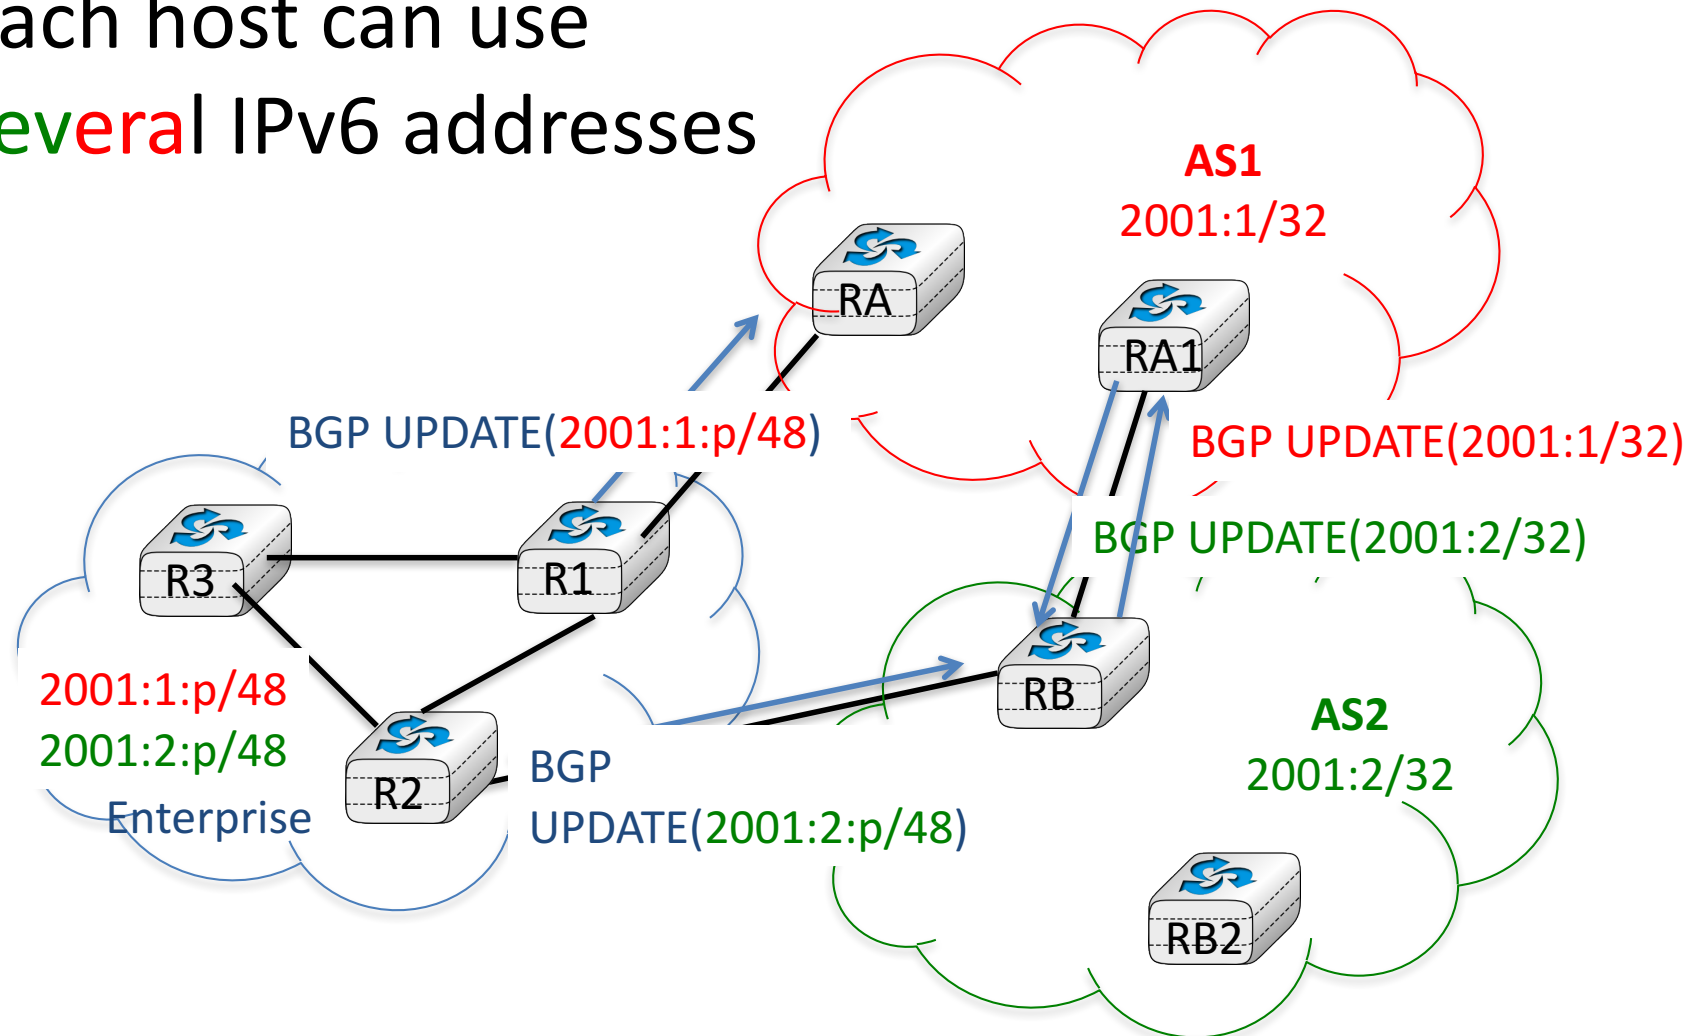


# Limitations of one prefix PA Multihoming

- Enterprise is stuck with the provider that allocated the prefix
  - Changing this provider requires renumbering
- Traffic engineering
  - Not really better
- Scalability
  - Not really better since AS1 also needs to advertise the enterprise prefix in addition to its own

# Multihoming with **several** PA prefixes

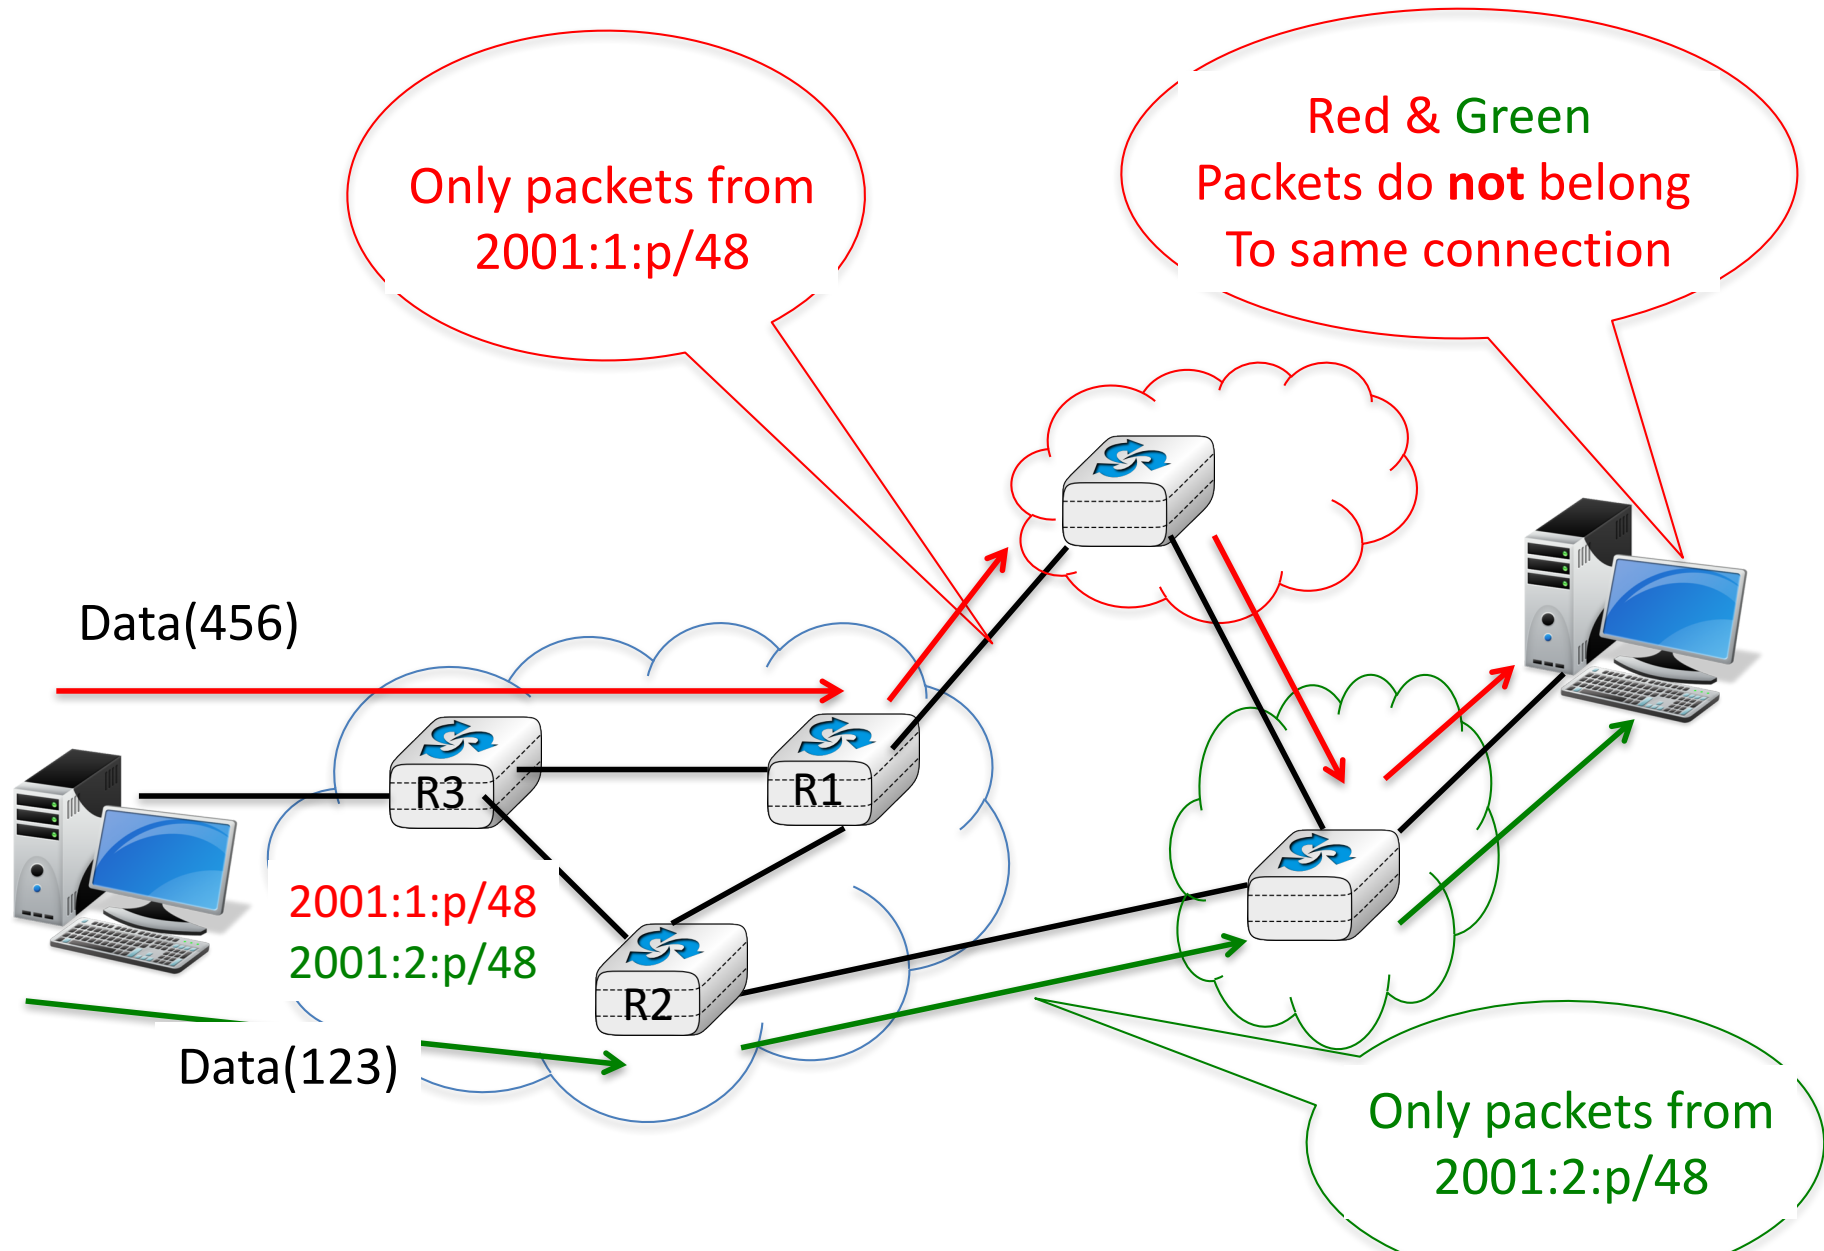
- Each host can use **several** IPv6 addresses



# Challenges for **several** PA Multihoming

- TCP support
  - TCP expects a single address on each connection
- Failures
  - How to cope with the failure of
    - A link to one provider, a router, An entire upstream provider
- Traffic engineering
  - How to control outgoing traffic ?
    - Selecting the best provider (delay, throughput, ...)
    - Load balancing
  - How to control incoming traffic ?
    - Selecting the best provider (delay, throughput, ...)
    - Load balancing

# TCP and BCP38



# Network Layer solution : shim6

Network Working Group  
Request for Comments: 5533  
Category: Standards Track

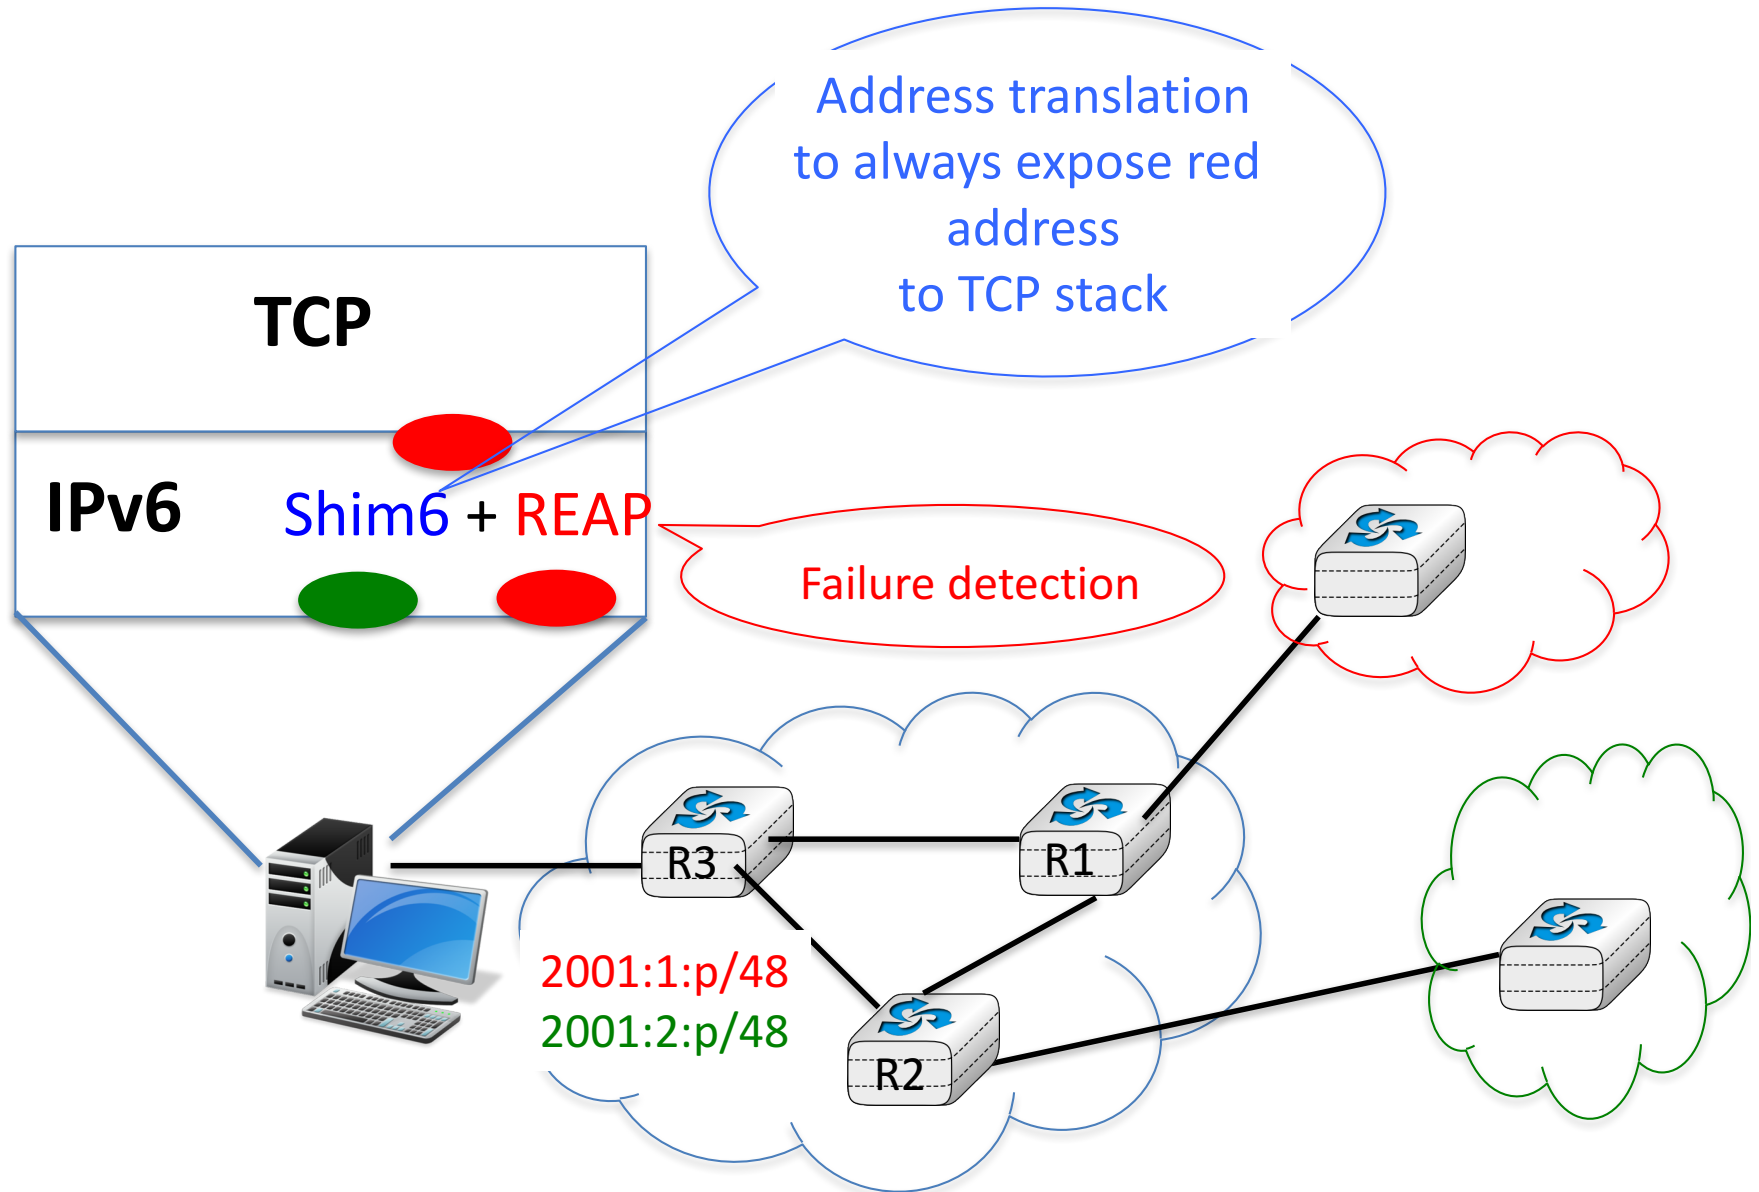
E. Nordmark  
Sun Microsystems  
M. Bagnulo  
UC3M  
June 2009

## **Shim6: Level 3 Multihoming Shim Protocol for IPv6**

### Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

# Shim6 in one slide



# shim6 in practice

- Implemented in the Linux kernel
  - But not in other stacks ...
- Failure detection, but no traffic engineering capability



Computer Communications  
Volume 34, Issue 14, 1 September 2011, Pages  
1685-1695



## Implementation and evaluation of the Shim6 protocol in the Linux kernel

S. Barré <sup>a</sup>  , J. Ronan <sup>b</sup> , O. Bonaventure <sup>a</sup> 

 [Show more](#)

<https://doi.org/10.1016/j.comcom.2011.03.005>

[Get rights and content](#)

### Abstract

In the changing landscape of the today's Internet, several solutions are under investigation to allow efficient, flexible and scalable multihoming. One of the proposals is shim6, a host-based multihoming solution based on the use of multiple IPv6 addresses on each host. In this work, we first describe the main



# One layer above

- Can we move multihoming support in the transport layer ?
  - Transport needs to support multiple IP addresses
    - TCP and UDP don't work as is
    - SCTP,  
Multipath TCP  
Multipath QUIC  
have all the  
required features

Internet Engineering Task Force (IETF)  
Request for Comments: 6824  
Category: Experimental  
ISSN: 2070-1721


A. Ford  
Cisco  
C. Raiciu  
U. Politechnica of Bucharest  
M. Handley  
U. College London  
O. Bonaventure  
U. catholique de Louvain  
January 2013

## **TCP Extensions for Multipath Operation with Multiple Addresses**


### **Abstract**

TCP/IP communication is currently restricted to a single path per connection, yet multiple paths often exist between peers. The simultaneous use of these multiple paths for a TCP/IP session would improve resource usage within the network and, thus, improve user experience through higher throughput and improved resilience to network failure.


# Can we do better ?



# BCP38 !

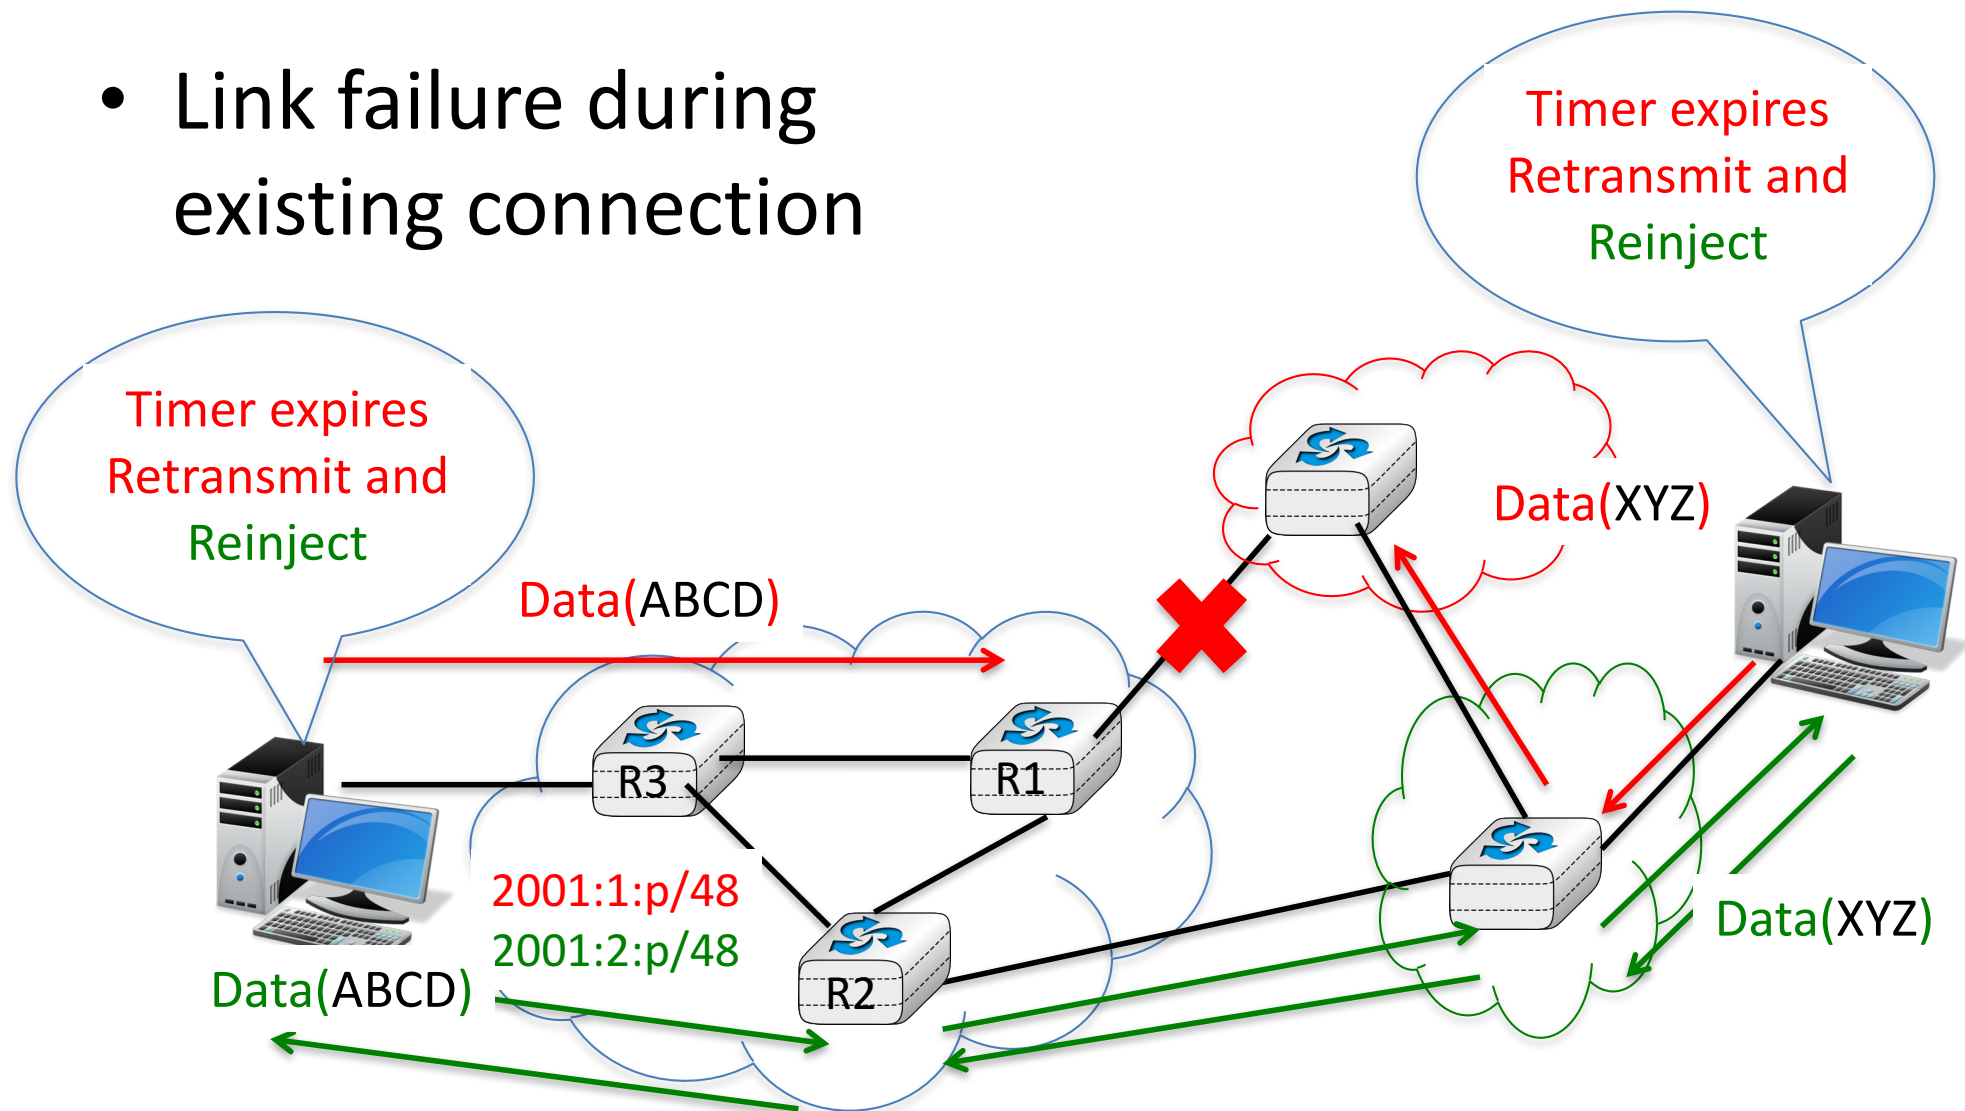


# Connection establishment with Multipath TCP



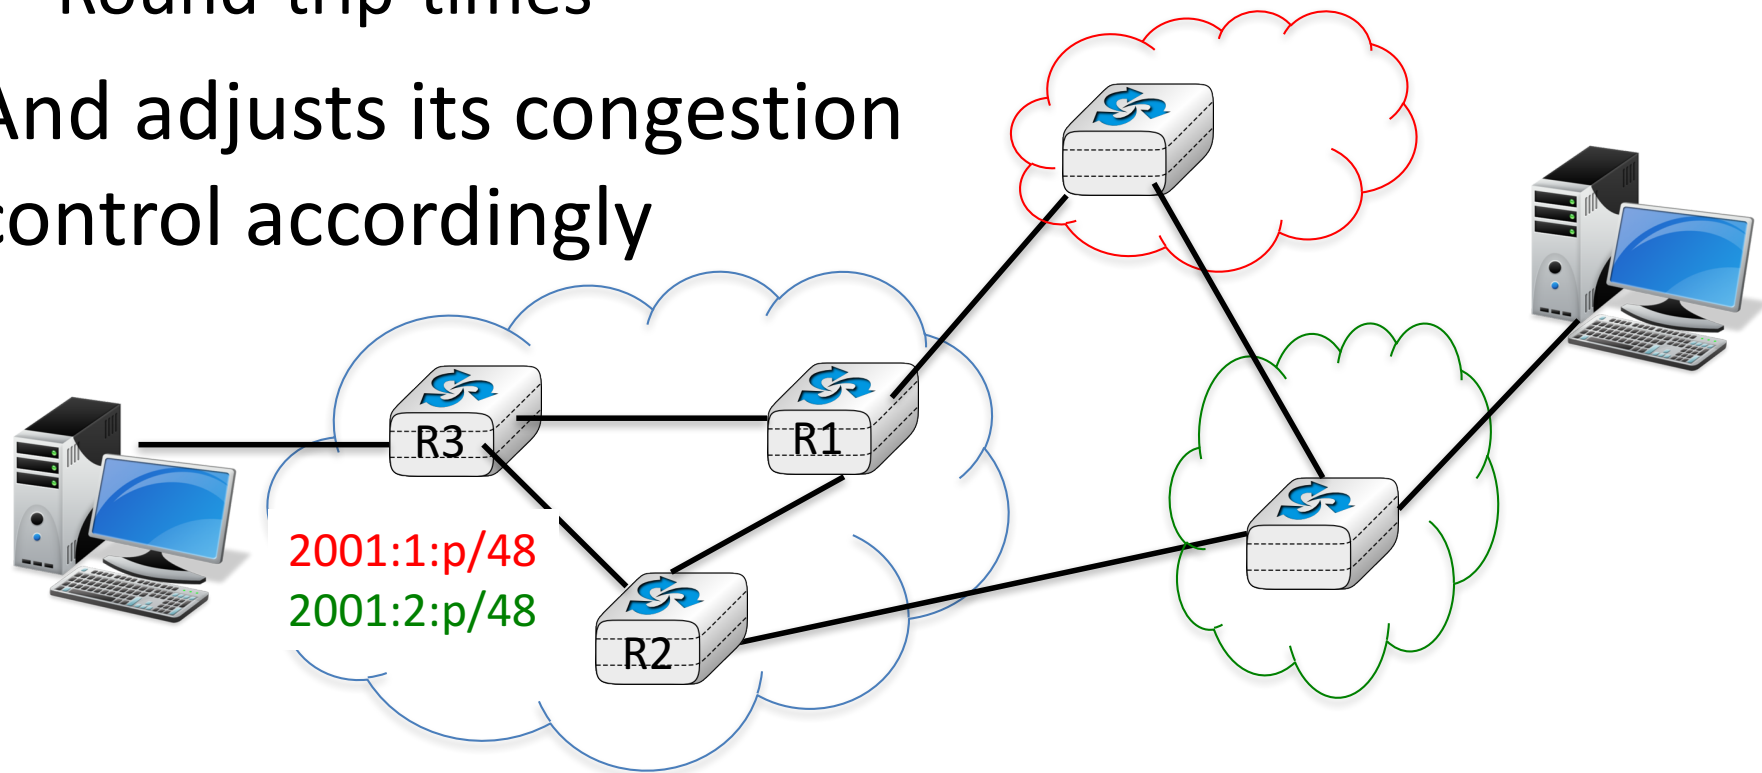
# How to cope with failures ?

- Link failure during existing connection



# Traffic engineering ?

- Multipath TCP naturally measures
  - Packet losses
  - Round-trip-times
- And adjusts its congestion control accordingly





# Multipath TCP makes Multihoming with **several** PA possible

- TCP support
  - Multipath TCP copes naturally with multiple addresses
- Failures
  - Multipath TCP detect failures via packet losses or ICMP and reacts by moving traffic away from failures
    - Within a few round-trip-times
- Traffic engineering
  - Controlling outgoing traffic
  - How to control incoming traffic ?
- Implementations
  - Linux : <http://www.multipath-tcp.org>
  - Apple iOS for all apps since iOS11

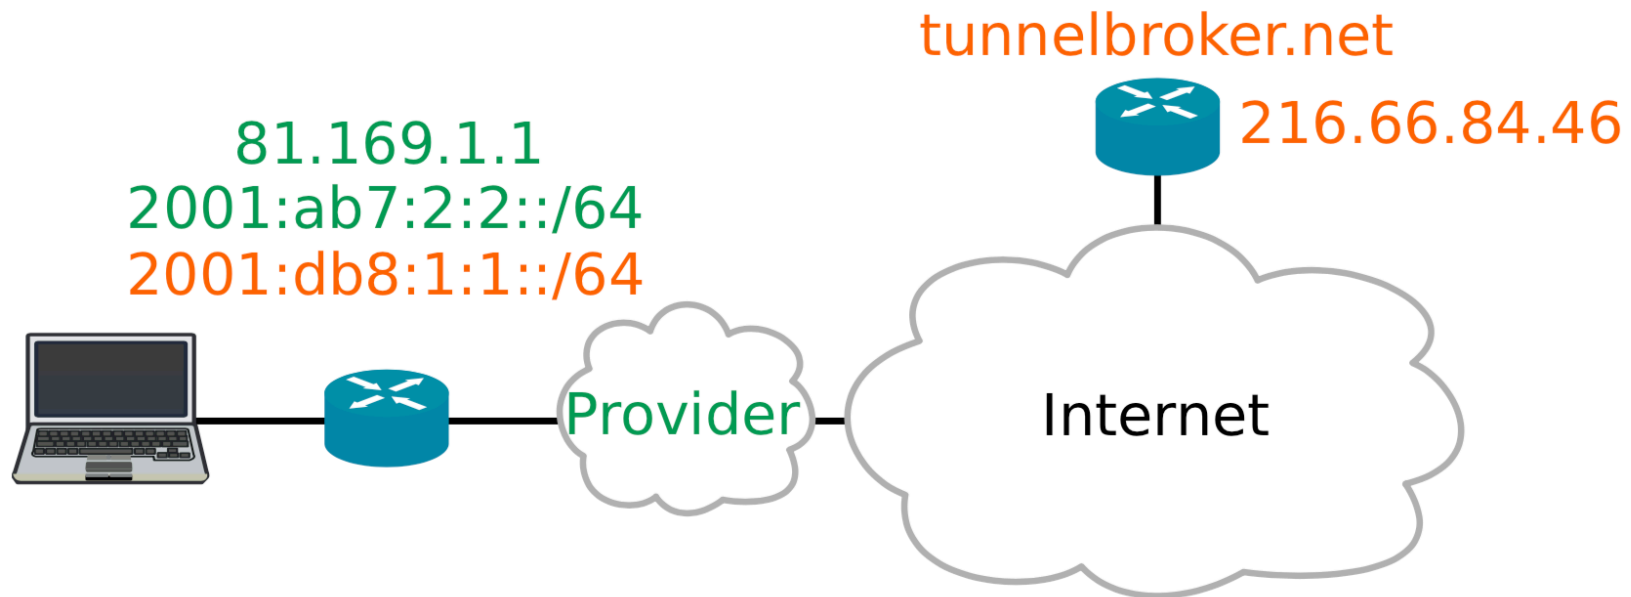
# Testing IPv6 Multihoming at home

- How to test IPv6 multihoming with a single link ?

81.169.1.1  
2001:ab7:2:2::/64




# Use 6 over 4 tunnels






# Default provider with native IPv6



# Second provider over tunnel



# Practical issues

- HE's tunnel broker might not be perfect
  - IPv4 address must be pingable
  - CPE router must allow protocol 41
- SIXXS.net would have been perfect



**SixXS** Main | About | Contact | News | User Home | PoPs | Presentations | FAQ | Forum | Tools

---

Summary | Rationale | Conclusion | Faq | Timeline

---

## Sunsetting SixXS


- Author: Pim van Pelt, Jeroen Massar
- Contact: <[staff@sixxs.net](mailto:staff@sixxs.net)>
- Date: March 2017
- Status: Draft | Review SixXS | Review Admins | Final | **Published**

### Summary

SixXS will be sunset in H1 2017. All services will be turned down on **2017-06-06**, after which the SixXS project will be retired. Users will no longer be able to use their IPv6 tunnels or subnets after this date, and are required to obtain IPv6 connectivity elsewhere, primarily with their Internet service provider.

# How to participate ?

- What you'll need



RaspberryPi



- We'll provide
  - Software images
  - Routing configuration
  - MPTCP support
  - Monitoring tools
- Contact: [mathieu.jadin@uclouvain.be](mailto:mathieu.jadin@uclouvain.be)