



**Euro-IX
route
server
activity
report**

RIPE-60

May 2010



Route servers: an update on Euro-IX activity

Nick Hilliard
CTO, INEX

Dogsbody, RS-VWG



**Euro-IX
route
server
activity
report**

RIPE-60

May 2010



Brief background

- Exchange growth has driven IXP users to use route servers in large numbers
- Quagga was the only viable option until mid-2009
- IXP funding was needed to develop required functionality in BGP stacks
- Four candidates were identified: Quagga, OpenBGPD, BIRD, BOGART
- Unfortunately, no progress on BOGART



**Euro-IX
route
server
activity
report**

RIPE-60

May 2010



Software Development

- BIRD
 - developed at Charles University, Prague; now funded by CZ.NIC Labs
 - Route server functionality performing well in live systems
 - Advanced configuration scripting and filtering options available
- OpenBGPD
 - Part of base O/S on OpenBGPD
 - Recent route server functionality funded by AMS-IX
 - Doesn't run well on "not-OpenBSD"
 - Problems with very large prefix filters



**Euro-IX
route
server
activity
report**

RIPE-60

May 2010



Software Development

- Quagga
 - Stability problems found when operating under high load
 - Euro-IX RS-VWG has co-ordinated funding for developer time
 - Euro-IX plans to feed all work back to the community under GPL
 - Initial development has focused on two problem areas:
 - stability
 - threading
 - some optimisation





**Euro-IX
route
server
activity
report**

RIPE-60

May 2010



Quagga Development

- Created threading framework
 - BGP sessions
 - CLI sessions
 - RIB engine
- Optimisation
 - session drop / route withdrawal
 - outgoing update handling
 - prefix filter handling
- BGP session handling re-designed
 - collision management
 - threaded, so keepalives not dropped
 - separation of transport / session



**Euro-IX
route
server
activity
report**

RIPE-60

May 2010



Development Results

- Starts up fine on large systems
- Doesn't fall over when spanked with lots of updates (200 clients, 1000 pfx)
- Reads multi-megabyte files very fast
- CLI always responsive
 - some commands respond asynchronously
 - some commands will respond even when when BGP engine is busy
- session collision problems substantially reduced
- still some memory leakage



**Euro-IX
route
server
activity
report**

RIPE-60

May 2010



Future Development

- integrate dev work into Quagga tree after testing / peer review / etc
- support for BGP ADD-PATH
- optimisation of memory management
- optimisation of CPU usage
 - scaling goal: 1000 clients, average 500 prefixes per client
 - i.e. up to 500m client updates on startup
- establishment of a test environment
 - updates
 - session creation and teardown
- test environment not limited to Quagga



Questions?

**Euro-IX
route
server
activity
report**

RIPE-60

May 2010

