

Quagga

a GPL licensed IPv4/IPv6 Routing Software Suite
- a very short introduction and overview -

AS4 support for Quagga bgpd
- how to deal with 4 byte AS numbers in quagga -

Juergen Kammer <j.kammer@eurodata.de>

Quagga

- > What is it, anyway?
- > Installation and configuration
- > Hints, born out of personal experiences
- > Resources

Quagga: what is it, anyway?

What is it, anyway?

> wikipedia:

The quagga is an extinct subspecies of the plains zebra.

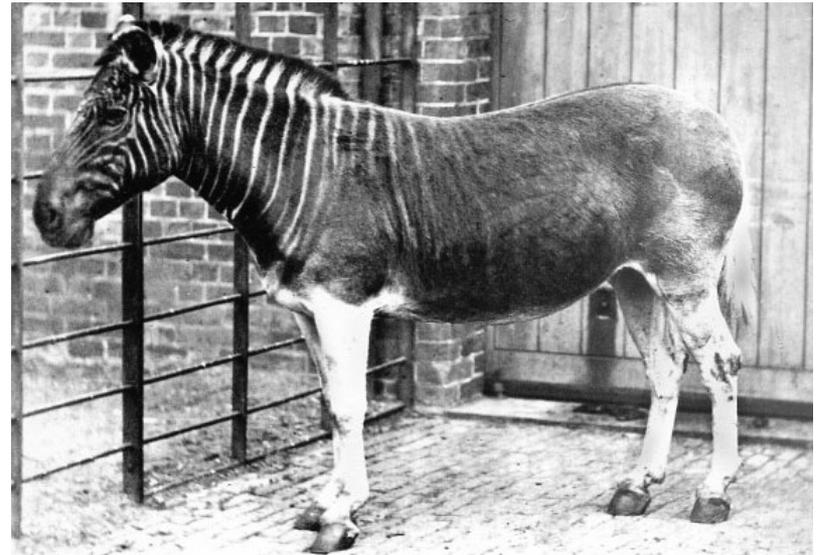
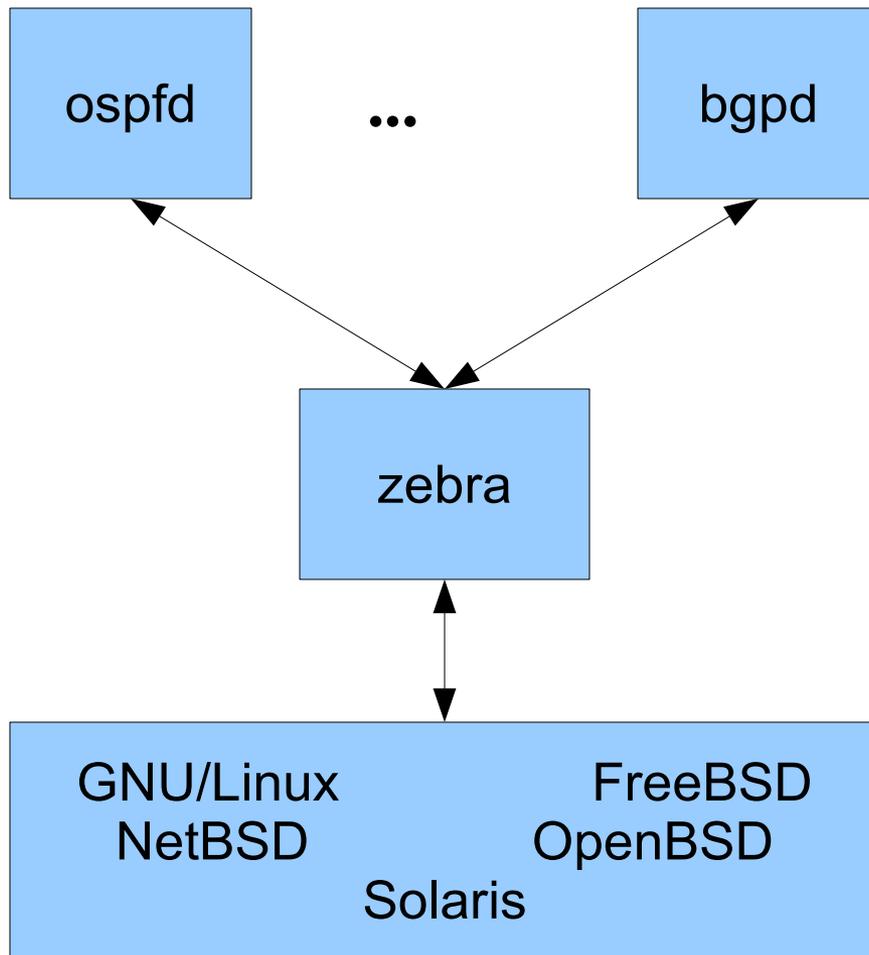


Photo F. York, London, Regent's Park ZOO, 1870
The only known photo of a living quagga.
(image in public domain, copyright expired)

What is it, anyway?

- > a fork of GNU zebra (written by Kunihiro Ishiguro)
- > not extinct 😊
- > initially “zebra-pj”, but forget that...
- > “The Quagga tree aims to build a more involved community around Quagga than the current centralised model of GNU Zebra” (from <http://www.quagga.net>)
- > there is a group of maintainers, and “the community”

Quagga: What is it, anyway?



- > one daemon per protocol
- > telnet interface to daemons
 - zebra (static routes & kernel interaction)
 - ospfd
 - bgpd
 - ospf6d
 - ripd
 - ripngd
 - isisd
- > vtysh
 - telnets to all daemons

Quagga

- > What is it, anyway?
- > **Installation and configuration**
- > Hints, born out of personal experiences
- > Resources

Installation

- > Linux
 - Debian packages, RPM packages, ...
- > Install from source
 - the usual “./configure, make, make install”
 - download via <http://www.quagga.net>
 - use releases, or cvs snapshots, or cvs directly
- > 0.99.7 latest unstable, 0.98.6 latest stable

Configuration

- > Configuration 90% like cisco
 - differences: trivial/specials/not implemented
 - command line completion like cisco
 - *if you know cisco CLI, you'll also manage quagga*

Quagga

- > What is it, anyway?
- > Installation and configuration
- > Hints, born out of personal experiences
- > Resources

Hints *(based on own experiences, YMMV)*

- > 0.99.<max> CVS Snapshots can be used productively
 - sometimes some features only available there
 - you will have luck, sometimes even bad such
 - *and that is why you should be as careful as you always are*

Quagga: Hints

- > Be aware that quagga itself *does not forward packets*, it just runs the routing protocols
 - reverse path filtering of your OS?
 - ip forwarding activated?
 - what about your firewall?
 - Linux: `/proc/sys/net/ipv4/netfilter_ip_conntrack_max`
 - Kernel - Zebra interaction problems?
 - Linux: `/proc/sys/net/core/rmem_{max,default}`
 - Linux: `/proc/sys/net/ipv4/igmp_max_memberships`
- > Number of neighbors in BGP near 1024?
 - `FD_SETSIZE` is 1024, quagga uses a select with it

Quagga: Hints

- > Relatively new features in 0.99.x
 - fast ospf (if you have to have it...)
 - ospf: passive interface default (problems when introduced)
- > Closing up to 1.0.0
 - only very few outstanding patches / known bugs
 - some autoconf-cleanups are just going into CVS
- > Latest IMHO hefty thingy resolved in CVS head as of 02/2007
 - bgpd - race condition on route flap, re-appearing route may be lost
- > 0.99.7 came out on 20.04.2007 - brand new!
- > Active community, very eager to squash bugs.

Quagga: Hints

- > *MD5 in quagga bgpd, on linux*
 - MD5 long time kernel patch by Hasso Tepper
 - md5qd
 - MD5 is in the newest linux kernels (>2.6.20)
 - http://www.solinno.co.uk/md5sig/quagga_linux-2.6.20_md5sig.diff
(Leigh Brown, December 2006)
 - not many experiences (seen one “works perfectly”)
 - work currently in progress
 - may be included in quagga 1.0.0

Quagga

- > What is it, anyway?
- > Installation and configuration
- > Hints, born out of personal experiences
- > **Resources**

Quagga: Hints

- > <http://www.quagga.net>
- > <http://www.dishone.st/pipermail/zebra-pj-commit/>
 - mailing list archive, commits on quagga cvs
- > <http://bugzilla.quagga.net/>
 - bugzilla, look whether *your* bug is in there
- > <http://www.quagga.net/lists.php>
 - quagga-users, quagga-dev
- > <http://wiki.quagga.net/index.php/Main/AddRes>
 - additional resources: patches etc.

AS4 support for quagga

- > Motivation
- > Status
- > Feedback received
- > Problems / Caveats
- > Conclusion

Motivation

- > RIPE-53
 - Henk was desperately looking for quagga AS4 support
- > No takers at that time, just some “if and when”
- > Basically a project of determinable size
- > Opportunity to support quagga, which we rely on

AS4 support for quagga

- > Motivation
- > Status
- > Feedback received
- > Problems / Caveats
- > Conclusion

Status

- > Patch at <http://quagga.ncc.eurodata.de> since November 2006
- > Basic AS4 handling looks good:
 - stable through 6 patch versions
 - new versions to follow changes in quagga CVS
- > Merge into mainline planned
 - some issues to resolve
 - more feedback would be nice

Contents of the patch

- > Basic AS4 handling
 - draft-ietf-idr-as4bytes-13.txt
- > MRT dumps
 - see GROW and draft-ietf-grow-mrt-04.txt (08.03.07)
 - Erik Romijn (RIPE NCC) did this, thanks!
- > Extended communities:
 - draft-rekhter-as4octet-ext-community-01.txt

Contents of the patch

- > New debug commands:
 - debug bgp as4
 - debug bgp as4 segment
 - *beware of this with a full routing table:
logs several lines per route received!*

AS4 support for quagga

- > Motivation
- > Status
- > Feedback received
- > Problems / Caveats
- > Conclusion

Feedback received

- > Erik Romijn, RIPE NCC
 - RIS is based completely on AS4 capable quagga since 28.03.2007, did you notice?
- > Geoff Huston
 - interaction with AS4 patched OpenBGPd works (all permutations)
 - used quagga for further real-life tests because OpenBGPd got problems

Feedback received

- > There have been more downloads than these two!
 - but not much more feedback
- > Not many ASes > 65535 have entered the DFZ yet
- > Where are you, people? Still testing around?

AS4 support for quagga

- > Motivation
- > Status
- > Feedback received
- > Problems / Caveats
- > Conclusion

Problems / Caveats

- > AS number format
- > MRT dumps
- > Extended communities

AS number format

- > [<higher16bitsindecimal>.<lower16bitsindecimal>
 - “**asdot**” ex. “3.1”, “65510”, “9188”
 - draft-michaelson-4byte-as-representation-02.txt (expired)
 - draft-uijterwaal-rpsl-4byteas-ext-01.txt
 - RIRs, IANA, at least one router vendor use this

AS number format

- > <32bitsindecimal>
 - “**asplain**” ex. “196609”, “65510”, “9188”
 - preferred by network operators (on NANOG), also in use!?

- > <higher16bitsindecimal>.<lower16bitsindecimal>
 - “**asdot+**” ex. “3.1”, “0.65510”, “0.9188”
 - draft: “once all systems are [..] using 4-byte AS numbers”

Choices, choices...

- > No clear consensus about which format to use now!
- > No impact *at all* on basic AS4 functionality
- > But implementation has to use something for input and output of AS numbers
- > As-path-accesslists work on “printed” aspaths, i.e. output!

Quagga AS4 patch choice:

- > Use **any** format on input (auto-detect)
 - no impact on semantics whether “0.3” or “3”, “1.0” or “65536” are supplied wherever as numbers are entered (*beware: as numbers, not regular expressions*)
- > Use configured format on output, **asdot** is default
 - CLI configuration command:
`bgp asnum format (asdot|asdot+|asplain|asip)`

Impact of choosable format

- > Hard bgp reset needed when changing it
- > User responsible for regular expressions
 - must be aligned with chosen format
- > Only one format will remain in the end!?
 - probably not entering mainline quagga before that

Problems / Caveats

- > AS number format
- > MRT dumps
- > Extended communities

MRT dumps (Erik Romijn, RIPE NCC)

- > There was no format defined which was able to encompass all the AS4 information.
- > Format derived in discussions on GROW, draft-ietf-grow-mrt-04 with it is out (08.03.2007).
- > If you use dumps natively (also from RIPE NCC's RIS) you will have to adapt, your old software will not understand this.
Regarding RIS you must have done that already...

Problems / Caveats

- > AS number format
- > MRT dumps
- > Extended communities

Extended communities

- > draft-rekhter-as4octet-ext-community-01.txt defines the format *on-the-line* (as it should)!
- > Nothing like draft-ietf-idr-as4bytes which defines all interaction with AS2 bgp partners, so what do we do?
 - *nothing, only give possibility to use it.*
- > No CLI template to copy a command syntax from, choose which?

Extended communities

- > First patch version tried: keep syntax, autodetect everything.
 - introduced first real bug
 - solution of bug (patch v03) voided “semantic is independent from AS input syntax”
- > New idea?

Extended communities

- > Solution (v05): use '+' in front of AS number to indicate ENCODE_AS4 for the extended community, i.e.

```
rt 100:1 +65536:2 +3.5:30 1.2:45 +100:7
```

ENCODE_AS

ENCODE_AS4

ENCODE_AS4

ENCODE_AS
converts to
23456
i.e. will use
AS_TRANS

ENCODE_AS4

AS4 support for quagga

- > Motivation
- > Status
- > Feedback received
- > Problems
- > Conclusion

Conclusion

- > AS4 is available, no need to stay in the AS2 corner!
 - *<http://quagga.ncc.eurodata.de>*
- > Quagga gives you the ability to try all textual AS number representations (and deal with the consequences), use it!
- > Be aware of new things in MRT and extended community handling!

AS4 support for quagga: Thank you!

Thank you!

?

?

?

?

...Questions...

?

Juergen Kammer <j.kammer@eurodata.de>