



**RIPE
NCC**

RIPE NCC Updates on the Routing Information Service

Colin Petrie
Global Information Infrastructure



Current status



- 12 active collectors
 - 1 multihop collector in Amsterdam
 - 11 local collectors at IXPs around the world
- Quagga-based
 - Store BGP updates in MRT format every 5 minutes
 - Store BGP table dumps in MRT format every 8 hours
 - Provide looking-glass query via RIPEstat
 - Data archived since 1999
- Has been static for some time
 - Most recent collector added in 2008

- Back-end replacement
 - Migration from old MySQL architecture
 - Scaling problems - needed a MySQL server per collector in some cases
 - Data retention - MySQL stored about 3 months per collector
 - Replaced with Hadoop
 - Horizontally-scalable processing and storage cluster
 - Map/Reduce performs data import, processing, and historical aggregations
 - HBase serves live queries from RIPEstat

- New collectors now possible
- In discussions with several interested parties
 - Currently developing RRC18 at CATNIX, Barcelona
 - Other locations currently in discussions with potential hosts and participating IXPs
- If you are interested in hosting an RRC, please contact us




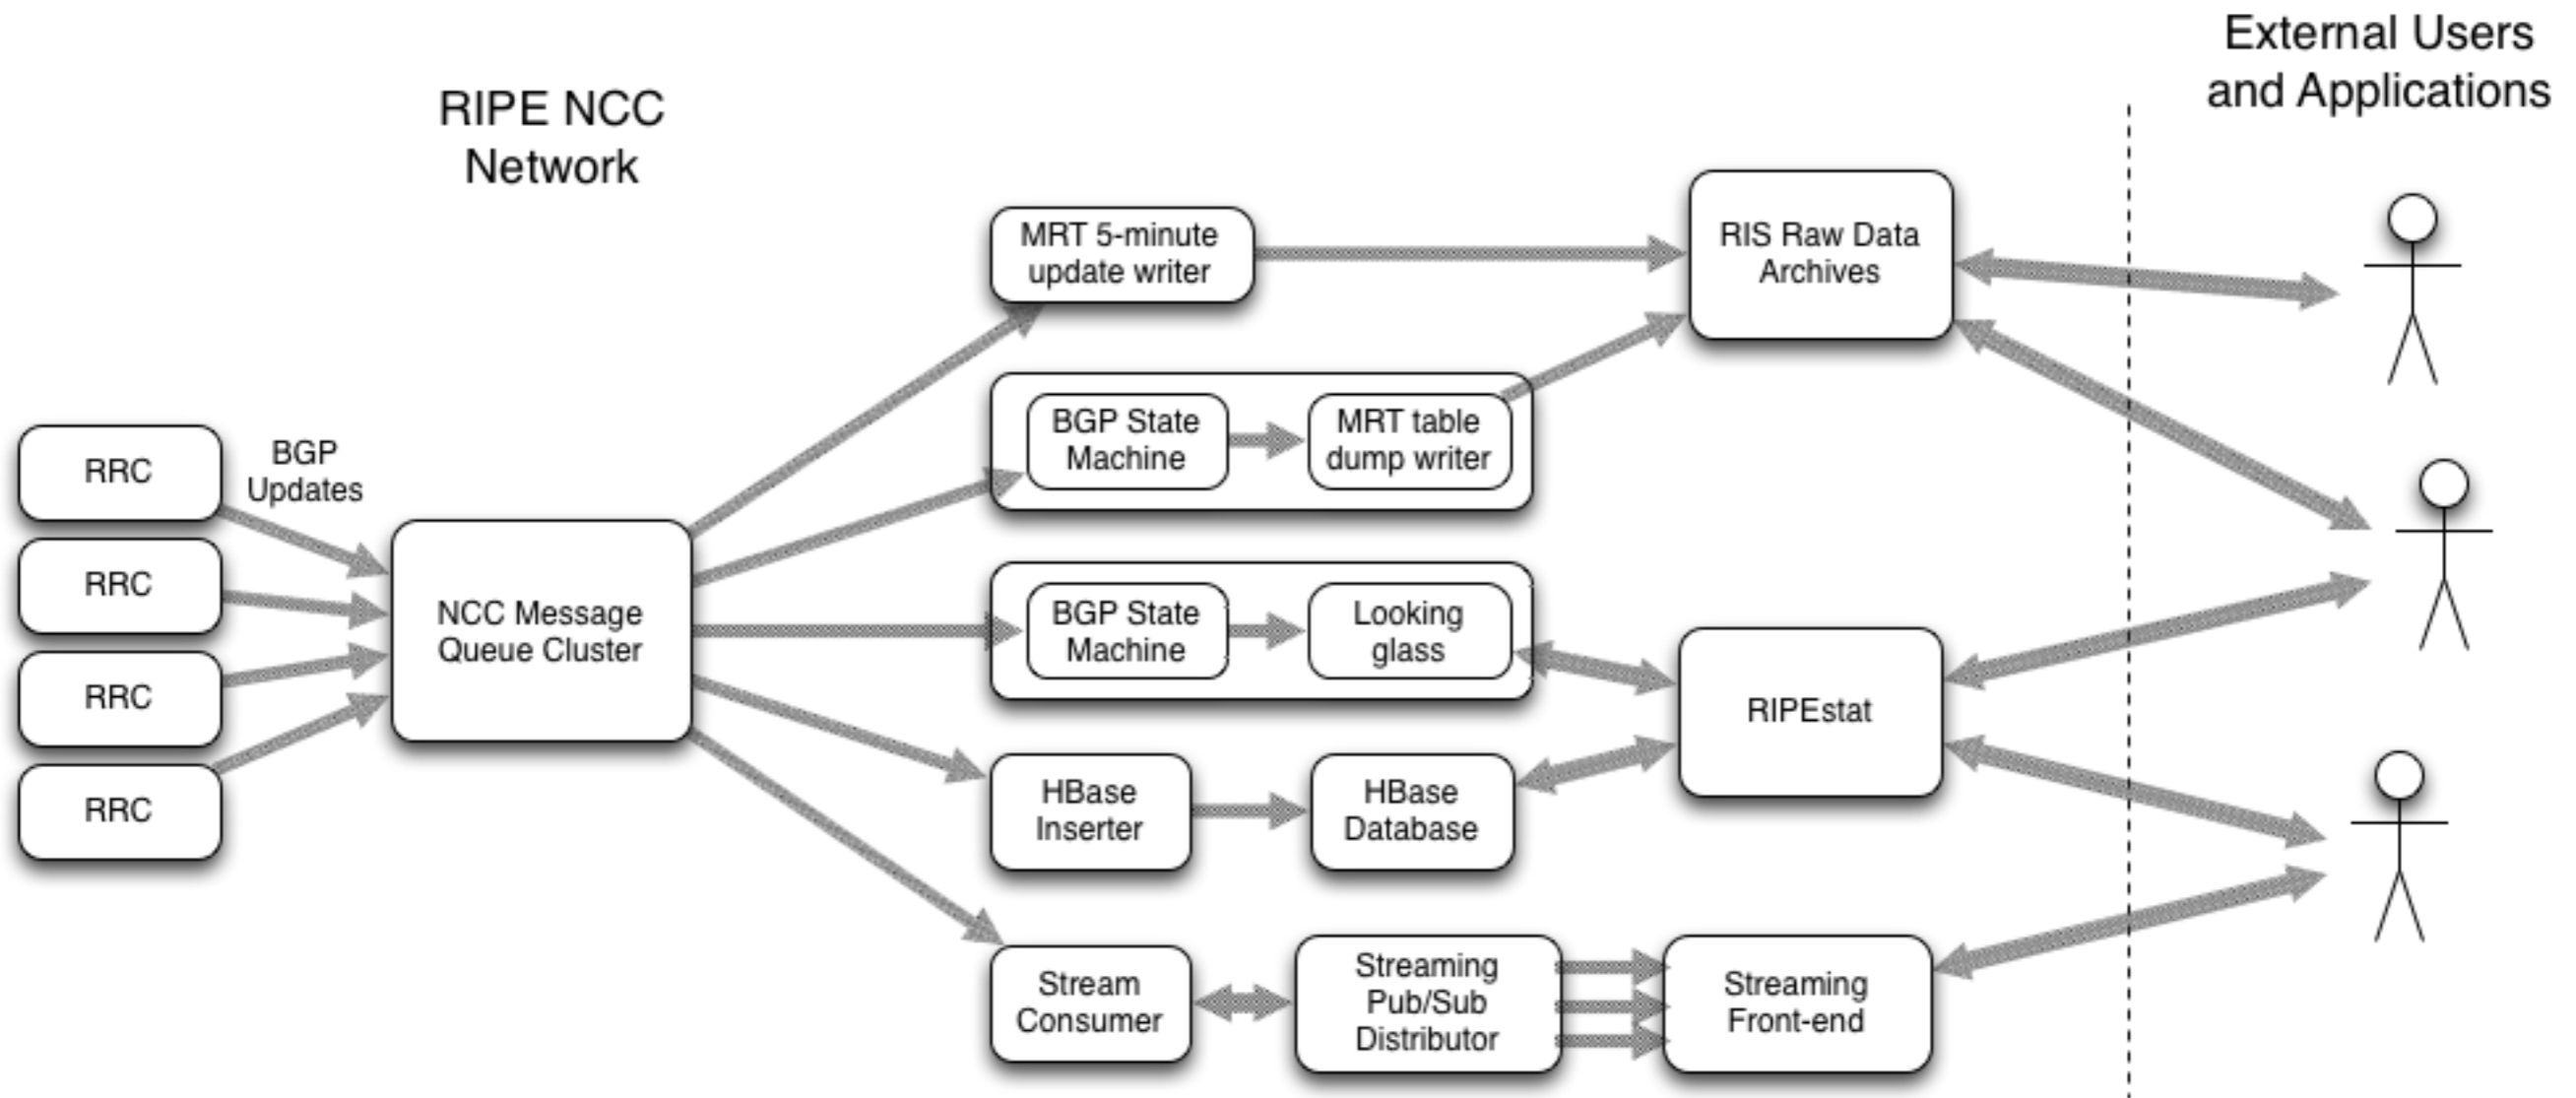
New developments



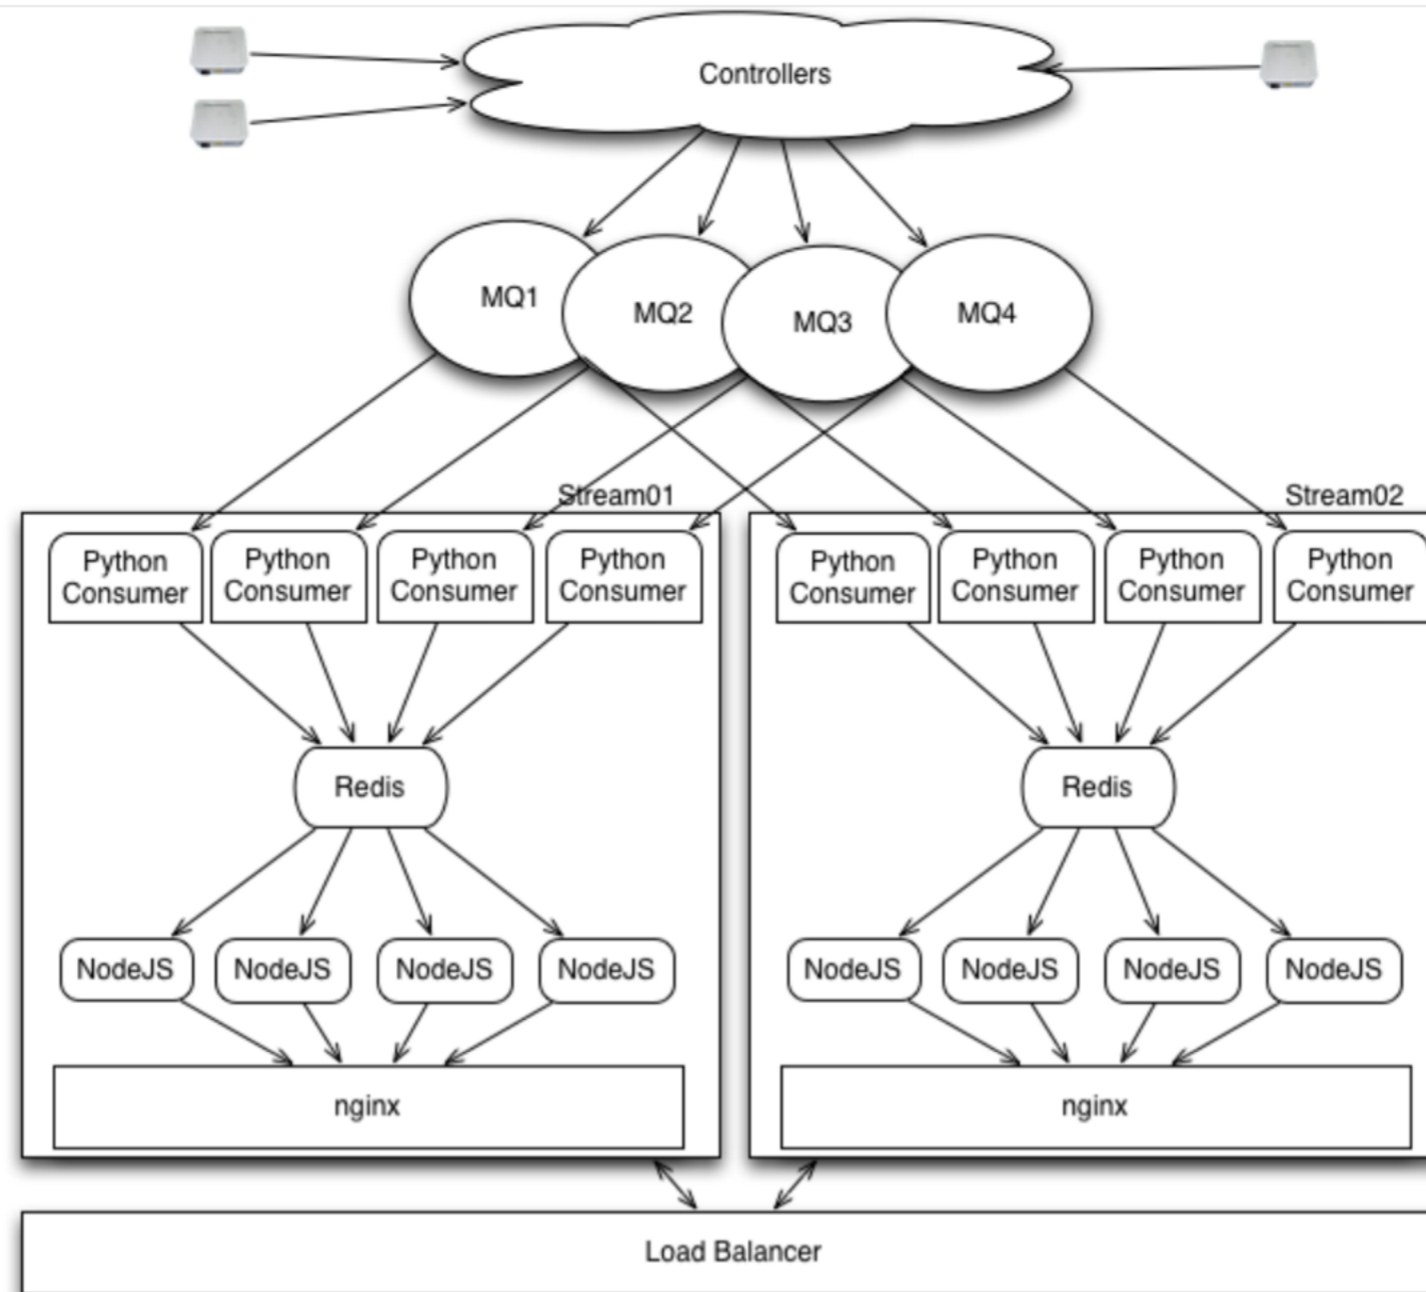
- Current Quagga implementation
- Single-threaded
 - Not as scalable on modern multi-core CPUs
- Locks updates during table-dump process
 - Requires that dump completes before the hold timer expires, or BGP session will drop
- Some data consistency issues
 - Sometimes updates are missing from the update dumps at the time of a table dump
 - This makes it difficult to accurately rebuild BGP state at an intermediate time, if updates are not reliable in-between

- This has been under design for 2 years
- Last year, a prototype developed by Wouter Miltenburg as part of his bachelor thesis at Hogeschool van Amsterdam
- Two RIPE Labs articles describing the architecture
 - https://labs.ripe.net/Members/wouter_miltenburg/researching-next-generation-ris-route-collectors
 - https://labs.ripe.net/Members/wouter_miltenburg/build-the-next-generation-ris-route-collectors
- Research paper also available (linked from the labs articles listed above)





- Same as the RIPE Atlas Streaming interface
- Diagram below from Massimo Candela's presentation on Monday Plenary session



Let's see if the live demo will work!



Other updates

BGPdump



- bgpdump v1.4.99.15 released
- Bugs fixed
 - Updated the maximum prefixes permitted in a single message to match the BGP RFC
 - Output buffering corruption
 - Crashes in int2str handling
- New feature
 - Added support for MRT type 5 messages, produced by earlier versions of RIS collectors before 2004
- We're now looking at importing historical information before 2004 into backend for RIPEstat



Other updates

BGP Additional Paths



- <https://tools.ietf.org/html/draft-ietf-idr-add-paths>
- Permits a prefix to be advertised more than once by a BGP neighbour without replacing the previous entry
- Allows a BGP peer to send not just their best paths, but all their available paths (or a subset)
- Allow more insight into provider networks in RIS?
- Previous tuple: {RRC, peer, prefix}
- New tuple: {RRC, peer, path-id, prefix}
- Requires MRT format update (new IETF draft?)

