



# GOLCONDE

*ce n'est pas un système de réplication*





---

# GOLCONDE

René Magritte, 1953





# GOLCONDE

---

- NOT REPLICATION\*
- DATA DISTRIBUTION SYSTEM
- ONE CANONICAL SOURCE
- ANY NUMBER OF TARGET RELATIONS
- QUEUE BASED

\* UNLESS YOU START FROM A CLEAN SLATE





# FEATURES

- LIGHT-WEIGHT DAEMON
- "AUTOSQL" AND CUSTOM HANDLERS
- TWO METHODS OF OPERATION
- REAL-TIME VISUALIZATION





# WHY GOLCONDE?

---

- GOOD REPLICATION TOOLS EXIST
  - SLONY, LOGICALSTANDBY, BRUARD, ETC
- NO "GOLDEN HAMMER"
- CURRENT TOOLS ARE RESOURCE INTENSIVE
  - LIVE ON THE DATABASE TIER
  - LIVE IN THE DATABASE





# WHY GOLCONDE?

---

- NEEDED A SCALE OUT SOLUTION
- CROSS DATA-CENTER IMPLEMENTATIONS
- MESSAGE BROKERS ARE MORE EFFICIENT FOR QUEUES
- DATABASE HARDWARE IS EXPENSIVE (\$\$\$)
- QUEUE AND CONSUMER GRADE HARDWARE ARE NOT
- NEED CUSTOMIZABLE WORKFLOW FOR DISPARATE RELATIONS





# REQUIREMENTS

---

- STOMP SUPPORTING MESSAGE BROKER
- POSTGRESQL 8.3\* OR HIGHER
- PYTHON 2.6
  - PYYAML
  - PSYCOPG2
  - STOMP.PY







---

[HTTP://ACTIVEMQ.APACHE.ORG/](http://activemq.apache.org/)







# STOMP

- 
- STREAMING TEXT ORIENTED MESSAGE PROTOCOL
  - INTEROPERABLE WIRE FORMAT
  - SUPPORTED BY MULTIPLE BROKERS

[HTTP://STOMP.CODEHAUS.ORG/](http://stomp.codehaus.org/)







# AMQP

---

- ADVANCED MESSAGE QUEUING PROTOCOL
- OPEN STANDARD WIRE PROTOCOL
- EMERGING STANDARD SUPPORTED BY MULTIPLE BROKERS

[HTTP://WWW.AMQP.ORG/](http://www.amqp.org/)







OpenAMQP

QPID

RedHat Enterprise MRG

[HTTP://WWW.RABBITMQ.COM/](http://www.rabbitmq.com/)





# ARCHITECTURE

---

- START OF ACTION LIFECYCLE
  - CLIENT-BASED: ENQUEUE FROM APPLICATION
  - TRIGGER-BASED: ENQUEUE FROM POSTGRESQL
- MULTI-THREADED PYTHON DAEMON





# LEXICON

---

- DESTINATIONS ARE THE CANONICAL RELATIONS
  - ARE ACTED UPON IN CLIENT-BASED WORKFLOWS
  - ENQUEUES DATA IN TRIGGER-BASED WORKFLOWS
- TARGETS ARE THE DISTRIBUTED RELATIONS





# ARCHITECTURE

---

- DATA IS PASSED IN JSON ENCODED PACKETS
- COMMANDS:
  - ADD (INSERT)
  - SET (UPSERT)
  - DELETE
  - UPDATE





# MESSAGE EXAMPLES

---

- {"ACTION": "ADD", "DATA": {"FRIEND\_ID": 47, "USER\_ID": 16, "STATUS\_ID": 1}}
- {"ACTION": "DELETE", "RESTRICTION": {"FRIEND\_ID": 11, "USER\_ID": 126}}
- {"ACTION": "SET", "DATA": {"FRIEND\_ID": 112, "USER\_ID": 111, "TIMESTAMP": "WED APR 1 13:56:54 2009", "STATUS\_ID": 1}}
- {"ACTION": "UPDATE", "RESTRICTION": {"FRIEND\_ID": 5, "USER\_ID": 41}, "DATA": {"TIMESTAMP": "WED APR 1 13:56:38 2009", "STATUS\_ID": 3}}





# AUTOSQL

---


- DEFAULT HANDLER TYPE FOR GOLCONDE
- EXAMINES PG\_CATALOG DATA FOR SCHEMA
- CACHES SCHEMA
- GENERATES SQL
- REQUIRES SAME SCHEMA





# CUSTOM HANDLERS

---



THAR BE DRAGONS

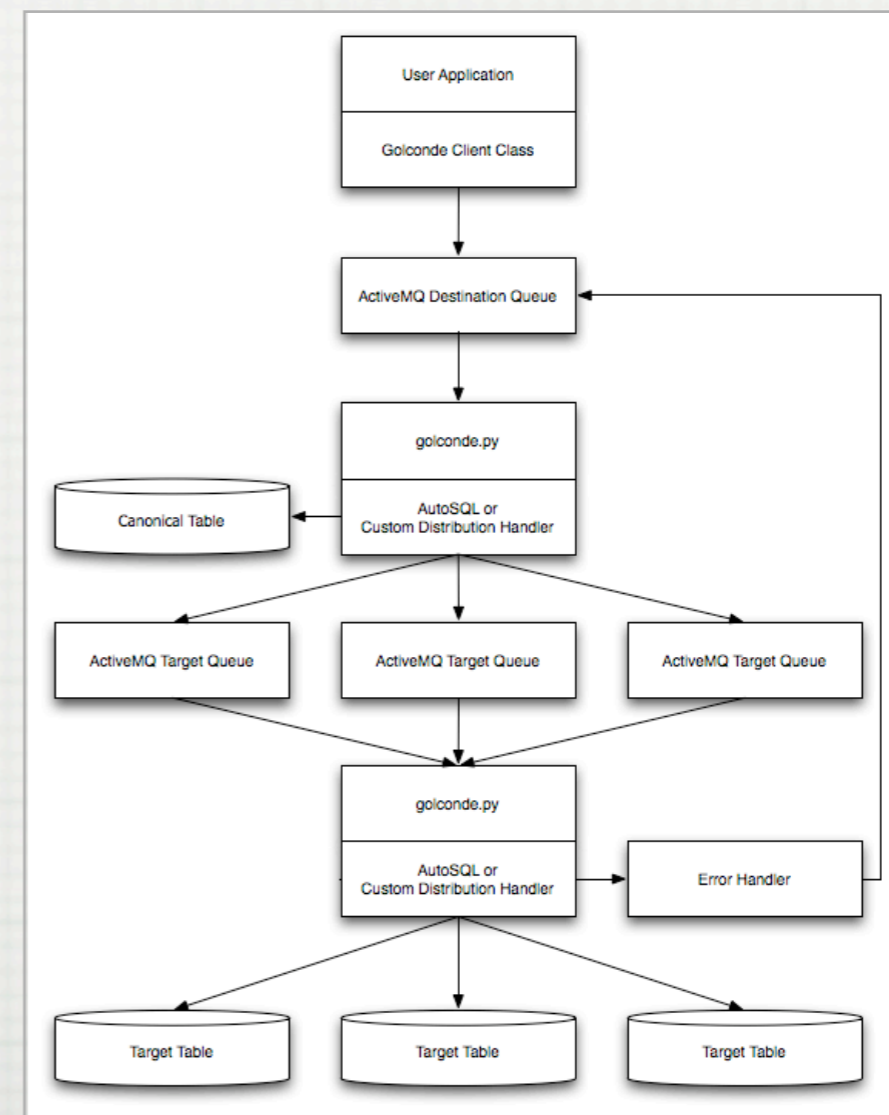
- ALLOWS FOR DIFFERENT SCHEMAS FOR THE SAME DATA
- EXAMPLE:
  - LEGACY SCHEMA VS NEW SCHEMA
  - DIFFERENT FOCUSED TARGET RELATIONS
- WARNING: NEED TO UNDERSTAND INTERNAL GOLCONDE FLOW





# CLIENT-BASED FLOW

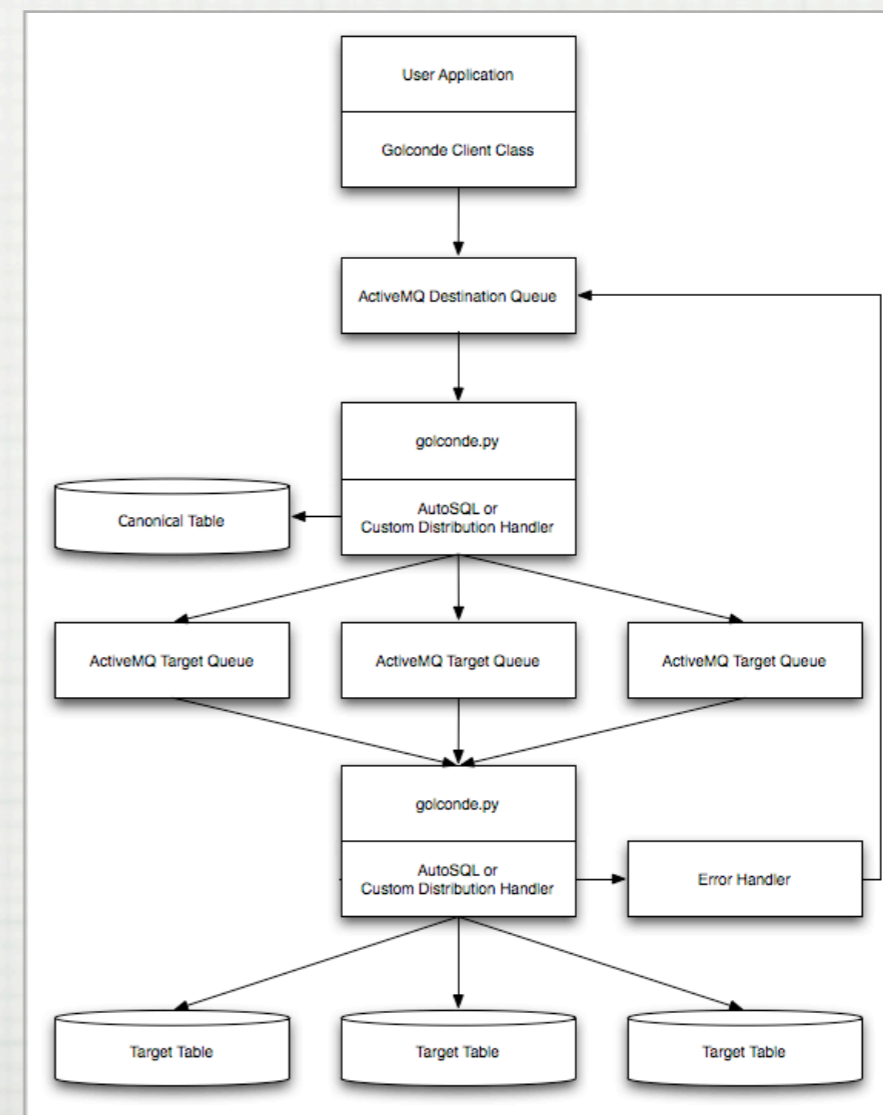
- CLIENT APPLICATION ENQUEUES
- GOLCONDE PERFORMS TRANSACTION ON CANONICAL DESTINATION RELATION
- SERIALY DISTRIBUTES TO TARGET QUEUES





# CLIENT-BASED FLOW

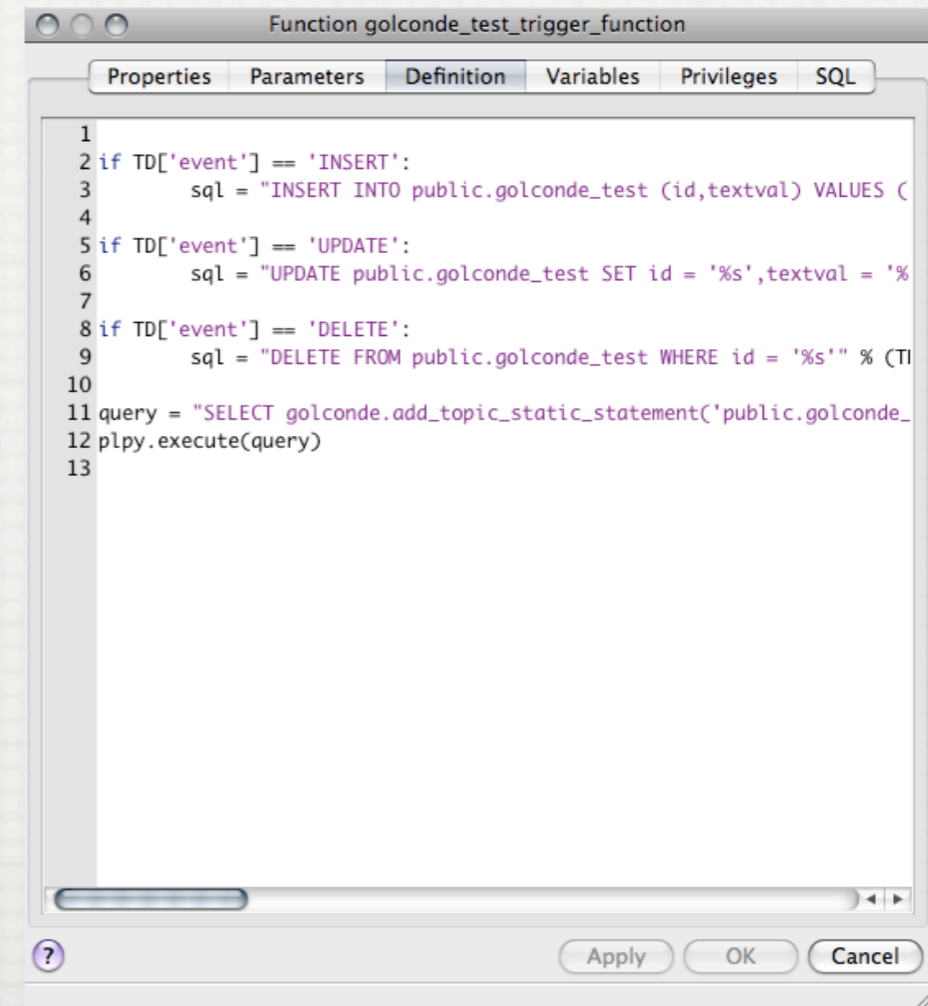
- TARGET HANDLER DEQUEUES
- PERFORMS TRANSACTION ON TARGET RELATIONS





# TRIGGER-BASED USAGE


- REQUIRES PL/PYTHON
- TRIGGER-UTIL.PY BUILDS AND INSTALLS TRIGGERS
- TRIGGER IS INSTALLED ON THE CANONICAL TABLE
- NO MODIFICATION OF CLIENT APPLICATION REQUIRED



The screenshot shows a window titled "Function golconde\_test\_trigger\_function" with tabs for Properties, Parameters, Definition, Variables, Privileges, and SQL. The "Definition" tab is active, displaying the following PL/Python code:

```
1
2 if TD['event'] == 'INSERT':
3     sql = "INSERT INTO public.golconde_test (id,textval) VALUES ("
4
5 if TD['event'] == 'UPDATE':
6     sql = "UPDATE public.golconde_test SET id = '%s',textval = '%"
7
8 if TD['event'] == 'DELETE':
9     sql = "DELETE FROM public.golconde_test WHERE id = '%s'" % (TI
10
11 query = "SELECT golconde.add_topic_static_statement('public.golconde_
12 plpy.execute(query)
13
```

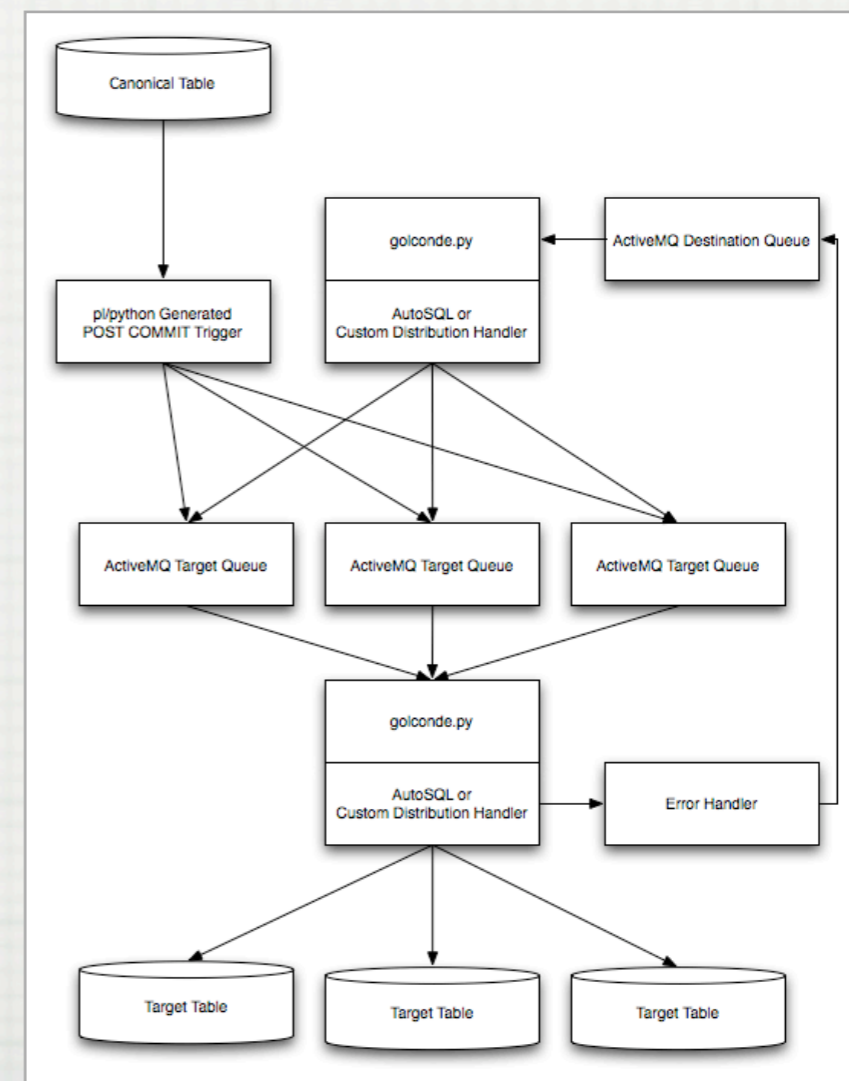
At the bottom of the window, there are buttons for "?", "Apply", "OK", and "Cancel".





# TRIGGER-BASED FLOW

- FIRES AFTER INSERT/UPDATE/DELETE
- ENQUEUES INTO TARGET QUEUES
- GOLCONDE PERFORMS TRANSACTIONS ON THE TARGET RELATIONS





# CONFIGURATION

- YAML BASED
- PARAMETERS FOR:
  - LOGGING
  - HTTP DAEMON FOR REALTIME MONITORING & VISUALIZATION
  - DESTINATION / TARGET GROUPS

[HTTP://WWW.YAML.ORG/](http://www.yaml.org/)

```
%YAML 1.2
---
Locations:
  base: /Users/gmr/Source/golconde
  logs: logs
Logging:
  filename: log.txt
  format: "%(levelname) -10s %(asctime)s %(message)s"
  # Valid values: debug, info, warning, error, critical
  level: debug
HTTPServer:
  enabled: True
  listen: 127.0.0.1
  port: 8000
Destinations:
  Test1:
    stomp: localhost:61613
    queue: /queue/golconde.test1
    pgsq: host=localhost user=gmr dbname=gmr
    # If you exclude the schema, it will default to public
    target: friendships
    function_name: AutoSQL
  Targets:
    target1:
      stomp: localhost:61613
      queue: /queue/golconde.test1.target1
      pgsq: host=localhost user=gmr dbname=destination
      # Schema.Relation format
      target: public.friendships
      function_name: AutoSQL
    target2:
      stomp: localhost:61613
      queue: /queue/golconde.test1.target2
      pgsq: host=localhost user=gmr dbname=destination_two
      target: public.friendships
      function_name: AutoSQL
```





# STATISTICS & VISUALIZATION

---

- BUILT-IN HTTP SERVER
- STATS REQUESTS RETURNS JSON
  - USE WITH STAPLR, NAGIOS, ETC
- INTERNAL REAL-TIME VISUALIZATION





# DEMONSTRATION

```
Terminal — Python — 160x48
DEBUG 2009-03-31 23:22:11,289 Attempting connection to host localhost, port 61613
DEBUG 2009-03-31 23:22:11,288 starting receiver loop
INFO 2009-03-31 23:22:11,289 Established connection to host localhost, port 61613
DEBUG 2009-03-31 23:22:11,289 Sent frame: type=CONNECT, headers={}, body=''
DEBUG 2009-03-31 23:22:11,290 Sent frame: type=SUBSCRIBE, headers={'ack': 'auto', 'destination': '/queue/golconde.test1.target2'}, body=''
DEBUG 2009-03-31 23:22:11,290 starting receiver loop
INFO 2009-03-31 23:22:11,290 Target Initialized
DEBUG 2009-03-31 23:22:11,291 Attempting connection to host localhost, port 61613
DEBUG 2009-03-31 23:22:11,291 starting receiver loop
DEBUG 2009-03-31 23:22:11,288 Sent frame: type=CONNECT, headers={}, body=''
DEBUG 2009-03-31 23:22:11,292 starting receiver loop
DEBUG 2009-03-31 23:22:11,292 Sent frame: type=CONNECT, headers={}, body=''
INFO 2009-03-31 23:22:11,292 Destination Initialized
DEBUG 2009-03-31 23:22:11,292 Attempting connection to host localhost, port 61613
INFO 2009-03-31 23:22:11,292 Established connection to host localhost, port 61613
DEBUG 2009-03-31 23:22:11,293 Sent frame: type=SUBSCRIBE, headers={'ack': 'auto', 'destination': '/queue/golconde.test2.target1'}, body=''
DEBUG 2009-03-31 23:22:11,294 Sent frame: type=CONNECT, headers={}, body=''
DEBUG 2009-03-31 23:22:11,295 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:46'}, body=''
DEBUG 2009-03-31 23:22:11,295 Sent frame: type=SUBSCRIBE, headers={'ack': 'auto', 'destination': '/queue/golconde.test1.target1'}, body=''
DEBUG 2009-03-31 23:22:11,296 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:48'}, body=''
DEBUG 2009-03-31 23:22:11,297 starting receiver loop
INFO 2009-03-31 23:22:11,295 Established connection to host localhost, port 61613
INFO 2009-03-31 23:22:11,297 Established connection to host localhost, port 61613
INFO 2009-03-31 23:22:11,298 Established connection to host localhost, port 61613
DEBUG 2009-03-31 23:22:11,298 Sent frame: type=CONNECT, headers={}, body=''
DEBUG 2009-03-31 23:22:11,300 starting receiver loop
DEBUG 2009-03-31 23:22:11,301 Sent frame: type=CONNECT, headers={}, body=''
DEBUG 2009-03-31 23:22:11,302 Sent frame: type=SUBSCRIBE, headers={'ack': 'auto', 'destination': '/queue/golconde.test1'}, body=''
DEBUG 2009-03-31 23:22:11,301 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:49'}, body=''
DEBUG 2009-03-31 23:22:11,301 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:51'}, body=''
DEBUG 2009-03-31 23:22:11,302 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:47'}, body=''
DEBUG 2009-03-31 23:22:11,302 starting receiver loop
DEBUG 2009-03-31 23:22:11,302 starting receiver loop
DEBUG 2009-03-31 23:22:11,302 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:52'}, body=''
DEBUG 2009-03-31 23:22:11,302 Sent frame: type=SUBSCRIBE, headers={'ack': 'auto', 'destination': '/queue/golconde.test2'}, body=''
DEBUG 2009-03-31 23:22:11,301 Sent frame: type=CONNECT, headers={}, body=''
DEBUG 2009-03-31 23:22:11,303 listener <_main_.TargetHandler object at 0xe60590> has no such method on_connected
DEBUG 2009-03-31 23:22:11,303 listener <_main_.TargetHandler object at 0xe60250> has no such method on_connected
DEBUG 2009-03-31 23:22:11,304 listener <_main_.TargetHandler object at 0xe60370> has no such method on_connected
DEBUG 2009-03-31 23:22:11,304 Sent frame: type=SUBSCRIBE, headers={'ack': 'auto', 'destination': '/queue/golconde.test2.target2'}, body=''
DEBUG 2009-03-31 23:22:11,306 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:50'}, body=''
DEBUG 2009-03-31 23:22:11,307 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:53'}, body=''
DEBUG 2009-03-31 23:22:11,307 listener <_main_.DestinationHandler object at 0xe60430> has no such method on_connected
DEBUG 2009-03-31 23:22:11,308 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:55'}, body=''
DEBUG 2009-03-31 23:22:11,308 listener <_main_.TargetHandler object at 0xe60550> has no such method on_connected
DEBUG 2009-03-31 23:22:11,308 Received frame: result='CONNECTED', headers={'session': 'ID:gmr-mbp.local-53423-1238425677001-5:54'}, body=''
DEBUG 2009-03-31 23:22:11,308 listener <_main_.DestinationHandler object at 0xe600b0> has no such method on_connected
```





# CURRENT STATUS

---

- 0.5 BETA RELEASE - 03/02
- 0.6 BETA COMING SOON
  - ADDS REAL-TIME STATS SERVER
  - BUG FIXES





# ROADMAP

---

- TWO-PHASE COMMIT LIKE BEHAVIOR
  - ALL TRANSACTIONS MUST COMMIT OR ROLLBACK
  - ROLLBACK IS DEFINED BY ROLLBACK DATA PACKETS IN AUTOSQL OR HANDLERS
- ADDITIONAL QUEUE SERVERS AND PROTOCOLS
- "NATIVE" CLIENT CLASSES IN PHP, PYTHON





# QUESTIONS?

---

CONTACT INFO:

[GMR@MYYEARBOOK.COM](mailto:GMR@MYYEARBOOK.COM)

TWITTER: @CRAD

[HTTP://GAVINROY.COM](http://GAVINROY.COM)

