

Comdb2: Bloomberg's Highly Available Relational Database System

Engineering

Bloomberg

25th Annual Tcl/Tk Conference
October 17, 2018

Joe Mistachkin
Senior Software Engineer

TechAtBloomberg.com

What is Comdb2?

- Comdb2 is a clustered relational database management system that uses optimistic concurrency control techniques.
- Comdb2 offers full transactional support, a standard relational model, and the expressivity of SQL.

Some Notable Features

- Provides multiple isolation levels, including snapshot and serializable isolation.
- Read/Write transactions run on any node, with the client library transparently negotiating connections to lowest cost (latency) node which is available.
- The client library provides transparent reconnect and resume.

Anything else?

- Stored procedures via a dialect of Lua.
- Built using a variety of (heavily customized) open source technologies, including but not limited to:
 - SQLite
 - BerkeleyDB

Anything else?

- Far more detailed information can be found in the VLDB 2016 paper, available here:

<http://www.vldb.org/pvldb/vol9/p1377-scotti.pdf>

What about Tcl?

- Tcl bindings are (now) included.
- One command will be added to the interpreter, named [cdb2].
- The [cdb2] command has approximately 20 sub-commands.

Server Installation

- The server can be installed using the instructions at: <https://bloomberg.github.io/comdb2/install.html>
- The examples that follow assume the initial test database is named “myDb”.

Bindings Example #1

```
# load the bindings...
package require tclcdb2
# open database "myDb"...
set db [cdb2 open myDb local]
# perform a query...
cdb2 run $db "SELECT now();"
# close database...
cdb2 close $db
```

Bindings Example #2, part 1

```
cdb2 unbind $db
```

```
cdb2 bind $db 1 cstring two
```

```
cdb2 run $db \  
"SELECT x FROM t1 WHERE x = ?;"
```

Bindings Example #2, part 2

```
# while there are more results...
while {[cdb2 next $db]} {
  # get column name, value, etc.
  puts [cdb2 colcount $db]
  puts [cdb2 colname $db 0]
  puts [cdb2 colvalue $db 0]
  puts [cdb2 coltype $db 0]
  puts [cdb2 colsize $db 0]
}
```

What about TDBC?

- The thinking was that a TDBC interface for Comdb2 can be built upon the low-level interface exposed via the [cdb2] command.

Bloomberg

Engineering

Questions?

[TechAtBloomberg.com](https://www.techatbloomberg.com)

© 2018 Bloomberg Finance L.P. All rights reserved.

Useful Links

- Source Code: <https://github.com/bloomberg/comdb2>
- Documentation: <http://bloomberg.github.io/comdb2>
- VLDB 2016 Paper: <http://www.vldb.org/pvldb/vol9/p1377-scotti.pdf>