



June 10, 2018

#### What Was Old Is New Again

#### YottaDB – https://yottadb.com



- A NoSQL database with a proven, mature code base that both scales up to enterprise-scale applications and scales down to the needs of embedded systems.
- Rock Solid. Lightning Fast. Secure. Pick any three.





- The Past
  - Where are we and how did we get here?
- Making What was Old New Again
- The Future
- Demo
  - Still a work in progress



#### The Past

#### Where are we and how did we get here?



# The Original Computer Database



- IBM Information Management System (IMS)
- Created to manage bill of materials & inventory of Saturn V & Apollo
  - Hierarchical data model a NoSQL database!
- First released 1966; latest release 2017
- Runs on mainframe ⇒ Expen\$ive

Yøtta<sup>DB</sup>



Massachusetts General Hospital, Boston



- Animal research laboratory circa 1966
  - Limited funding for computing
- Minicomputers spare DEC PDP-7
- Accessible talent across the river, in Cambridge
  - Massachusetts Institute of Technology
  - Bolt, Beranek and Newman





- <u>Massachusetts General Hospital Utility Multi-</u> <u>Programming System</u>
  - Operating system + hierarchical database file system + user interface + programming language + ...
  - First used 1966/67
  - Ecosystem culture user driven development; users and developers work closely together ⇒ pragmatic software without deep Computer Science theory
    - Not respected by CS academia

#### Key-Value Tuples



["Capital", "Belgium", "Brussels"]
["Capital", "Thailand", "Bangkok"]
["Capital", "USA", "Washington, DC"]

Key

Value

Always sorted – MUMPS means you never have to say you're sorting...

#### Schemaless



["Capital", "Belgium", "Brussels"] ["Capital", "Thailand", "Bangkok"] ["Capital", "USA", "Washington, DC"] ["Population", "Belgium", 13670000] ["Population", "Thailand", 84140000] ["Population", "USA", 325737000]

Default order for each key:

- Empty string ("")
- · Canonical numbers in numeric order
- Strings (blobs) in lexical order

Schema determined entirely by application – MUMPS assigns no meaning

Numbers and strings (blobs) can be freely intermixed in values and keys except first

#### Mix Key Sizes



```
["Capital", "Belgium", "Brussels"]
["Capital", "Thailand", "Bangkok"]
["Capital", "USA", "Washington, DC"]
["Population", "Belgium", 13670000]
["Population", "Thailand", 84140000]
["Population", "USA", 325737000]
["Population", "USA", 17900802, 3929326]
["Population", "USA", 18000804, 5308483]
```

["Population", "USA", 20100401, 308745538]

yyyymmdd

"Population" + 1 more key means value is latest population

"Population" + 2 more keys means value is population on date represented by last key

#### Keys ↔ Array References



POPULATION Population("Belgium")=13670000 TREE Population("Thailand")=84140000 Population Population("USA")=325737000 Population("USA",17900802)=3929326 "Thailand" "Belaium" Population("USA", 18000804)=5308483 84140000 13670000 ... Population("USA", 20100401)=308745538 17900802 3929326 First key is Other keys are variable name subscripts

Array references are a familiar programming paradign



Any JSON structure is representable as a tree, but not vice versa

# Sharing and Persistence – Database Access Yotta<sup>DB</sup>

• Process private, available only for lifetime of process

Population("Belgium")
Population("Thailand
Population("USA")

"local" variables

Shared across processes, persistent beyond lifetime of any process
 ^Population("Belgium")
 ^Population("Thailand")
 \*global" variables

Spot the difference?

#### Universal NoSQL



- Satisfies common major NoSQL use cases
  - http://mgateway.com/docs/universalNoSQL.pdf
- NoSQL means "Not only SQL"

#### Noteworthy Features



- Tight binding of database to language
- Direct source code execution (initial implementation)
- Dynamic linking
- Multitasking
- Interactive / incremental usage
- Hierarchical locks (traffic light semantics)

#### Noteworthy Contempories



- C
- SQL
- TCP/IP
- UNIX

#### Evolution ... 1



- 1970s
  - Language+database separate from operating system
- 1980s (GT.M forerunner to YottaDB)
  - Programs are just text files in the file system
    - Compiled to object code for execution
    - While maintaining interactive / incremental usage

#### Evolution ... 2



- 1990s
  - ACID transactions (GT.M)
  - Vendor consolidation
    - Just two commercial implementations left
- 2000s
  - GT.M/Linux moves to free / open source license

#### ACID Transactions



- Atomic it all happens or none of it happens
- Consistent logic inside a transaction cannot see internal state of another transaction
- Isolated no other logic can see inside this transaction
- Durable once committed, state change is permanent

ACID Transactions in GT.M/YottaDB



- Ensuring Consistency & Isolation with high concurrency is hard
- Optimistic Concurrency Control
  - http://daslab.seas.harvard.edu/reading-group/paper s/kung.pdf
- Achieves high levels of concurrency & scalability
  - At the cost of a pathological case that application code must avoid

#### GT.M/YottaDB Today



- At the heart of mission-critical applications the largest real-time core-banking and patient-centric healthcare systems in the world
- But not widely used in general purpose computing



- Consequences of direct execution of source code
  - Needed to save memory and run fast
  - Single letter abbreviations of commands, short names

```
hello
write "Hello, World!",!
quit
```

```
hello w "Hello, World!",! q
```



- Consequences of direct execution of source code
- Enterprise-scale applications on small computers
  - Expert friendly code, e.g. S %P1=\$S(\$L(%P1)>8:\$E(%P1,1,8)-1700000\_"."\_\$E(%P1,9,14),1:%P1-1700000) ;%P1 is now in FM format I %P1[".",+\$P(%P1,".",2)=0 S %P1=\$\$FMADD(+%P1,-1)\_".24" ;If HL7 tz and local tz are the same I %P2["L",%TZ=%LTZ S %P2="" I (%P2["U")!(%P2["L"),%P1["." D ;Build UCT from dat . S %=\$TR(%TZ,"+-","-+") ;Reverse the sign . S %H=\$E(%,1,3),%M=\$E(%,1)\_\$E(%,4,5) . S %P1=\$\$FMADD(%P1,,%H,%M) Q



- Consequences of direct execution of source code
- Enterprise-scale applications on small computers
- Successful applications have long lives
  - Code written in the 1970s and 1980s was written to different standards of readability than code today
  - Application consistency for maintainability means coding style lags best practices for readability



- Direct execution of source code
- Enterprise-scale applications on small computers
- Successful applications have long lives
- Vendor consolidation ended language evolution & standardization
  - One vendor able to acquire all implementations except GT.M



- Direct execution of source code
- Enterprise-scale applications on small computers
- Successful applications have long lives
- Vendor consolidation ended language evolution & standardization
- Cultural issues inside and outside community



# Making What was Old New Again



#### The Diamond is the Database



- Mature, proven code
  - "Rock Solid. Lighning Fast. Secure. Pick any three."

# The Language is What it is



- You either love it or you hate it
  - Like anchovies on your pizza
  - or like emacs vs. vi[m] vs. ...
  - or like your religion vs. the other guy's religion
  - or...

#### YottaDB Strategy



- Build on what works well
- Accommodate what's new



Public domain from Wikimedia Commons



Photos are almost 100 years apart

By GT1976 [CC BY-SA 4.0 (https://creativecommons.org/licenses/by-sa/4.0)], from Wikimedia Commons



#### From GT.M to YottaDB

Building on Strengths and Accommodating What's New



# Tight Database Binding is a Strength

- Y∕tta<sup>DB</sup>
- Create tight binding from database to C, just like the tight binding from database to the MUMPS language
- Make it as easy to use as any other library source /usr/local/lib/yottadb/ydb\_env\_set #include "libyottadb.h" gcc -I \$ydb\_dist -L \$ydb\_dist -o myprog myprog.c -lyottadb ./myprog

#### Simple API – Key Functions



ydb\_data\_s() - determine whether node and/or subtree exist ydb\_delete\_s() - delete node or both node & subtree ydb\_delete\_excl\_s() - delete all local variables (optionally except specified) ydb\_get\_s() - get a value from a local or global variable node ydb\_node\_next\_s() - get next node (depth-first order) ydb\_node\_previous\_s() - get previous node ydb\_set\_s() - set the value at a node ydb\_subscript\_next\_s() - get next subscript at deepest level (breadth-first order) ydb\_subscript\_previous\_s() - get previous subscript at deepest level ydb\_subscript\_previous\_s() - get previous subscript at deepest level

# Extend to Other Languages



- C is the *lingua franca* of computer languages
  - JavaScript wrapper next major version to use Simple API
    - https://github.com/dlwicksell/nodem
  - Go is coming
  - SQL for reporting and to leverage other tools
    - Expert-friendly FOSS SQL/JDBC engine exists; make it user friendly
  - C++, Python, Rust, Java

Driven by user input and funding

#### Extend Platforms



- Linux on 32-bit ARM
  - ARMv7-A (e.g., Raspberry Pi 3, BeagleBone Black) added 2017
  - ARMv6 added 2018 (e.g., Raspberry Pi Zero)
- Linux on 64-bit ARM
  - Anticipated late 2018



#### The Future



#### "YottaDB Everywhere"



- Footprint fits in embedded systems
- Scales up to manage very large databases
- And everything in-between

Everywhere

• "Rock solid. Lightning fast. Secure. Pick any three."

# YottaDB Initial Targets



- Traditional GT.M applications
  - Including core banking and electronic health records
- "Big data"
  - e.g., alternative to Hadoop
- Internet of Things (IoT)
  - One database for the complete stack

#### Traditional – A Better GT.M ... 1



- Stay upward compatible
  - GT.M enhancements & fixes merged into code base



#### Close Relationship to GT.M is a Strength



### Ensuring Upward Compatibility





More than 20 years experience working together on GT.M

#### Traditional – A Better GT.M ... 2



- GT.M enhancements & fixes merged into code base
- Performance & scalability specific to YottaDB
- Fixes specific to YottaDB

# Big Data Approach



- Create Hadoop system from real-time banking
   application
  - Real time data feed replication stream  $\rightarrow$  Hadoop
  - Write queries for working application
- Replace Hadoop database & application with YottaDB
- Incrementally expand functionality to more Hadoop applications and perhaps others (e.g., R)









#### A Picture is Worth 1E3 Words ... 3





46



#### Demo

### Still a work in progress







• Demonstrate YottaDB as a single database used on the edge and in the cloud



#### Internet of Things Demo – EEG Sensor





#### Internet of Things Demo – Chernoff Faces



Y⊚tta<sup>DB</sup>

#### Chernoff Face Reading My Mind







# Demo Technology – All FOSS



https://github.com/YottaDB/YottaDBDemos/tree/master/mindwave







- Web site https://yottadb.com
- Quick start –

https://docs.yottadb.com/MultiLangProgGuide/Mul tiLangProgGuide.html#quick-start

- User documentation –
   https://yottadb.com/resources/documentation/
- Blog https://yottadb.com/blog/
- Contact K.S. Bhaskar / bhaskar@yottadb.com





Thank You!

yottadb.com