



FUSIONDB

Unconference - XML Prague
2019-02-07

Adam Retter

 adam@evolvedbinary.com

 [@adamretter](https://twitter.com/adamretter)



EVOLVED BINARY

This talk...

- Trying not to pre-empt. tomorrows conference talk!
- Demo some pre-beta version features
- AMA - Ask Me Anything!



What is FusionDB?

- **Multi-model database**
 - Document
 - XML and Binary
 - JSON native - work in progress
 - Key/Value
- **ACID**
 - MVCC
 - Snapshot Isolation
- **Written in C++ 14 and Java**
- **100% eXist-db API Compatible**
- **Will be Open Sourced...**



Fusion Studio

- **IDE**

- Not designed to replace comprehensive tools like Oxygen!
- Create and Edit Documents in the database
- Execute XQuery against the database
- Serialization Settings and Paging

- **Management Tool**

- Connect to Multiple servers
- Browse Documents in the database
- Upload/Copy/Move/Delete Documents
- Modify Document Properties - permissions, type, metadata, etc.
- Create/Edit/Delete Indexes
- Create/Edit/Delete Users



Fusion Studio

- **Dashboard (in progress...)**

- Overview of system health and load
- Detailed statistics of system operation
- Access to logs

- **Status**

- Open Source

- <https://github.com/evolvedbinary/fusion-studio-extension> (GPL v3)
- <https://github.com/evolvedbinary/fusion-studio-api> (AGPL v3)
<https://app.swaggerhub.com/apis/evolvedbinary/pebble-api/1.1.0>
- <https://github.com/evolvedbinary/fusion-studio> (GPL v3)

- Theia - TypeScript

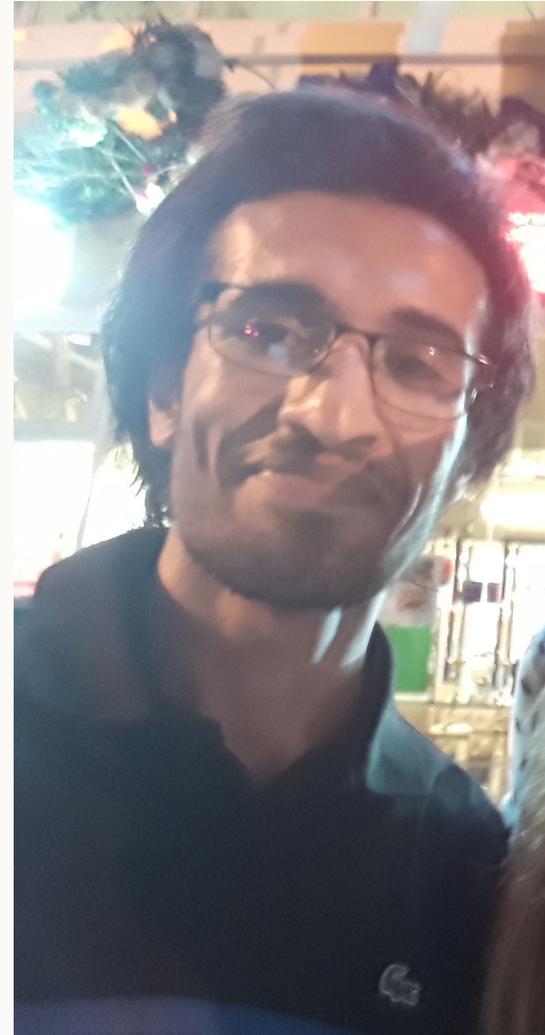
- Looking for more contributors - some paid opportunities
- Compatible with eXist-db (currently)



Behind Fusion Studio

- **Charaf Eddine Cheraa**

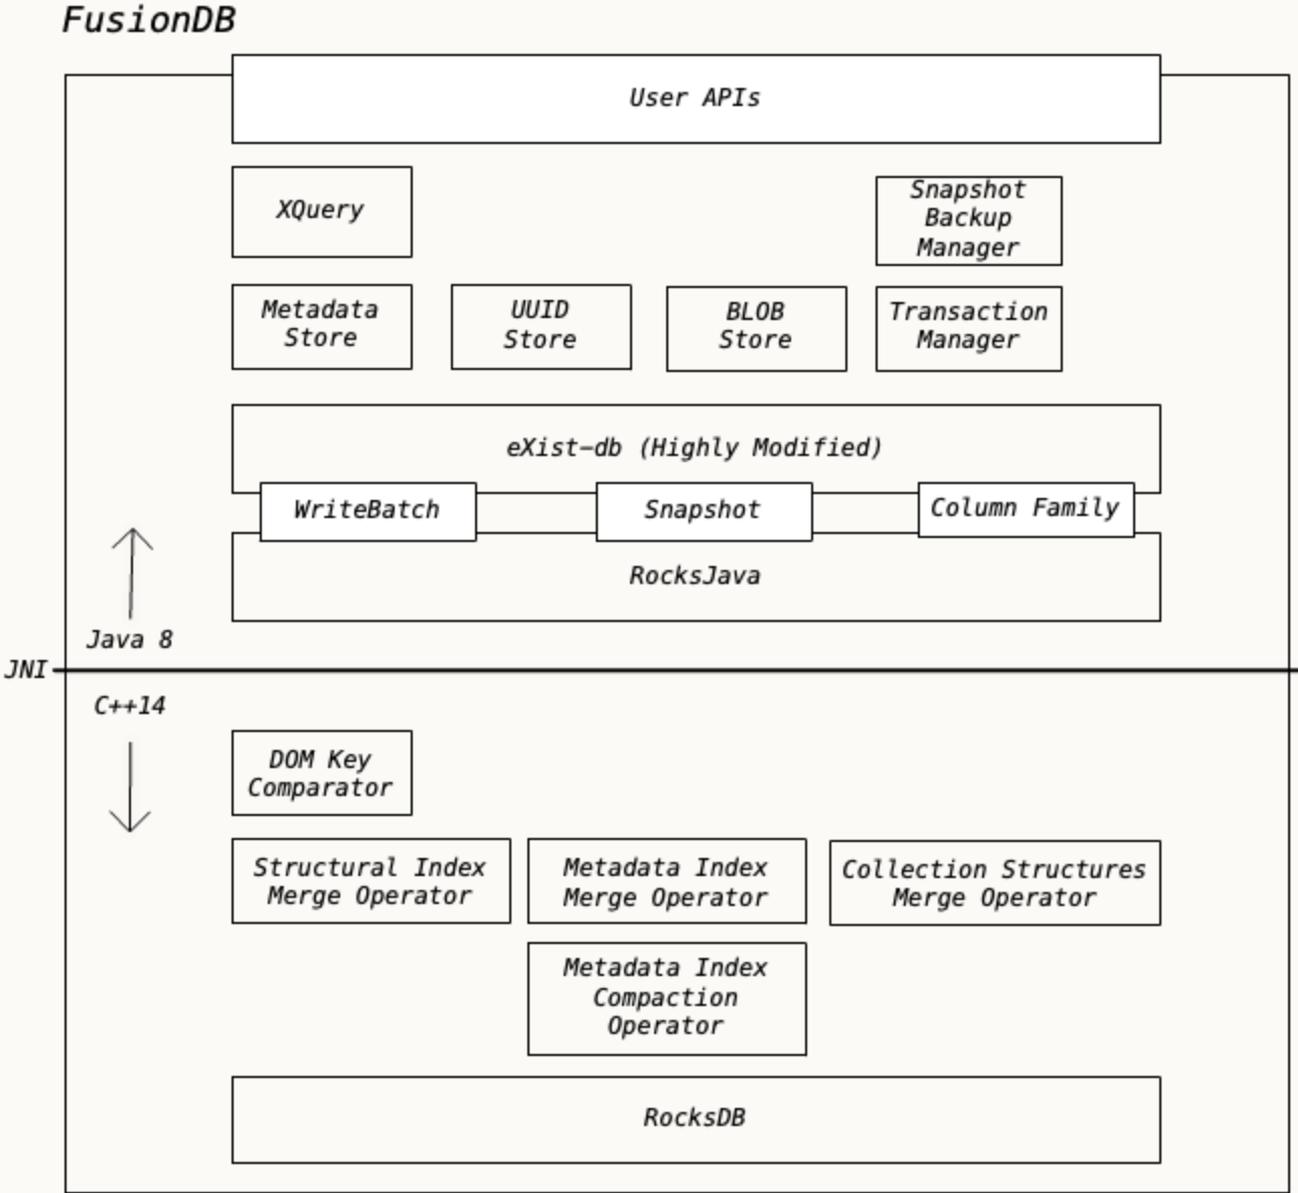
- Took the TypeScript holiday - <https://blog.adamretter.org.uk/typescript-holiday/>
- From Algeria, on GitHub - @ccheraa
- Sadly, he could not be here today!
- Thank you Charaf!



Demo - Fusion Studio



FusionDB Architecture



Demo - What's on the disk?



UUID Locators

- **Universally Unique Identifier**
 - v1 with random MAC
 - Very fast generation (> 5 Million / sec)
- **Assigned to every resource e.g. Documents and Collections**
- **Resources retrieval by UUID**
- **Immutable**
 - Immune to move/rename/copy
- **Persist across backup/restore**
 - Restore - Collision Protection
- **e.g. Provenance / External Systems / Archival**



Demo - UUID Locators



Key/Value Metadata

- **Uses the key/value Model**
 - Just the tip of the iceberg!
- **Metadata can be assigned to any resource e.g. Documents and Collections**
- **Resources Queryable by Key or, Key and Value**
- **Index Backed for Performance**
- **XQuery Metadata Module**
 - Metadata CRUD operations
 - Search by Metadata



Demo - Key/Value Metadata in XQuery



Backup Facilities

- **eXist-db style Data Export**
 - Snapshot Consistent
 - Lock-free!
- **Checkpoint Backup**
 - Exploits immutable nature of SST files
 - Full Backup! Could add differential in future
 - Almost Instantaneous!
 - Hard links on same filesystem
 - Simple file copy if across filesystems



Demo - Backup Facilities



ACID Transactions and XQuery

- **Snapshot Isolation**
- **Each XQuery is a distinct transaction**
 - auto-abort if error is raised
 - auto-commit on successful completion
- **try/catch**
 - Creates a Sub-Transaction
 - auto-abort on entering catch clause
 - auto-commit on try body completion
 - Allows the user to control transactional updates
 - Nesting try/catch... nests sub-transactions



Demo - Transactions and XQuery



BLOB Store

- **Deduplication**

- BLOB identity established through BLAKE256 digest
- Each BLOB maintains a reference count
- Vacuum process hoovers up any BLOBs with count=0
 - Unless they are recycled first!

- **Move is a zero I/O operation**

- **Copy is an almost zero I/O operation**

- Uses just a few bytes more for each copy

- **Backup**

- eXist style data export can also deuplicate BLOBs



Demo - BLOB Store



Thank You!
Ask Me Anything...

