

TRANSACTIONS AND ABSTRACTIONS OVER HBASE

Andreas Neumann
@anew68

Continuity

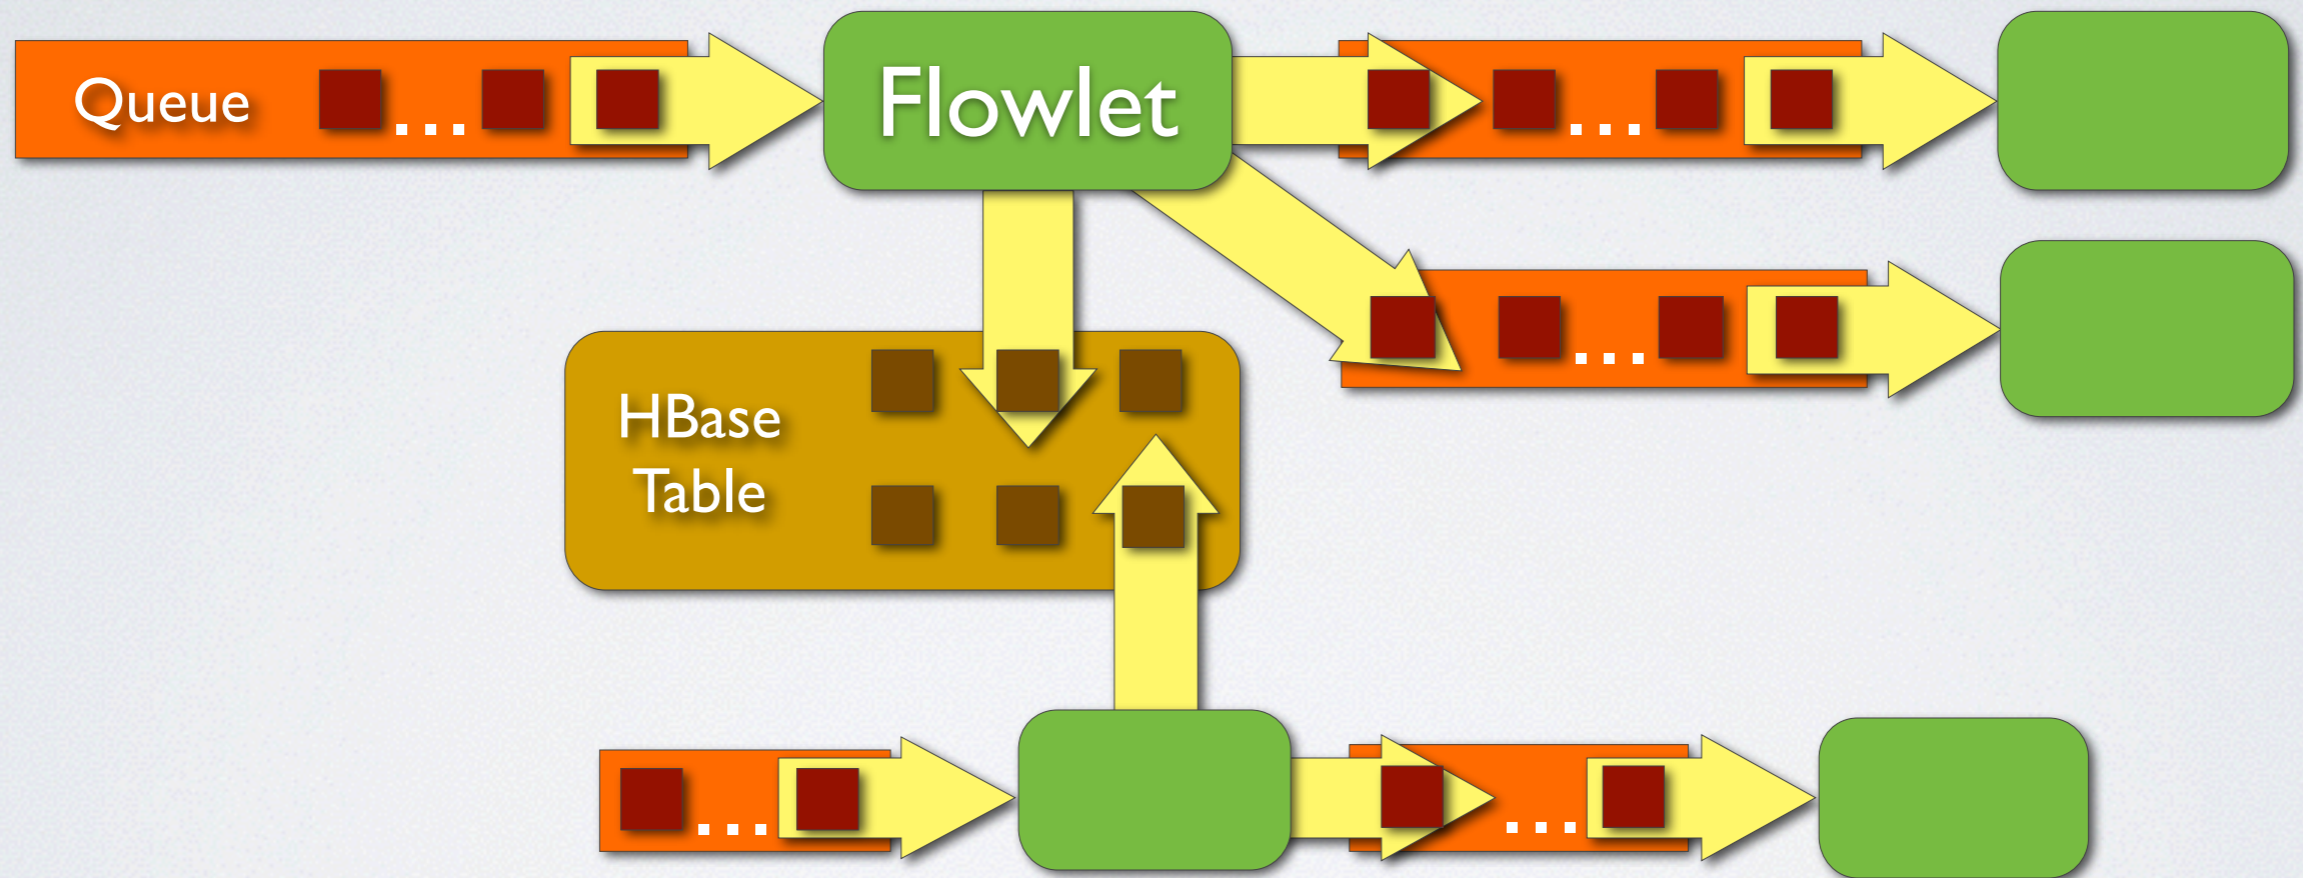
AGENDA

- Transactions over HBase: **Why? What?**
- Implementation: **How?**
 - The approach
 - Transaction Manager
- Abstractions
- Future

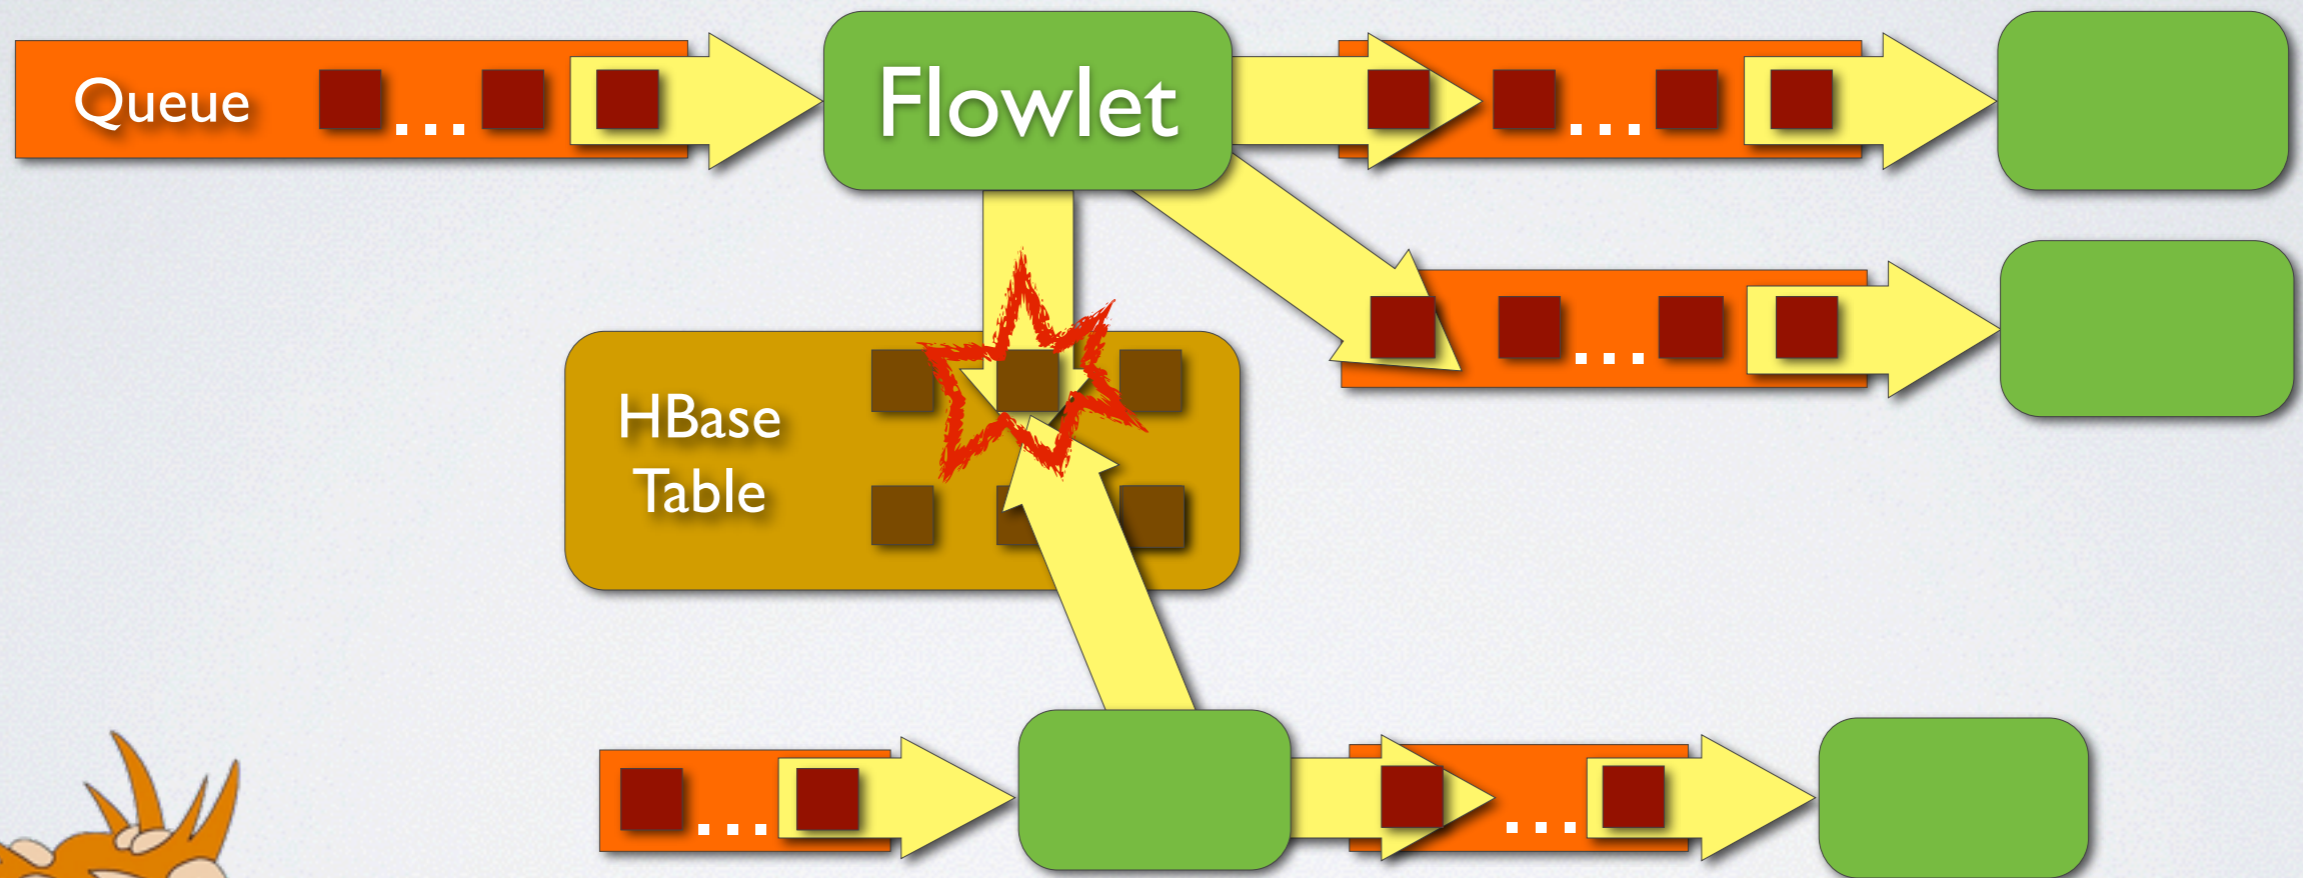
WHO WE ARE

- Simple Access to Powerful Technology
- **Continuity Reactor**: the world's first scale-out application server for Hadoop
 - Fast, easy development, deployment and management of **Hadoop** and **HBase** apps
 - Collect, **Process**, **Store**, and Query data
- **Real-time** Stream Processing


REAL-TIME STREAM PROCESSING



REAL-TIME STREAM PROCESSING



WRAP IN TRANSACTION!



TRANSACTIONS: WHAT?

- **Atomic** - Entire transaction is committed as one
- **Consistent** - No partial state change due to failure
- **Isolated** - No dirty reads, transaction is only visible after commit
- **Durable** - Once committed, data is persisted reliably

HBASE: QUICK “WHAT?”

- Open-source non-relational **distributed** column-oriented **database** modeled after Google’s BigTable
- Named Tables
 - Row key x Column key x *timestamp* -> Value
- Massive Scale
 - Key space partitioned into Regions.

WHAT ABOUT HBASE?

- **Atomic operations on cell value:**
checkAndPut, checkAndDelete, increment, append
- **Atomic batch of operations** on rows **within region**
- **No** cross region atomic operations support
- **No** cross table atomic operations support
- **No** multi-RPC atomic operations support

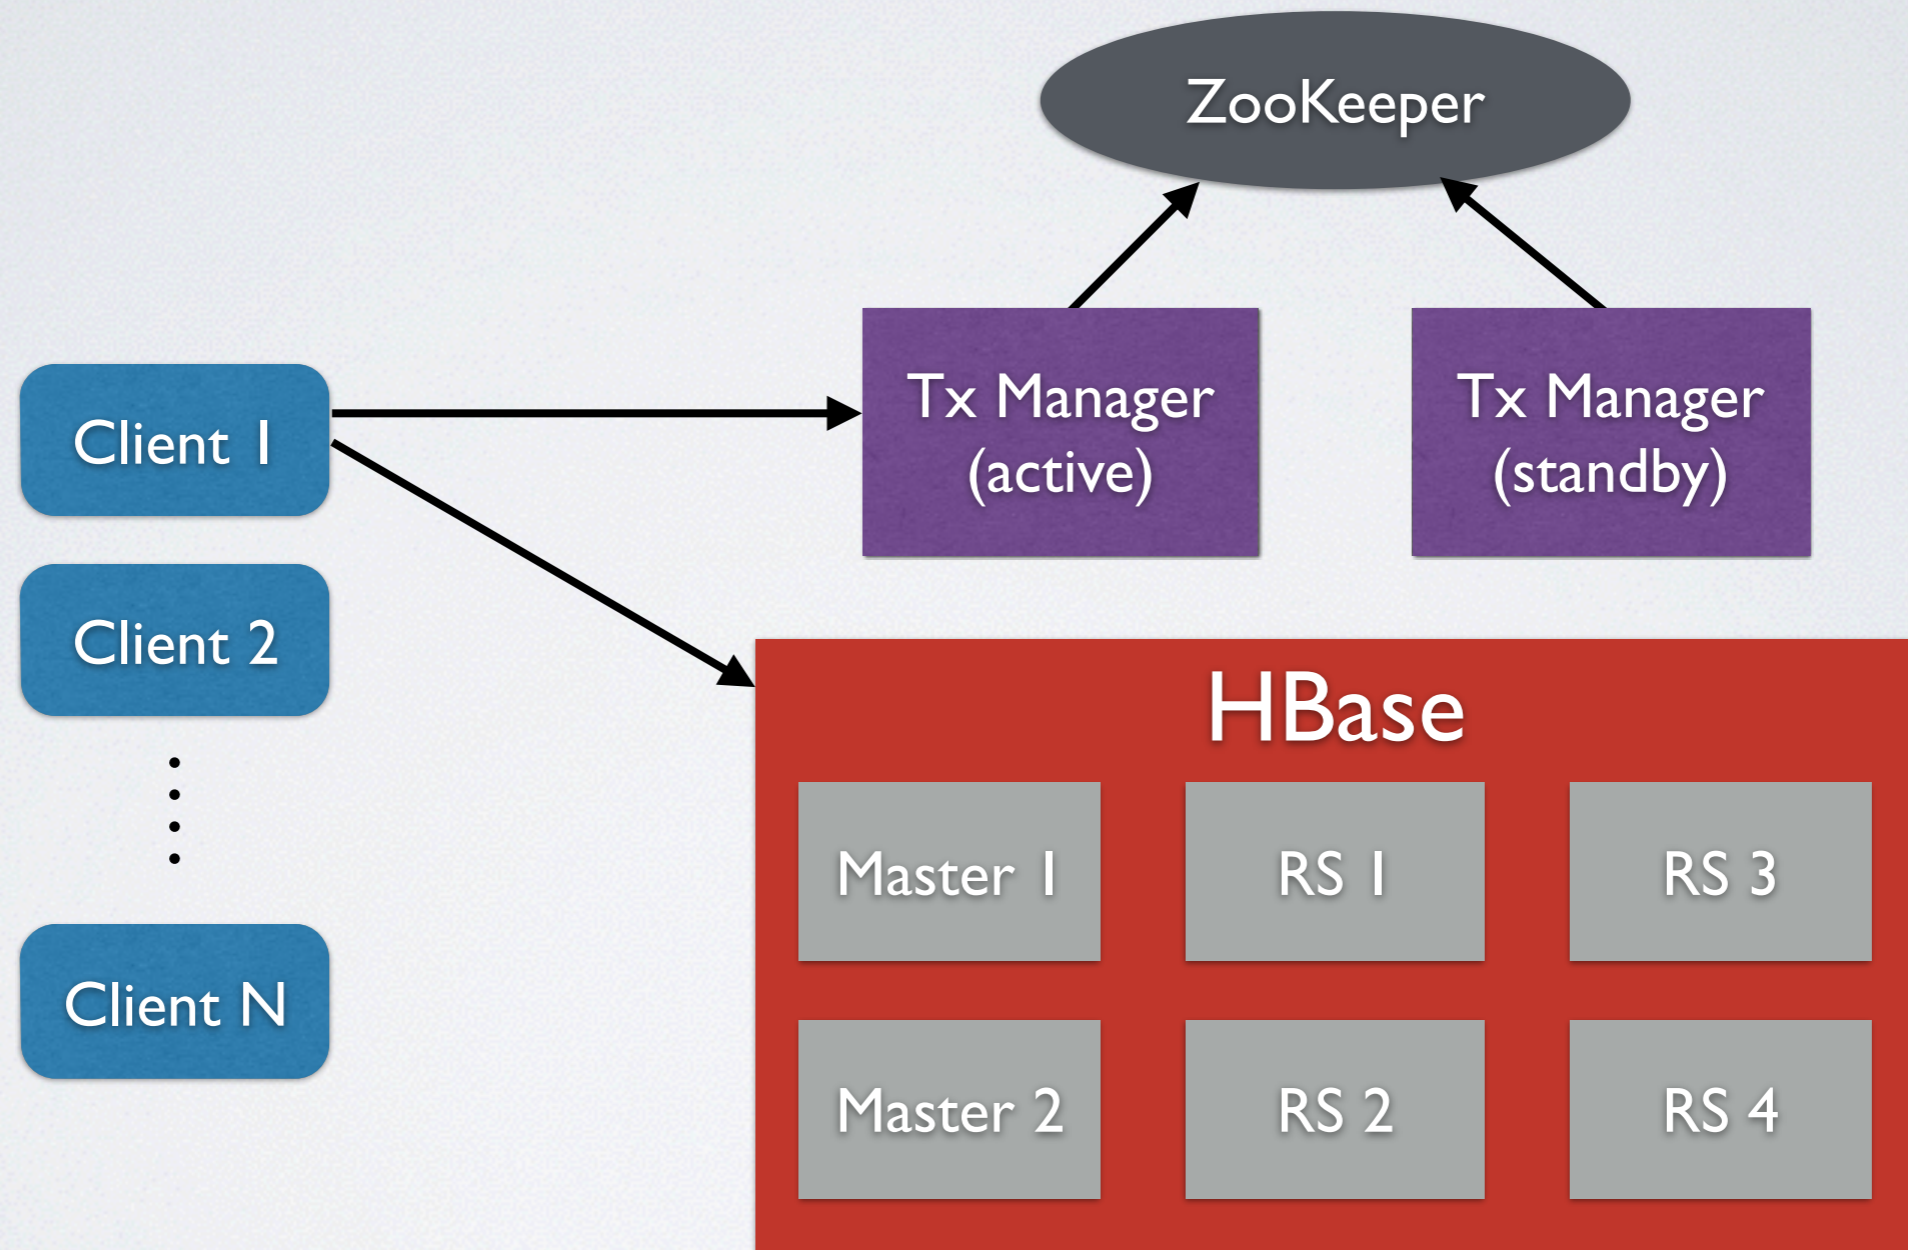
IMPLEMENTATION

- “OMID” style **Snapshot Isolation**
- **Multi-Version Concurrency Control**
 - **Cell version (timestamp) = transaction ID**
 - Reads exclude uncommitted transactions (for isolation)
- **Optimistic Concurrency Control**
 - Conflict detection at commit of transaction
 - Rollback in case of conflict (whichever commits later)

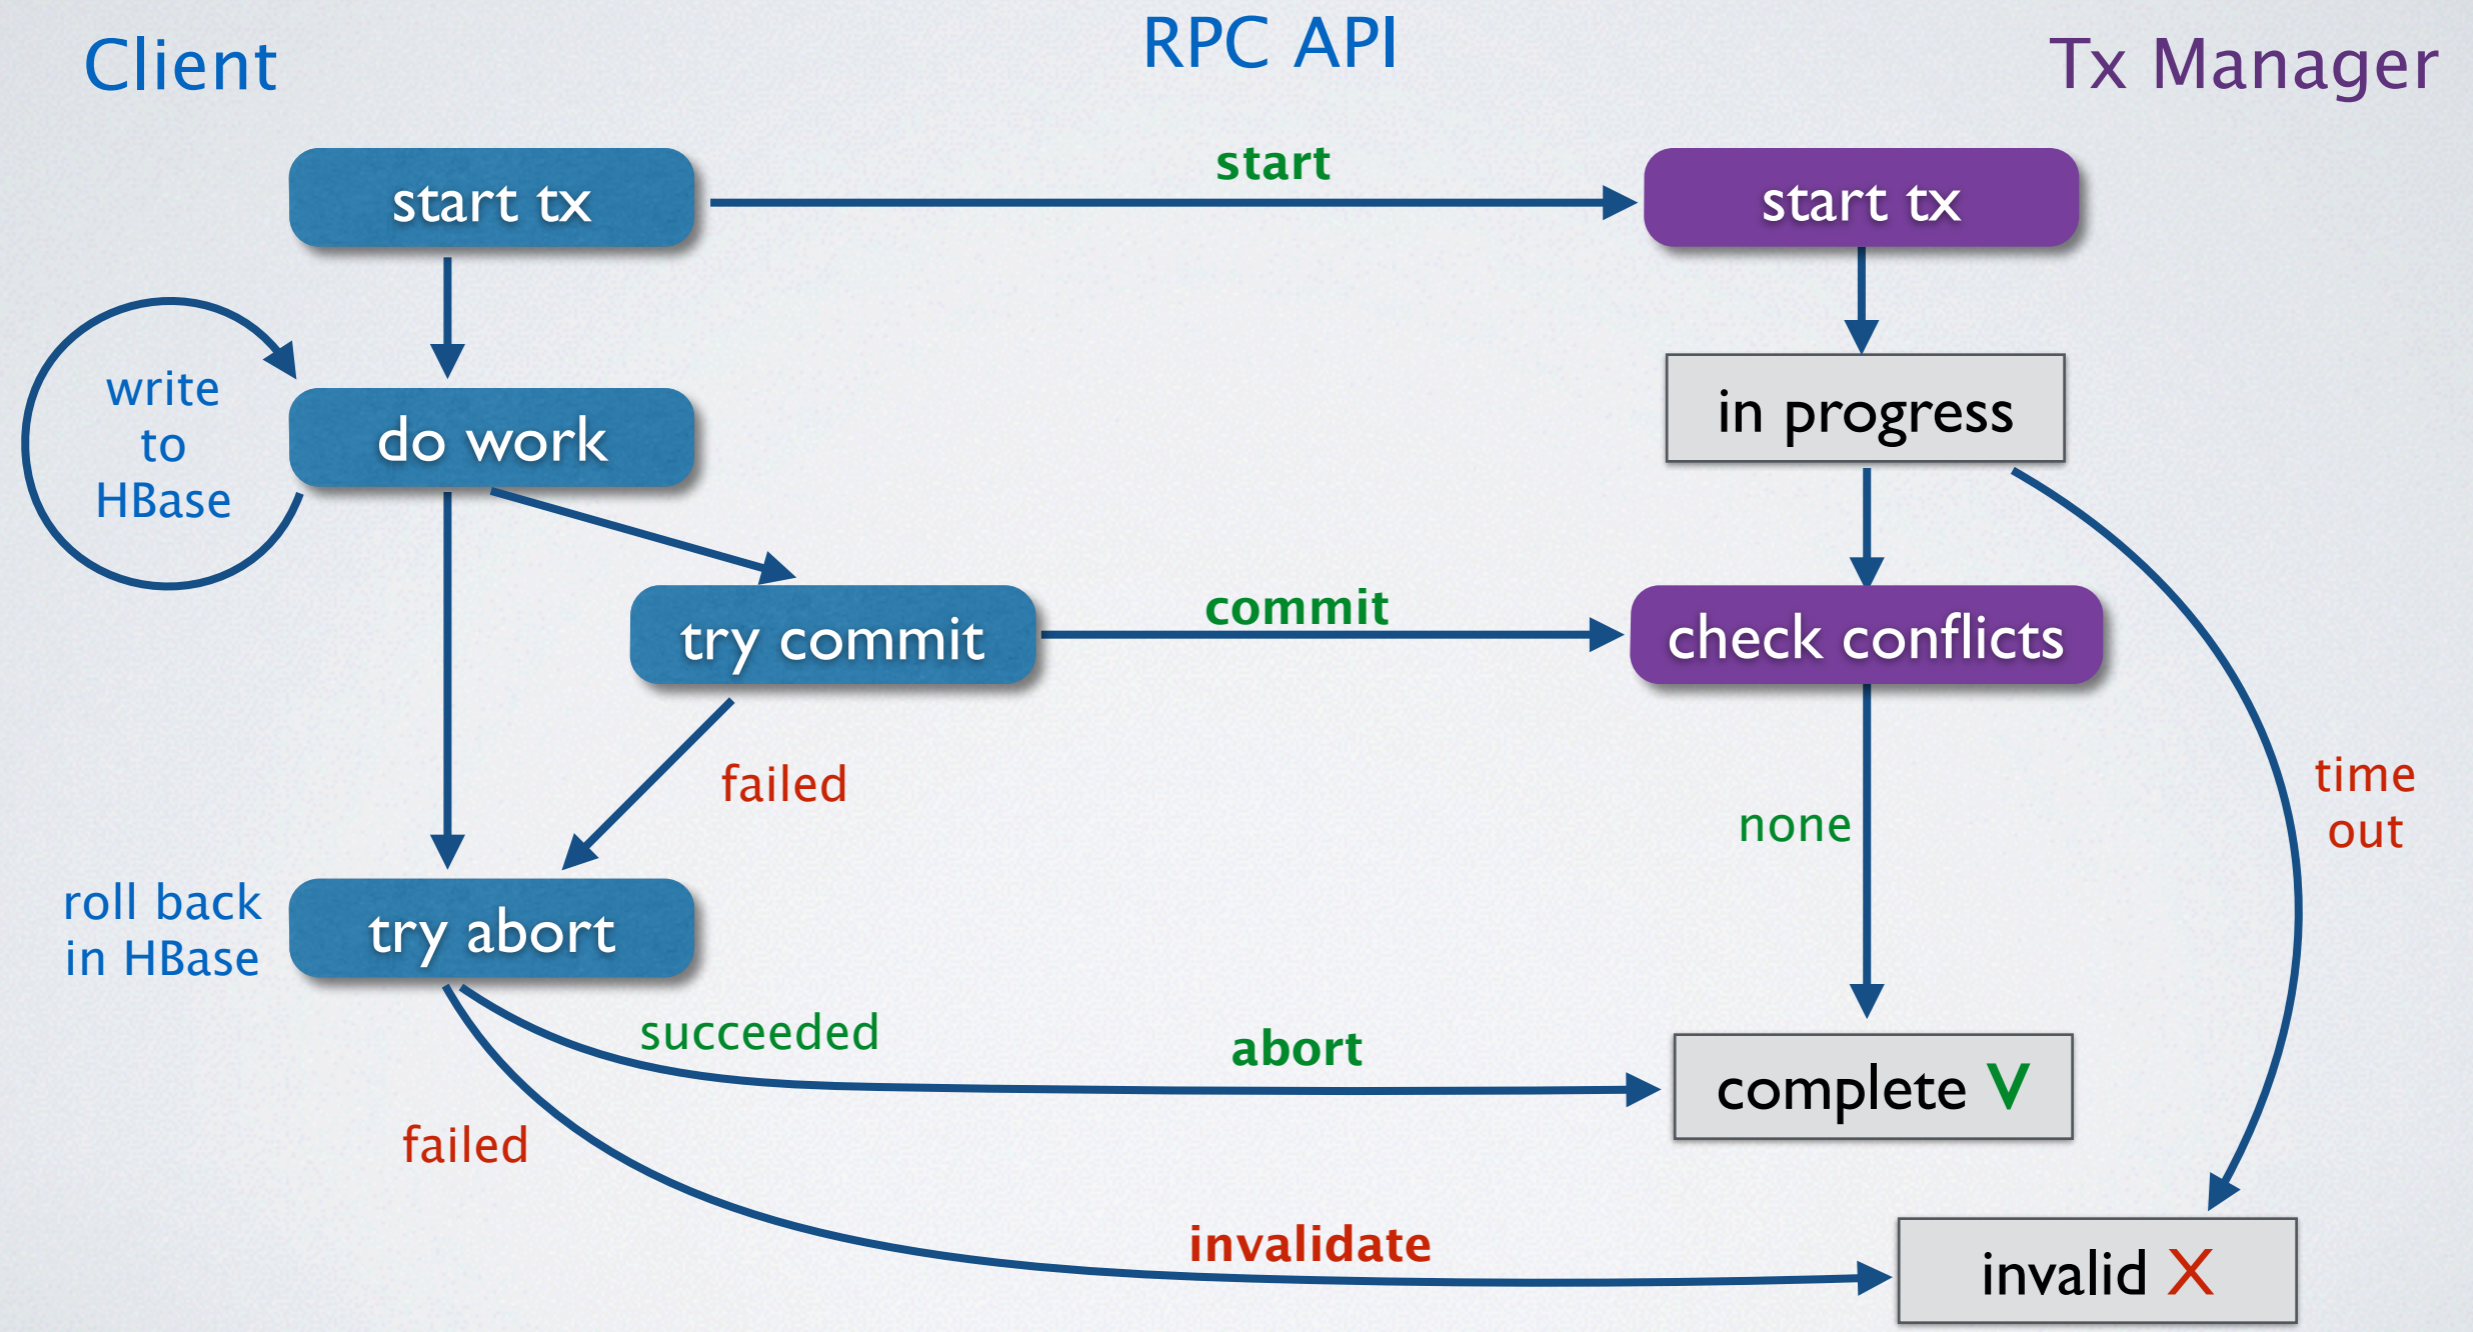
OPTIMISTIC CONCURRENCY CONTROL

- **Avoids cost of locking** rows and tables
- **No deadlocks** or lock escalations
- **Cost of conflict** detection and possible rollback is higher
- **Good if conflicts are rare:** short transaction, disjoint partitioning of work

TRANSACTIONS IN CONTEXT



TRANSACTION LIFE CYCLE



TRANSACTION MANAGER

- Create new transactions
 - Provides **monotonically increasing** write pointers
- Maintains all in-progress, committed, and invalid transactions
- Detect conflicts
- Transaction =

Write Pointer: Timestamp for HBase writes

Read pointer: Upper bound timestamp for reads

Excludes: List of timestamps to exclude from reads

TRANSACTION MANAGER

- Simple & Fast
 - All required state is in-memory
- Single point of failure?
 - Periodically persist **snapshot** of all state
 - **Write-ahead log** for all changes since last snapshot
 - Secondary Tx Manager watches for failure of Primary
 - Failover can happen quickly

TRANSACTION CLEANUP

- Some transactions time out or fail to roll back
 - Invalid transactions must be excluded from reads
 - Exclude list can get large over time
- Old versions may not be visible to any transaction
- TTL (time-to-live) expires old versions

DATA JANITOR

- RegionObserver coprocessor
- Maintains in-memory snapshot of recent invalid & in-progress sets
- Periodically updates from transaction snapshot in HDFS
- Purges data from invalid transactions and older versions on flush & compaction



ABSTRACTION

- Dataset implements TransactionAware interface

```
void startTx(Transaction tx);
```

```
Collection<byte[]> getTxChanges();
```

```
boolean commitTx() throws Exception;
```

```
boolean rollbackTx() throws Exception;
```

- Other Data Stores than HBase
- HBase, LevelDB, HyperSQL, In-Memory, ...

TRANSACTION AWARE

- Modulation of ACID by dataset implementation
 - Granularity of keys - row, column, ...
 - In-memory caching of writes
 - Skip rollback
 - ...

WHAT'S NEXT?

- Continue Scaling Tx Manager
 - Transaction Groups?
- Integration across other transactional stores
- Open Source
 - <http://continuity.com/blog>



Looking for the chance to work with a team that is defining
a new category within Big Data?

We are hiring!

<http://continuuity.com/careers>

[careers\[at\]continuuity.com](mailto:careers@continuuity.com)

Andreas Neumann @anew68 [andreas\[at\]continuuity.com](mailto:andreas@continuuity.com)