

\$

# Spirit

Online Schema  
Change Utility  
for MySQL 8.0

Morgan Tocker  
MySQL Belgian Days 2024



# Agenda.

1. **State of DDL**
2. How Spirit works
3. Optimizations
4. Feature requests for MySQL DDL

# ALTER TABLE t1 \_

**\* Simplified**

ADD COLUMN, DROP COLUMN, RENAME COLUMN,  
CHANGE DEFAULT VALUE, CHANGE ENUM/SET, RENAME TABLE

**INSTANT**

DROP INDEX, RENAME INDEX, ADD INDEX, ADD PARTITION, DROP  
PARTITION, TRUNCATE PARTITION, CHANGE NULL/NOT NULL,  
INCREASE VARCHAR (same ptr len)

**INPLACE**

CHANGE DATA TYPE, ADD PRIMARY KEY,  
DROP PRIMARY KEY, CHANGE CHARACTER SET

**COPY**

# INSTANT

- New to MySQL 8.0
- Works 99% as described
- Nuanced exceptions:
  - Max 64 instant ADD/DROP columns before a fully copy
  - Compressed tables
  - Temporary tables
  - Tables with FULLTEXT indexes



Assert with:

```
ALTER TABLE t1 ADD newcol  
INT, ALGORITHM=INSTANT;
```

# INPLACE

Time varies between operations

Blocks InnoDB Purge



	Statement	Algorithm	Time
Drop Index	DROP INDEX b	INPLACE	0.00s
Extend Varchar (same ptr len)	MODIFY b VARCHAR(61)	INPLACE	0.00s
Add Column	ADD COLUMN d INT	INSTANT	0.08s
Add Index	ADD INDEX(b)	INPLACE	8.49s
Extend Varchar (larger ptr len)	MODIFY b VARCHAR(70)	COPY	16.23s
Set NULL	MODIFY b VARCHAR(60) NULL	INPLACE	18.74s
Set NOT NULL	MODIFY b VARCHAR(60) NOT NULL	INPLACE	18.83s
NULL ALTER	ENGINE=InnoDB	COPY	18.88s

## COPY

- Rebuilds the whole table
- Some operations are also blocking:  
“Permits Concurrent DML: No”

```
ALTER TABLE t1  
ALGORITHM=copy, LOCK=none;
```



# How do you make schema changes then?

1. Once upon a time:  
replication flipping  
(aka blue green)
  - Not safe with Row-based Replication
2. Use an external tool:
  - Pt-online-schema-change
  - Gh-ost
  - Spirit (**new!**)

# Agenda.

1. State of DDL
2. **How Spirit works**
3. Optimizations
4. Feature requests for MySQL DDL

# \$ Basic Algorithm

- Copy table definition to a new table
- Alter the new table
- Start subscribing to replication changes
- Copy rows from origin to new
- Apply changes from replication
- Cutover

# Nuanced Algorithm

- **Attempt INSTANT DDL**
- **Parse DDL and apply INPLACE if only DROP/RENAME index**
- Copy table definition to a new table
- Alter the new table
- Start subscribing to replication changes
- Copy rows from origin to new
- Apply changes from replication
- **Checksum**
- Cutover

# Agenda.

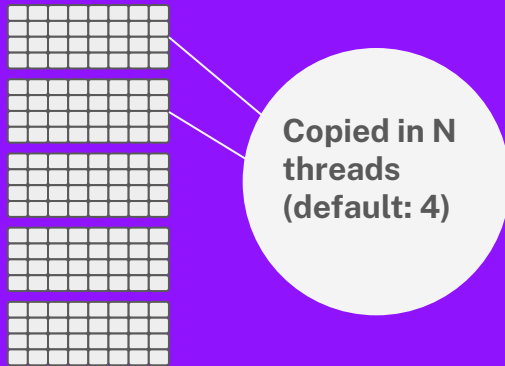
1. State of DDL
2. How Spirit works
- 3. Optimizations**
4. Feature requests for MySQL DDL

# Limitations (for some)

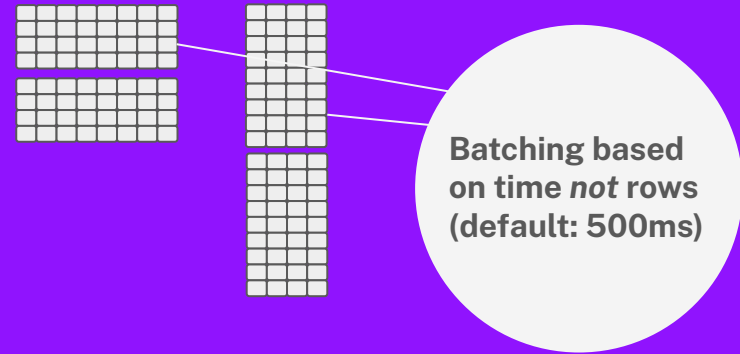
- Only MySQL 8.0
- Must have a Primary Key
- No FOREIGN KEY support
- No RENAME column
- No Triggers
- No lossy conversions
- Replication Throttling to minimum of 10s fidelity

# Optimizations: Copying Rows

## Parallel Copy

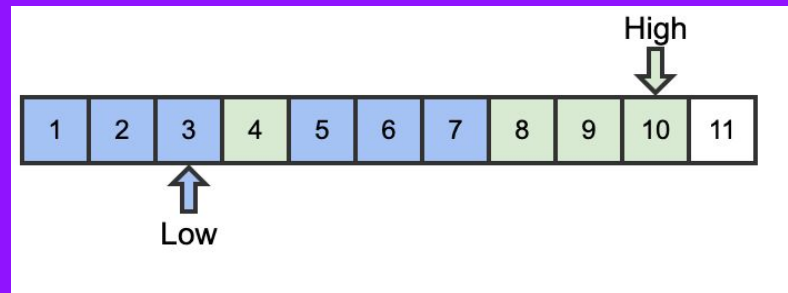


## Dynamic Chunking



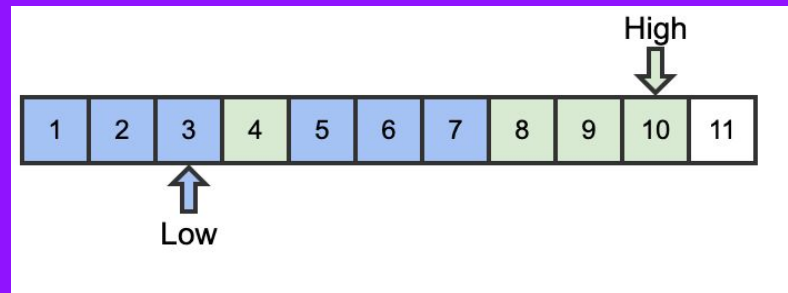
# Optimizations: Apply Changes

- Delta Map / Parallel Apply
  - Requires memory comparable Primary Key
- Key Above Watermark
  - Currently requires auto-increment Primary Key



# Features: Checkpointing

- Every ~1m save point the copy/apply has safely completed
- Allows for resume of changes



# Agenda.

1. State of DDL
2. How Spirit works
3. Optimizations
4. **Feature requests for MySQL DDL**

# BUG #113355

DROP/RENAME INDEX as INSTANT

ADD COLUMN, DROP COLUMN, RENAME COLUMN,  
CHANGE DEFAULT VALUE, CHANGE ENUM/SET, RENAME TABLE

**INSTANT**

DROP INDEX, RENAME INDEX, ADD INDEX, ADD PARTITION, DROP  
PARTITION, TRUNCATE PARTITION, CHANGE NULL/NOT NULL,  
INCREASE VARCHAR (same ptr len)

**INPLACE**

CHANGE DATA TYPE, ADD PRIMARY KEY,  
DROP PRIMARY KEY, CHANGE CHARACTER SET

**COPY**

# BUG #113476

Long running INPLACE grows HLL

ADD COLUMN, DROP COLUMN, RENAME COLUMN,  
CHANGE DEFAULT VALUE, CHANGE ENUM/SET, RENAME TABLE

**INSTANT**

DROP INDEX, RENAME INDEX, **ADD INDEX**, ADD PARTITION, DROP  
PARTITION, TRUNCATE PARTITION, CHANGE NULL/NOT NULL,  
INCREASE VARCHAR (same ptr len)

**INPLACE**

CHANGE DATA TYPE, ADD PRIMARY KEY,  
DROP PRIMARY KEY, CHANGE CHARACTER SET

**COPY**

# BUG #113355

Support INSTANT for some change type

ADD COLUMN, DROP COLUMN, RENAME COLUMN,  
CHANGE DEFAULT VALUE, CHANGE ENUM/SET, RENAME TABLE

**INSTANT**

DROP INDEX, RENAME INDEX, ADD INDEX, ADD PARTITION, DROP  
PARTITION, TRUNCATE PARTITION, CHANGE NULL/NOT NULL,  
INCREASE VARCHAR (same ptr len)

**INPLACE**

**CHANGE DATA TYPE**, ADD PRIMARY KEY,  
DROP PRIMARY KEY, CHANGE CHARACTER SET

**COPY**



# The Potential Future

vs OUR MOST COMMON DDLs

ADD COLUMN, DROP COLUMN, DROP INDEX, CHANGE DATA TYPE (expansion), RENAME COLUMN, CHANGE DEFAULT VALUE, CHANGE ENUM/SET

**INSTANT**

ADD INDEX

**INPLACE**

CHANGE DATA TYPE (other), CHANGE NULL/NOT NULL

**SPIRIT**

## Conclusion

### **Spirit is symbiotic**

It attempts MySQL DDL  
first, falls back as required

As more changes become  
instant, Spirit will take  
advantage

### **Instant DDL is awesome**

Large schema changes #1  
blocker to larger databases

[github.com/cashapp/spirit](https://github.com/cashapp/spirit)  
Apache 2.0 License