



# QuestDB

Time series data, faster



**Tom GRAUGNARD**


# Sommaire

- Qu'est ce que c'est ? 3
- Comment ça marche ? 5
- Fonctionnalités 10
- Démonstration 15
- Conclusion 16


# Qu'est ce que c'est ?



Built for performance



Optimized for time series



Implemented with SQL 3

# Performances

| Opération | 64-bit double  | 32-bit int     |
|-----------|----------------|----------------|
| Ecriture  | 120 Million /s | 240 Million /s |
| Lecture   | 240 Million /s | 480 Million /s |


→ Par thread

16 threads sur  
96 disponible



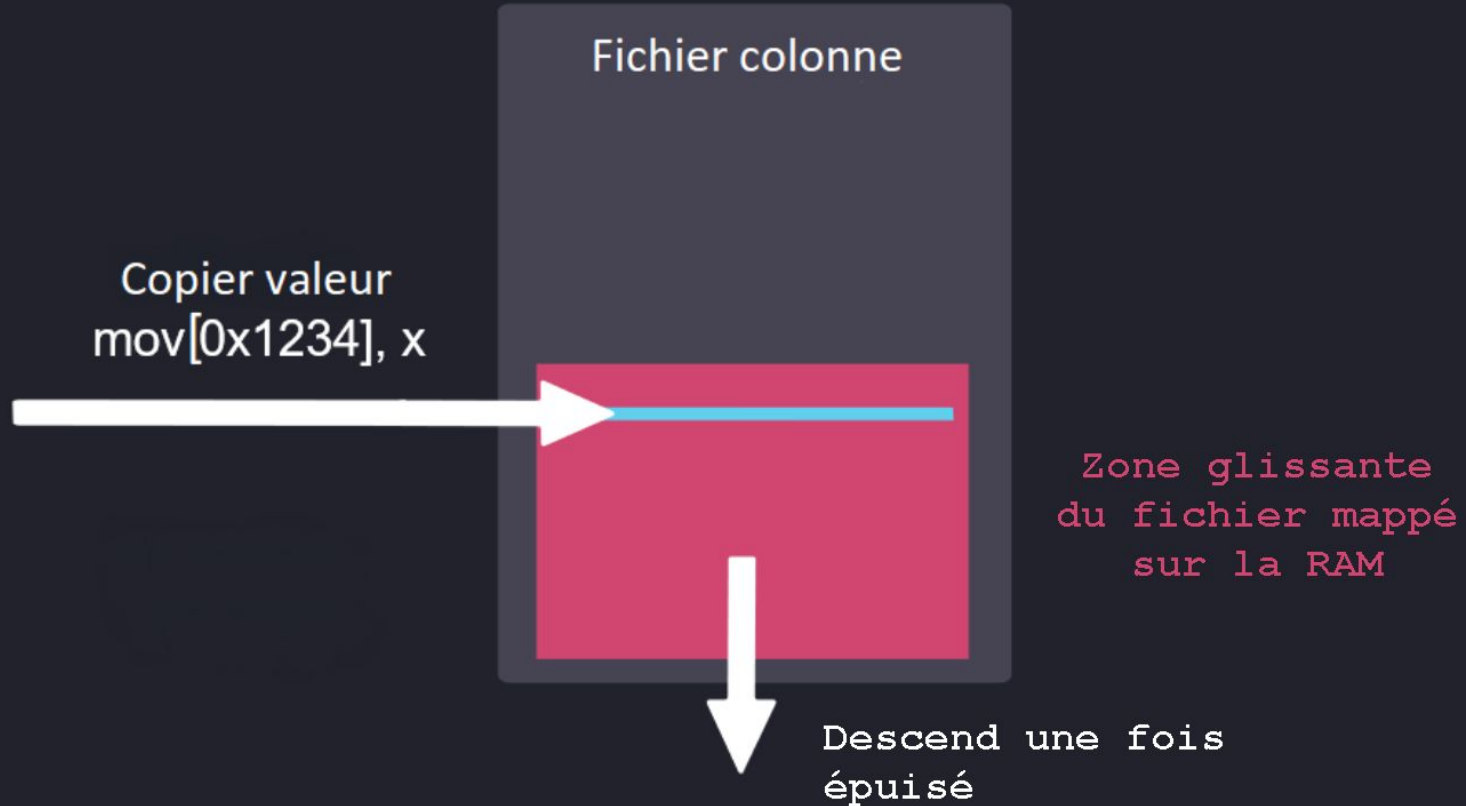
| Requête  | Temps d'exécution |
|--|-------------------|
| SELECT sum(double) FROM 1bn                                | 0.061 secs        |
| SELECT tag, sum(double) FROM 1bn                           | 0.179 secs        |
| SELECT tag, sum(double) FROM 1bn<br>WHERE timestamp='2019' | 0.05 secs         |

**Comment ça marche ?**




**Ajout à la fin de fichier**

# Ecriture

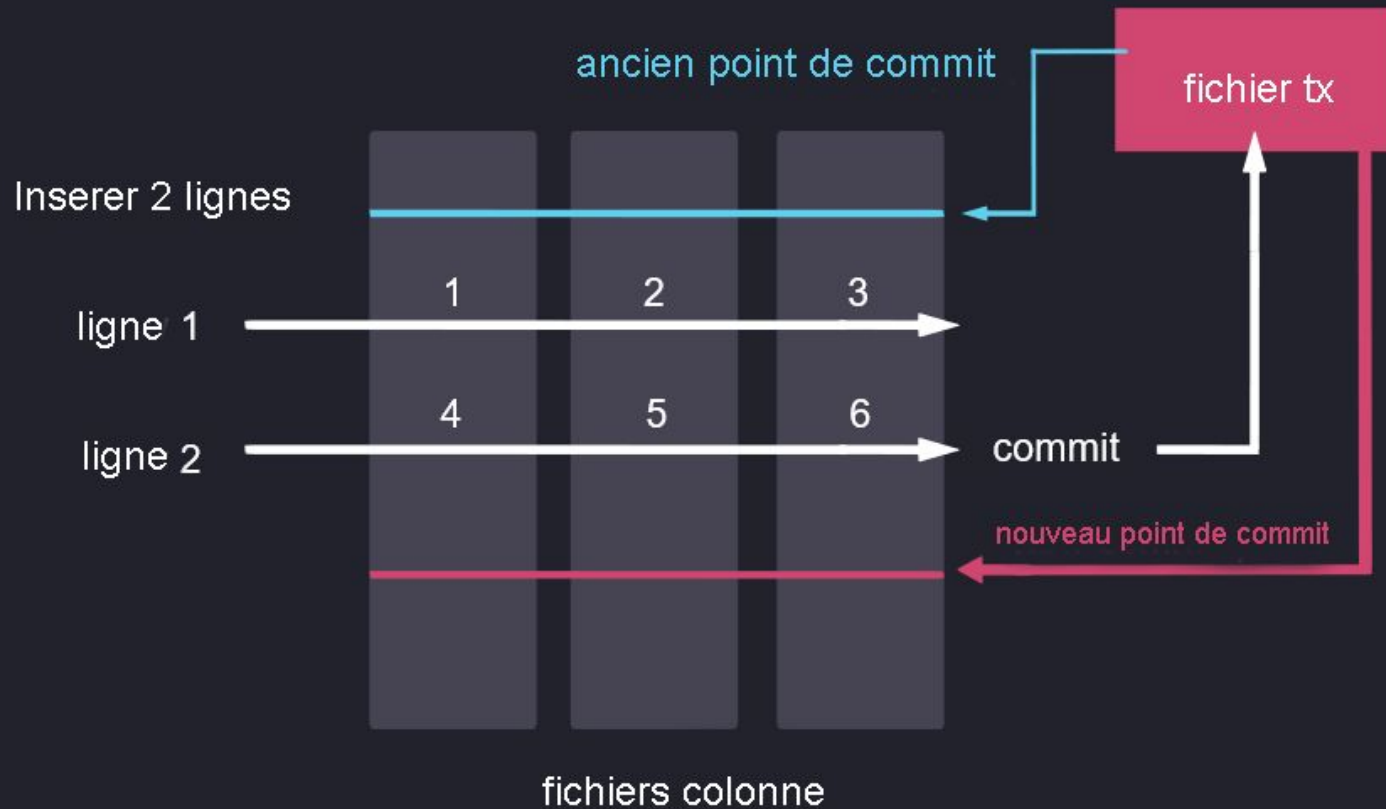


# Lecture



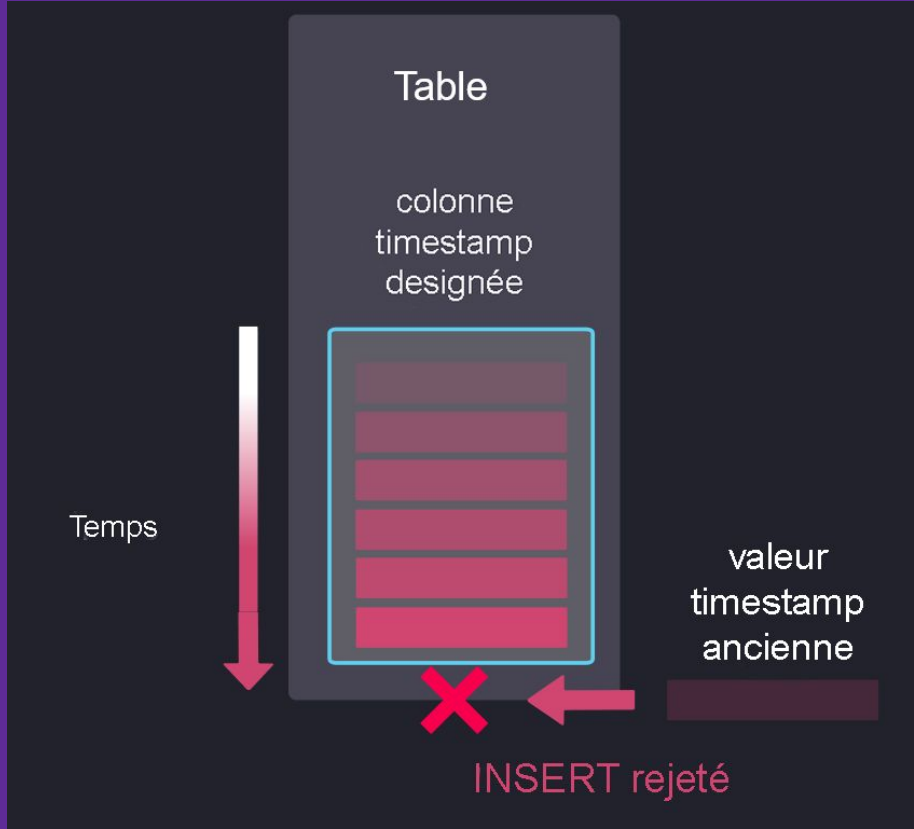
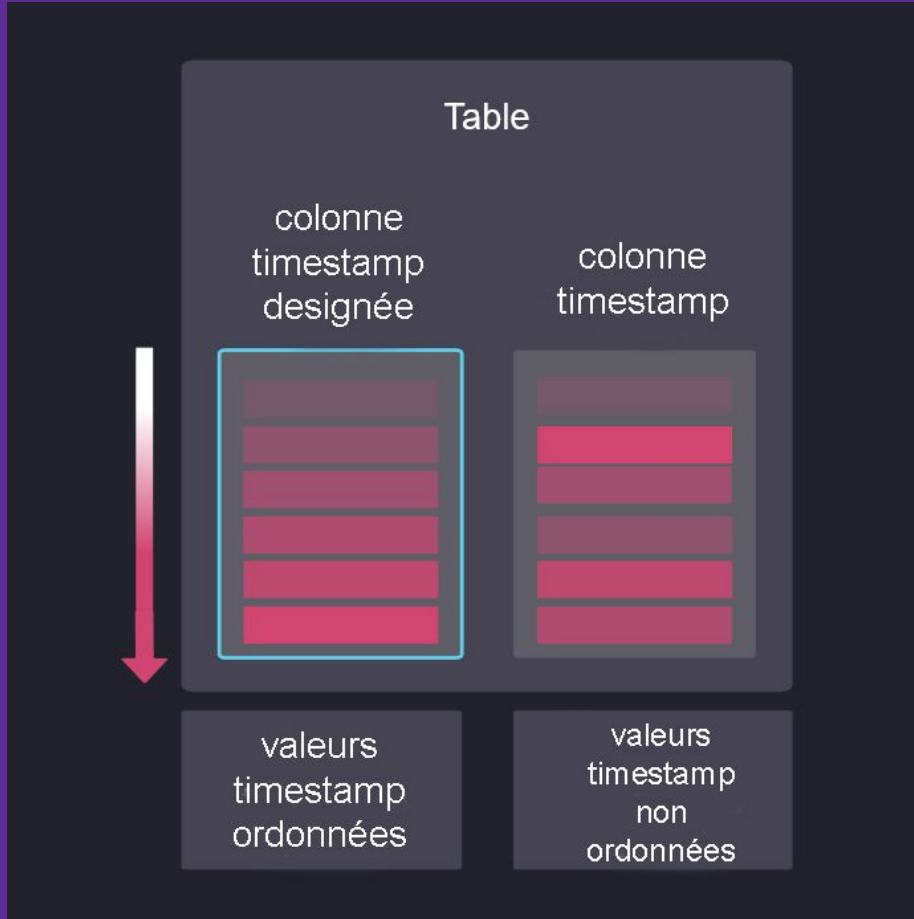


# Atomicité

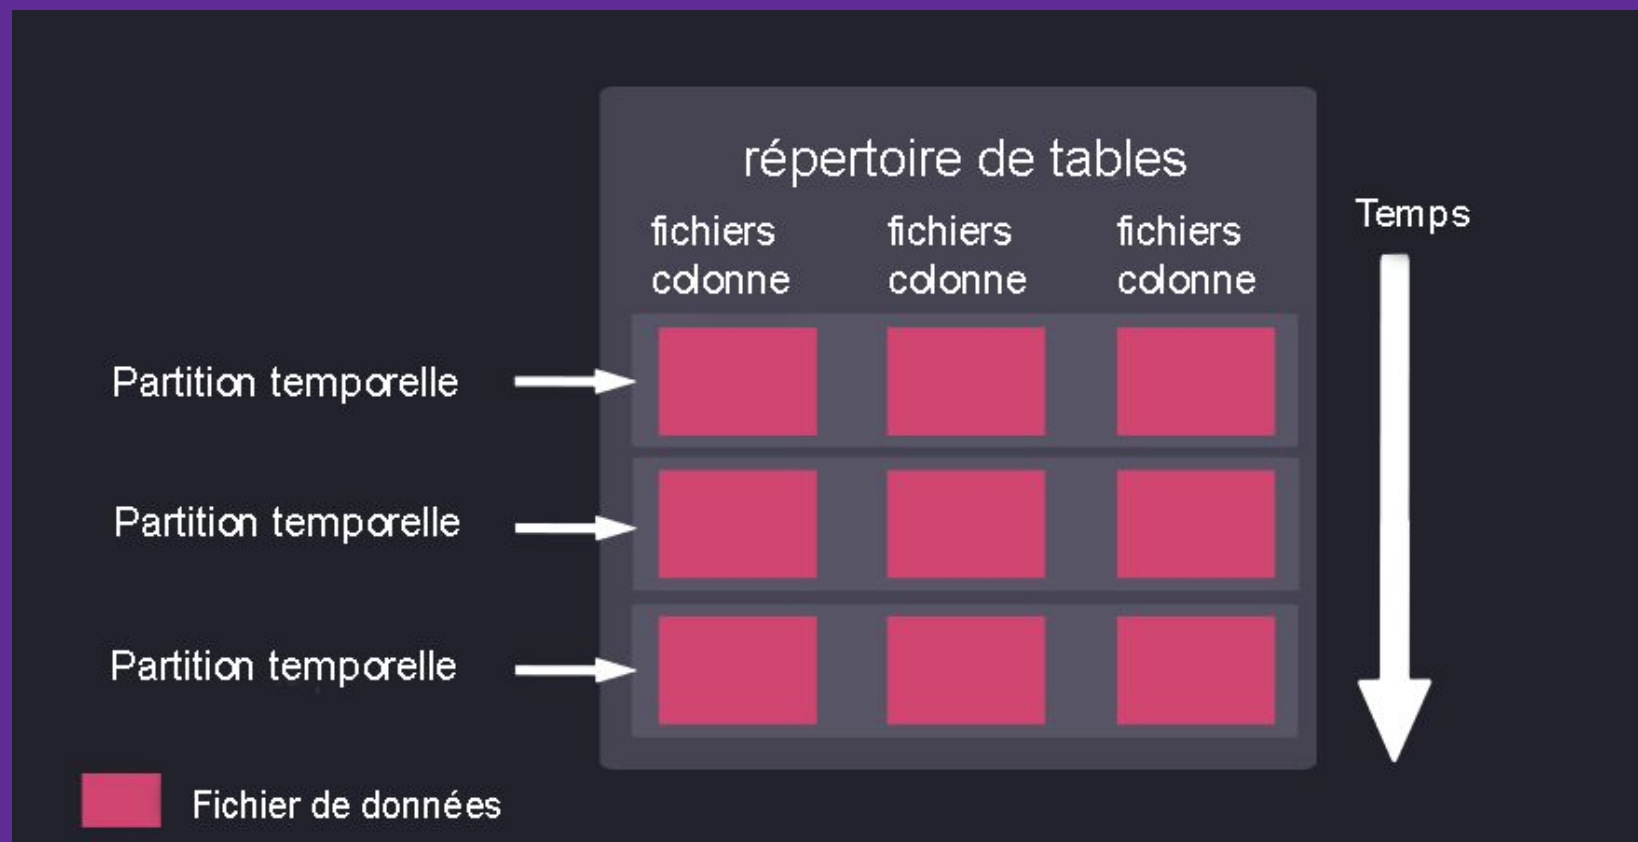


# Fonctionnalités

# designated timestamp



# Partition



# Symboles

| ID | Symbol |
|----|--------|
| 1  | Alpha  |
| 2  | Beta   |
| 3  | Omega  |

# Symboles

Table

| Row ID | Symbol | Value | ID | Symbol |
|--------|--------|-------|----|--------|
| 1      | 1      | 1     | 1  | Alpha  |
| 2      | 1      | 0     | 2  | Beta   |
| 3      | 2      | 1     | 3  | Omega  |
| 4      | 1      | 1     |    |        |
| 5      | 3      | 0     |    |        |

# Index

| Table  |        |       | Index  |         |
|--------|--------|-------|--------|---------|
| Row ID | Symbol | Value | Symbol | Row IDs |
| 1      | A      | 1     | A      | 1, 2, 4 |
| 2      | A      | 0     | B      | 3       |
| 3      | B      | 1     | C      | 5       |
| 4      | A      | 1     |        |         |
| 5      | C      | 0     |        |         |

# Index

| Table  |        |       | Index  |         |
|--------|--------|-------|--------|---------|
| Row ID | Symbol | Value | Symbol | Row IDs |
| 1      | A      | 1     | A      | 1, 2, 4 |
| 2      | A      | 0     | B      | 3, 6    |
| 3      | B      | 1     | C      | 5       |
| 4      | A      | 1     |        |         |
| 5      | C      | 0     |        |         |
| 6      | B      | 1     |        |         |



# Démonstration

# Conclusion

## Inconvénients :

- Contraintes, check, triggers non supportés
- Postgres pas encore 100% supporté
- Bugs pas encore corrigés

**Merci pour votre attention !**

**Questions ?**