QuestDB

Time series data, faster

Tom GRAUGNARD

Sommaire

Qu’est ce que c’est ?

Built for performance

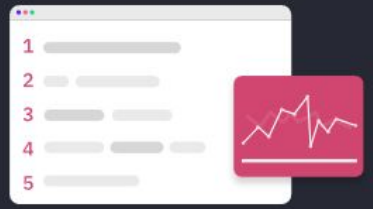

Optimized for time series

Implemented with SQL

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 WHERE timestamp='2019'	0.05 secs

Comment ça marche ?

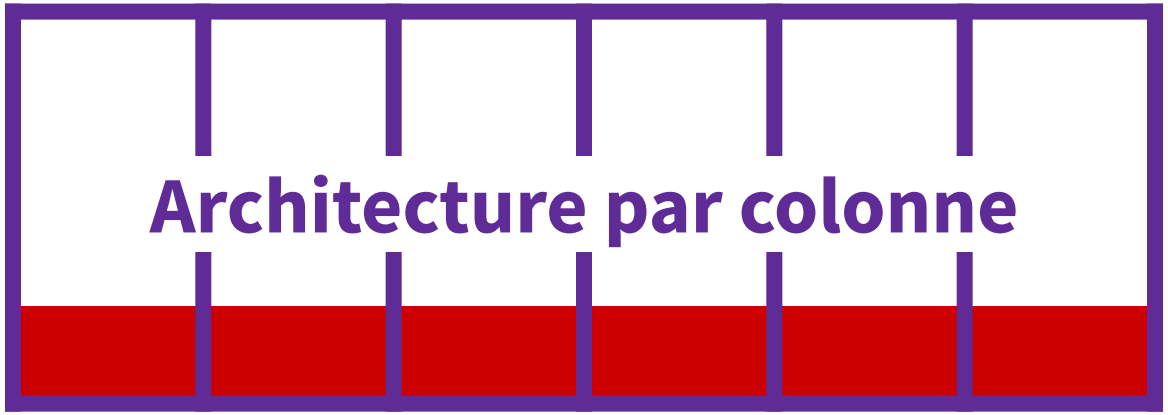

Ajout à la fin de fichier

Ecriture

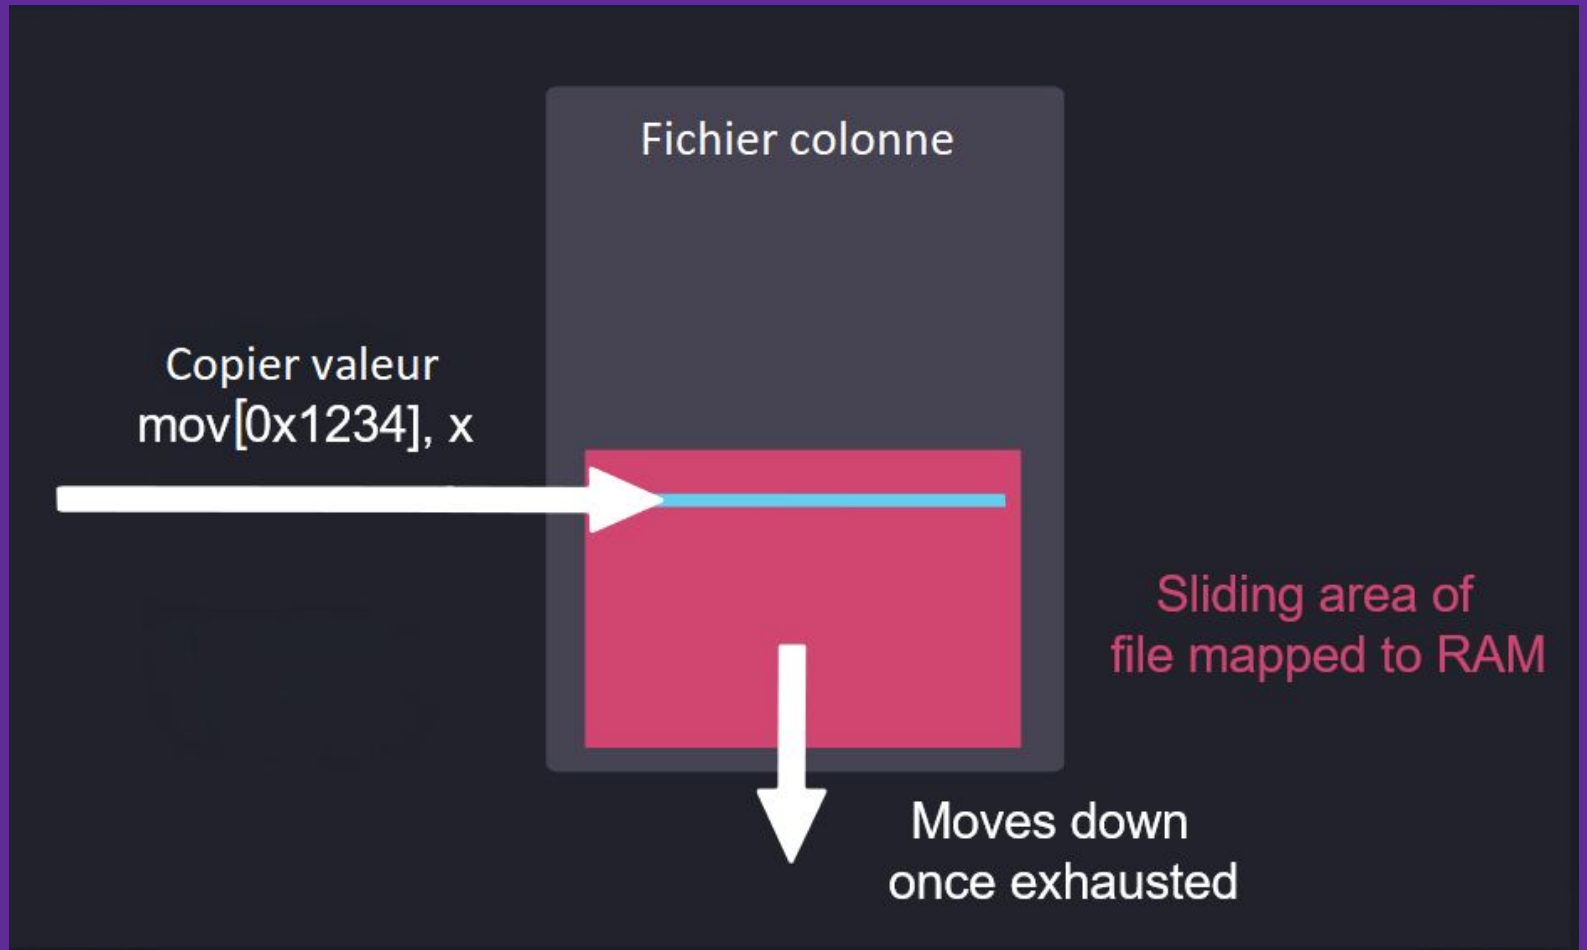
Fichier colonne
Copier valeur
mov[0x1234], x
Sliding area of
file mapped to RAM
Moves down
once exhausted

Lecture

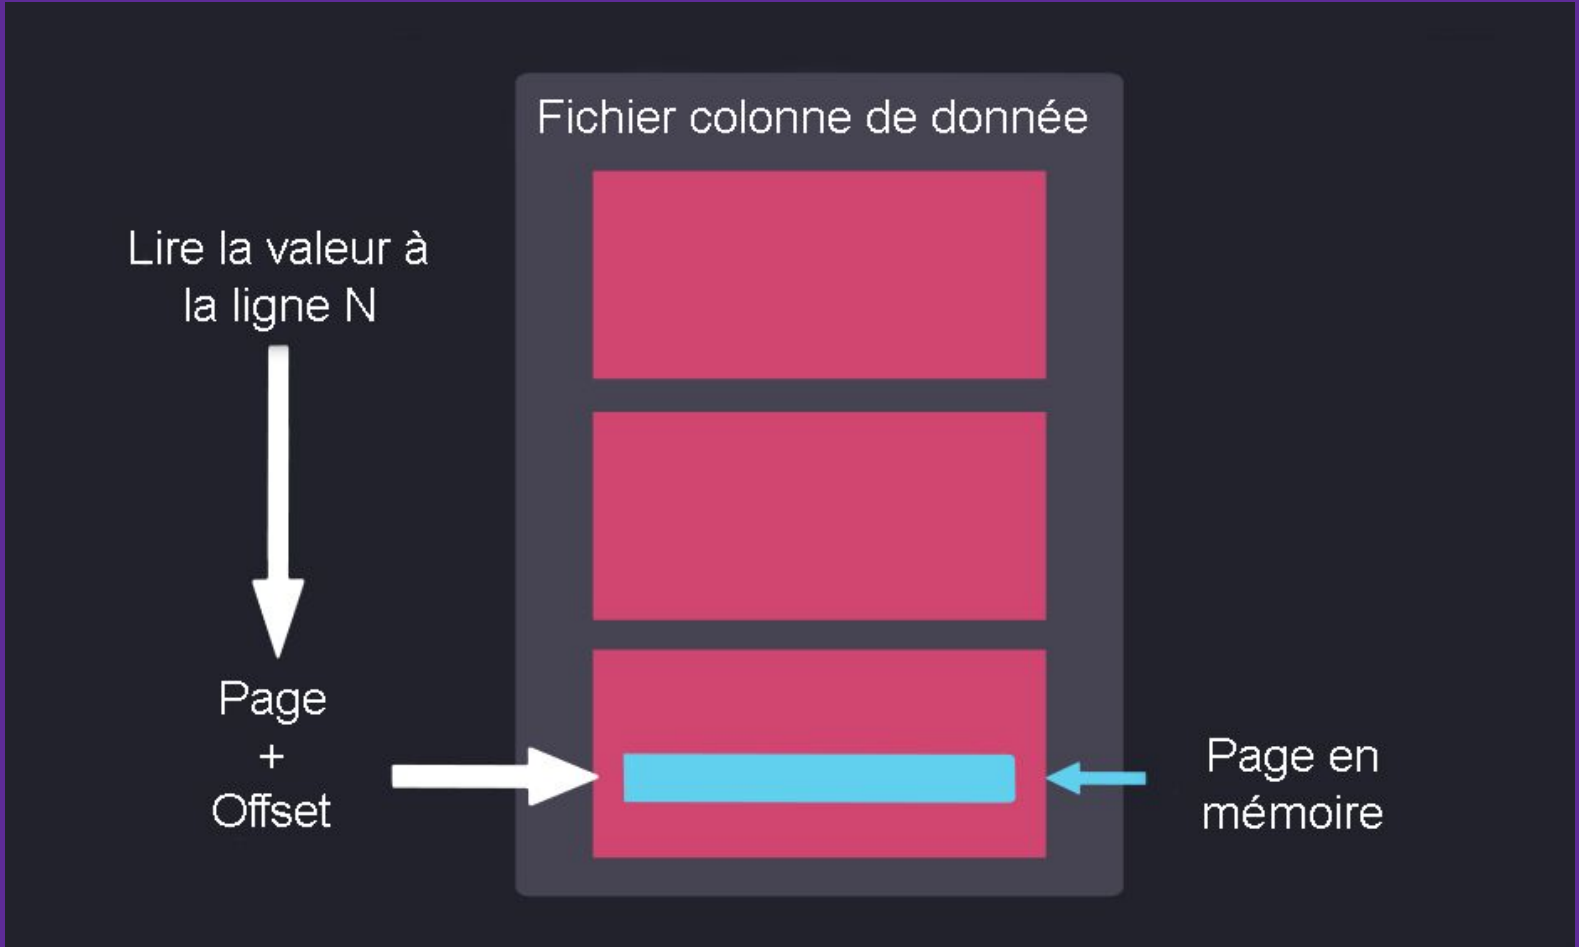
Fichier colonne de donnée
Lire la valeur à la ligne N
Page + Offset
Page en mémoire

Atomicité

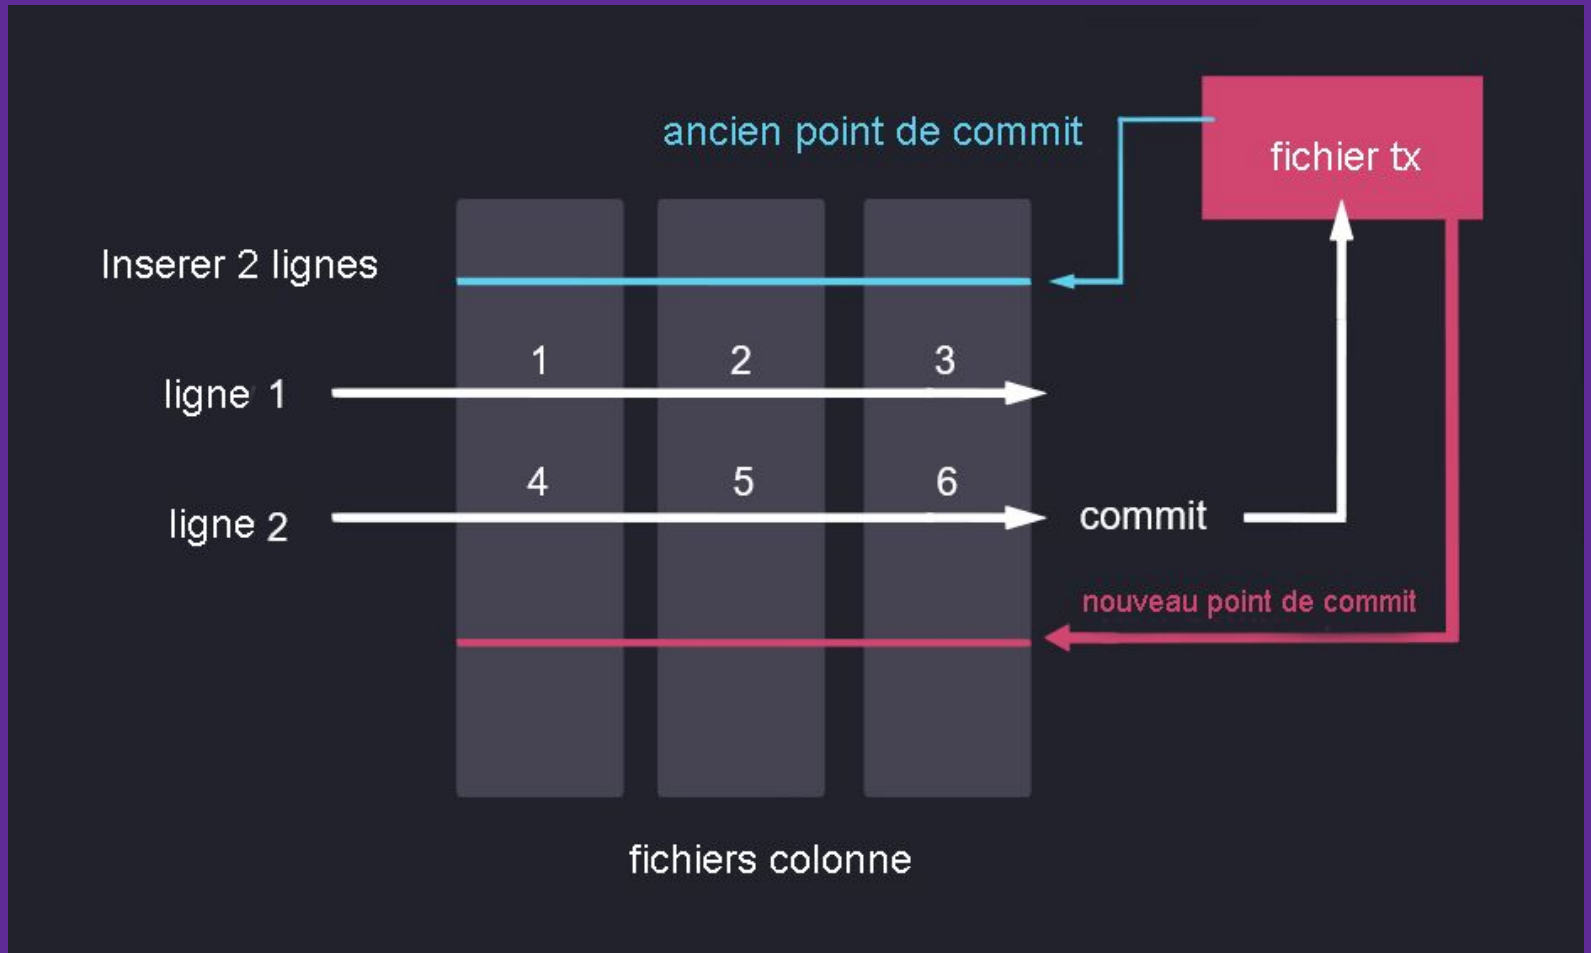
ancien point de commit
fichier tx
Inserer 2 lignes
ligne 1
1
2
3
ligne 2
4
5
6
commit
nouveau point de commit
fichiers colonne

Fonctionnalités

designated timestamp

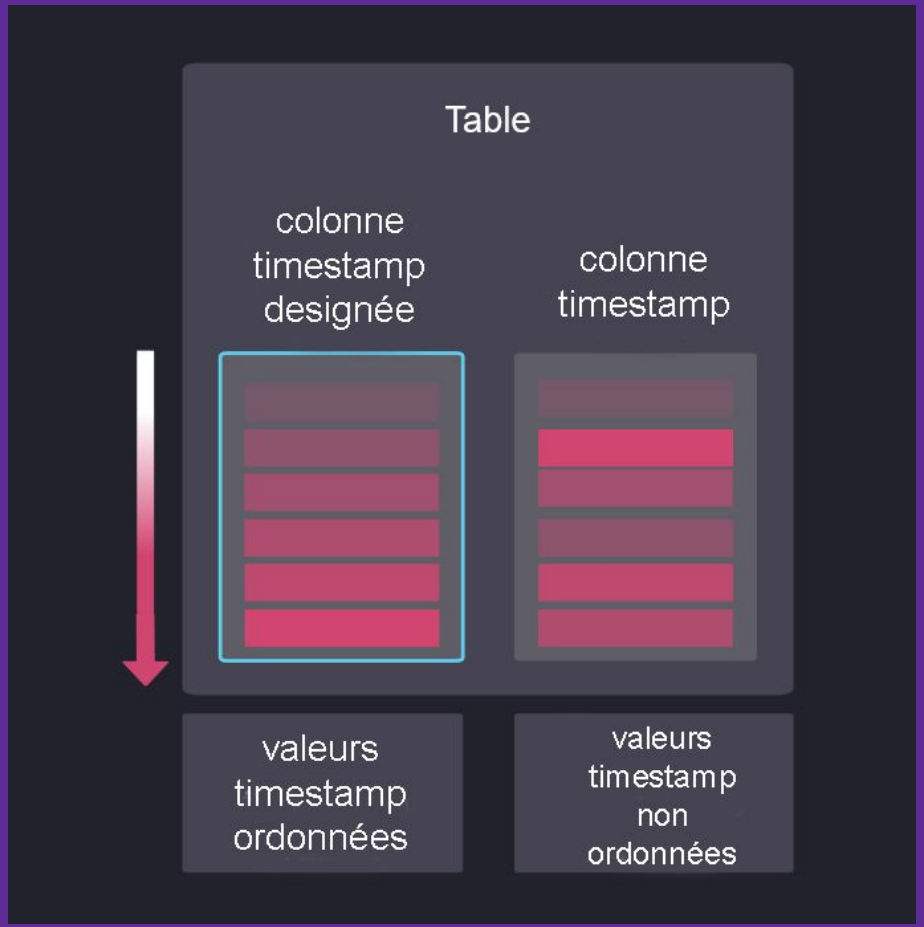

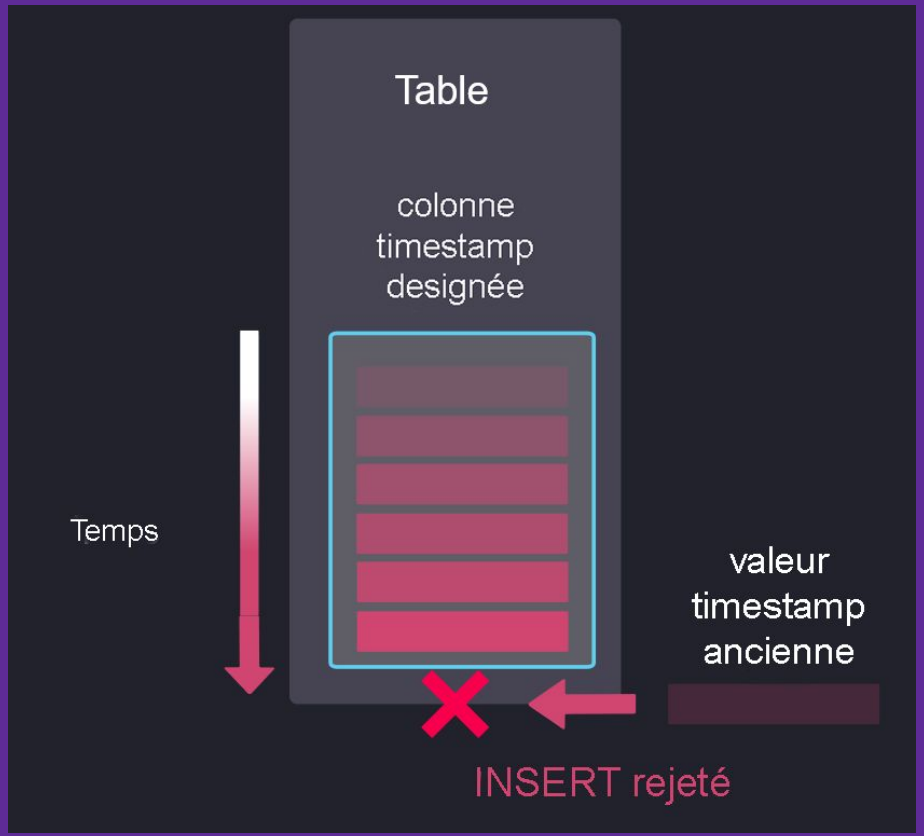

Partition

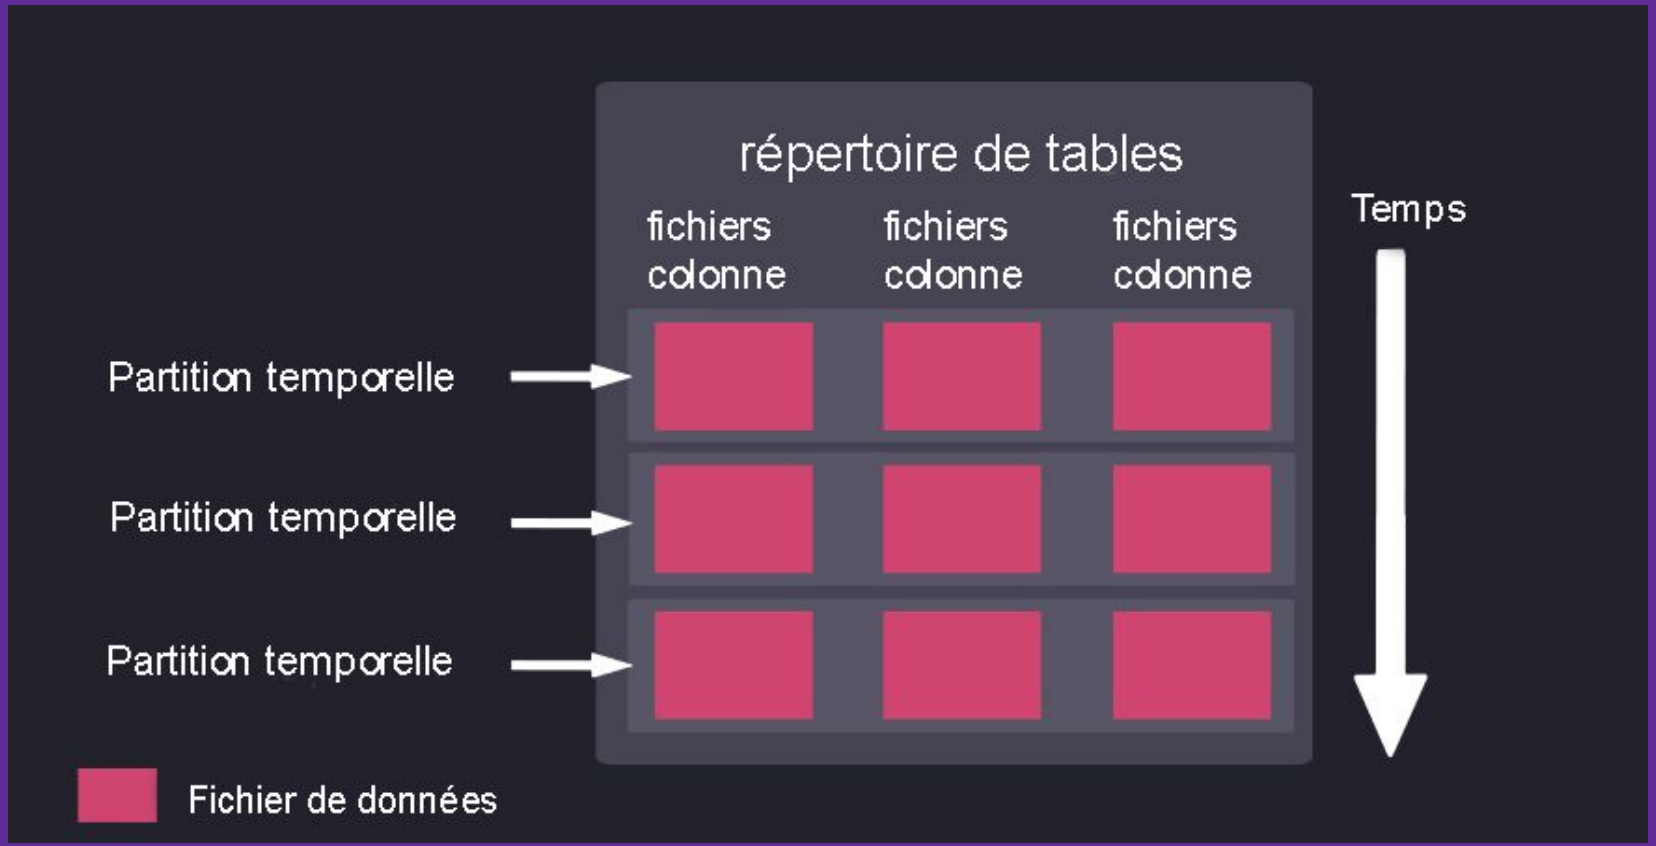
répertoire de tables
fichiers colonne
fichiers colonne
fichiers colonne
Temps
Partition temporelle
Partition temporelle
Partition temporelle
Fichier de données

Symboles et Index

Table

Row ID	Symbol	Value
1	A	1
2	A	0
3	B	1
4	A	1
5	C	0
6	B	1

Index

Symbol	Row IDs
A	1, 2, 4
B	3, 6
C	5

Symboles et Index

Table

Row ID	Symbol	Value
1	A	1
2	A	0
3	B	1
4	A	1
5	C	0
6	B	1

Index

Symbol	Row IDs
A	1, 2, 4
B	3, 6
C	5

Démonstration

Conclusion

Inconvénients :

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

Merci pour votre attention !

Questions ?