

# PostgreSQL Concurrency control



# Perspective

User 1



# Perspective



User 2



# Perspective

User 3



# Perspective





# Derk van Veen

## Database engineer

## Adyen

**adyen**

engineered  
for ambition

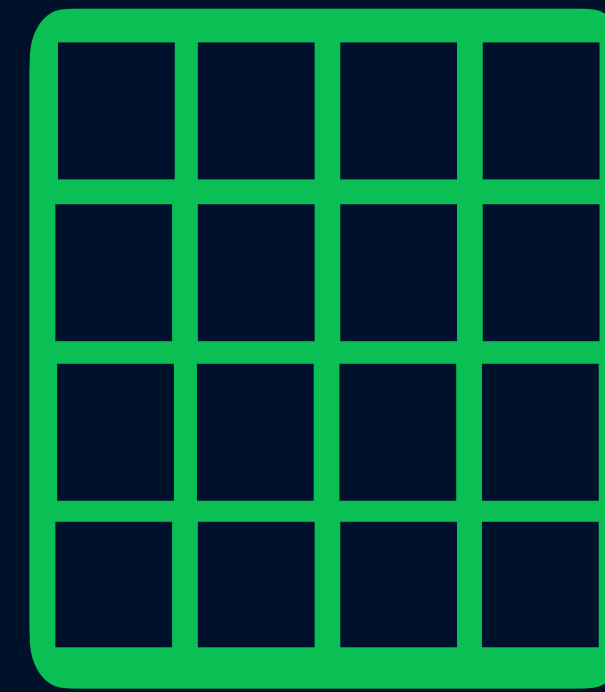
# Adyen ♥ Postgres



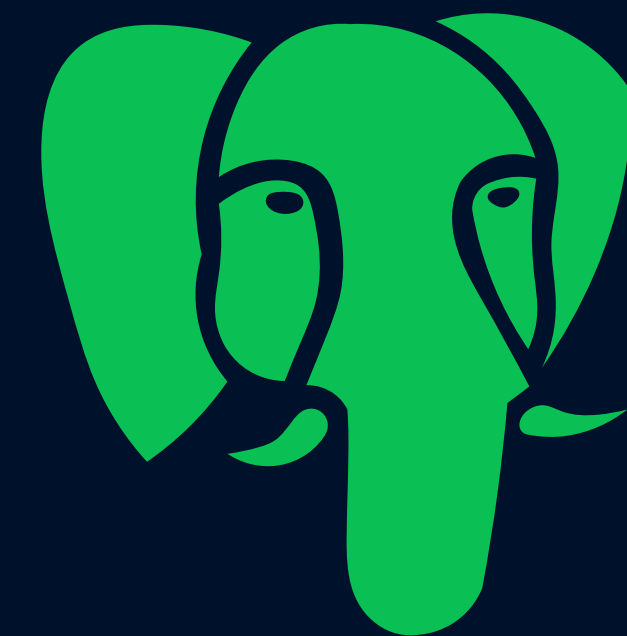
10 y



100s



1000



Many > 100TB

Possibly the biggest critical Postgres shop in the world!



# Agenda

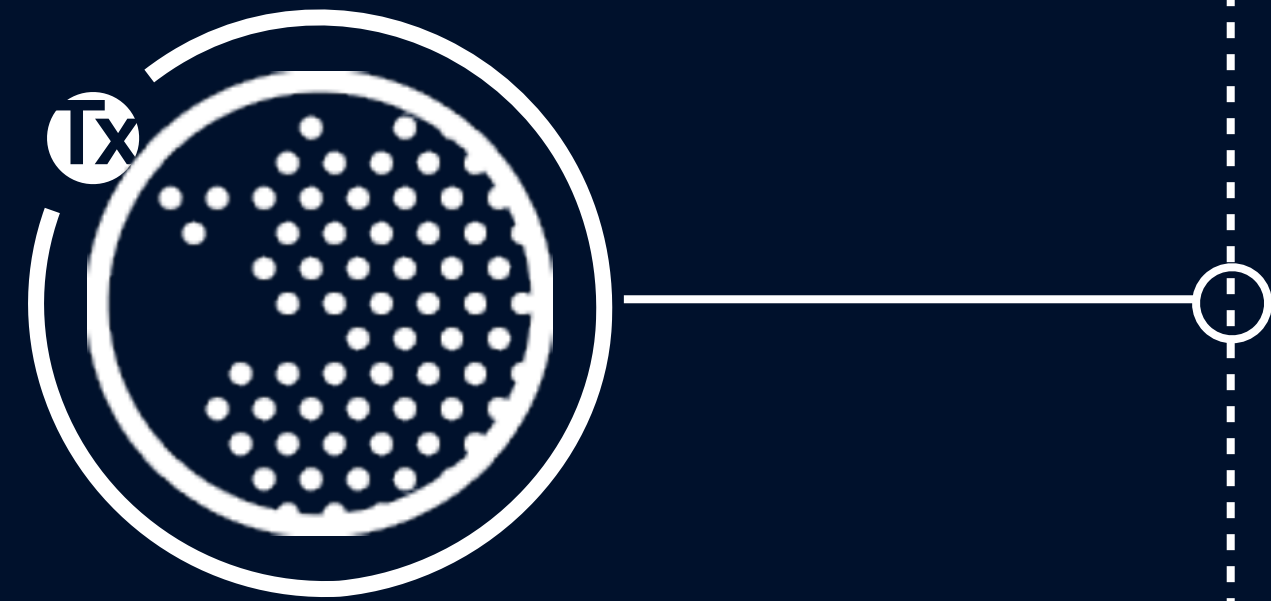




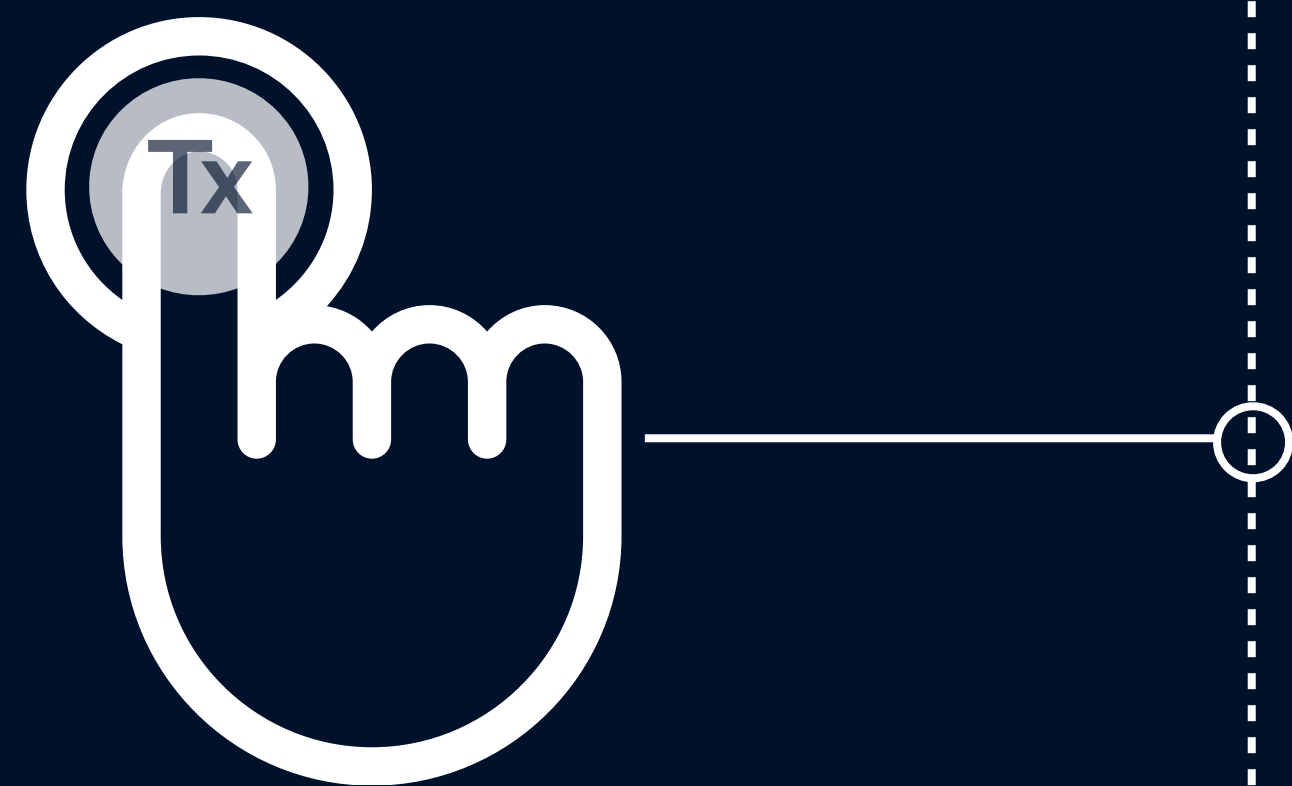
# Transaction definition



# Individual transactions



# Transaction space



# Using transactions



**Locks**



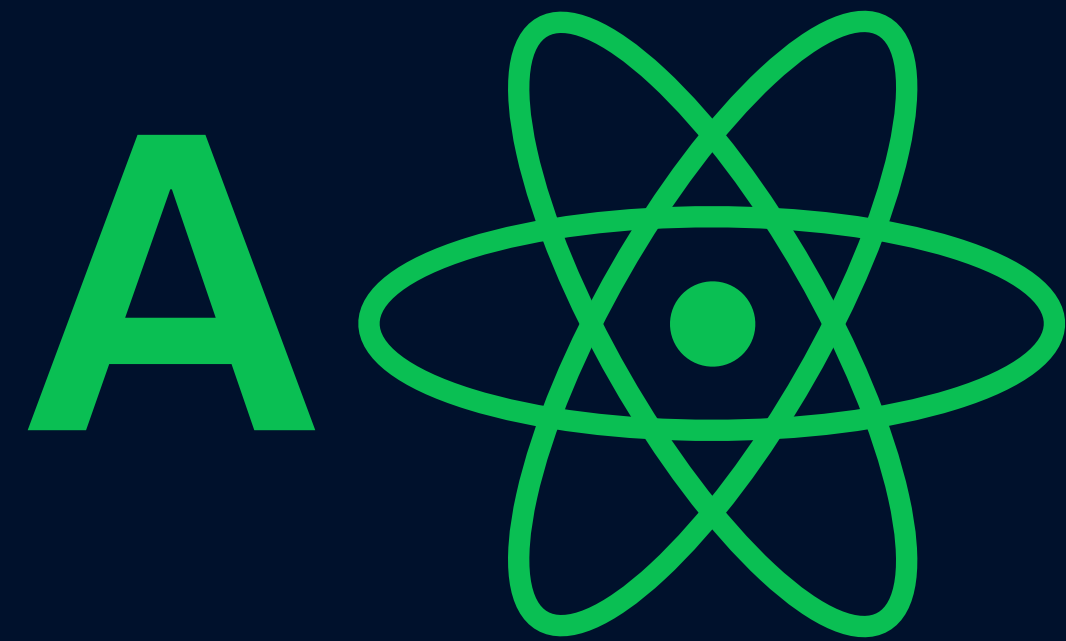
# Isolation levels



# Transaction definition



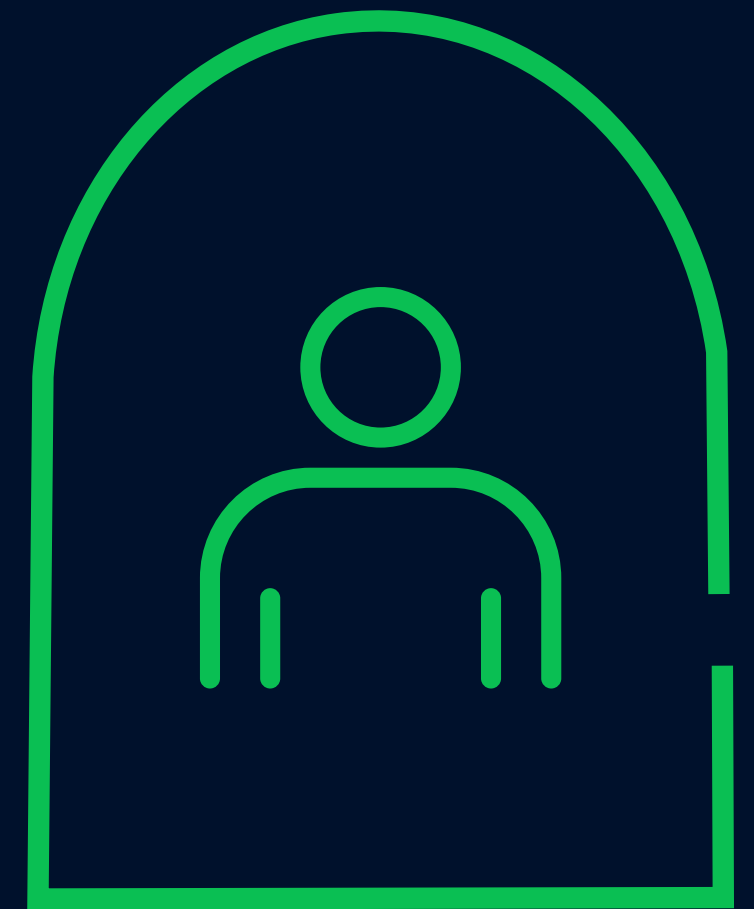
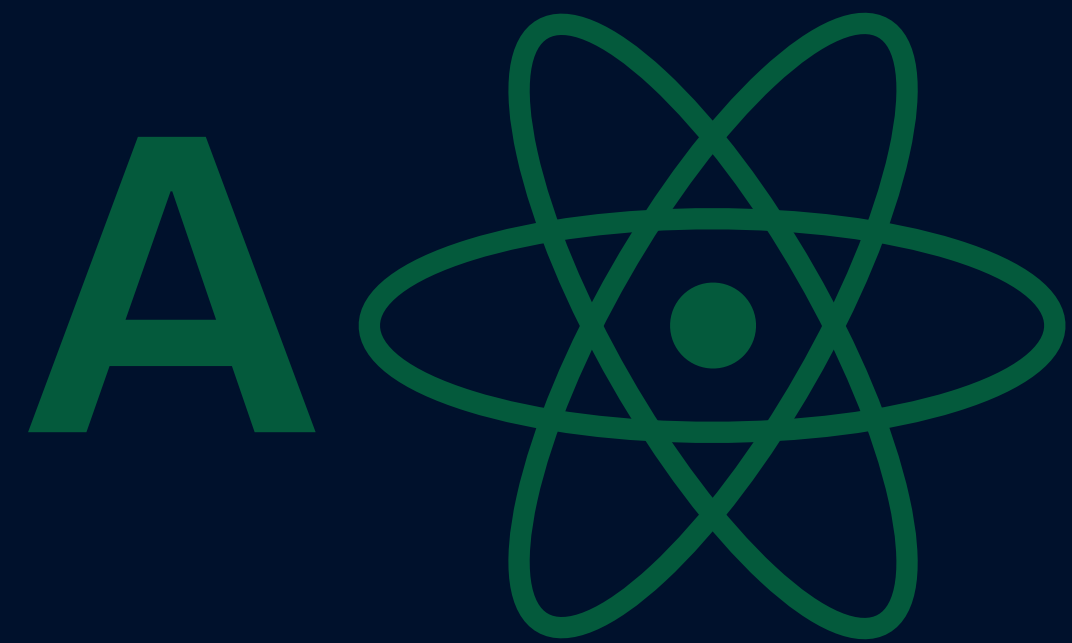
# Transaction definition



# Transaction definition



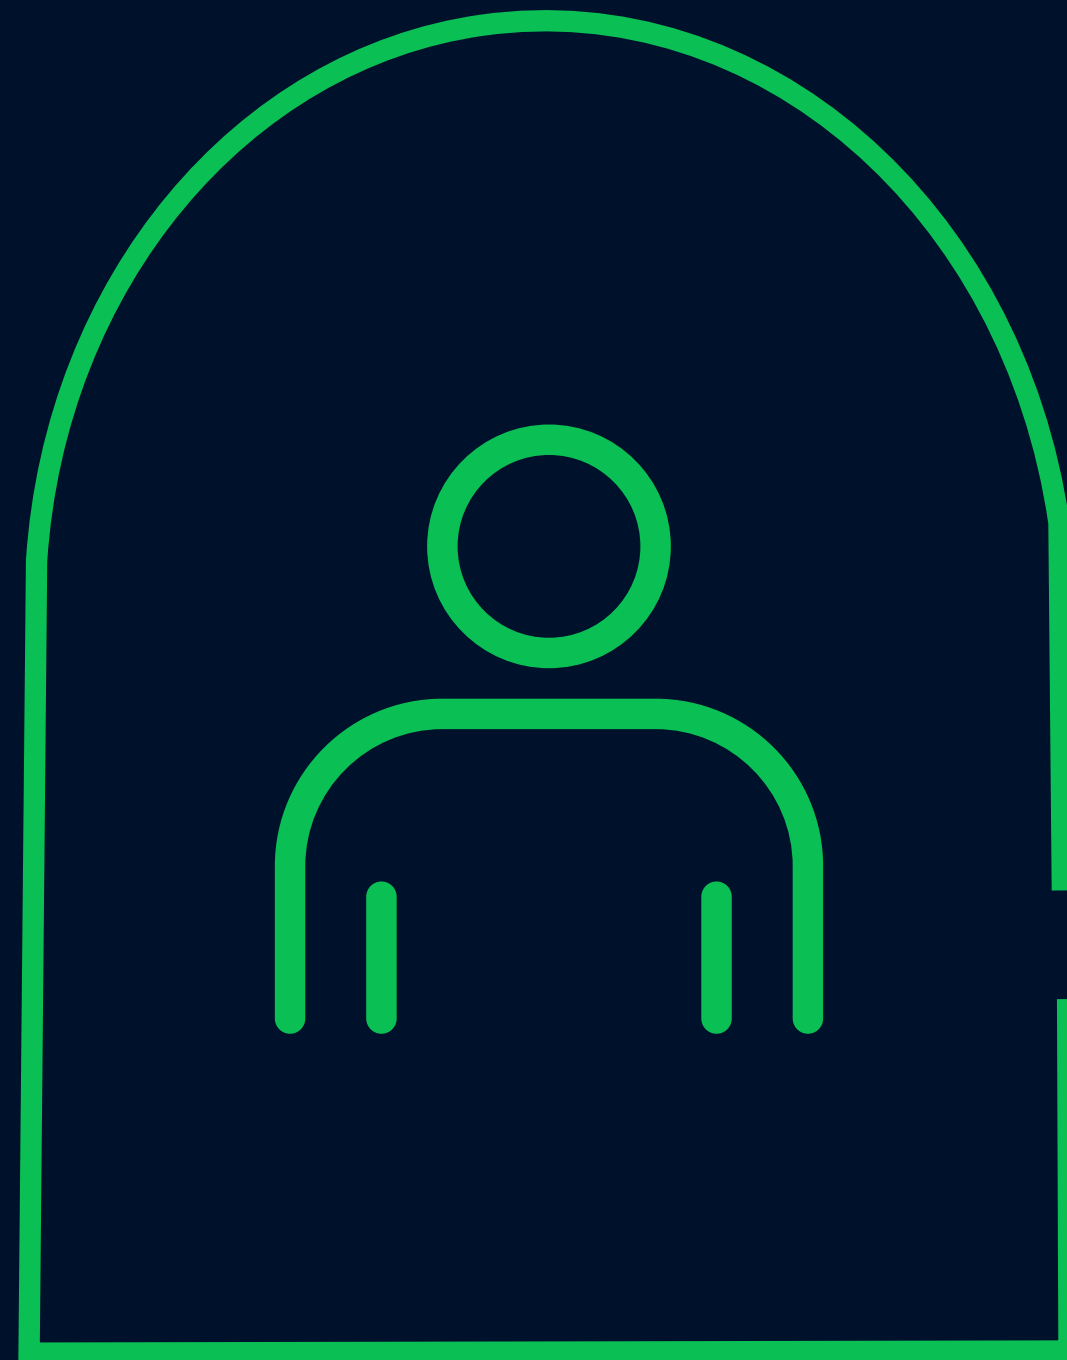
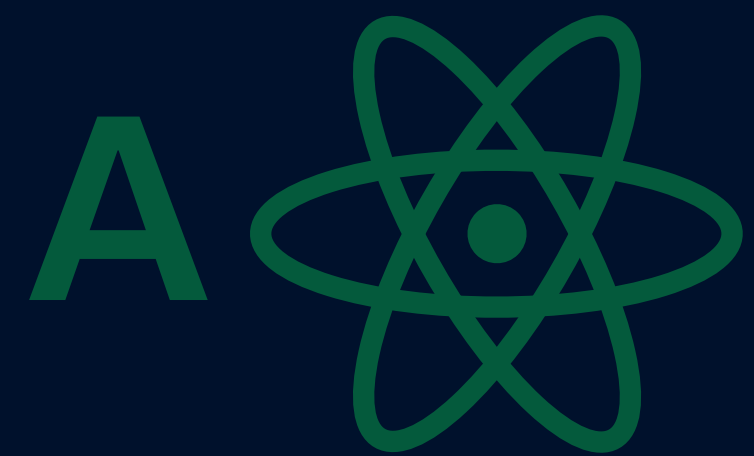
# Transaction definition



# Transaction definition



# Transaction definition





# Individual transactions

# Insert

```
Insert into rollercoaster  
values ('baron')
```



---

t=101

# Insert



t=101

SQL



# Insert



t=101

SQL

# Delete

```
delete from rollercoaster  
( 'baron' )
```



t=101

t=105

SQL

# Delete



t=101

t=105

SQL

SQL

SQL

# Update



Delete



Insert

# Update



t=103

SQL



# Update



t=103

SQL

# Overview



t=101

t=105

t=103



# Page level

Page header



$t_{\min}=101, t_{\max}=105$



$t_{\min}=10, t_{\max}=103$



$t_{\min}=103, t_{\max}=\dots$

Page footer



# Page level

Page header



$t_{\min}=101, t_{\max}=105$



$t_{\min}=10, t_{\max}=103$



$t_{\min}=103, t_{\max}=\dots$

$t=100$

select \* from  
rollercoaster

Page footer

# Page level

Page header



$t_{\min}=101, t_{\max}=105$



$t_{\min}=10, t_{\max}=103$



$t_{\min}=103, t_{\max}=\dots$

$t=102$

`select * from  
rollercoaster`

Page footer

# Page level

Page header



t\_min=101, t\_max=105



t\_min=10, t\_max=103



t\_min=103, t\_max=

t=106

select \* from  
rollercoaster

Page footer



# Transaction space



# Transaction space



oldest not committed



first available



active list

# Transaction snapshot

First not committed transaction

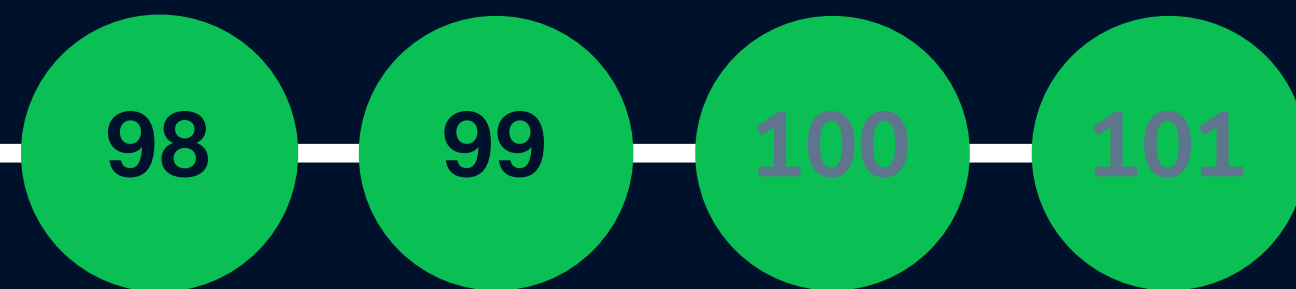
100:110:104,107 List of active transactions

First available transaction id



# Transaction snapshot

100:100:

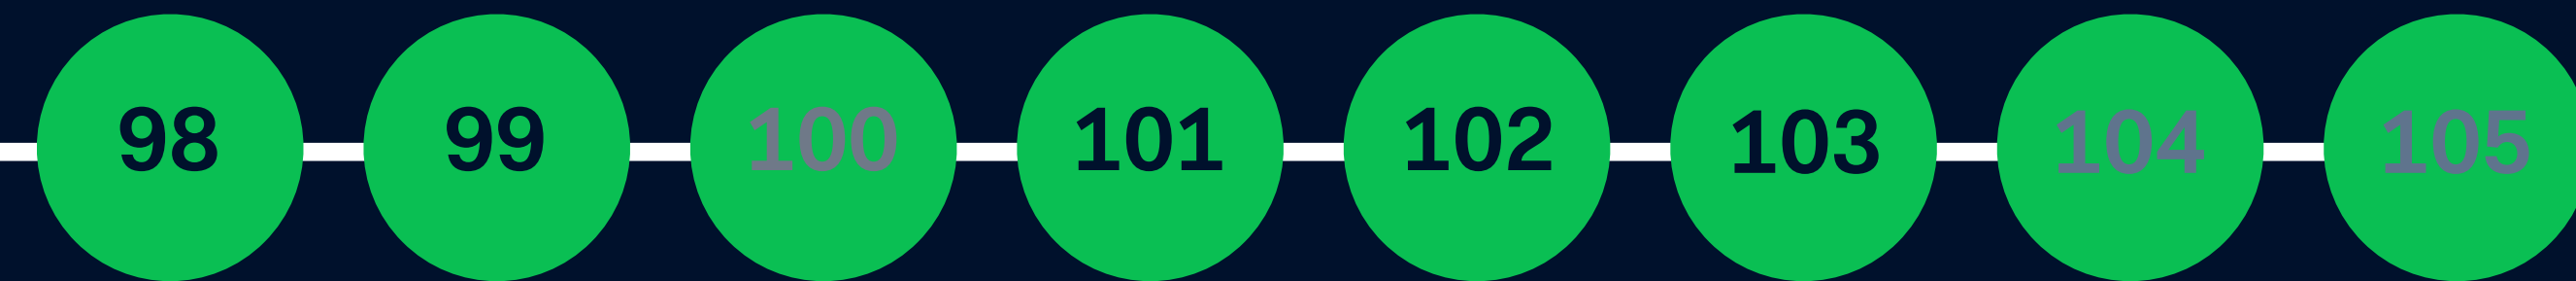


# Transaction snapshot

100:104:

Past

Future



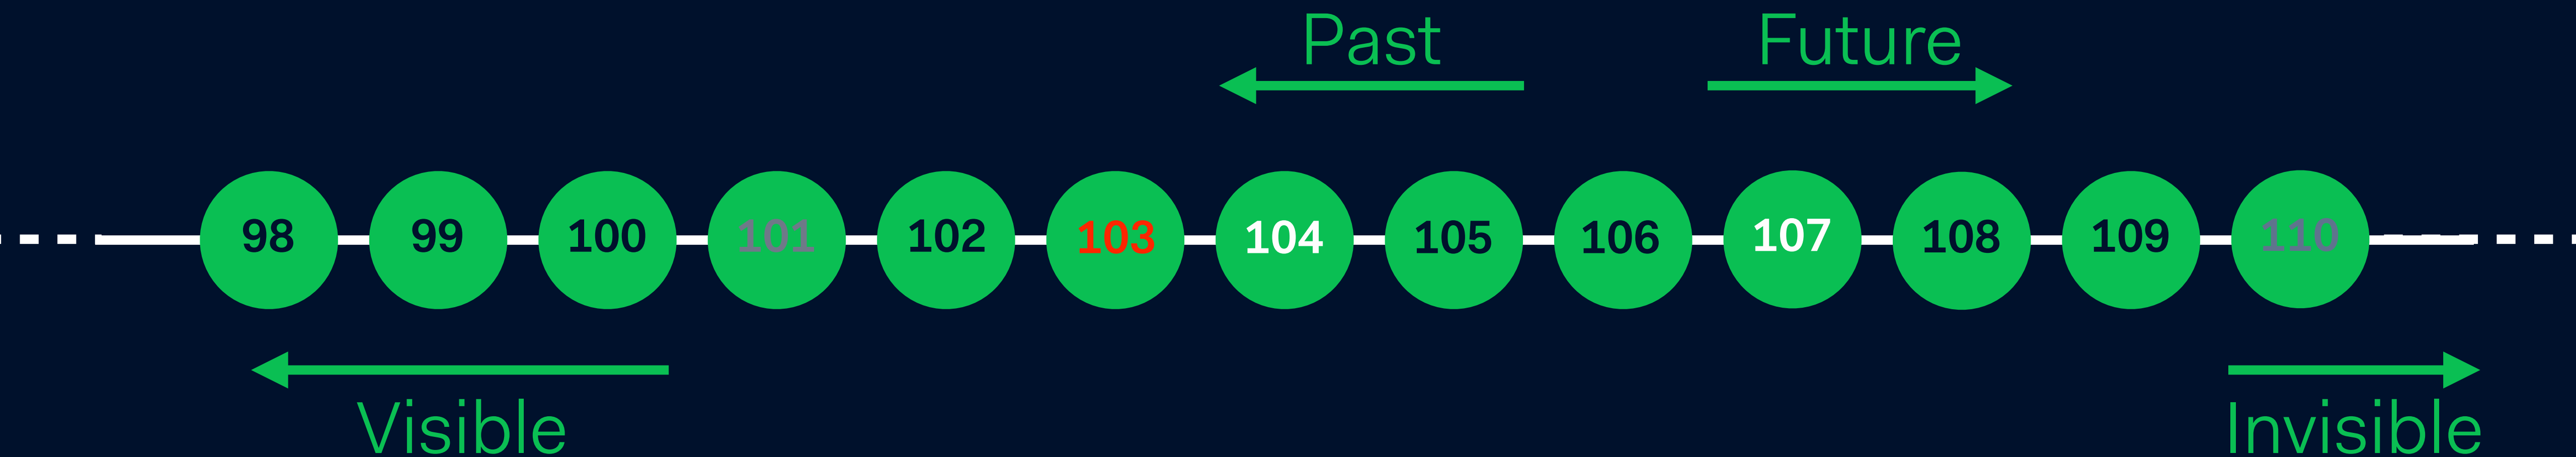
Visible

Invisible

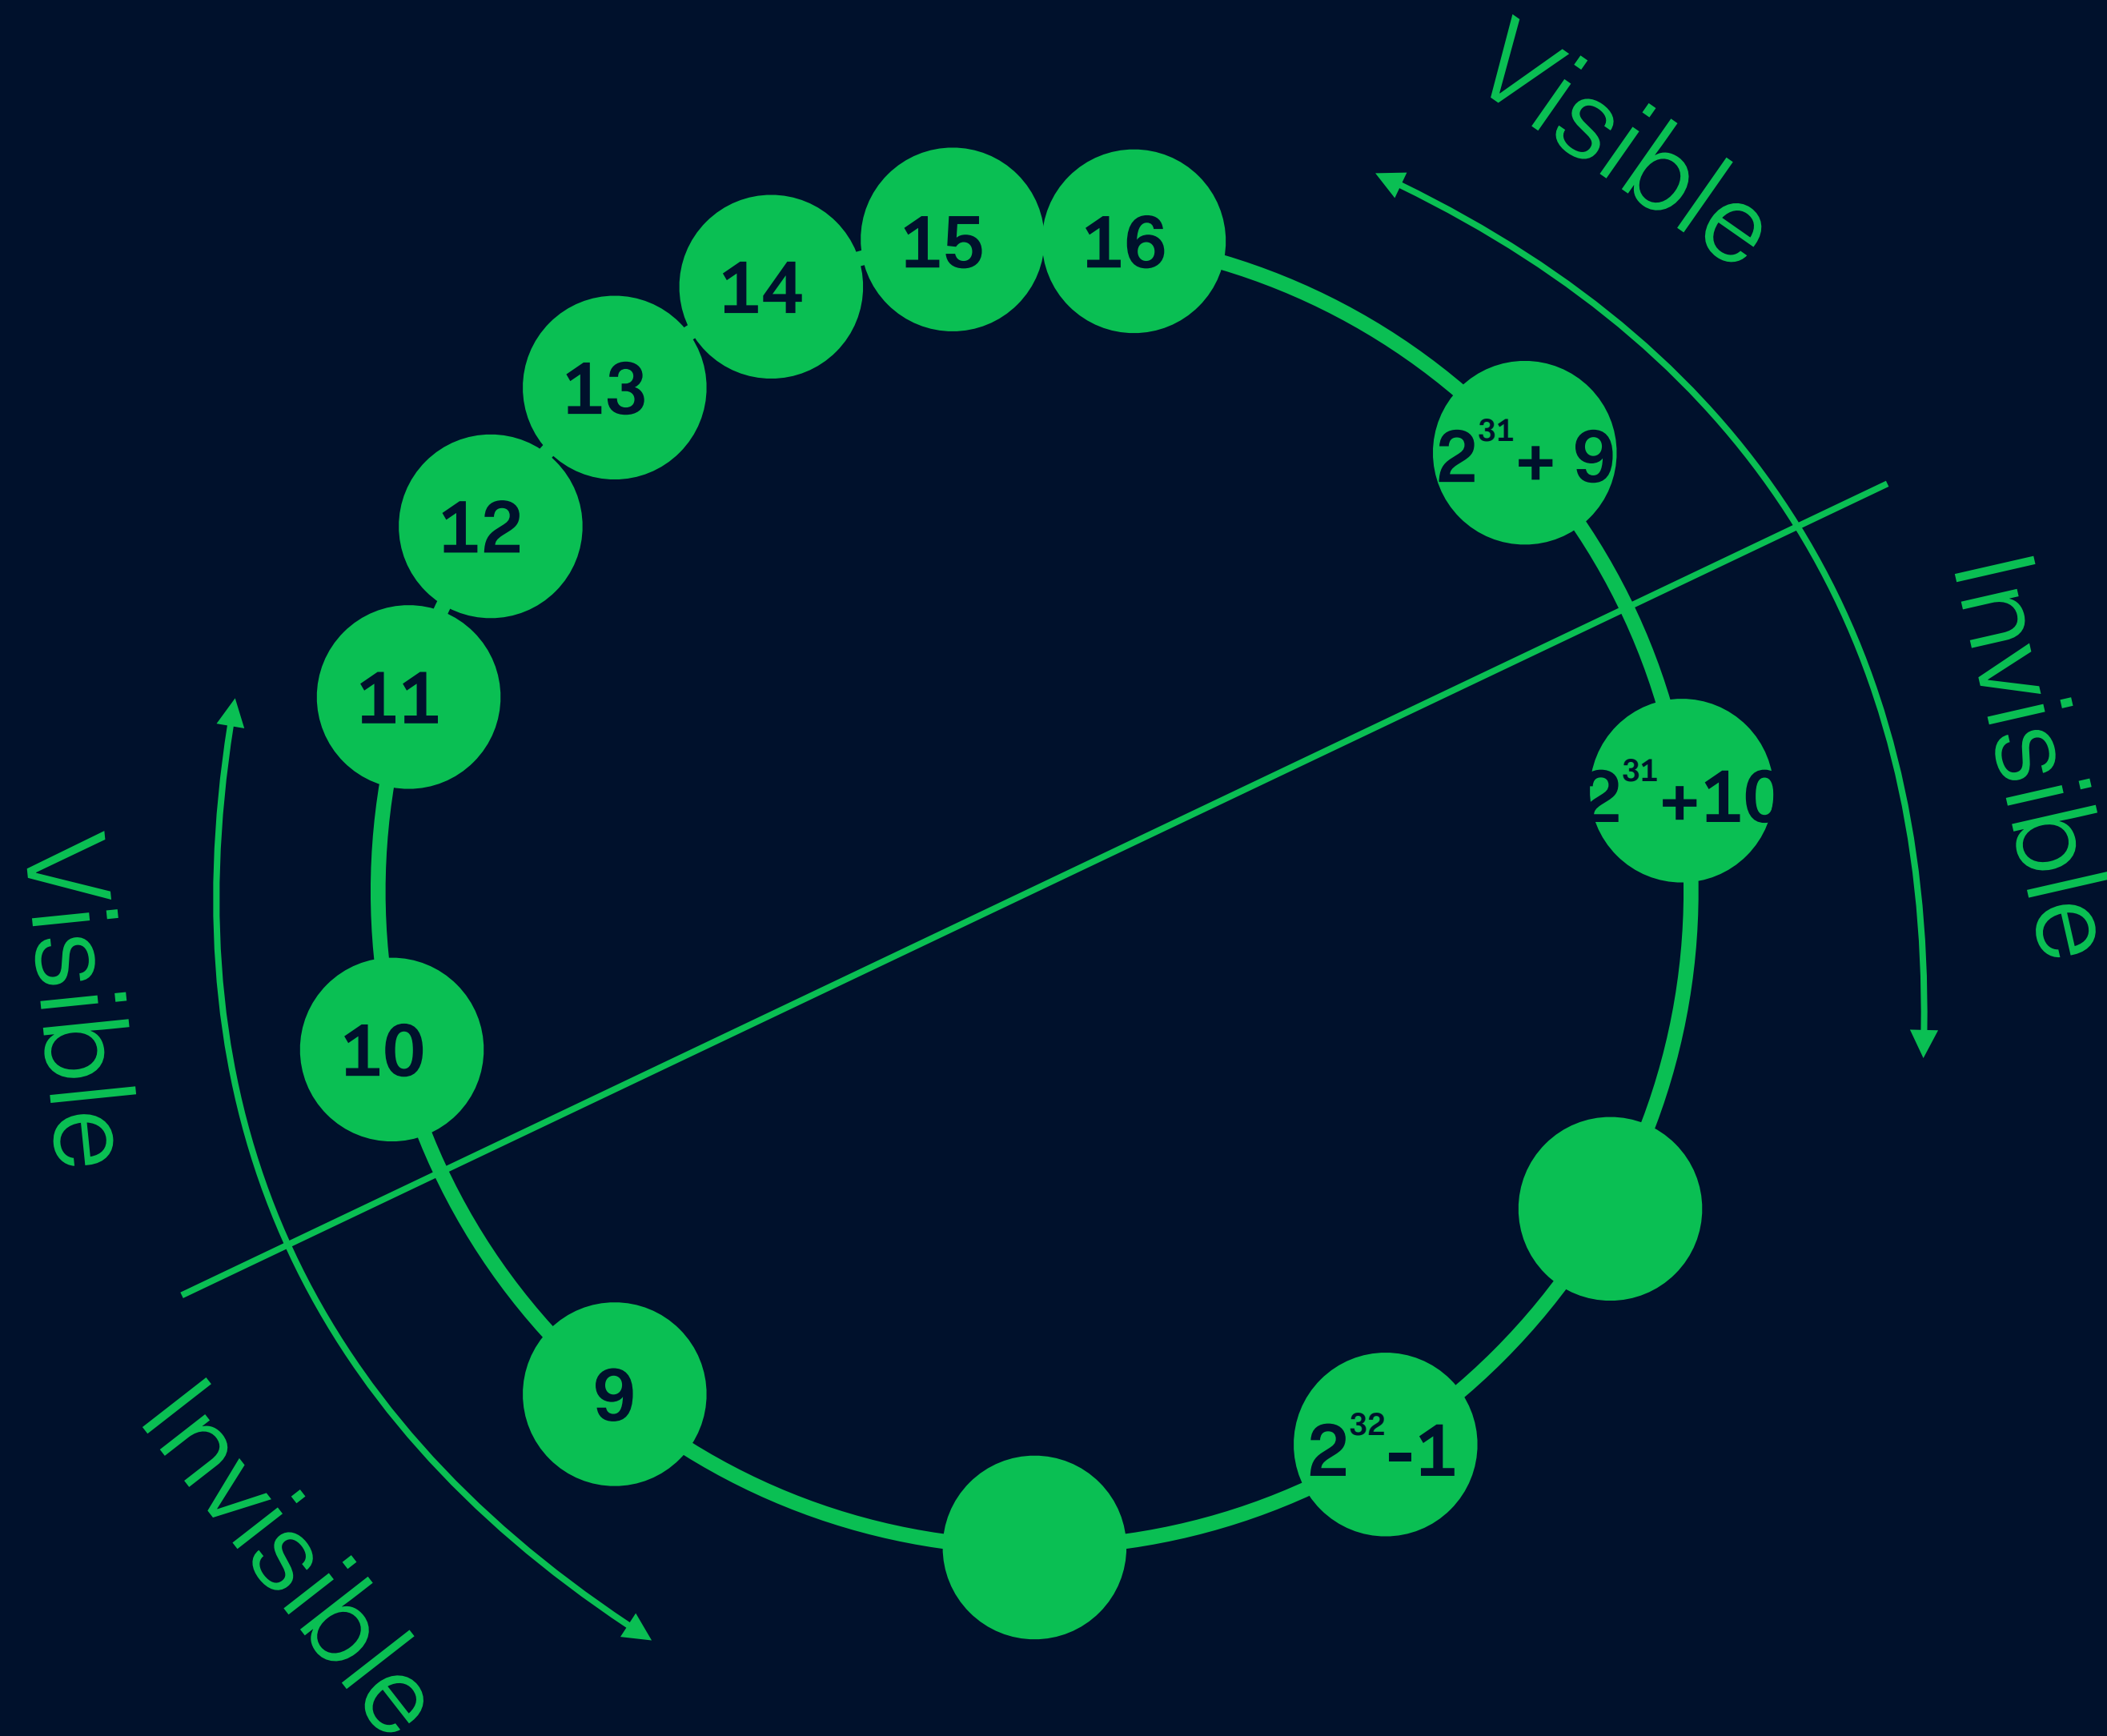


# Transaction snapshot

101:110:104,107



# Transaction snapshot



# Page level

Page header



$t_{\min}=101, t_{\max}=105$



$t_{\min}=10, t_{\max}=102$



$t_{\min}=102, t_{\max}=\dots$

Page footer

# Page level

Page header



$x_{\min}=101, x_{\max}=105$



$x_{\min}=10, x_{\max}=102$



$x_{\min}=102, x_{\max}=\dots$

Page footer

# Page level

Page header



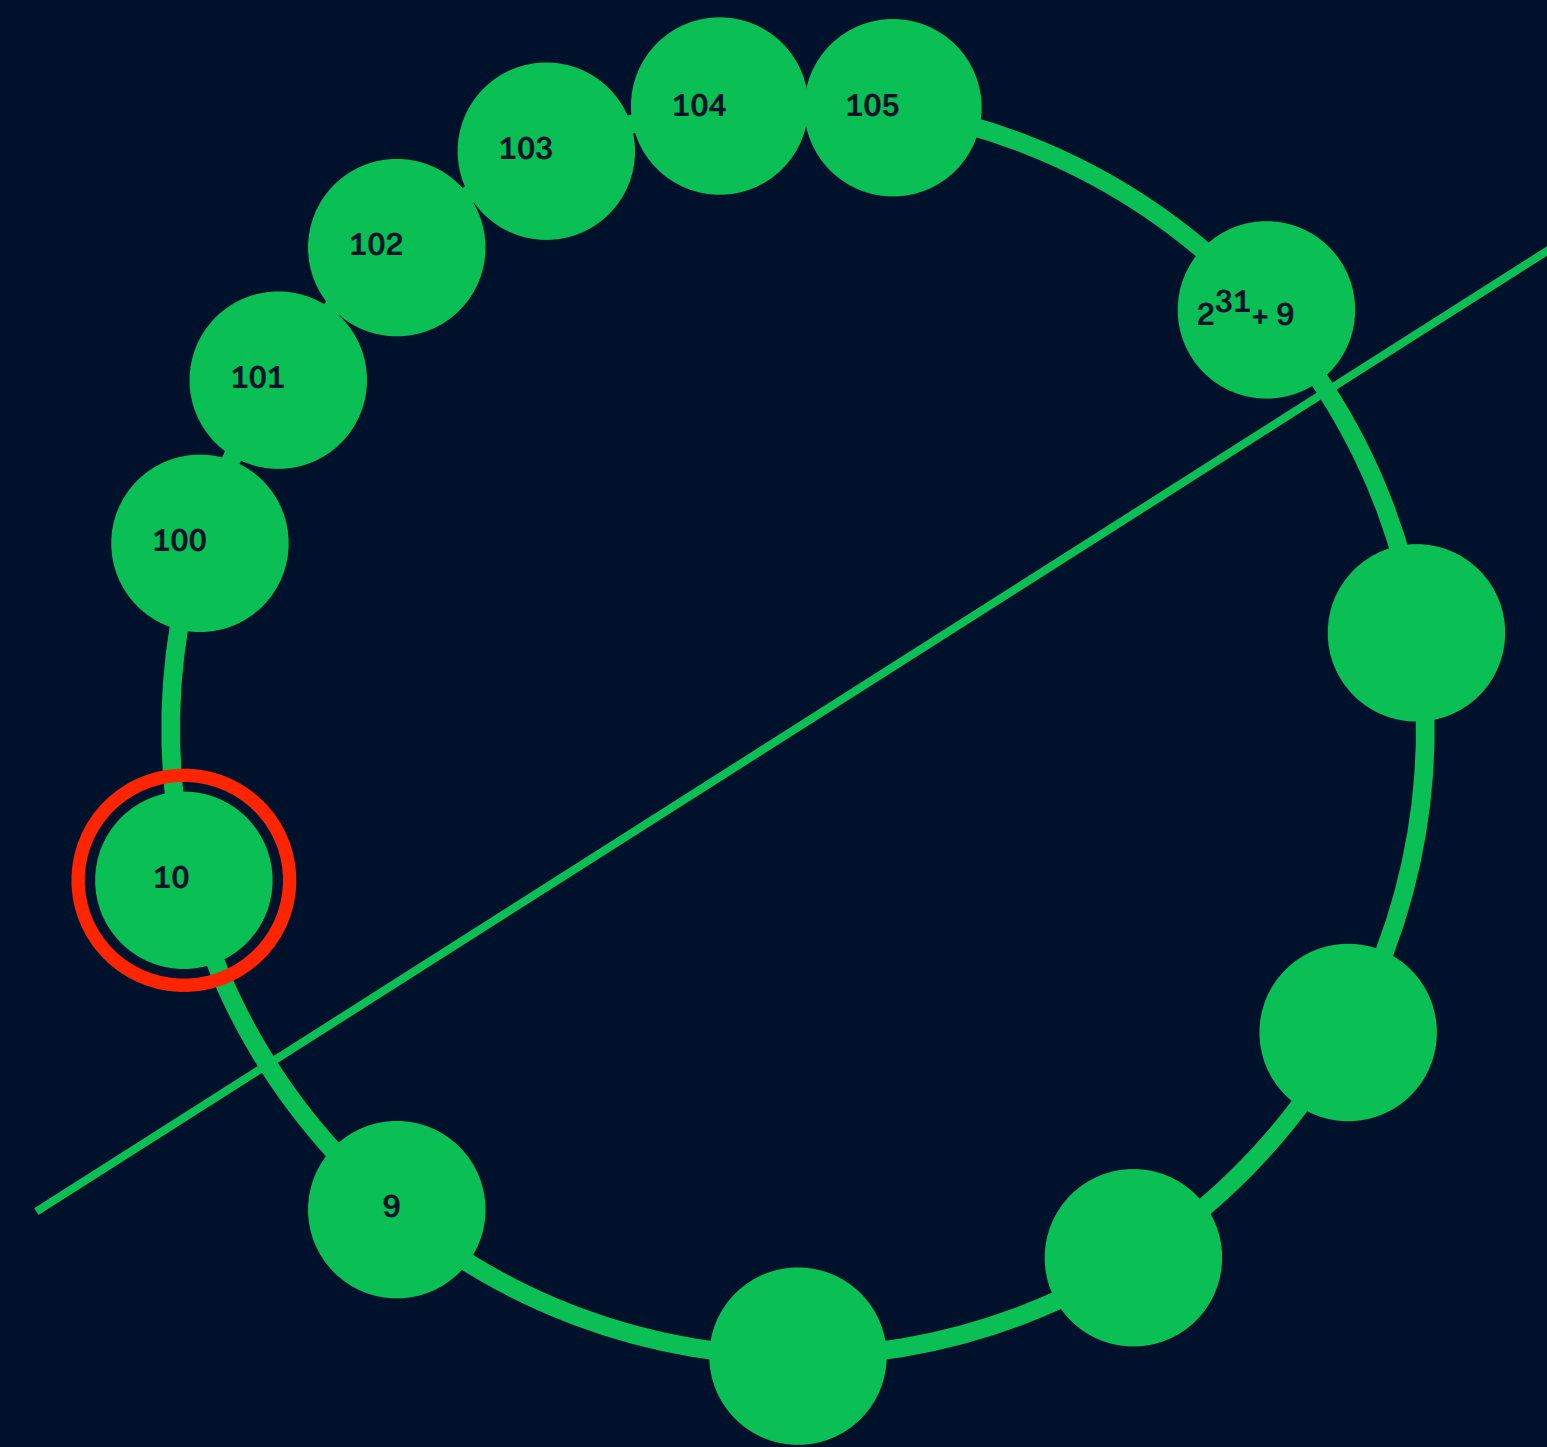
$x_{\min}=101, x_{\max}=105$



$x_{\min}=10, x_{\max}=102$



$x_{\min}=102, x_{\max}=\dots$



Page footer

# Page level

Page header



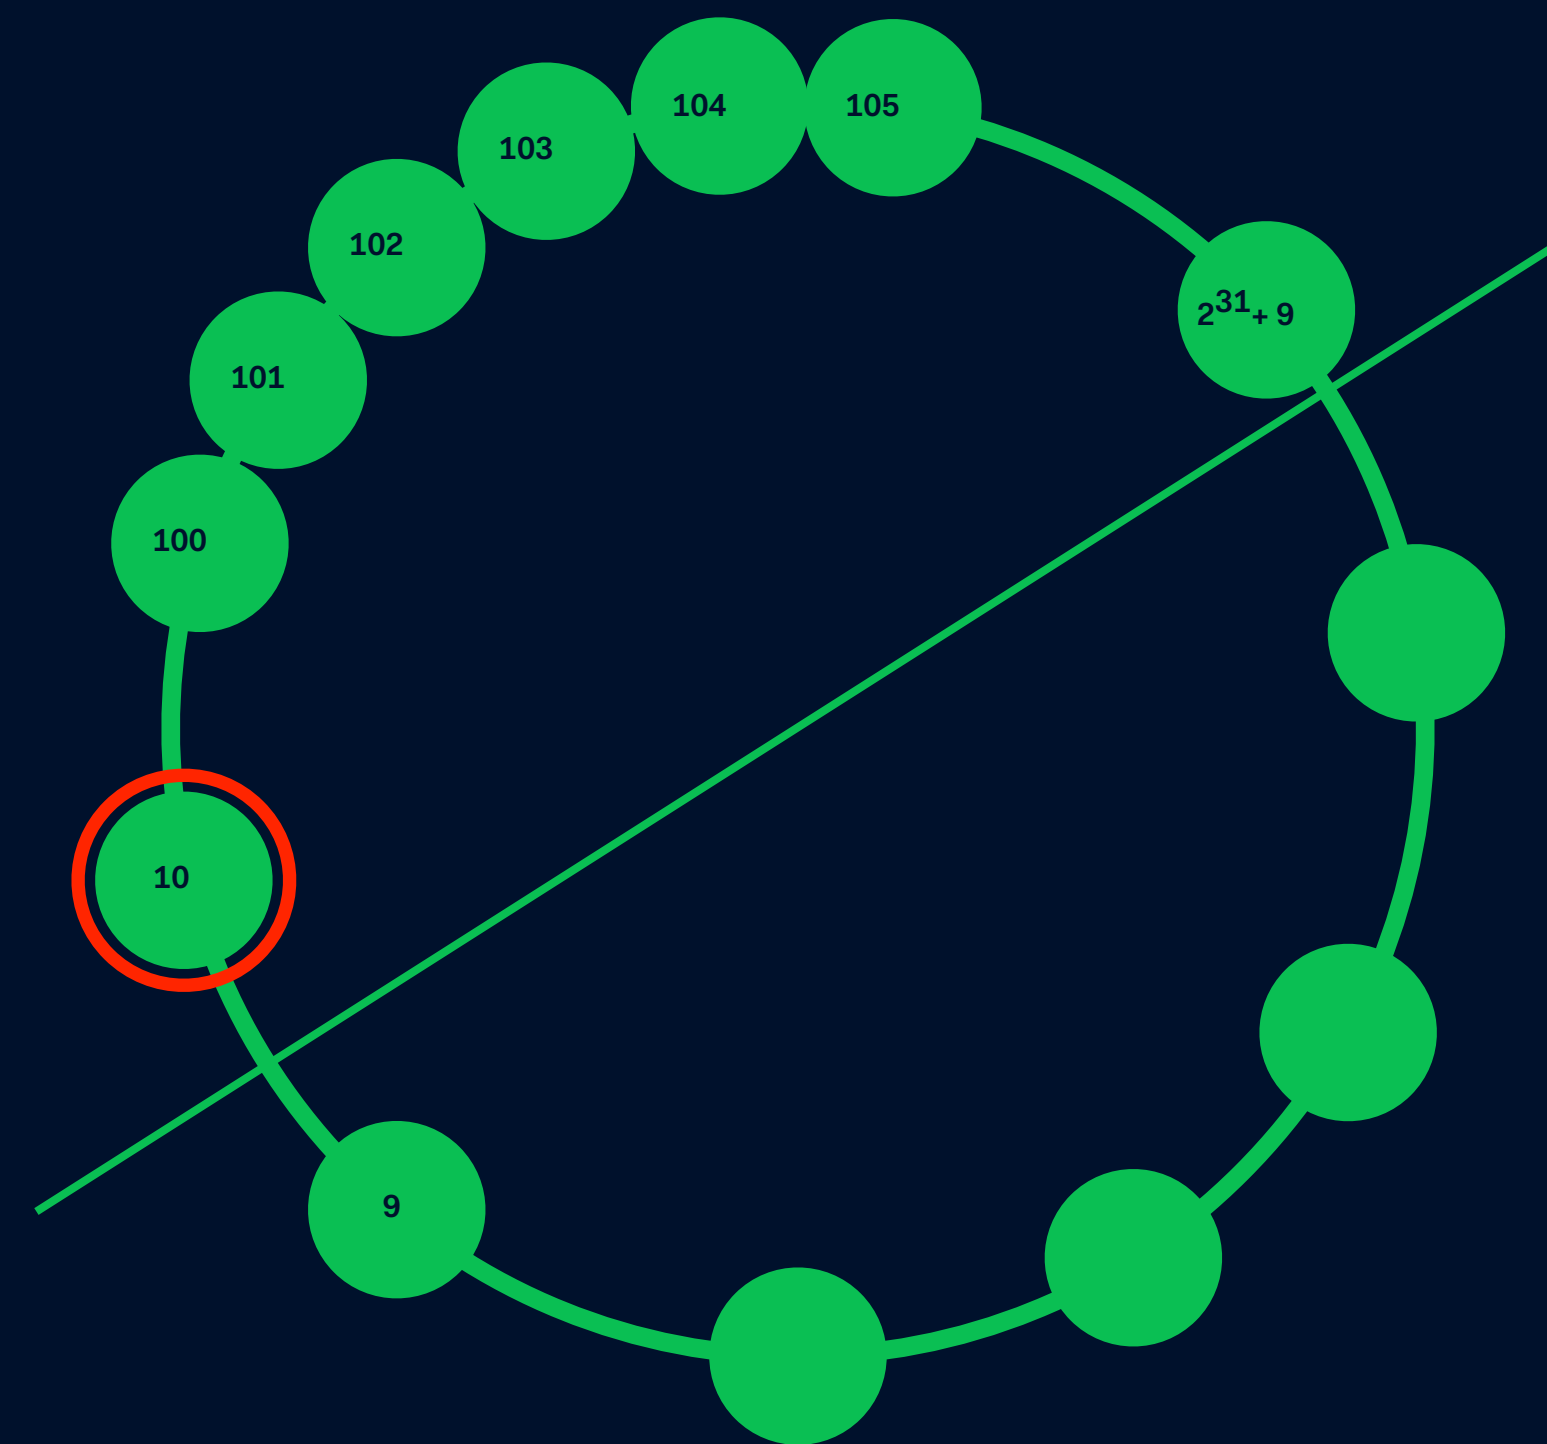
$x_{\min}=101, x_{\max}=105$



$x_{\min}=10, x_{\max}=102$



$x_{\min}=102, x_{\max}=\dots$



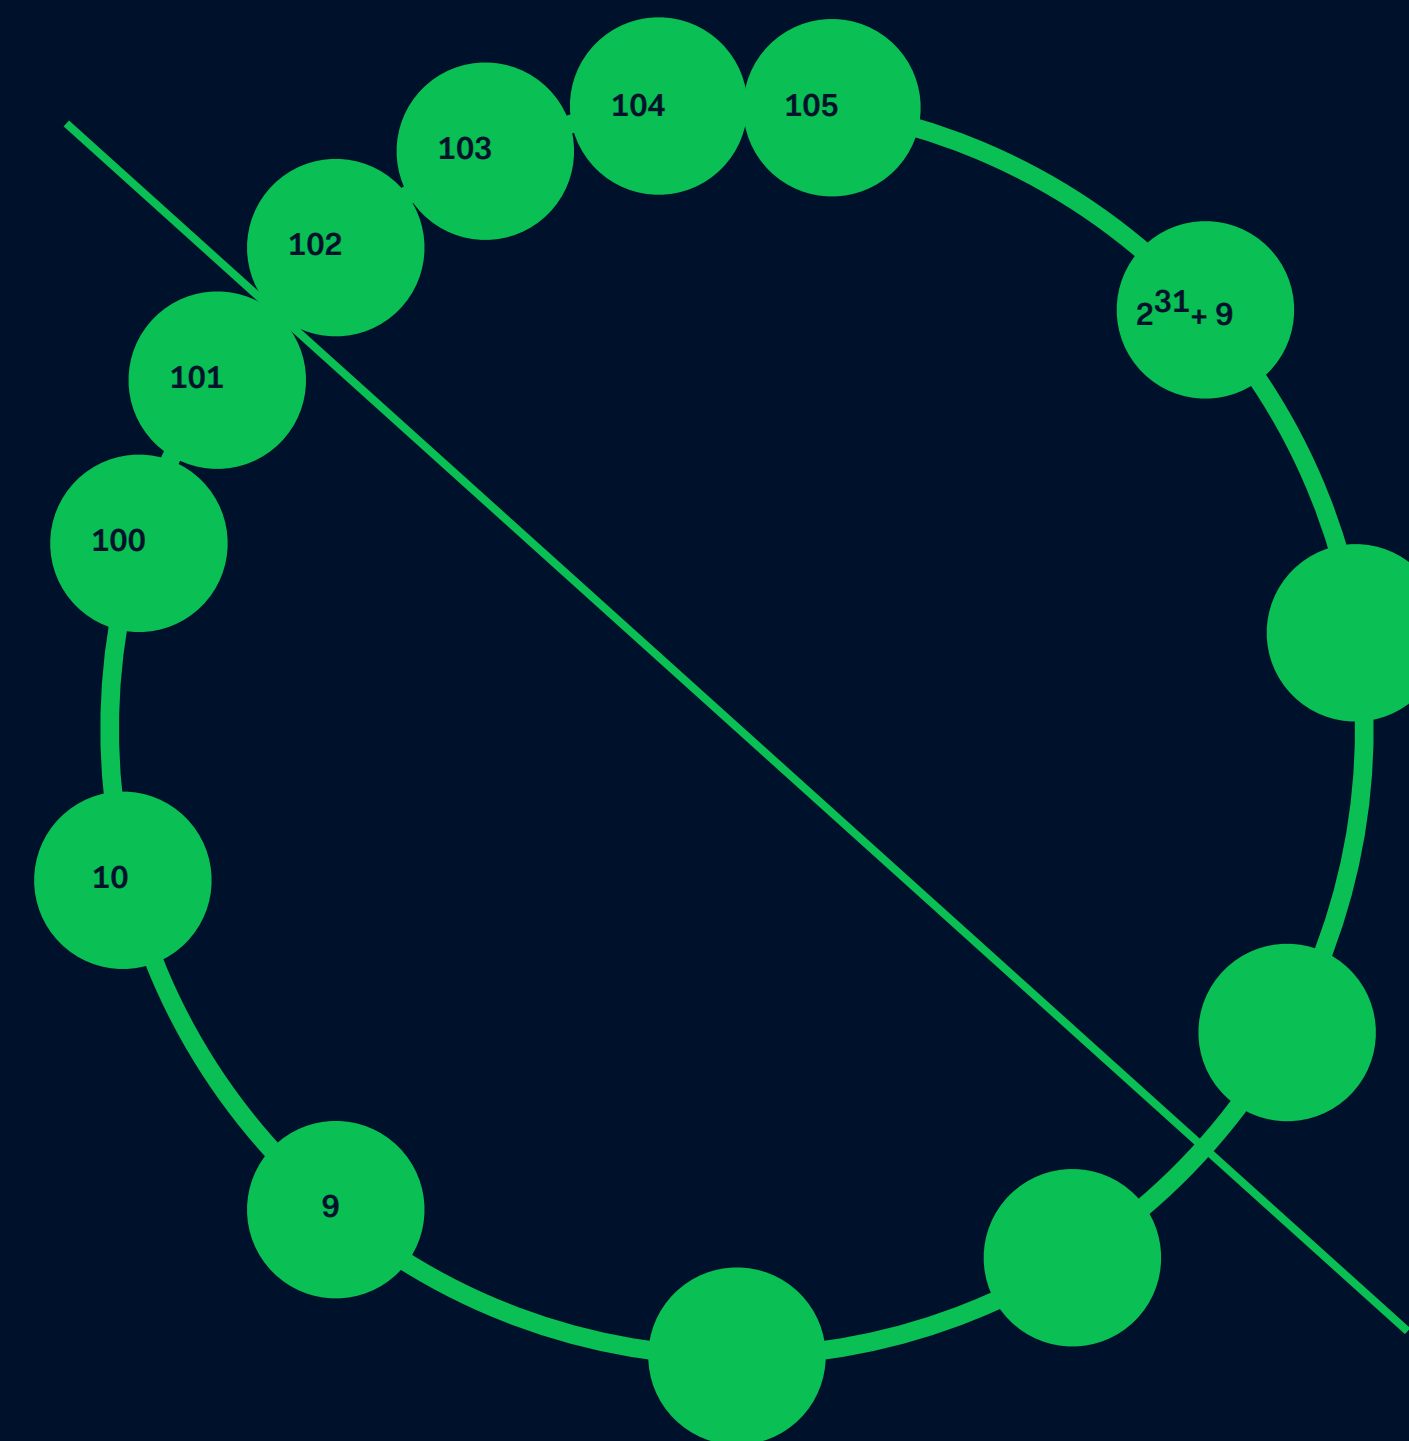
Page footer

# Page level

Page header



$x_{\min}=102, x_{\max}=\dots$



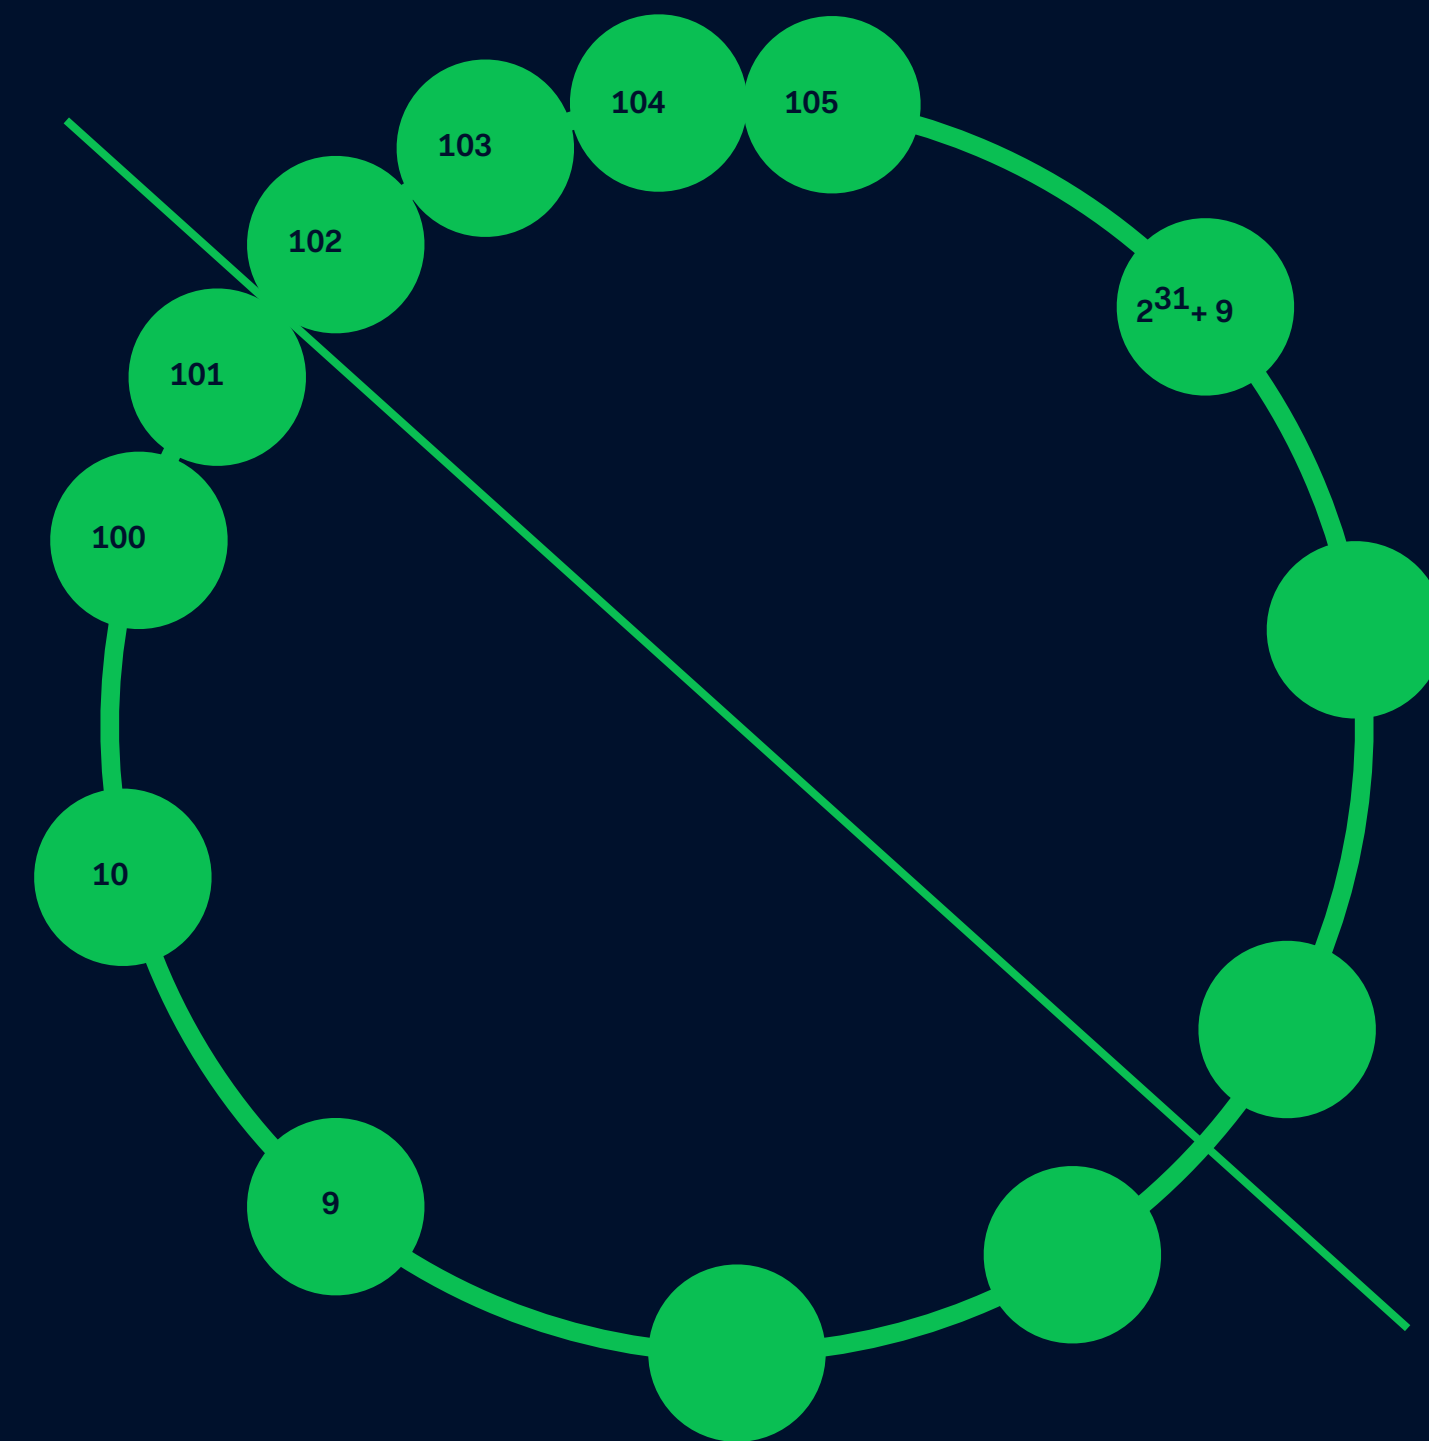
Page footer

# Page level

Page header



$x_{\min}=102, x_{\max}=\text{frozen}$



Page footer

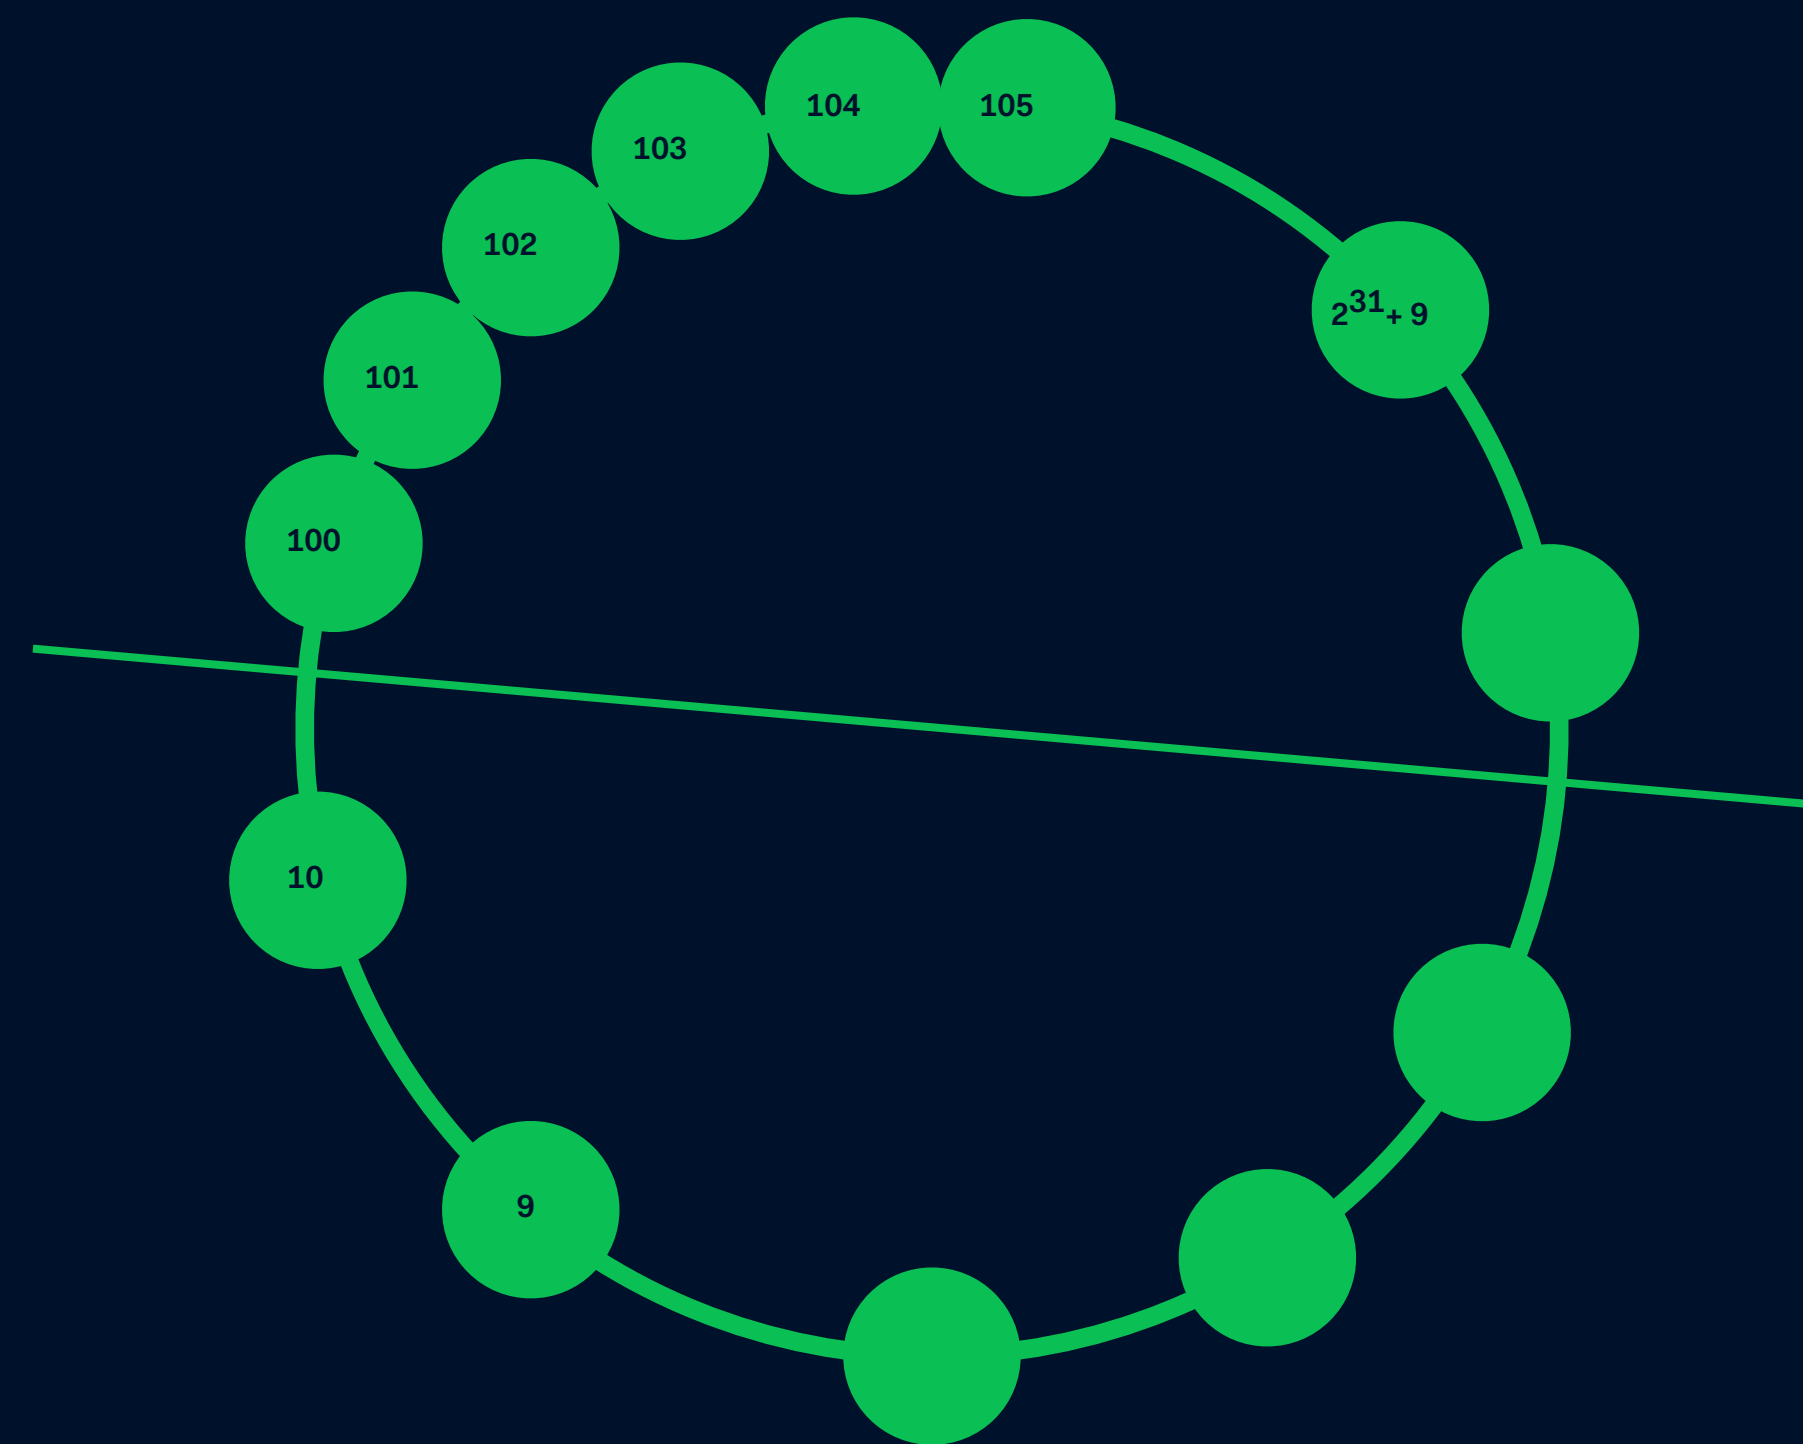


# Page level

Page header



$x_{\min}=102, x_{\max}=\text{frozen}$



Page footer



# Using transactions



# Using transactions

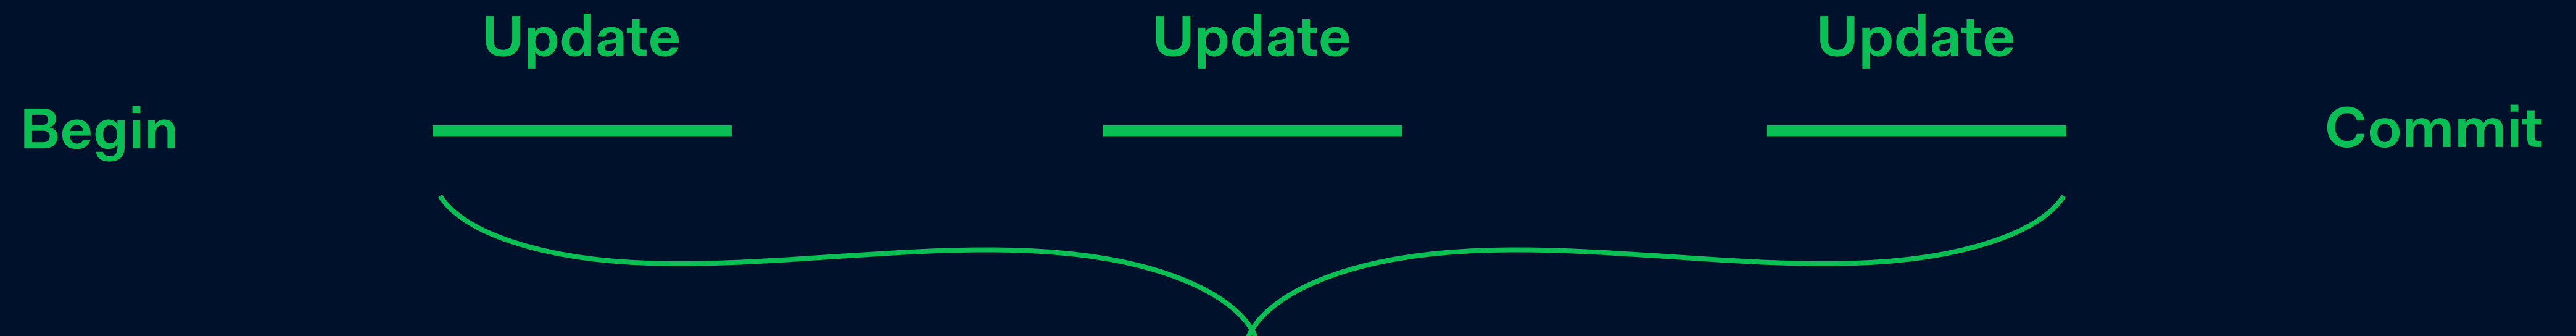
begin

# Using transactions

```
CREATE OR REPLACE PROCEDURE test_tx(v_commit boolean)
FOR i IN 0..1000000 LOOP
  PERFORM * FROM t1;
  IF (v_commit) THEN
    commit;
  END IF;
END LOOP;
```

Call test_tx(false)	<b>00:03:068</b>
Call test_tx(true)	<b>00:05:455</b>

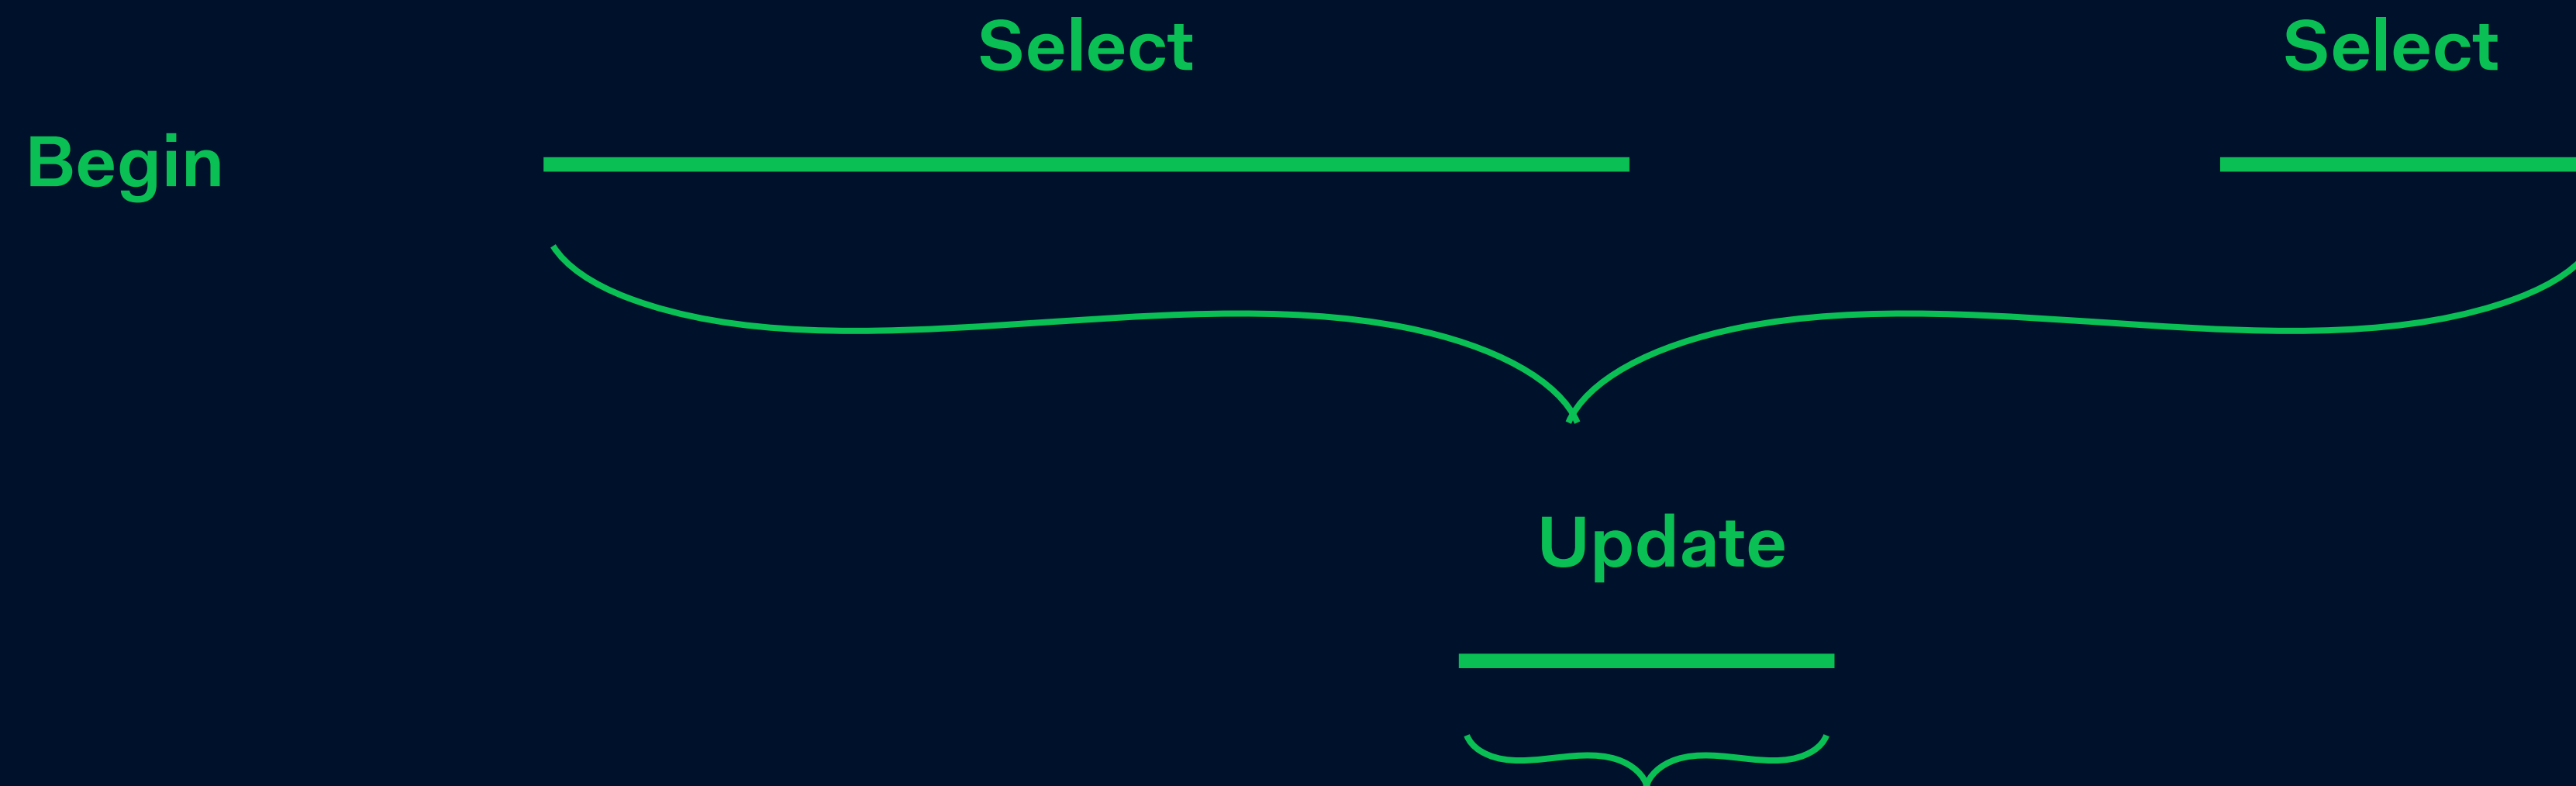
# Using transactions



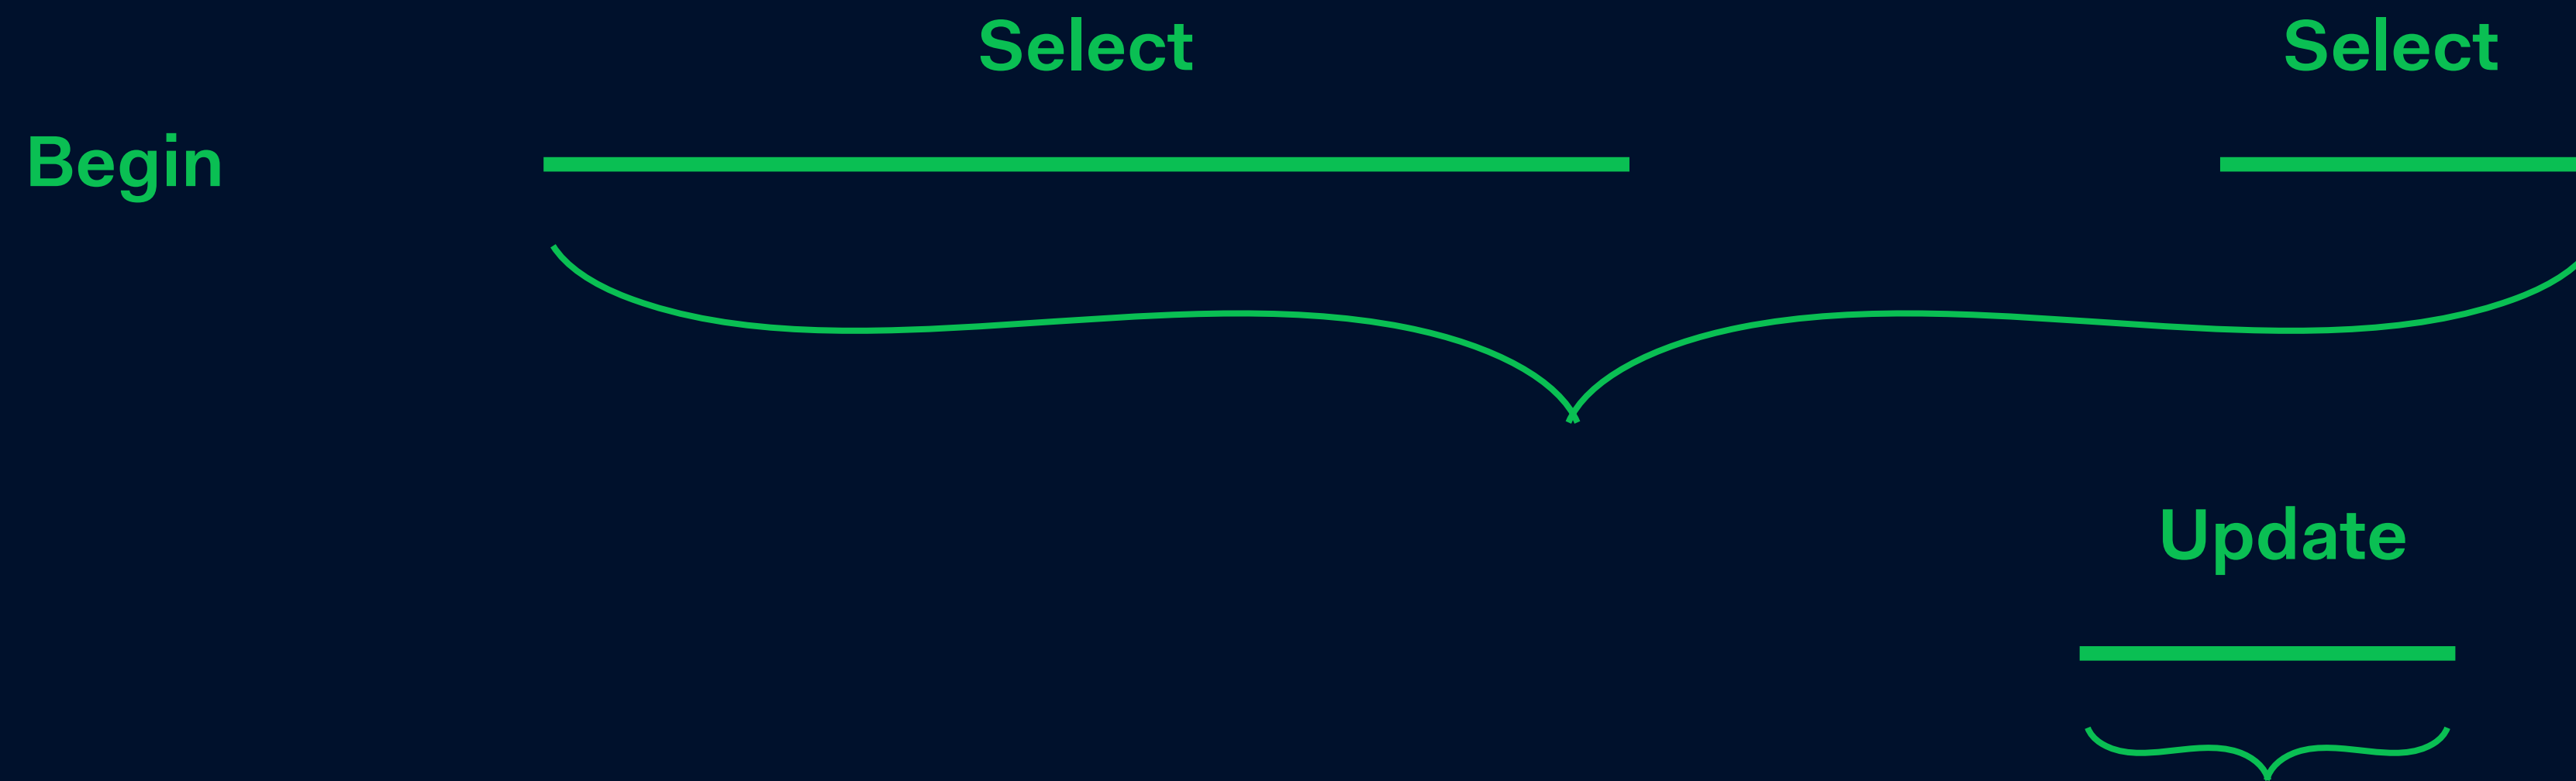
# Using transactions



# Using transactions



# Using transactions

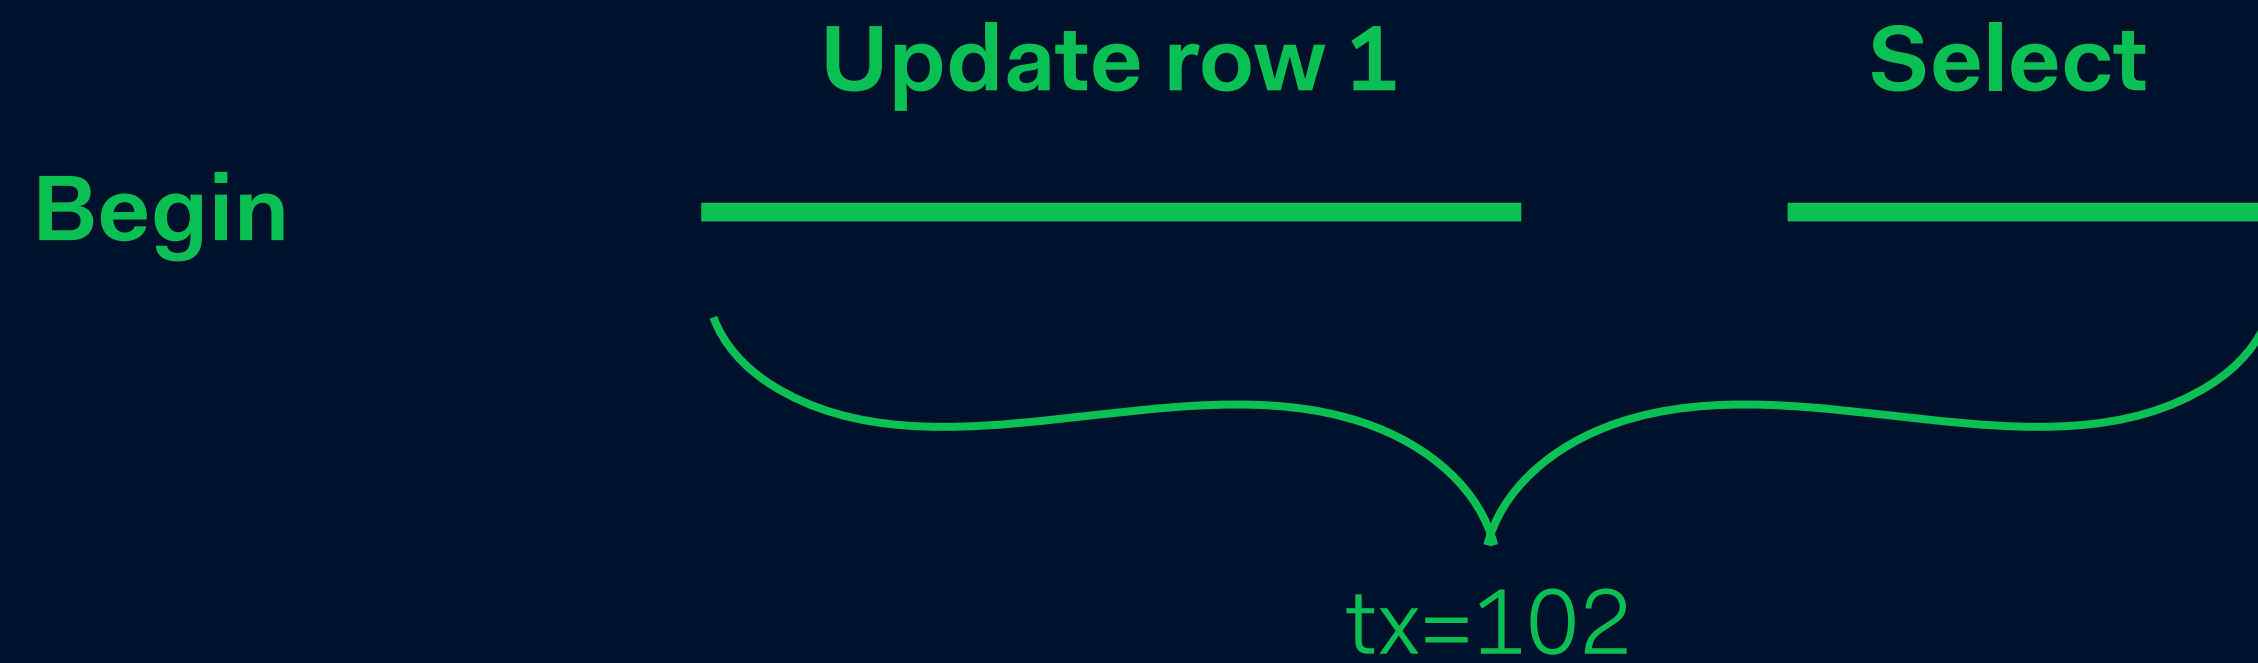




# Using transactions



# Using transactions



# Using transactions



# Page level

Page header



$x_{\min}=10, x_{\max}=102$



$x_{\min}=102, x_{\max}=102$



$x_{\min}=102, x_{\max}=\text{=}$

Page footer



**Locking**



# Locking

tx=100



Update

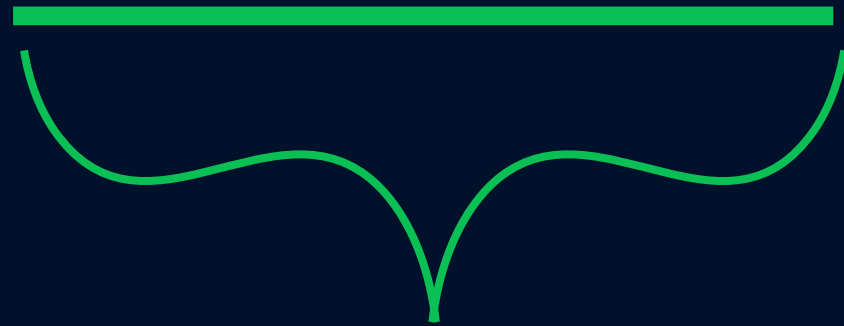
tx=102



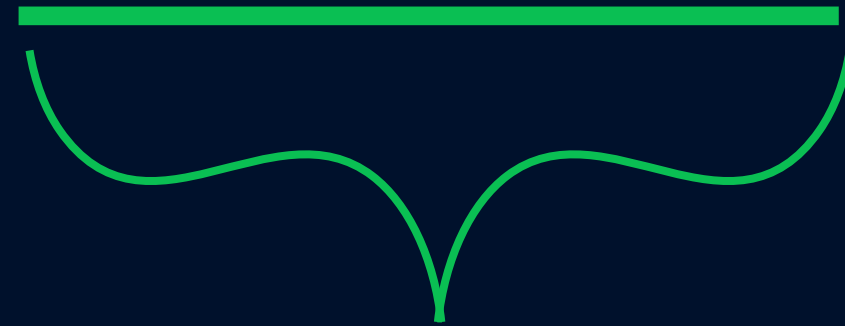
# Locking



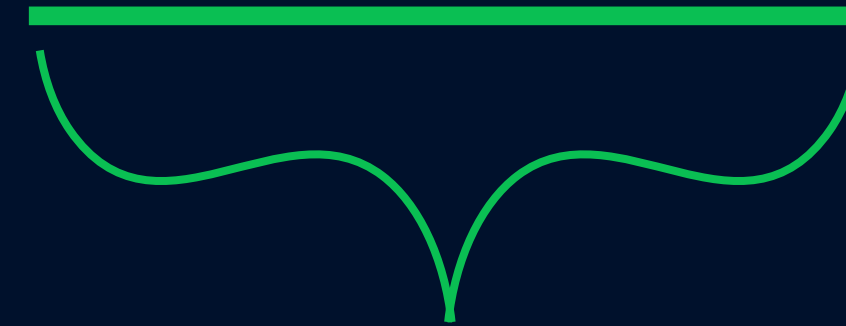
Select



Select



Select



# Locking

Begin

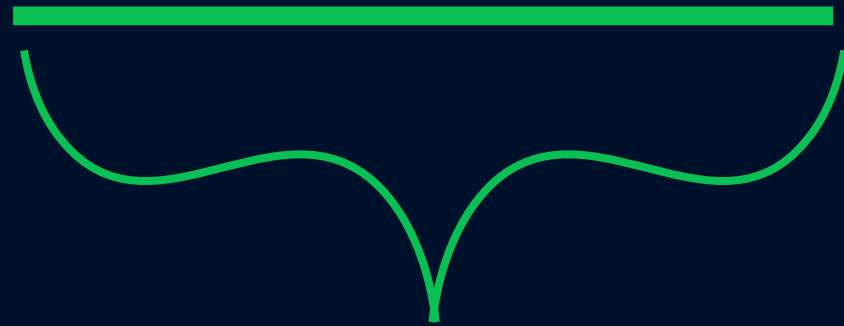


# Locking



Begin

Select



# Locking



# Locking



# Locking



# Locking



# Locking



# Locking



# Locking

```
Select * from table where flag is false for update;  
  
Update table where flag is false;
```



# Locking



Update selection

# Locking

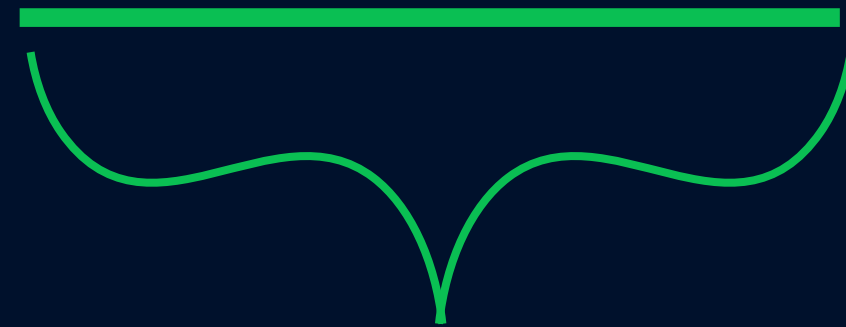
```
Begin;  
Select * from table where flag is false for update;  
  
Update table where flag is false;  
Commit;
```

# Locking

Begin



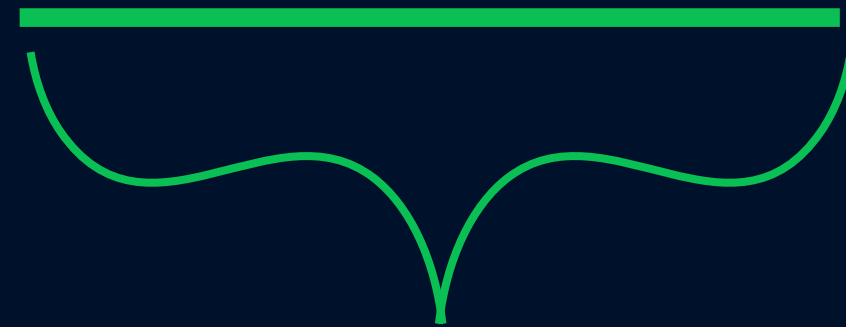
Select for update



# Locking

Begin;

Select for update



```
SELECT * FROM table FOR UPDATE;
```

# Locking



# Locking

tx=100



Update

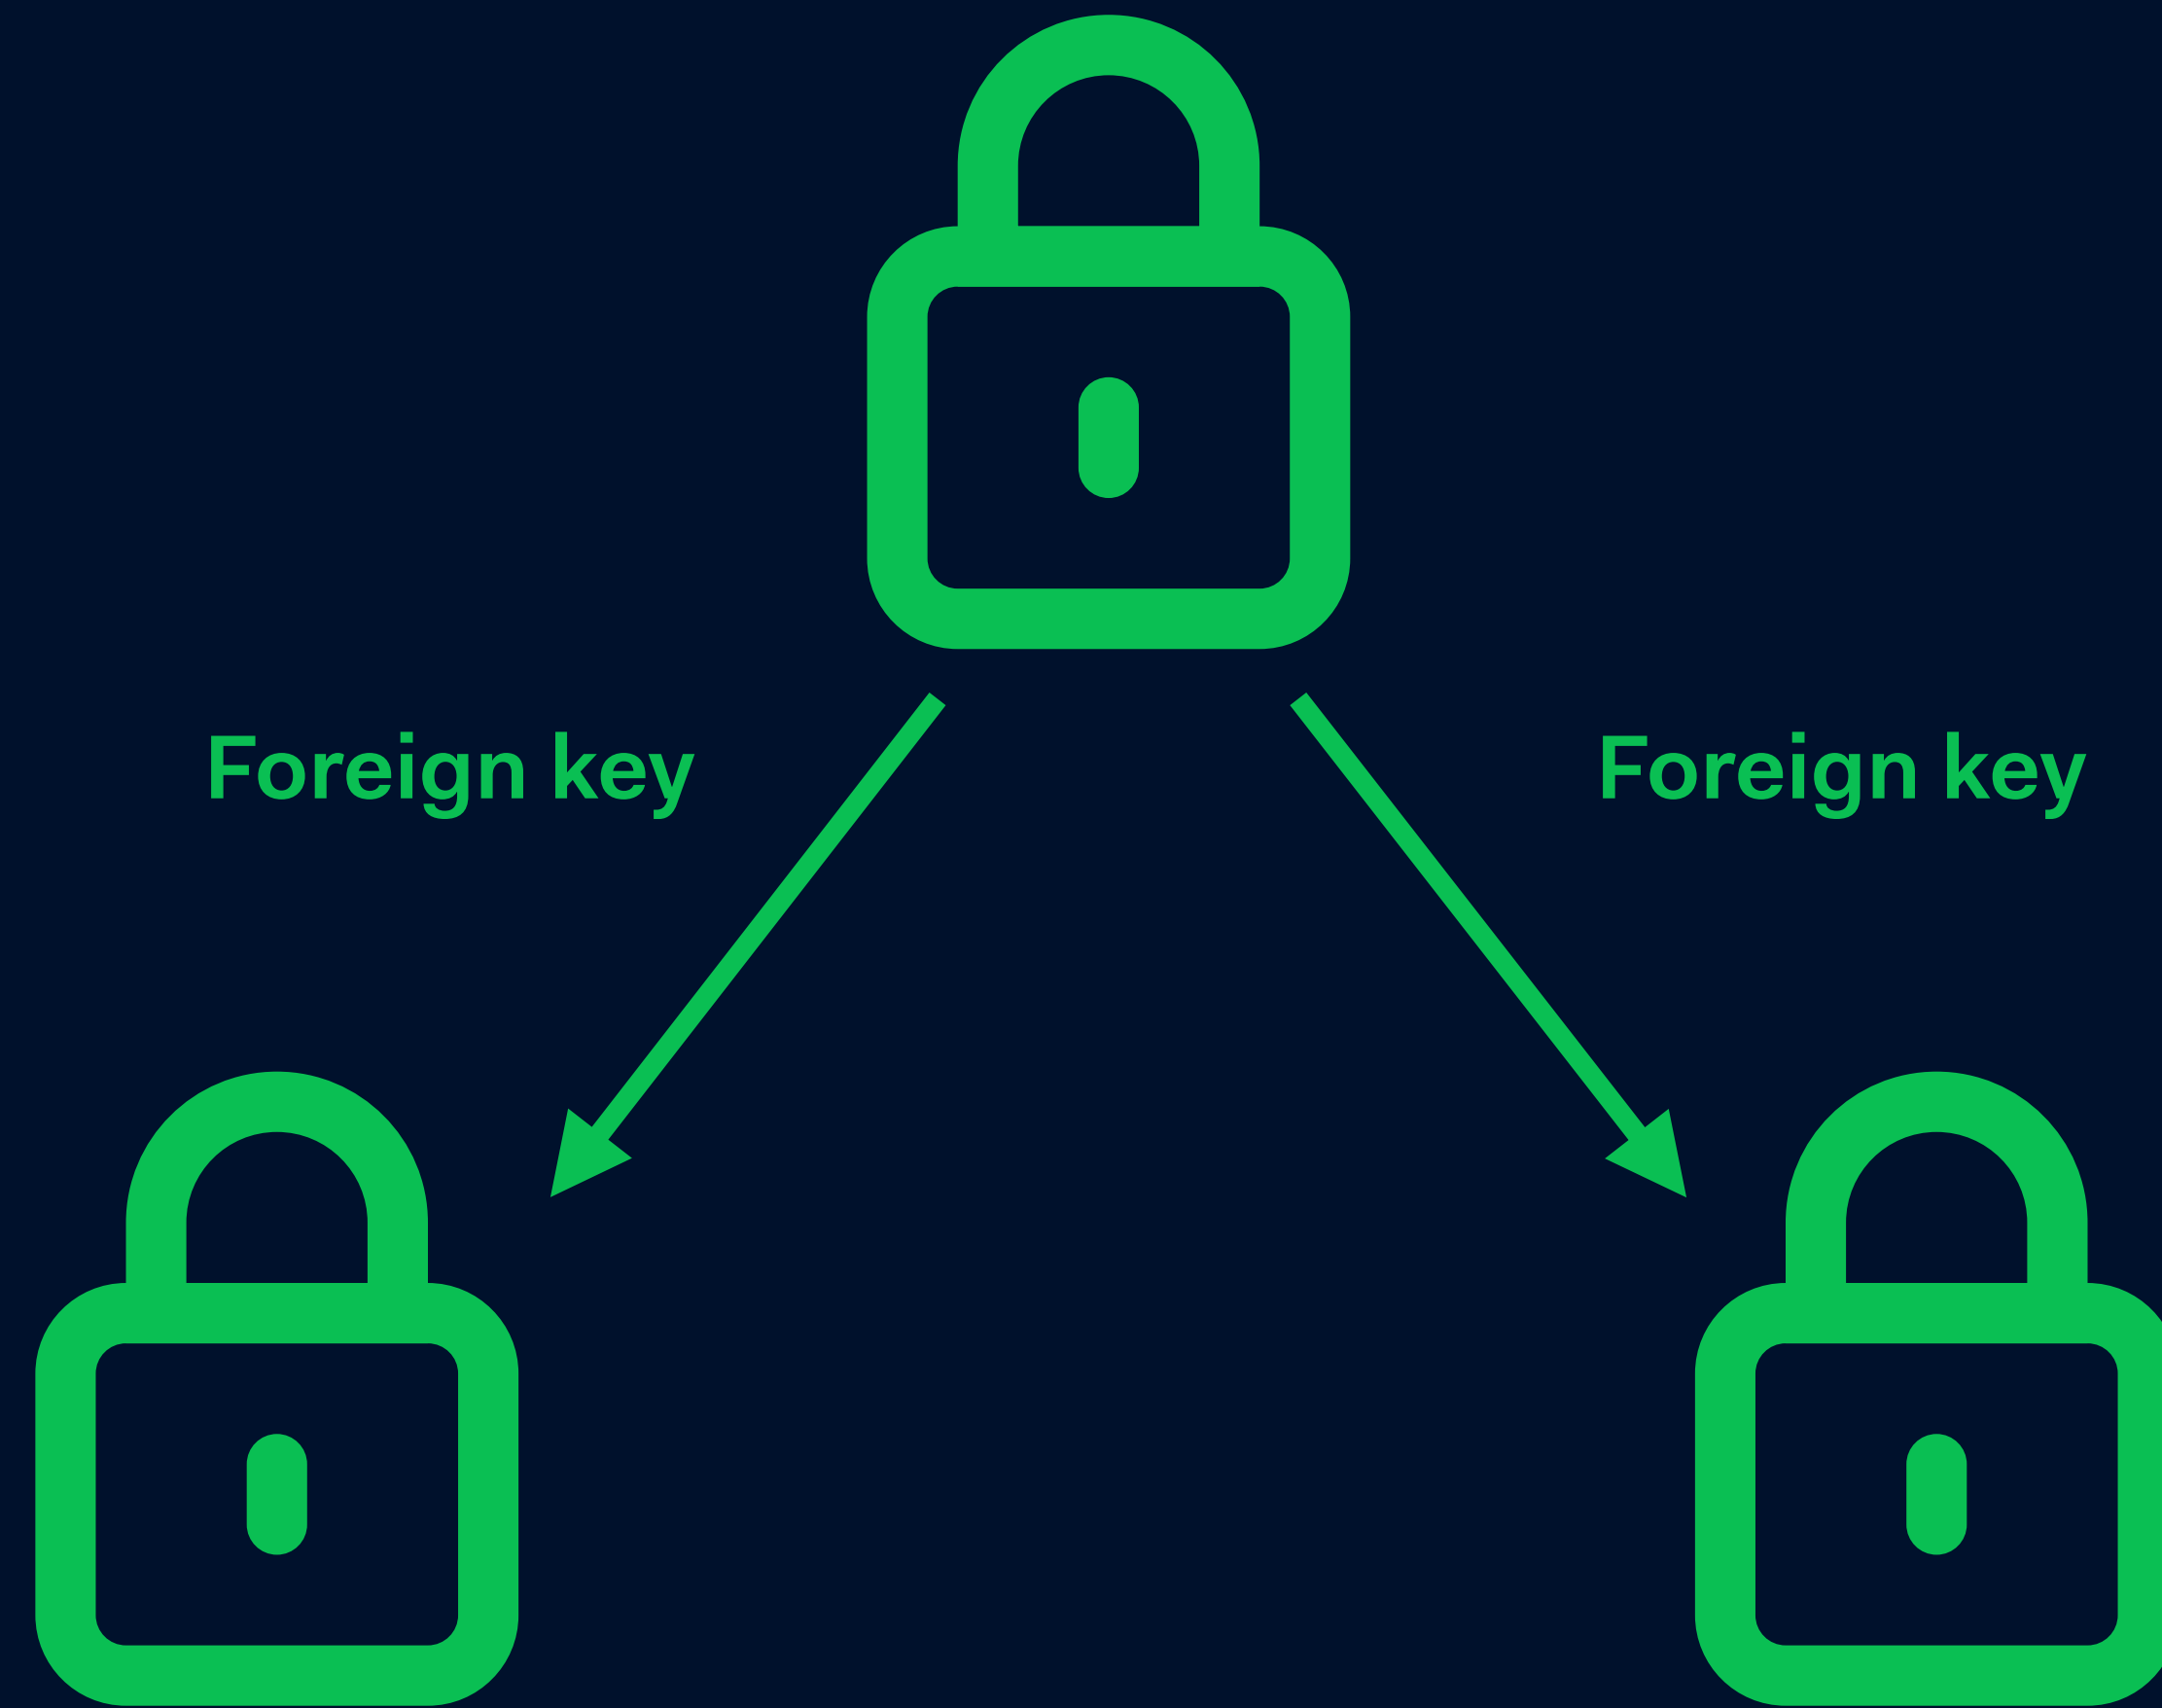
tx=102



# Locking

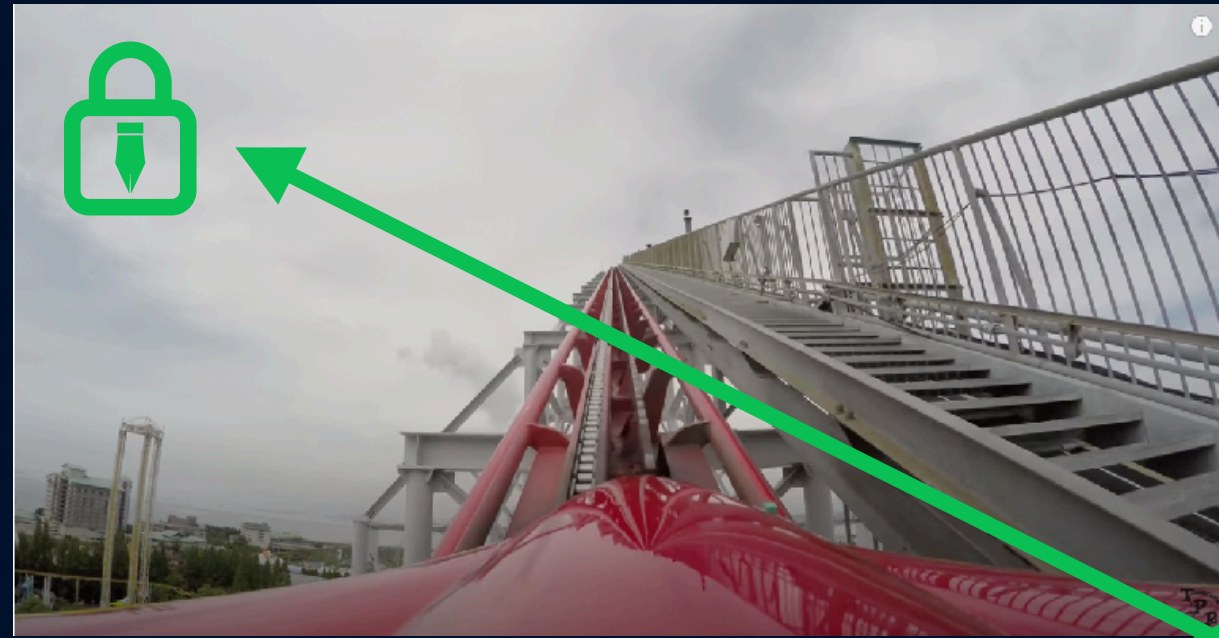


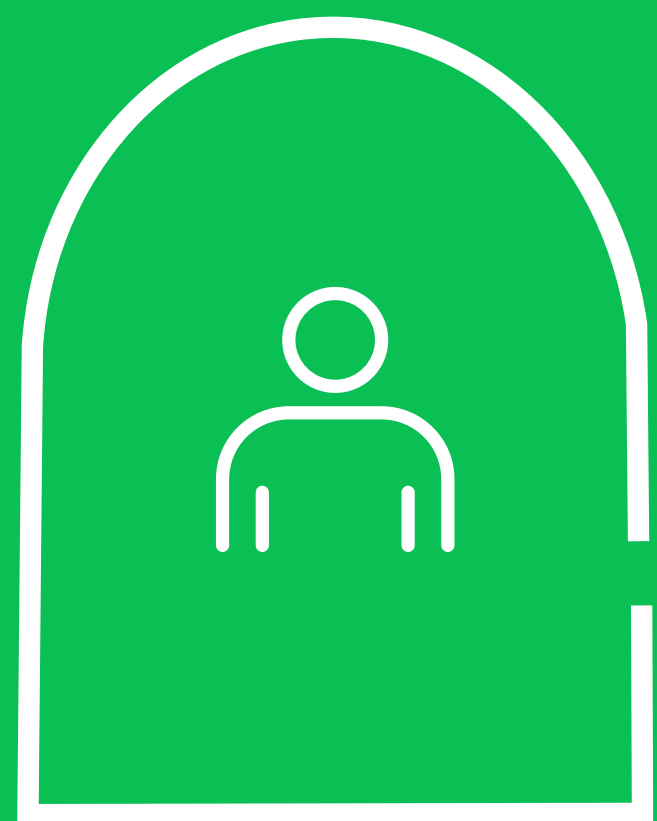
# Locking





# Locking



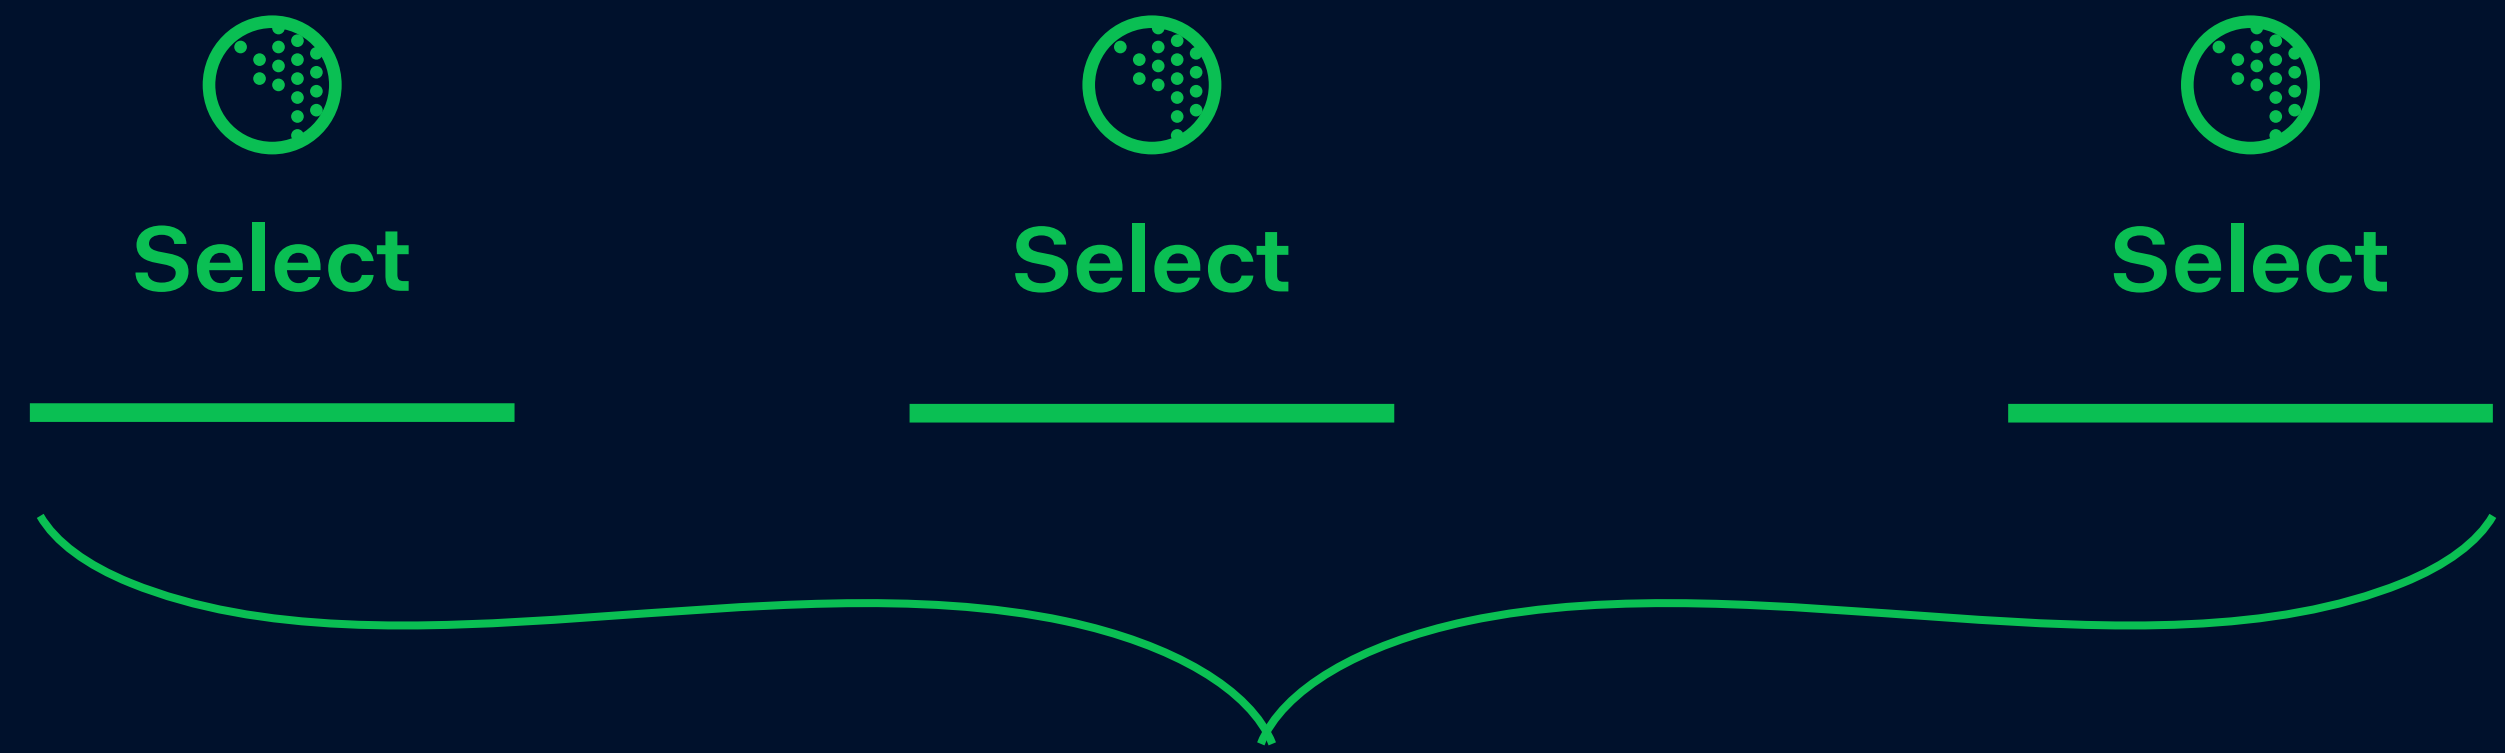


# Isolation levels

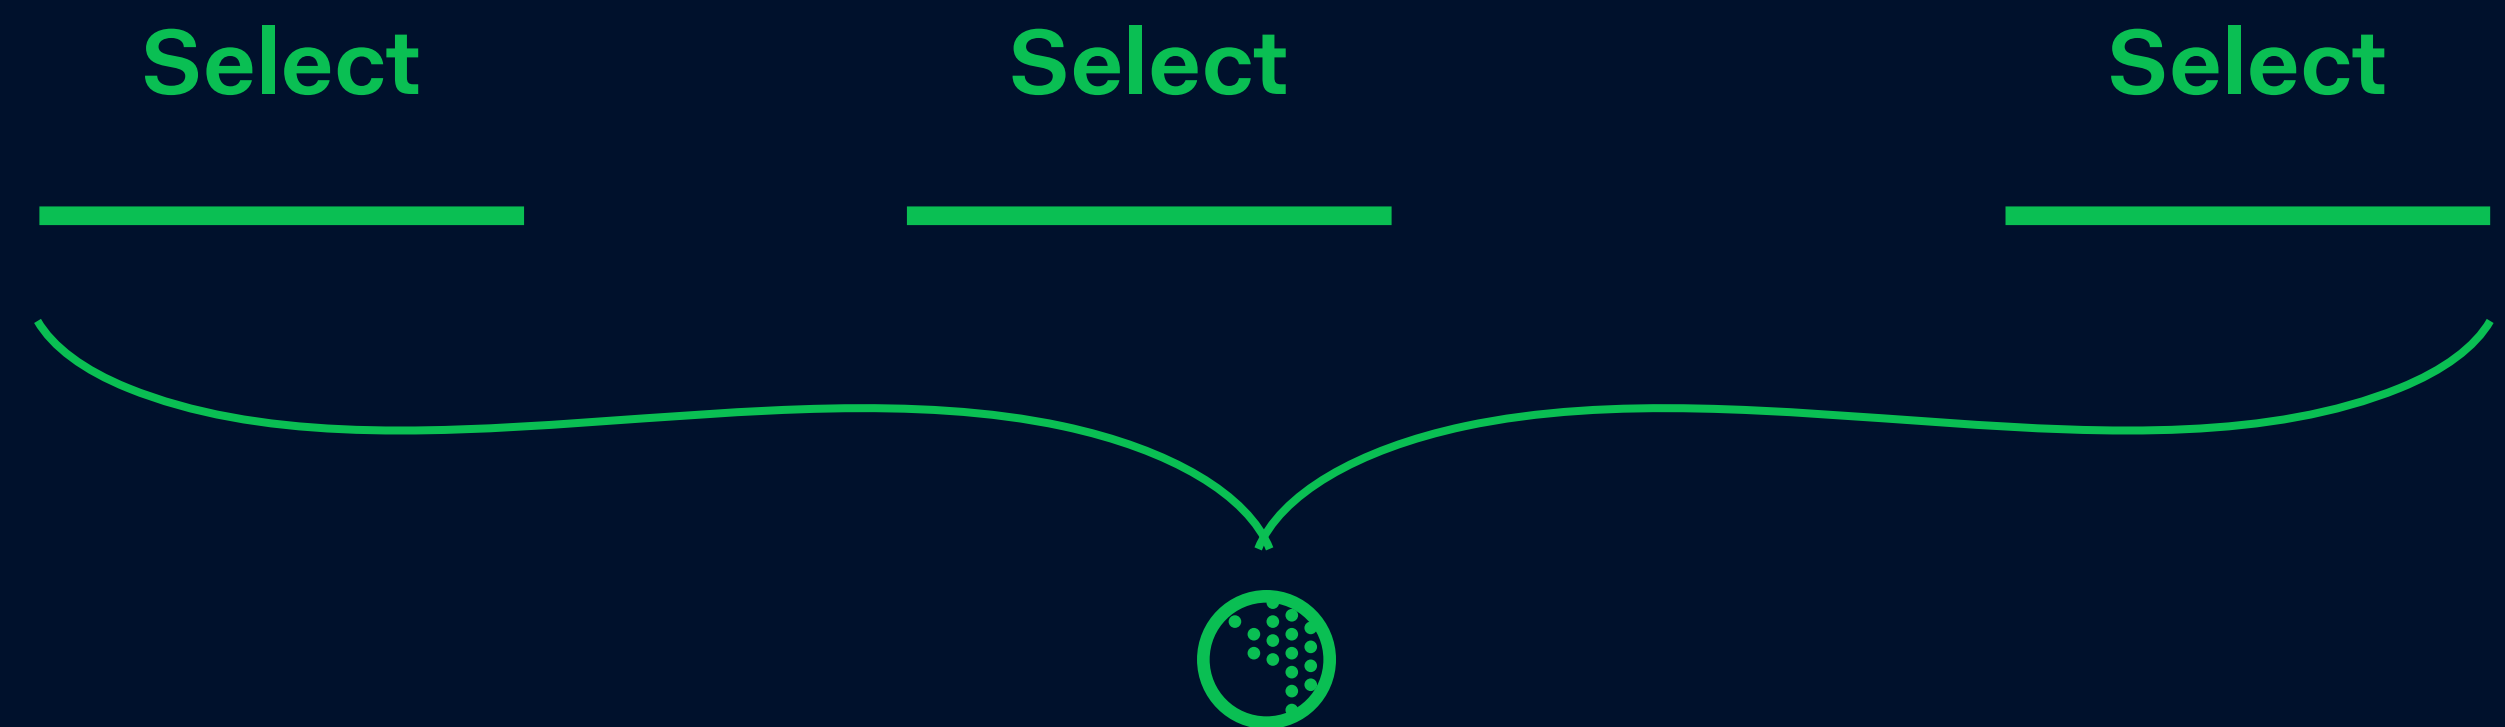


# Isolation levels

# Isolation levels



**Read committed**



**Repeatable read**

# Isolation levels

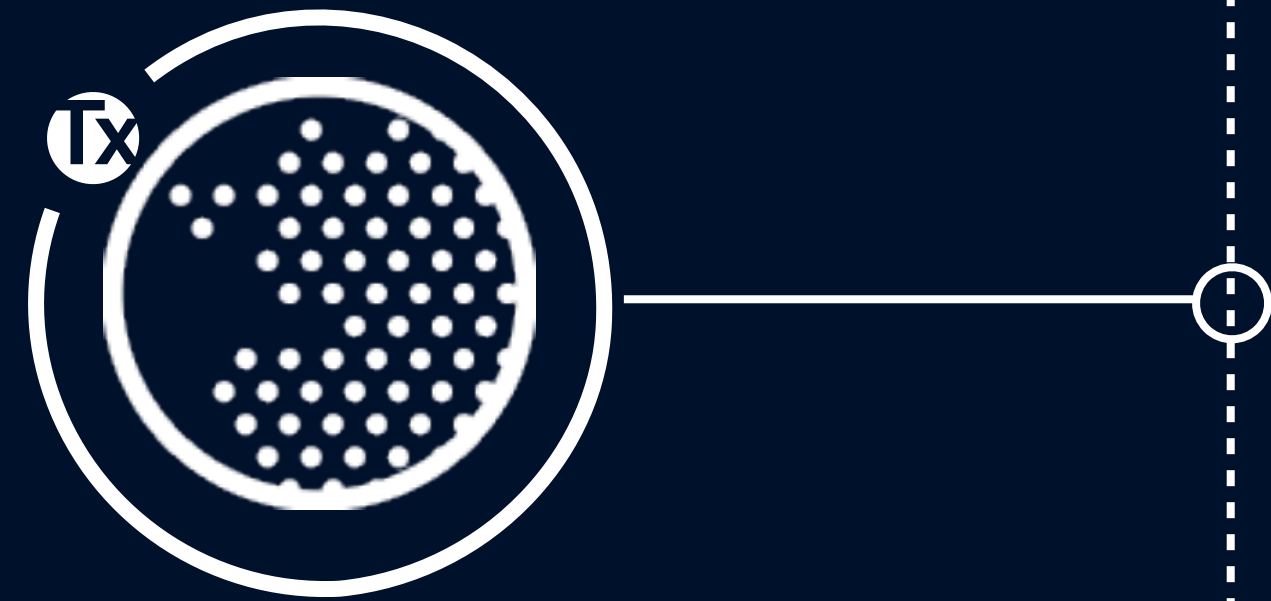


**Read committed**

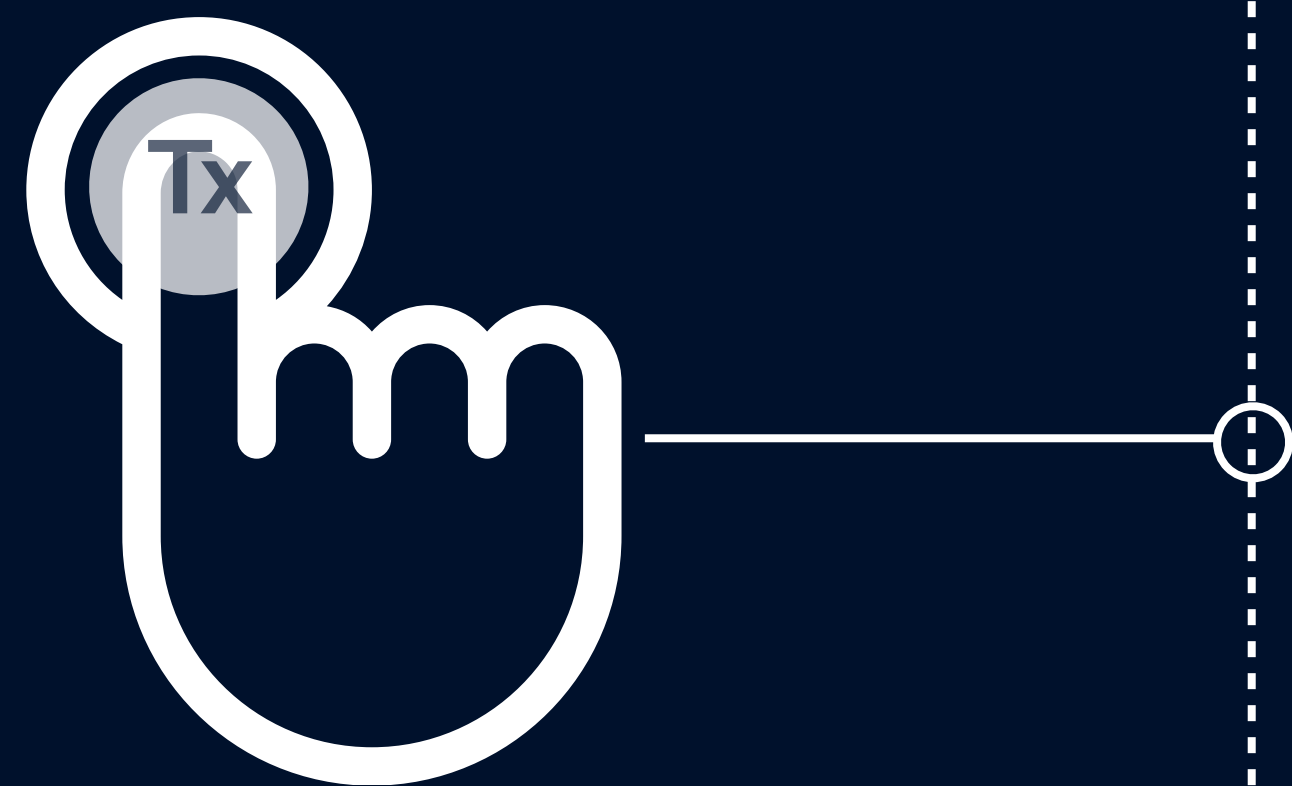
**Repeatable read**

# Learnings





# Transaction space



# Using transactions





**Locks**



# Isolation levels

Questions

```
def withdraw(amt, user_id):  
beginTxn()  
    bal = readBalance(user_id)  
    if (bal >= amt):  
        writeBalance(bal - amt, user_id)  
commit()
```