



Deep dive into the PostgreSQL Frontend/Backend protocol

Xavier Fischer
Swiss PGDay 2026

Who is using **psql**?

How on Earth is this possible?

you type a query text, and...

```
dvdrental=# select * from actor limit 10;
```

actor_id	first_name	last_name	last_update
1	Penelope	Guinness	2013-05-26 14:47:57.62
2	Nick	Wahlberg	2013-05-26 14:47:57.62
3	Ed	Chase	2013-05-26 14:47:57.62
4	Jennifer	Davis	2013-05-26 14:47:57.62
5	Johnny	Lollobrigida	2013-05-26 14:47:57.62
6	Bette	Nicholson	2013-05-26 14:47:57.62

... you get data back !!



psql does NOT know beforehand

- if the query syntax is OK
- how many columns the results will have
- the name, type for each column value

```
dvdrental=# select * from actor limit 10;
 actor_id | first_name | last_name | last_update
-----+-----+-----+-----
      1 | Penelope  | Guinness | 2013-05-26 14:47:57.62
      2 | Nick      | Wahlberg | 2013-05-26 14:47:57.62
      3 | Ed        | Chase    | 2013-05-26 14:47:57.62
      4 | Jennifer  | Davis    | 2013-05-26 14:47:57.62
      5 | Johnny    | Lollobrigida | 2013-05-26 14:47:57.62
      6 | Bette     | Nicholson | 2013-05-26 14:47:57.62
```

This is handled by libpq (the “**connector**”) after exchanging several messages with PostgreSQL

Who am I

Xavier Fischer

- Principal Software Engineer at EDB since 2023
- Working on EDB .NET Connector
- 20+ years of experience in .NET / database development
- Hailing in Aix-en-Provence, France

Who am I

Xavier Fischer

- Principal Software Engineer at EDB since 2018
- Working on EDB .NET Connector
- 20+ years of experience in .NET / data
- Hailing in Aix-en-Provence, France



What is a **connector**?

it's a **database driver**

available in **your programming language**
will **translate** calls to database language using
the **wire protocol**
aka the Frontend/Backend protocol



what's the **wire protocol**?

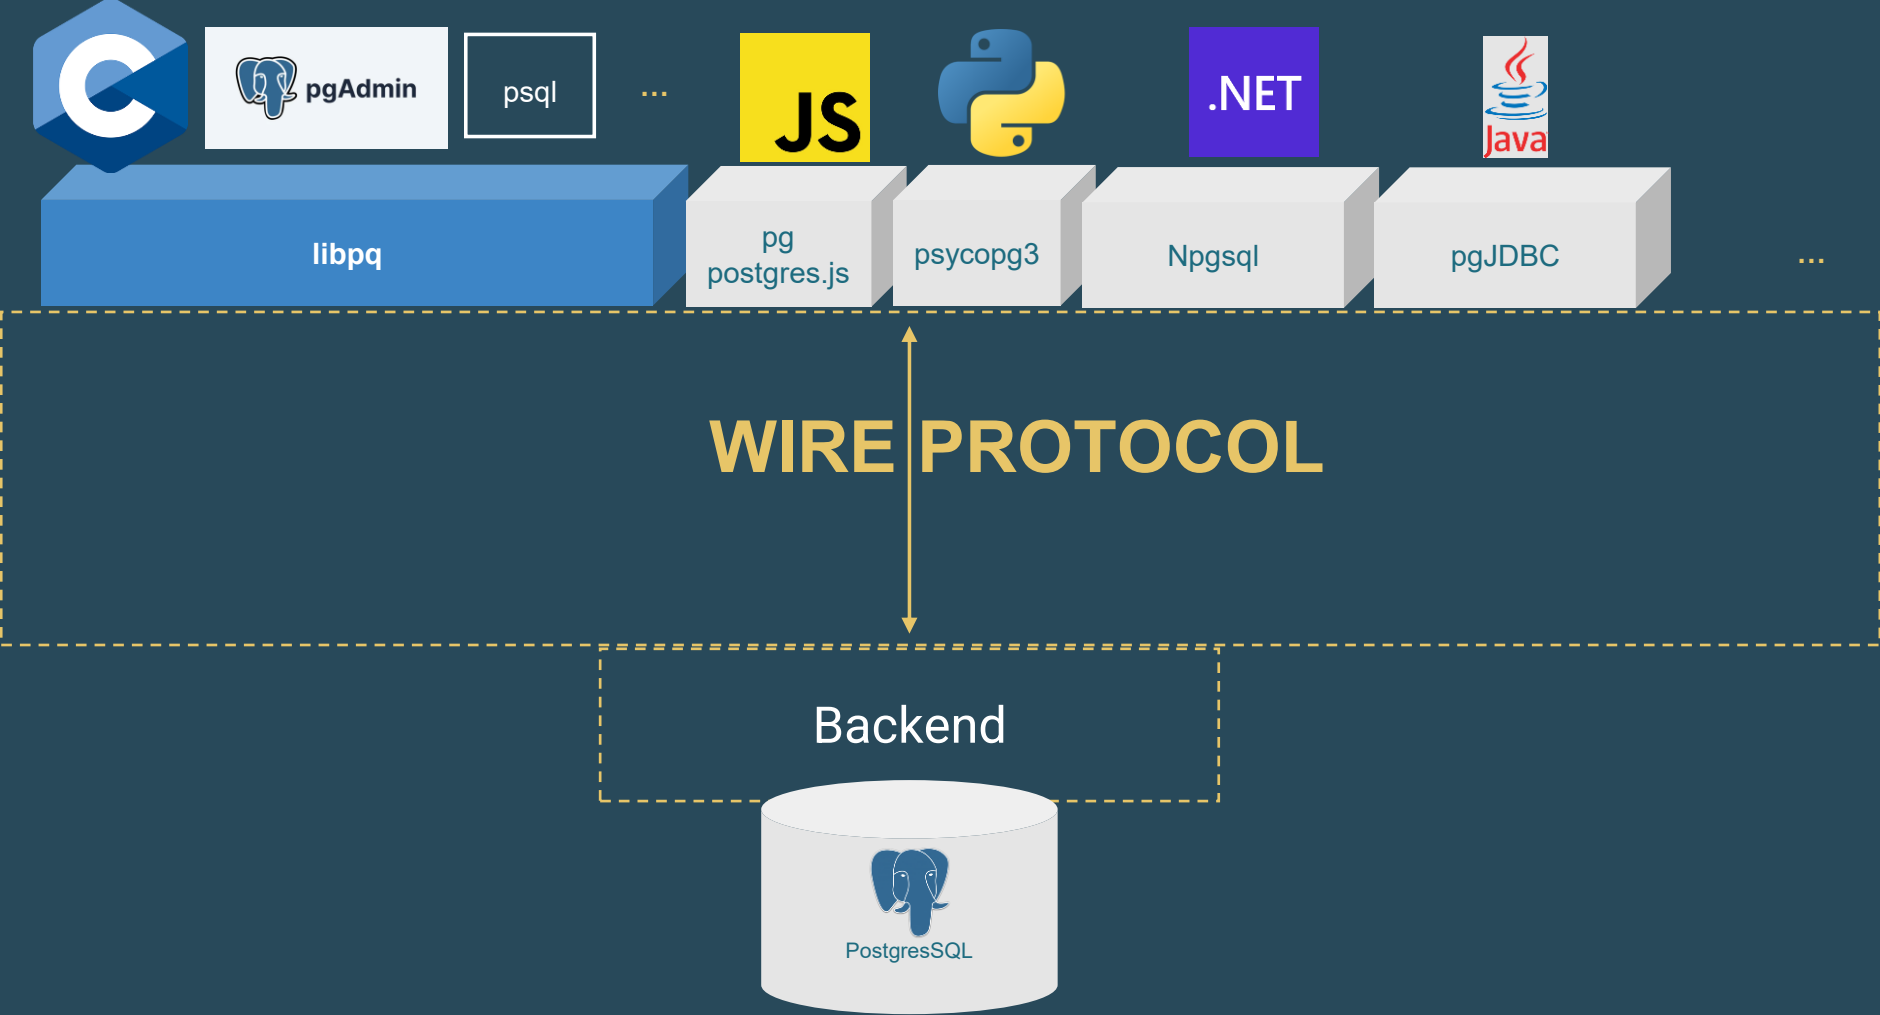
It's the PostgreSQL **message-based**
protocol between
frontends and backends
(clients, servers, poolers, tools, ...)
over TCP/IP or Unix sockets



Protocol history




Postgres Connector landscape



protocol messages

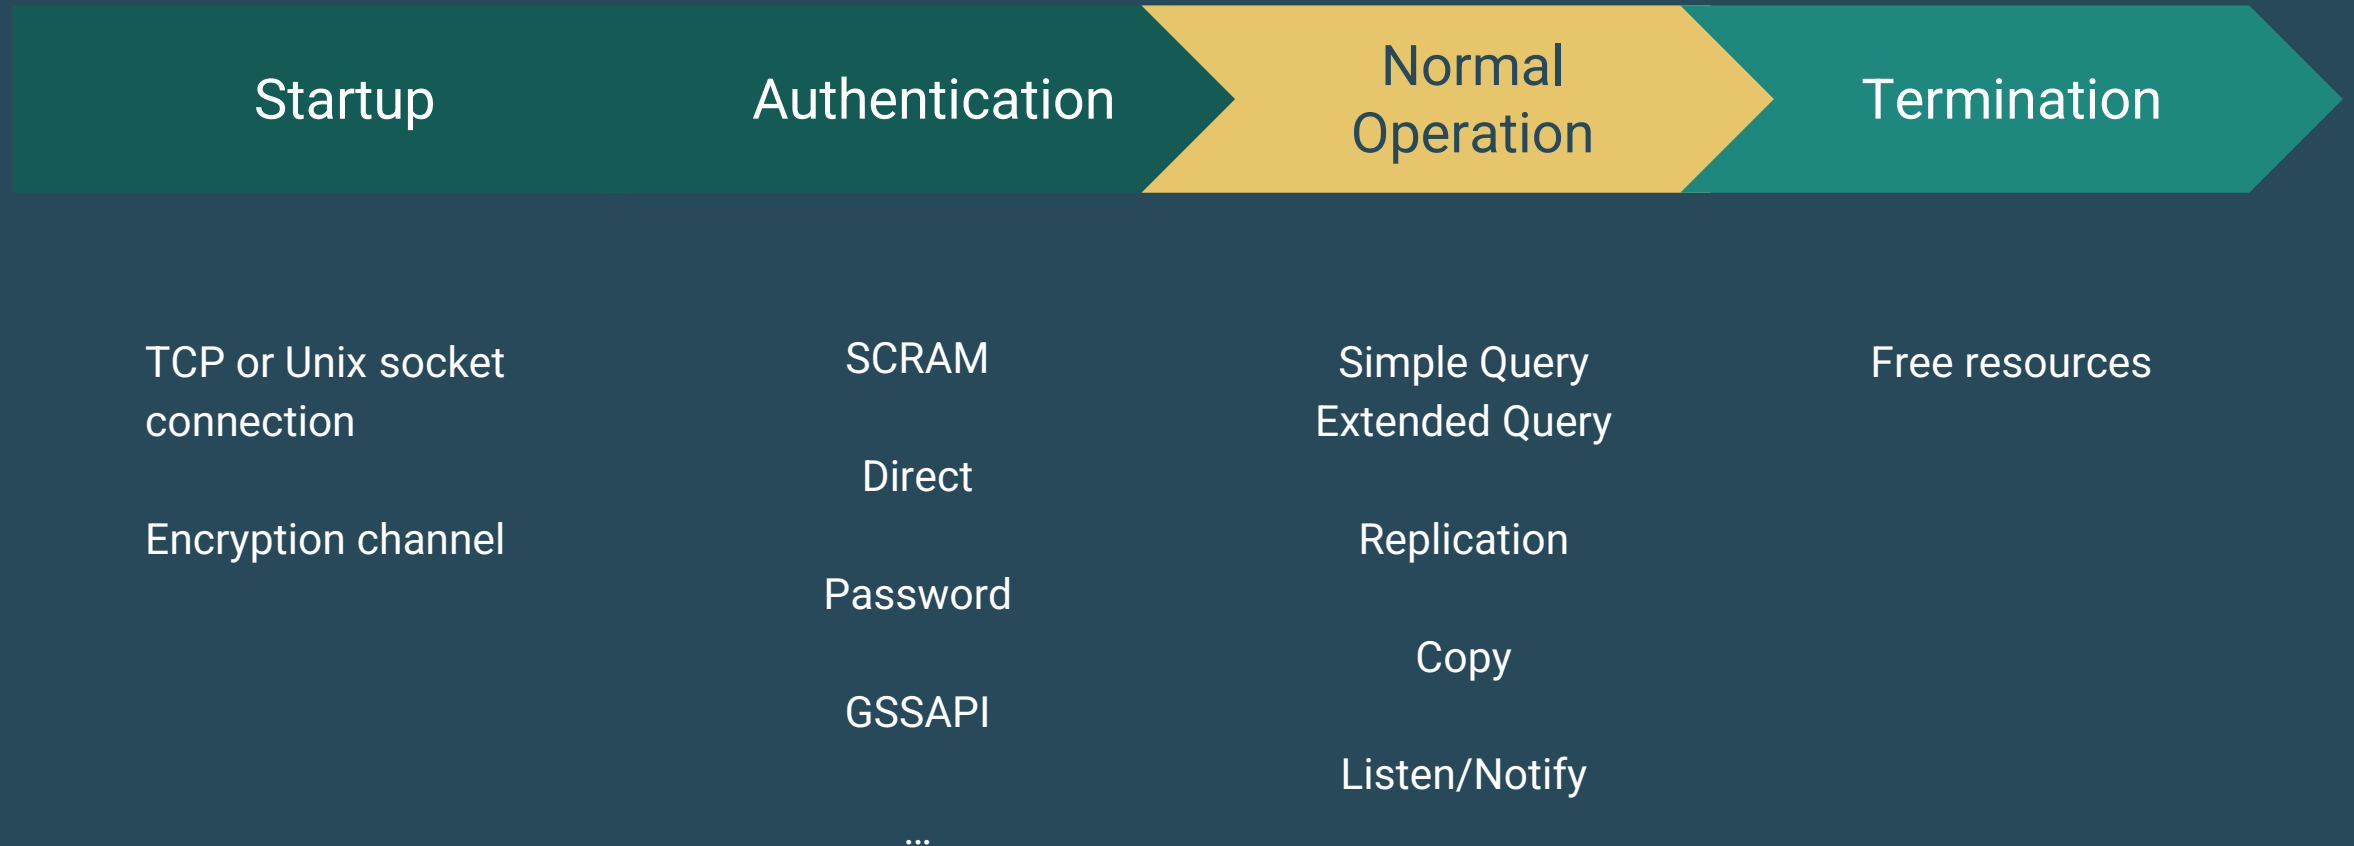
Protocol message format 101

Messages are usually Type-Length-Value (TLV)



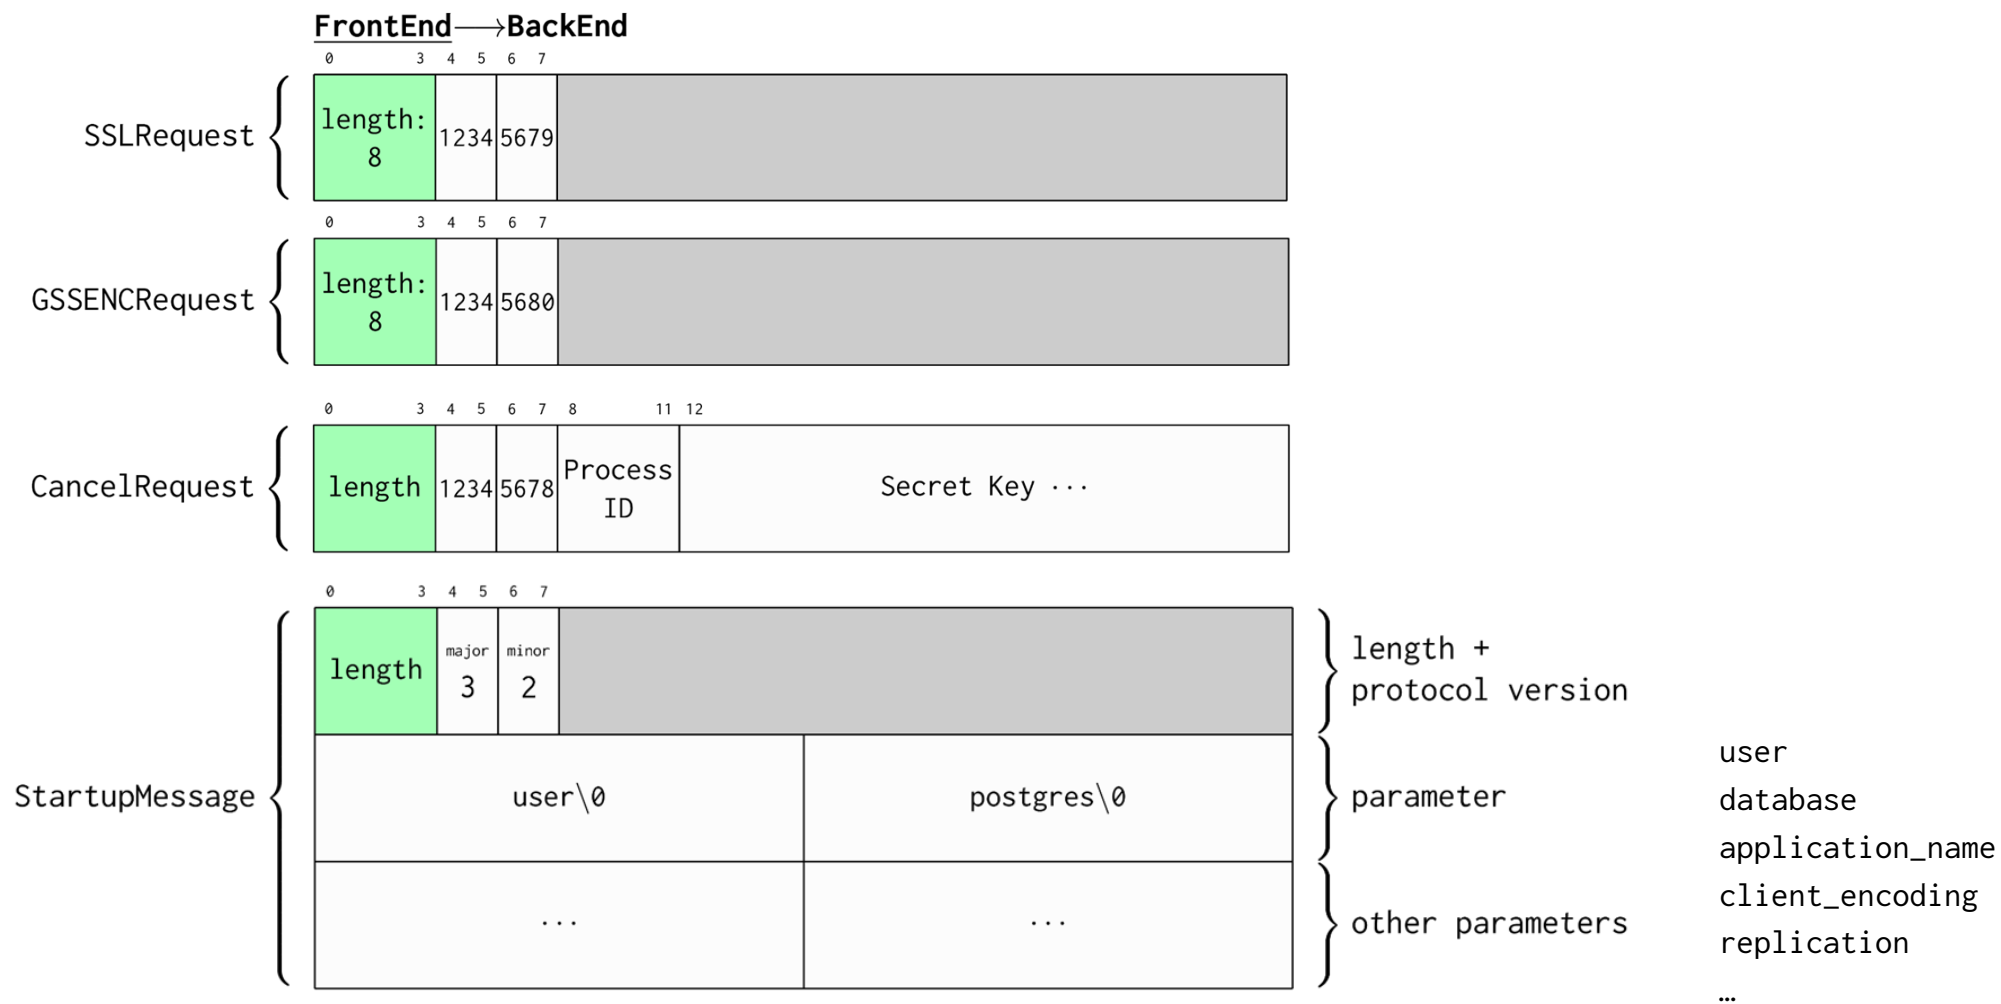
=> If a new protocol is released, unknown messages can be skipped easily

Protocol conversation 101



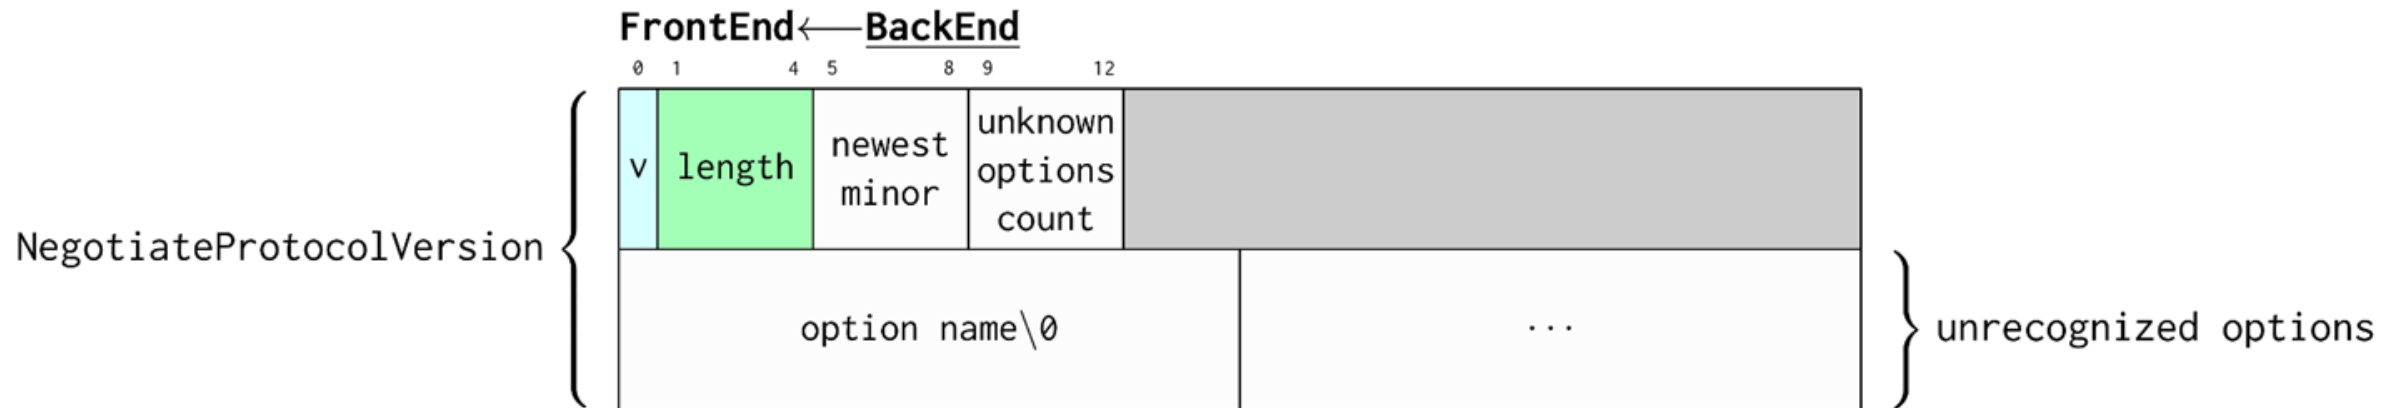
Startup Message

First message sent from the frontend is one of those



Startup Message

When backend doesn't support minor protocol version



Frontend should close the connection if it does not support this minor version.

With PG18 and V3.2 this could happen if frontend expects V3.2 and server only has V3.0

The frontend can usually continue if it can live without the unsupported options or stick to server version

Authentication

Method picked in `pg_hba.conf`, backend sends the Authentication* “challenge” message

- Messages exchanged vary for each method
- Final step: backend sends `Error` or `AuthenticationOk`

PG18 introduced OAUTH, deprecated MD5

SCRAM and OAUTH are recommended since PG18

you're (nearly) in!

how is the **backend** process
created?

Startup

Authentication

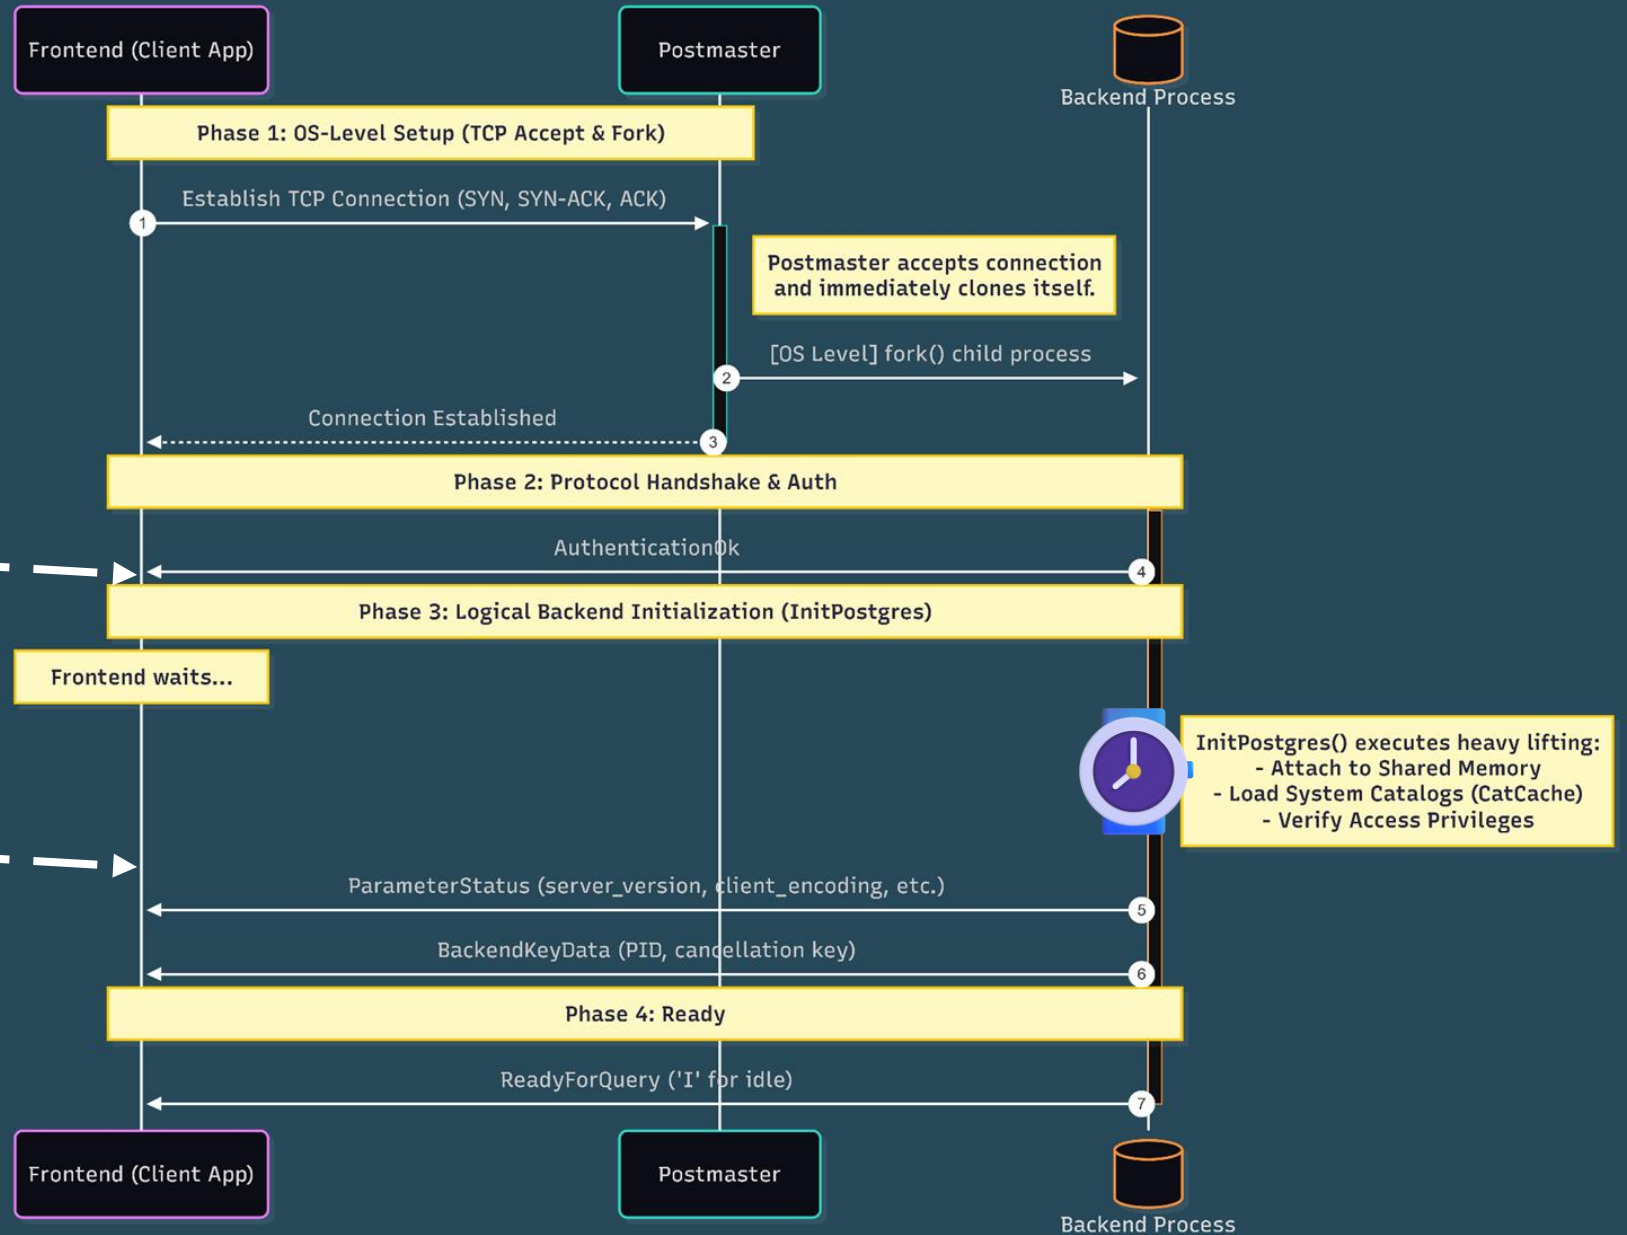
Normal Operation

Termination

Backend spawn

You are here

We are going there



Server parameters

- ParameterStatus (a lot)
- Backend Key Data: PID + key
(needed by CancelRequest on a fresh new connection)

- ReadyForQuery (🎉)

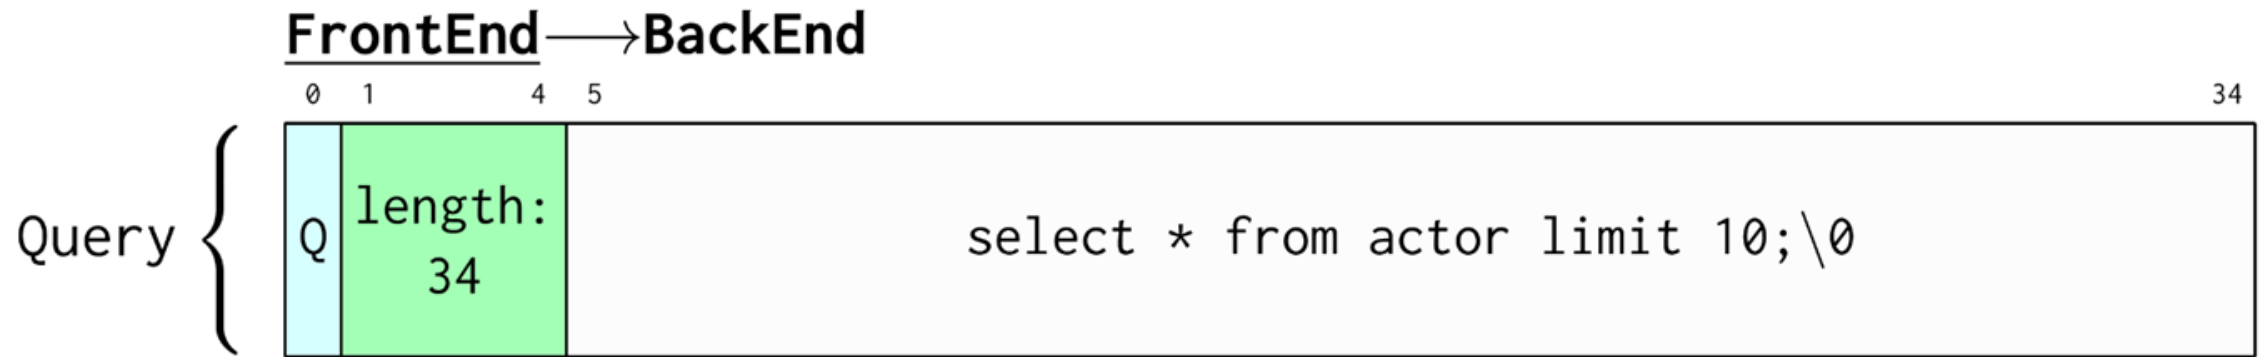
	FrontEnd	BackEnd	
AuthenticationOK	R	length: 8	type: 0
ParameterStatus	S	length: 23	in_hot_standby off
ParameterStatus	S	length: 25	integer_datetimes on
ParameterStatus	S	length: 26	TimeZone Europe/Paris
ParameterStatus	S	length: 27	IntervalStyle postgres
ParameterStatus	S	length: 32	search_path "\$user", public
ParameterStatus	S	length: 20	is_superuser on
ParameterStatus	S	length: 26	application_name psql
ParameterStatus	S	length: 38	default_transaction_read_only off
ParameterStatus	S	length: 26	scram_iterations 4096
ParameterStatus	S	length: 23	DateStyle ISO, MDY
ParameterStatus	S	length: 35	standard_conforming_strings on
ParameterStatus	S	length: 35	session_authorization postgres
ParameterStatus	S	length: 28	client_encoding WIN1252
ParameterStatus	S	length: 24	server_version 18.3
ParameterStatus	S	length: 25	server_encoding UTF8
BackendKeyData	K	length: 12	PID 19480 Secret 6796488
ReadyForQuery	Z	length: 5	Idle

Simple query explained


```
dvdrental=# select * from actor limit 10;
```

actor_id	first_name	last_name	last_update
1	Penelope	Guinness	2013-05-26 14:47:57.62
2	Nick	Wahlberg	2013-05-26 14:47:57.62
3	Ed	Chase	2013-05-26 14:47:57.62
4	Jennifer	Davis	2013-05-26 14:47:57.62
5	Johnny	Lollobrigida	2013-05-26 14:47:57.62
6	Bette	Nicholson	2013-05-26 14:47:57.62

Simple query explained: >q



Things can go wrong



Simple query explained: >Q <T

FrontEnd ← BackEnd

0 1 4 5

31

T	length: 120	fields 4					
name: actor_id		table oid: 1467366	index: 1	type oid: 23	length: 4	type modifier: -1	Text
name: first_name		table oid: 1467366	index: 2	type oid: 1043	length: -1	type modifier: 49	Text
name: last_name		table oid: 1467366	index: 3	type oid: 1043	length: -1	type modifier: 49	Text
name: last_update		table oid: 1467366	index: 4	type oid: 1114	length: 8	type modifier: -1	Text

Row Description


column 1
actor_id

column 2
first_name

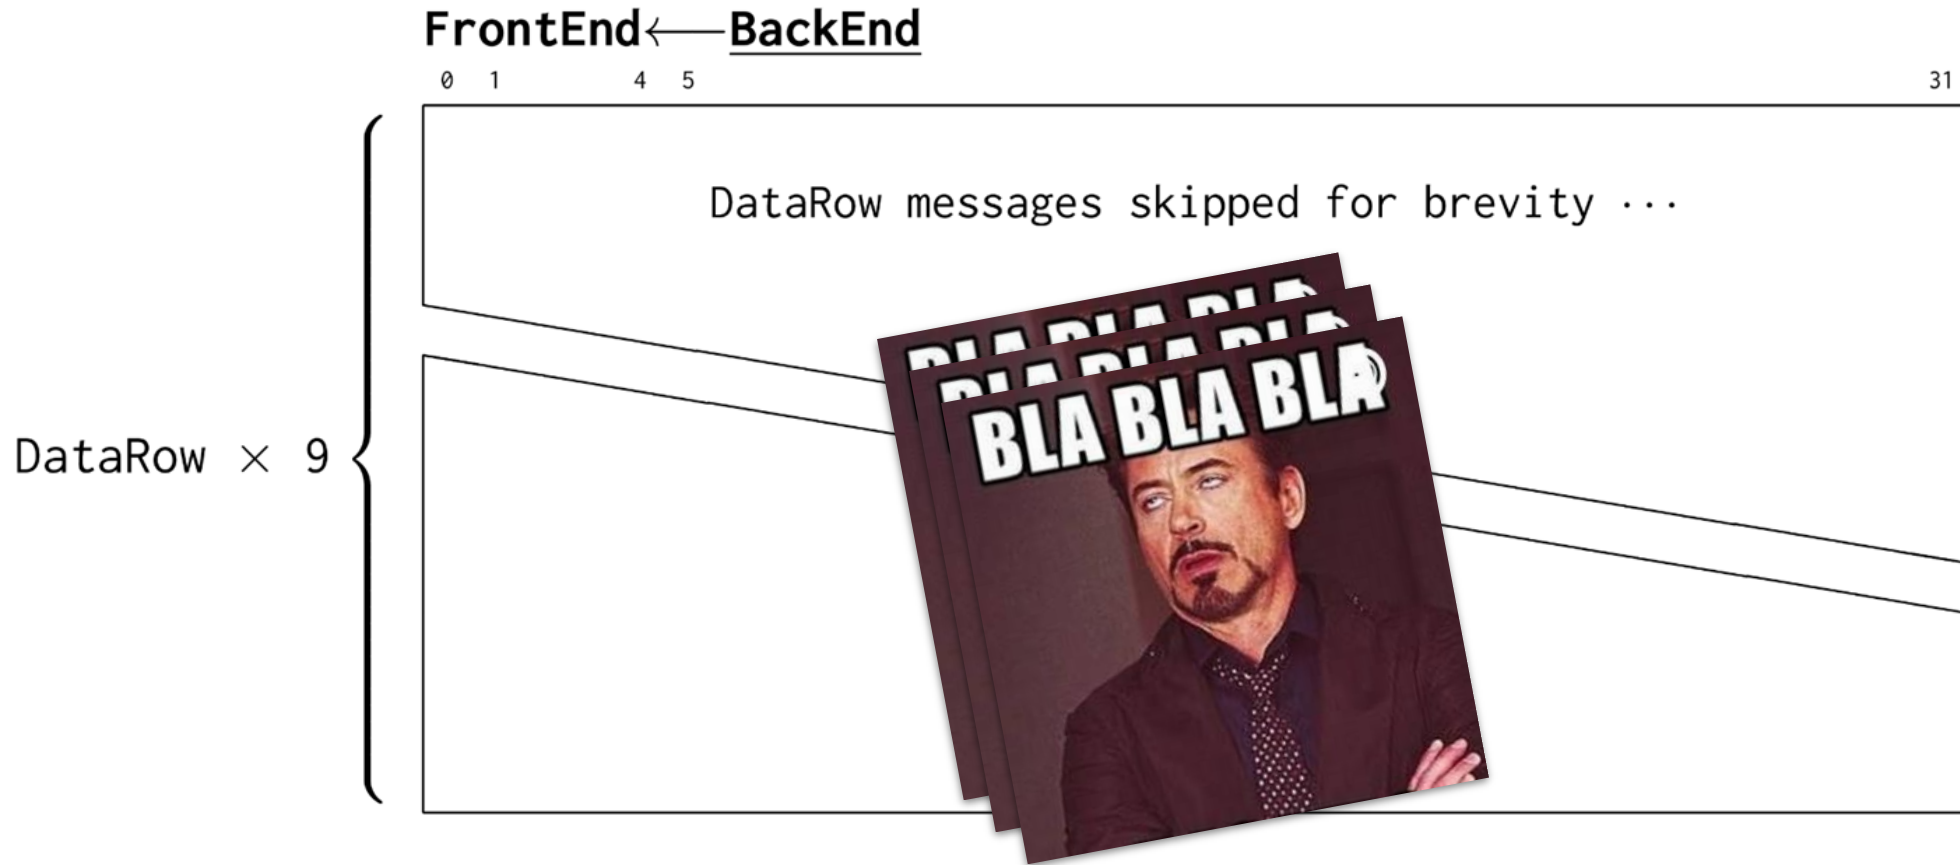
column 3
last_name

column 4
last_update


Simple query explained: >Q <T/D




Simple query explained: >Q <T/D/



Simple query explained: >Q <T/D/D/D/D/D/D/D/D/D/D/D/c



Simple query explained: >Q <T/D/D/D/D/D/D/D/D/D/D/D/C/Z



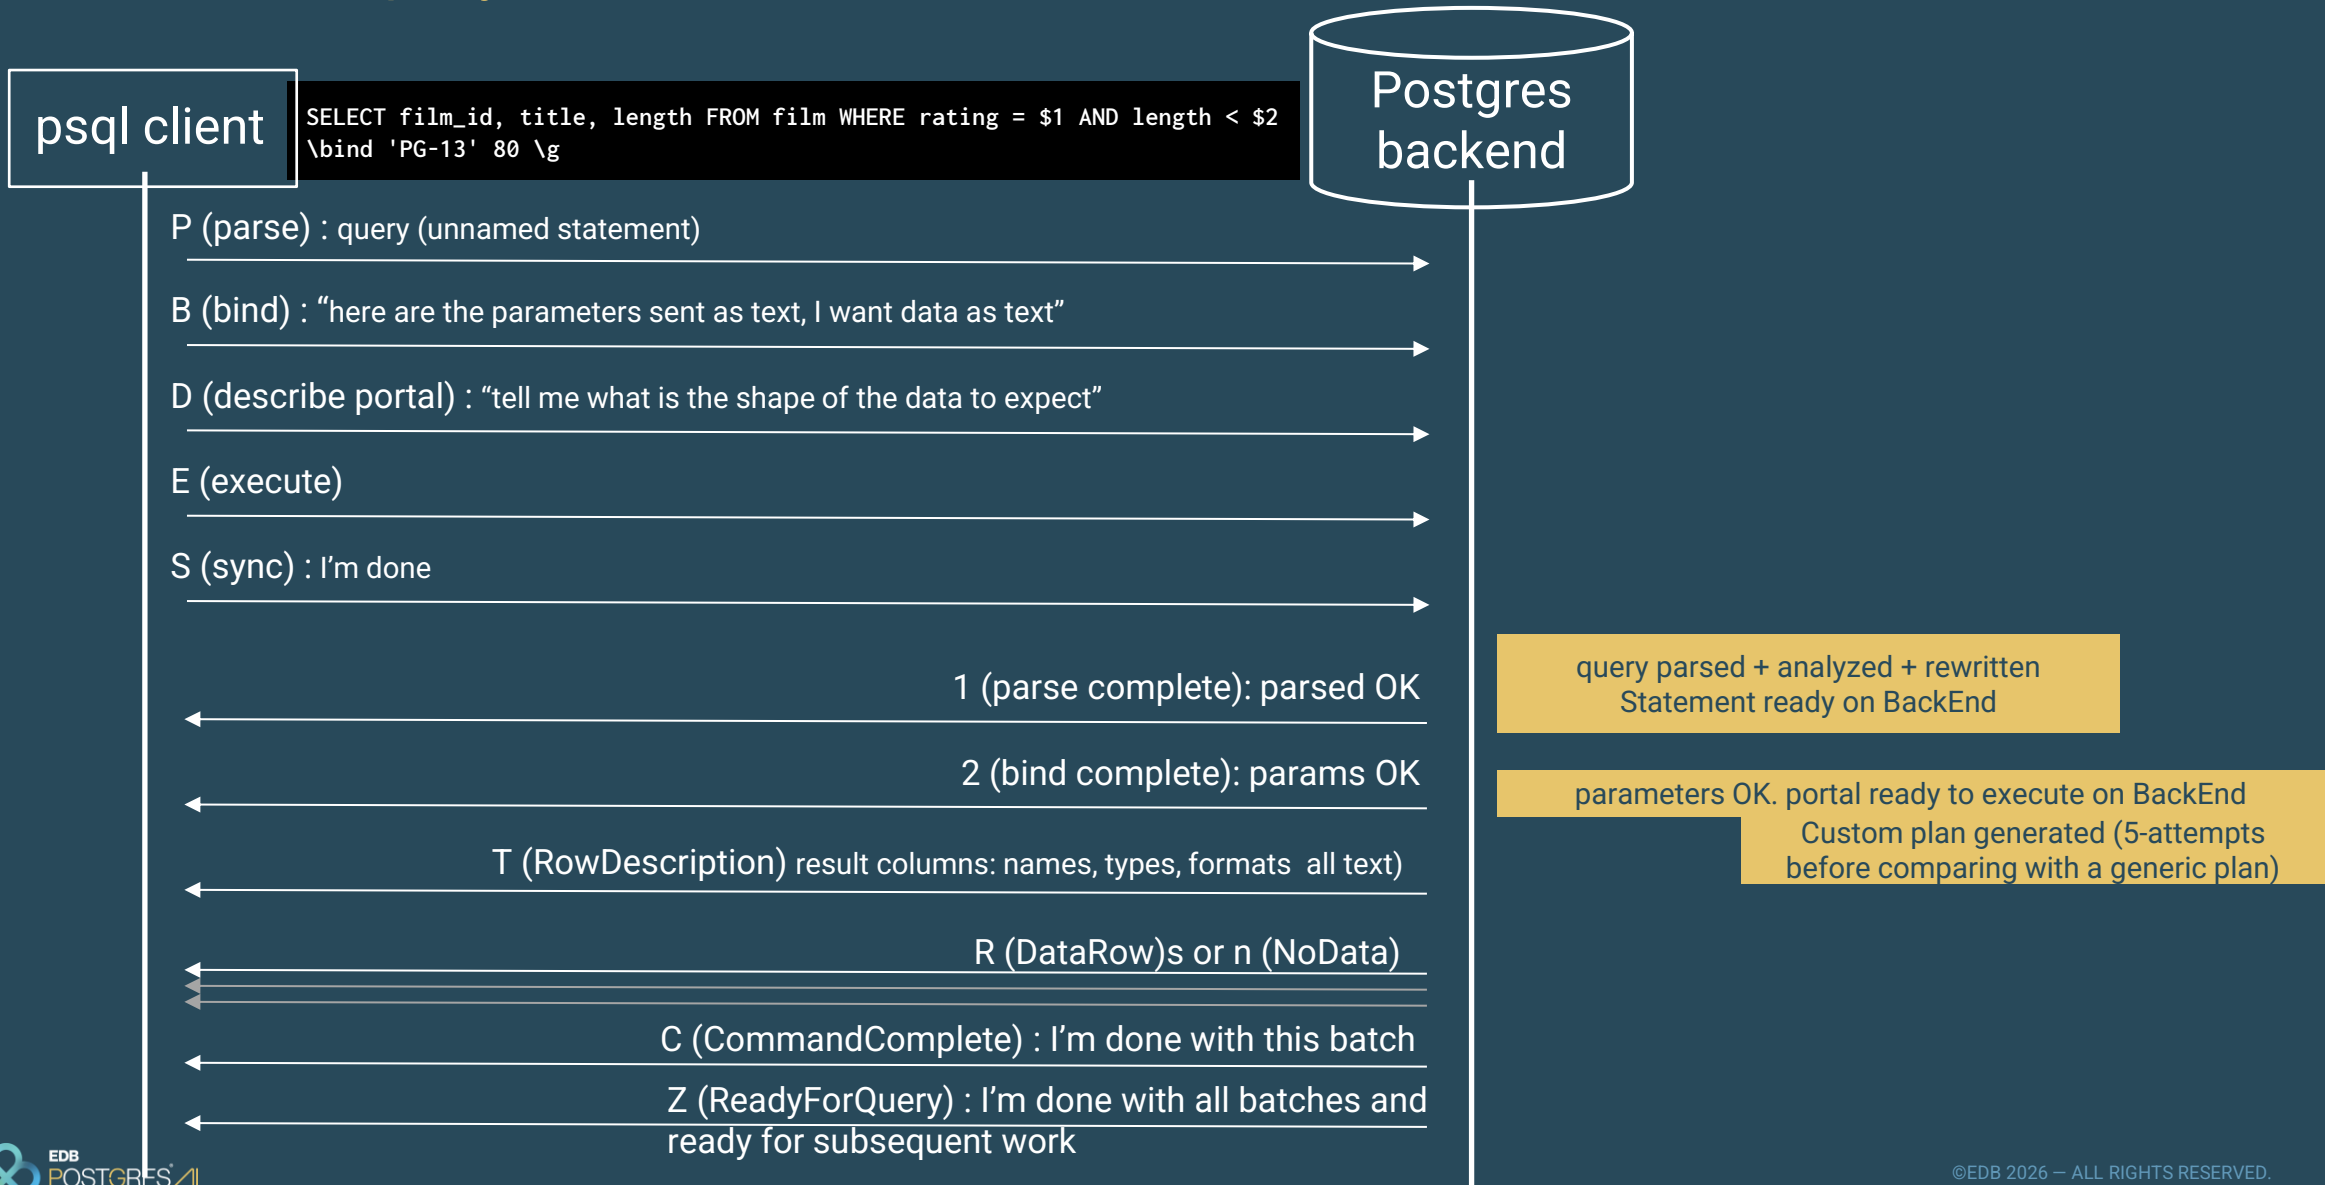
Simple query explained: >Q <T/D/D/D/D/D/D/D/D/D/D/C/Z



extended query

Extended query: simple example

parameterized query

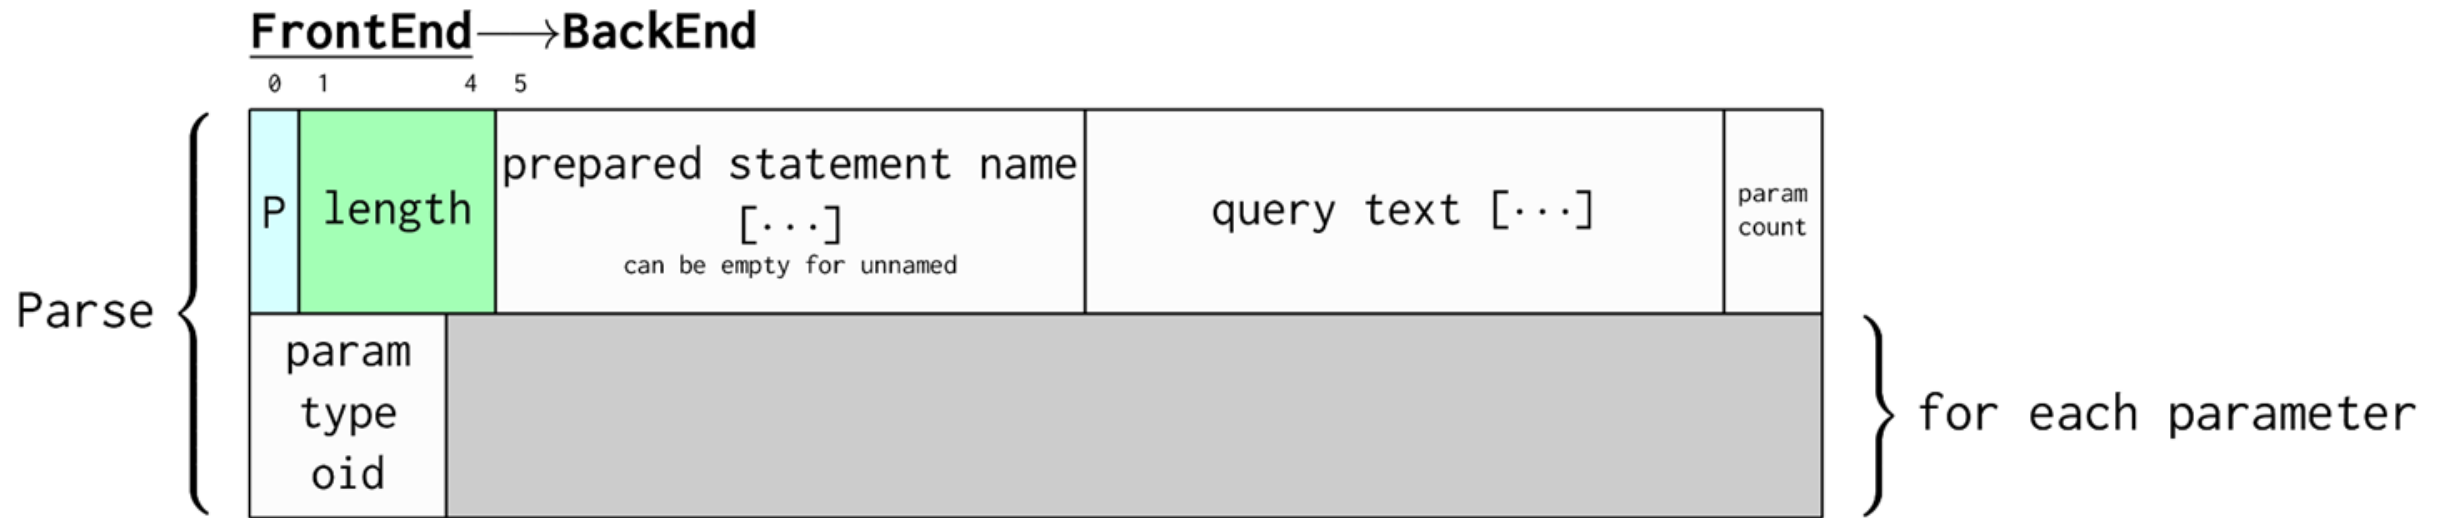


Statements / portals


- **Statement** is created with Parse. It's a reusable parameterized query "template"
=> kip parsing stage.
- **Portals** are created with Bind. It's statements with parameters, planned and ready to be executed.

Parse → statement → Bind → portal → Execute (all or fetch)

Extended query explained: Parse



=> Parse complete expected



Extended query explained: Bind

FrontEnd → BackEnd

0	1	4	5				
B	length	portal name [...] <small>(empty for unnamed)</small>		statement name [...] <small>(empty for unnamed)</small>			

0 or 1	No parameters (0), all text (0), all binary (1)						
count	format 1	[...]	format n	Format specified for each parameter			

} Parameter formats alternatives


count	param1 length	param1 value [...]	...	param n length	param n value [...]	
-------	---------------	--------------------	-----	----------------	---------------------	--

} Parameter values
unspecified (OUT params):
length=-1, no value


0 or 1	All results as text (0) or binary (1)						
count	format 1	[...]	format n	Format specified for each result			

} Result formats alternatives

Extended query explained: Describe (statement/portal)



=> Statement variant returns ParameterDescription, usually called before Bind



=> both variants return RowDescription

Extended query : prepared statement

Named statement
(in psql since PG18)

psql client

Postgres backend

```
SELECT film_id, title, length FROM film WHERE rating = $1 AND length < $2
\parse film_search
```

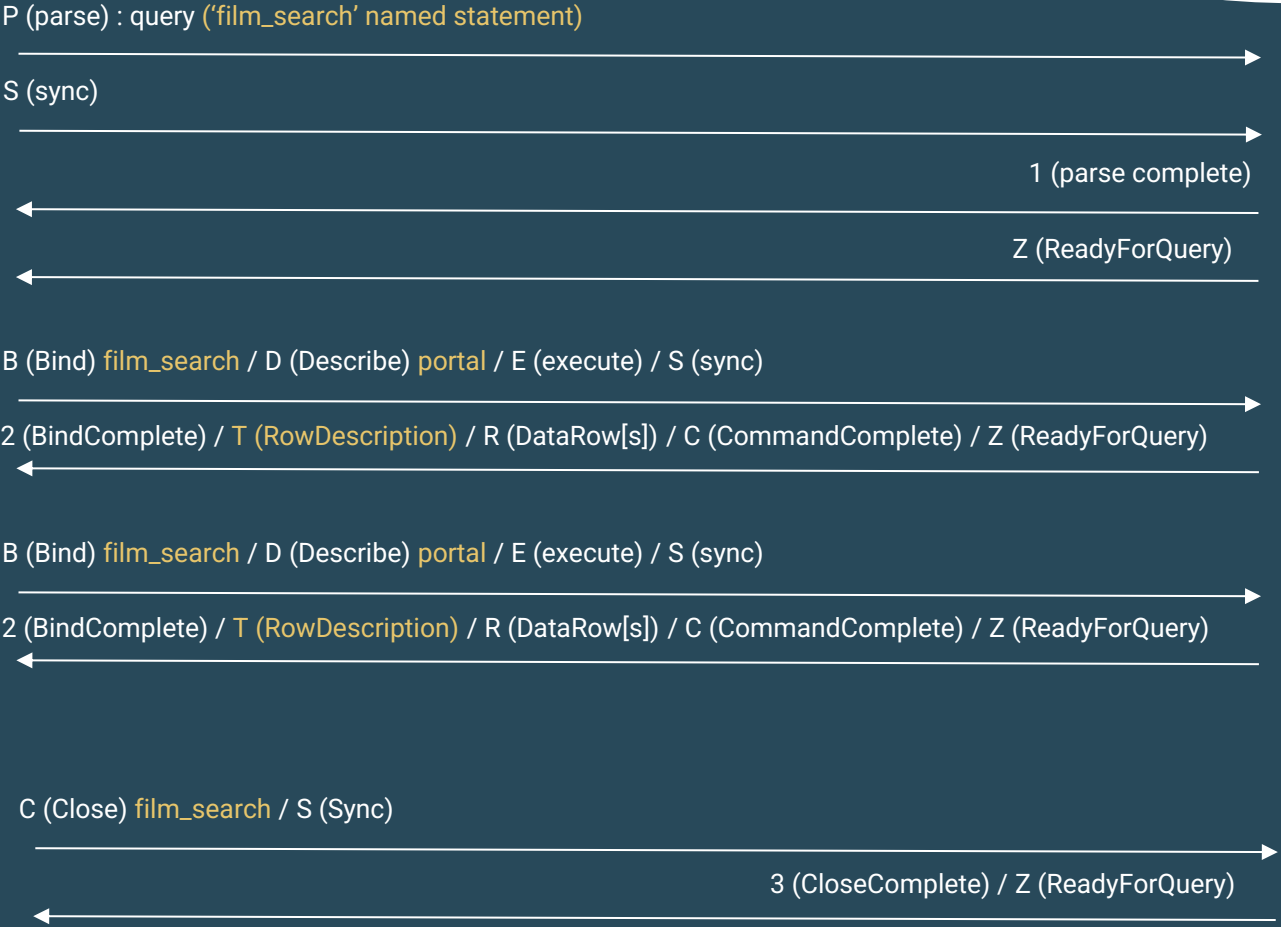
```
\bind_named film_search 'PG-13' 80
EXECUTE film_search;
```

Describe {	0	1	4	5	statement or portal name [...] <small>(empty string for unnamed)</small>
	D	length	S or P		

```
\bind_named film_search 'PG-13' 120
EXECUTE film_search;
```

...

```
\close_prepared film_search
```

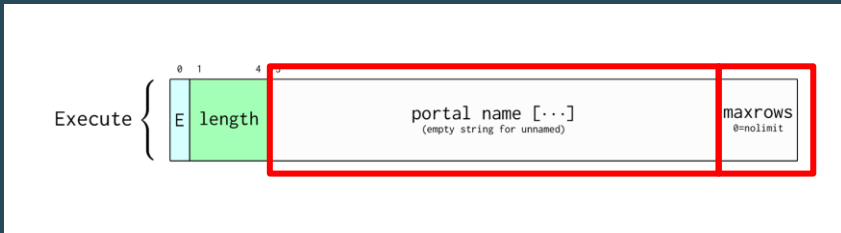


Extended query : cursors/portals

Producer/Consumer with slow consumer

psql client

Postgres backend



Slow processing...



END

P (parse) 'film_search' / D (Describe) S 'film_search' / S (Sync)

1 (ParseComplete) / t (ParameterDescription) / D (RowDescription) / Z

B (Bind) values + film_search + film_portal / E (execute) film_portal 100 rows / S (sync)

2 (BindComplete) / R (100 DataRow[s]) / s (PortalSuspended) / Z (ReadyForQuery)

E (execute) film_portal 100 rows / S (sync)

R (100 DataRow[s]) / s (PortalSuspended) / Z (ReadyForQuery)

R (20 DataRow[s]) / C (CommandComplete)

C (Close) film_portal / S (sync)

3 (CloseComplete) / Z (ReadyForQuery)

220 rows
in total

Extended query / wrap up

Pros

- Prevents SQL injection
- Binary data support
- Plan reuse with prepared statements
- Many ways to improve performance

Cons

- Slower for one-off queries
- Single statement only

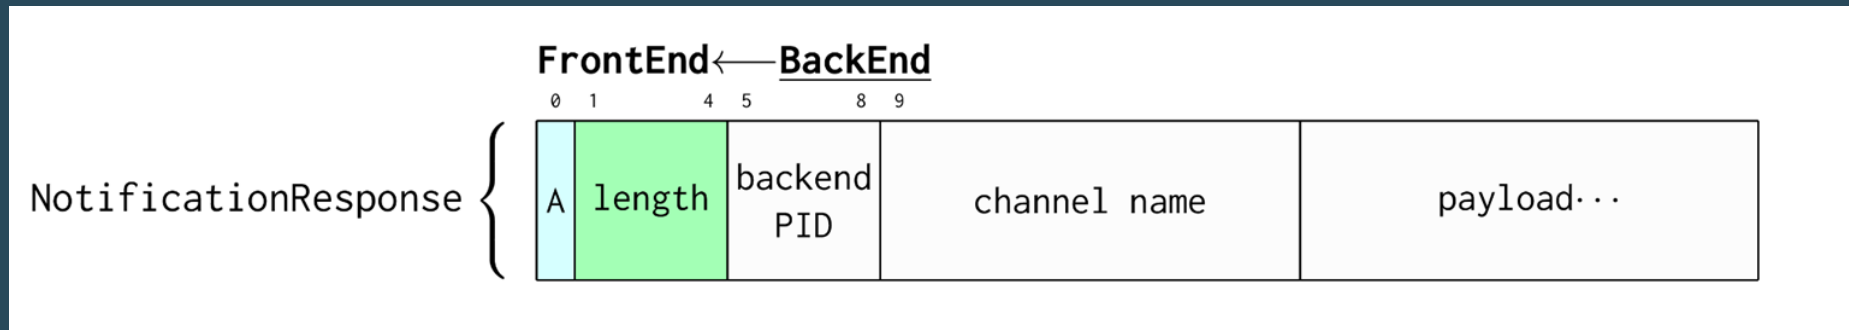
server-pushed messages

Listen / notify

Subscription model


- Any frontend subscribes via simple query *LISTEN* <channel name>
- Any backend can *NOTIFY* <channel name> payload
- Listeners receive a NotificationResponse message when connection is

idle



GUC changes

- Notifies frontends when any GUC sent after Startup changes
- Backend sends *ParameterStatus* message before the next *ReadyForQuery* for each changed GUC
- Needed for drivers to interpret wire data correctly and consistently



copy

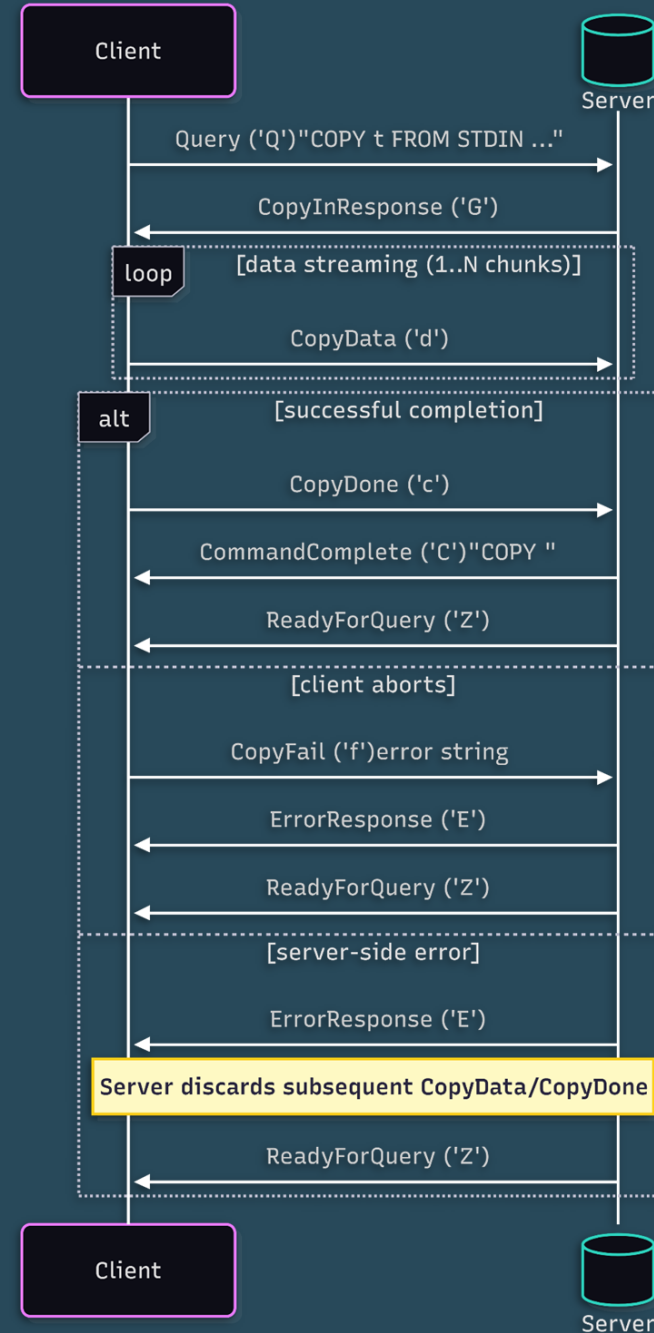
Bulk import/export

Copy is a sub protocol

Formats

- Text (default)
- CSV
- Binary (10-100x faster than INSERTs)

Used by shp2pgsql, pg_dump, pglogical



replication

Replication

Initiated with the **replication** parameter in the Startup packet.
It's a special mode of the protocol with its own grammar
(IDENTIFY_SYSTEM, START_REPLICATION...)

- **replication=true** physical replication
- **replication=<database>** logical replication

Replication

Physical replication works only with same:

- server major version
- Same on-disk binary format

Logical replication is “cross platform”

- Does not replicate DDL (functions, procedures, catalog)

Replication

Great talks from PG Conf EU 2023

“The journey towards active-active replication in PostgreSQL” by

Jonathan S. Katz

<https://www.postgresql.eu/events/pgconfeu2023/schedule/session/4783-the-journey-towards-active-active-replication-in-postgresql/>



performance

Performance considerations : connections

Connections are expensive (5 to 10MB, 5-50ms to open)

- use pooling (implemented in some drivers, or with pgBouncer, pgPool, ...)

Performance considerations : queries

Repetitive parameterized queries benefit from prepared statements

- use binary format
- 5-executions_trap: make sure data is well distributed with param variation; Postgres may switch to a generic plan whose selectivity assumptions are wrong for skewed data

Performance considerations : quick-wins

Avoid N+1 problem (fetch N rows, then perform a query per row)

- Postgres will do it better than you!

Avoid unnecessary roundtrips

- Use batches (INSERT, arrays)

Trust Postgres, trust the drivers, measure, and read the docs!

diagnosing


libpq has PQtrace()

If you use/wrap libpq, you can activate protocol tracing (improved in PG14)

```
PQtrace(conn, trace_file);|
void PQtrace(PGconn *conn, FILE *debug_port)
=== in fe-trace.c ===
Describe
```

```
F      33      Query      "select * from actor limit 10"
B      120     RowDescription 4 "actor_id" 1467366 1 23 4 -1 0 "first_name" 1467366 2 1043 65535 49
0 "last_name" 1467366 3 1043 65535 49 0 "last_update" 1467366 4 1114 8 -1 0
B      60      DataRow      4 1 '1' 8 'Penelope' 7 'Guinness' 22 '2013-05-26 14:47:57.62'
B      57      DataRow      4 1 '2' 4 'Nick' 8 'Wahlberg' 22 '2013-05-26 14:47:57.62'
B      52      DataRow      4 1 '3' 2 'Ed' 5 'Chase' 22 '2013-05-26 14:47:57.62'
B      58      DataRow      4 1 '4' 8 'Jennifer' 5 'Davis' 22 '2013-05-26 14:47:57.62'
B      63      DataRow      4 1 '5' 6 'Johnny' 12 'Lollobrigida' 22 '2013-05-26 14:47:57.62'
B      59      DataRow      4 1 '6' 5 'Bette' 9 'Nicholson' 22 '2013-05-26 14:47:57.62'
B      56      DataRow      4 1 '7' 5 'Grace' 6 'Mostel' 22 '2013-05-26 14:47:57.62'
B      61      DataRow      4 1 '8' 7 'Matthew' 9 'Johansson' 22 '2013-05-26 14:47:57.62'
B      53      DataRow      4 1 '9' 3 'Joe' 5 'Swank' 22 '2013-05-26 14:47:57.62'
B      60      DataRow      4 2 '10' 9 'Christian' 5 'Gable' 22 '2013-05-26 14:47:57.62'
B      14      CommandComplete "SELECT 10"
B      5      ReadyForQuery I
F      59      Parse      "actor_by_id" "select * from actor where actor_id = $" 0
F      4      Sync
B      4      ParseComplete
B      5      ReadyForQuery I
F      30      Bind      "" "actor_by_id" 0 1 1 '1' 1 1
F      6      Describe  P ""
F      9      Execute   "" 0
F      4      Sync
B      4      BindComplete
B      120     RowDescription 4 "actor_id" 1467366 1 23 4 -1 1 "first_name" 1467366 2 1043 65535 49
1 "last_name" 1467366 3 1043 65535 49 1 "last_update" 1467366 4 1114 8 -1 1
B      49      DataRow      4 4 '\x00\x00\x00\x01' 8 'Penelope' 7 'Guinness' 8
'\x00\x01\x80\xe2\xf57 '
```

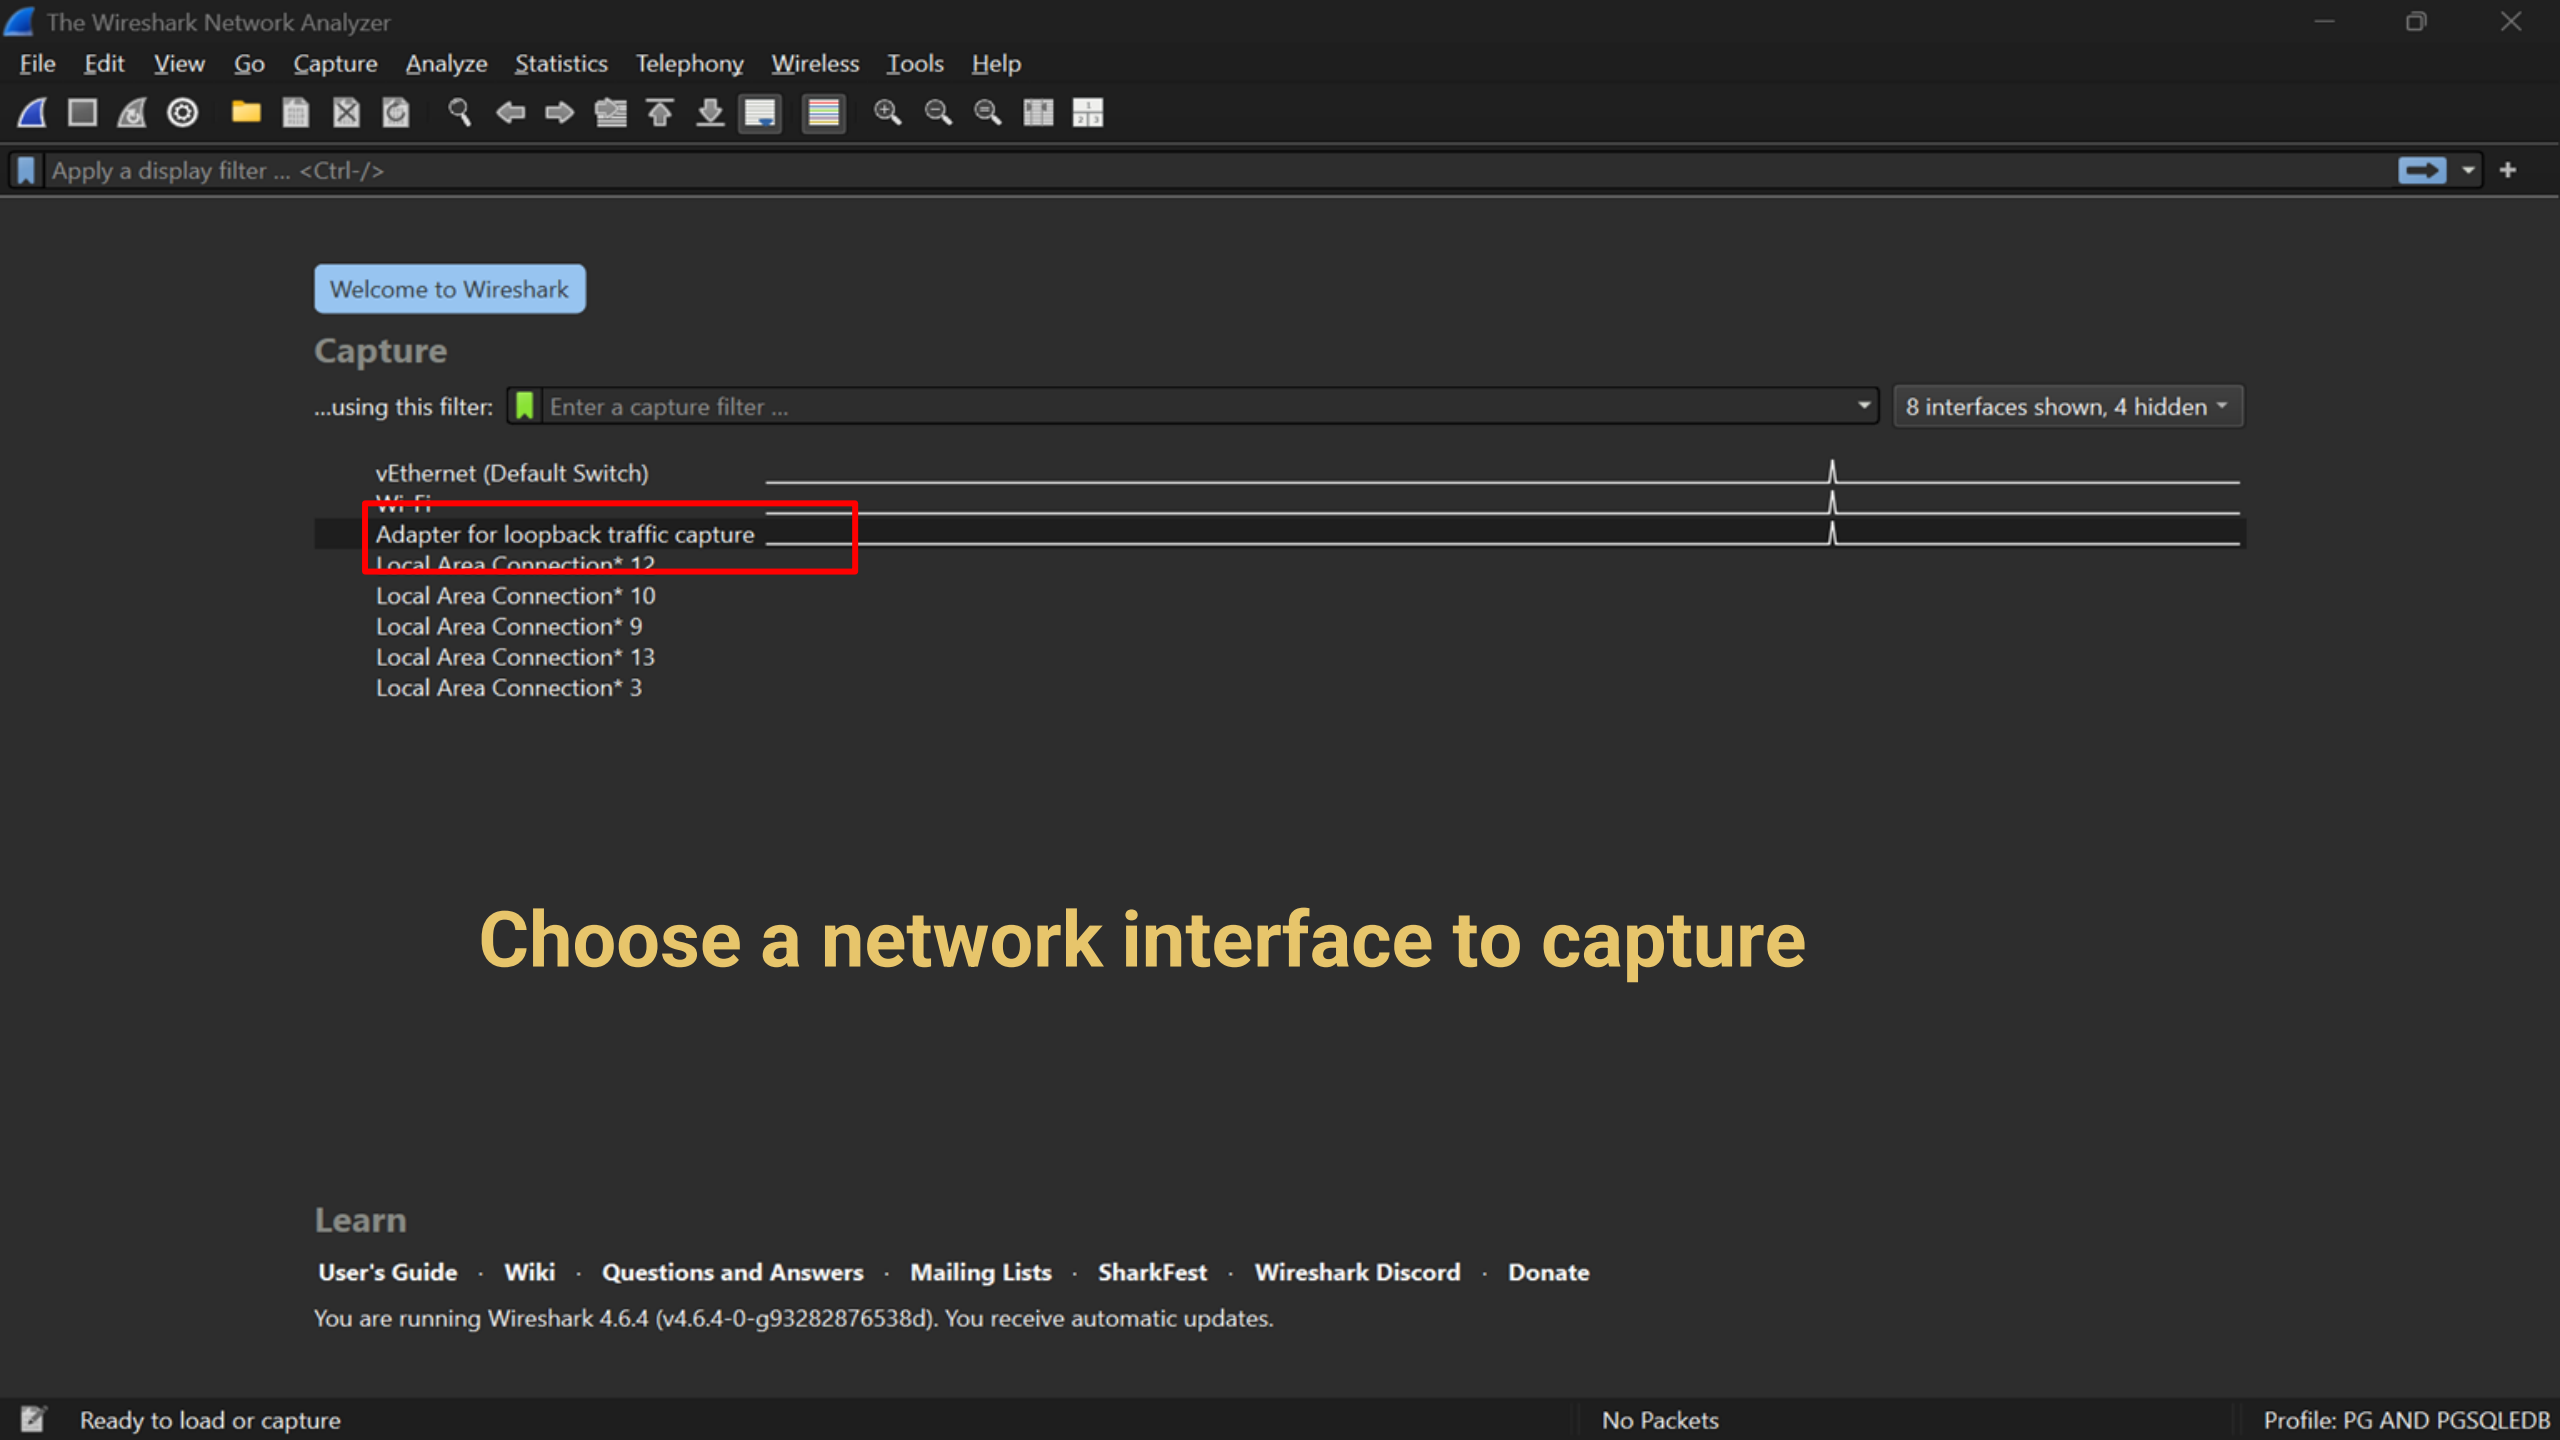
diagnosing issues with wireshark



wireshark allows to trace network messages

wireshark dissectors allows presenting the messages in a **understandable way** for a given **protocol**





Choose a network interface to capture

Learn

[User's Guide](#) · [Wiki](#) · [Questions and Answers](#) · [Mailing Lists](#) · [SharkFest](#) · [Wireshark Discord](#) · [Donate](#)

You are running Wireshark 4.6.4 (v4.6.4-0-g93282876538d). You receive automatic updates.

here comes




pg_protoexport

an open source tool do
document the protocol


```
> pg_protoexport demo
```

TEX

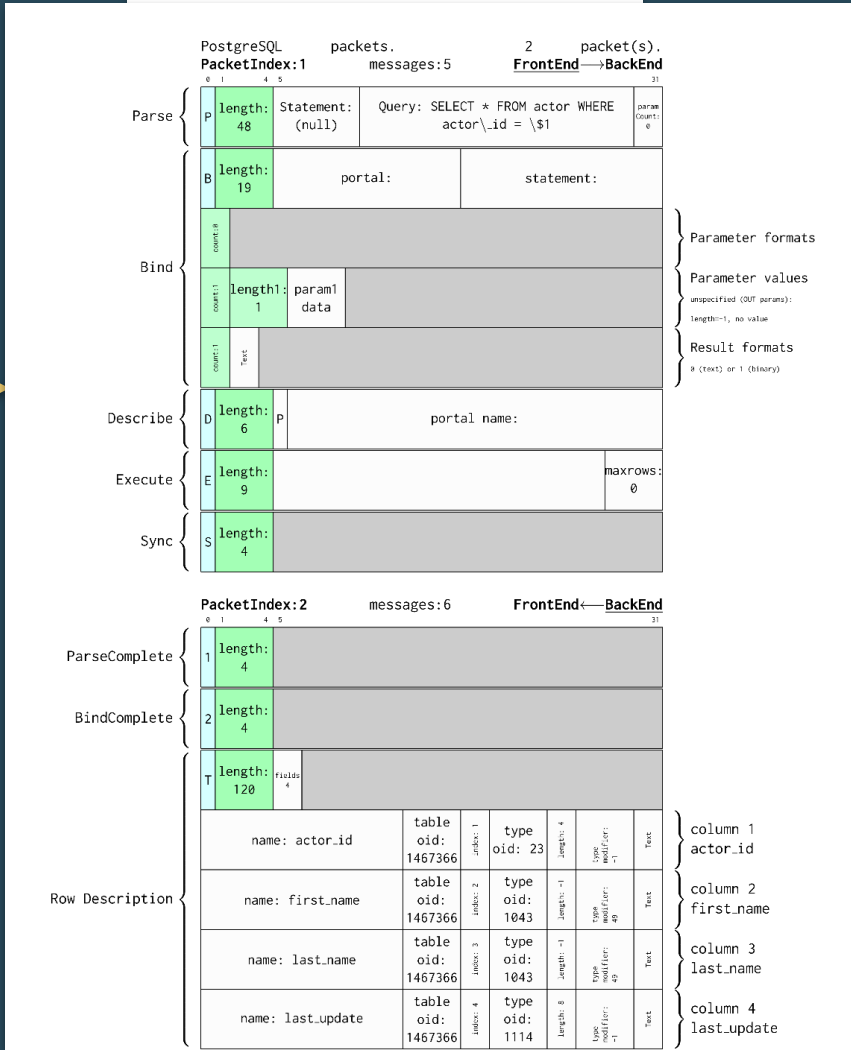


Wireshark

capture file

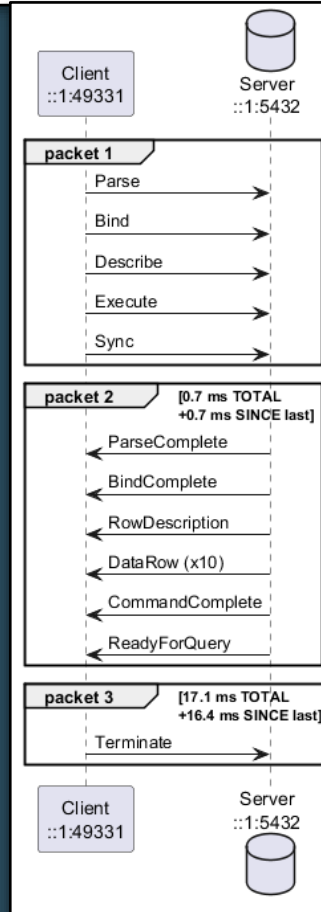
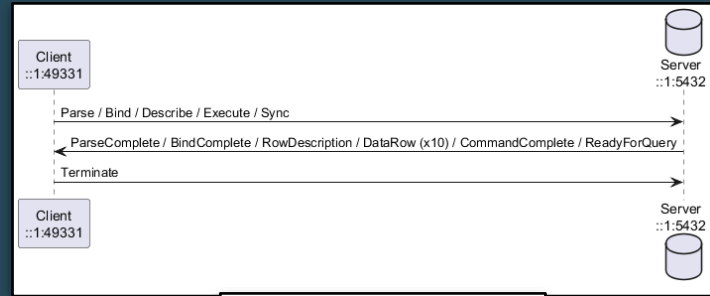
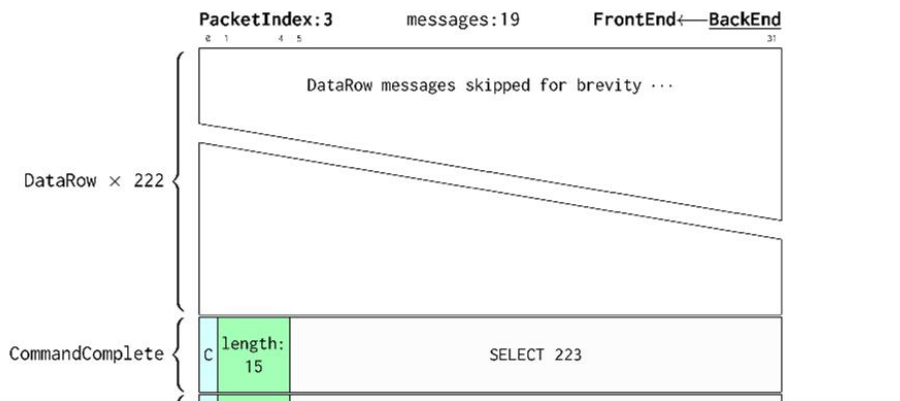
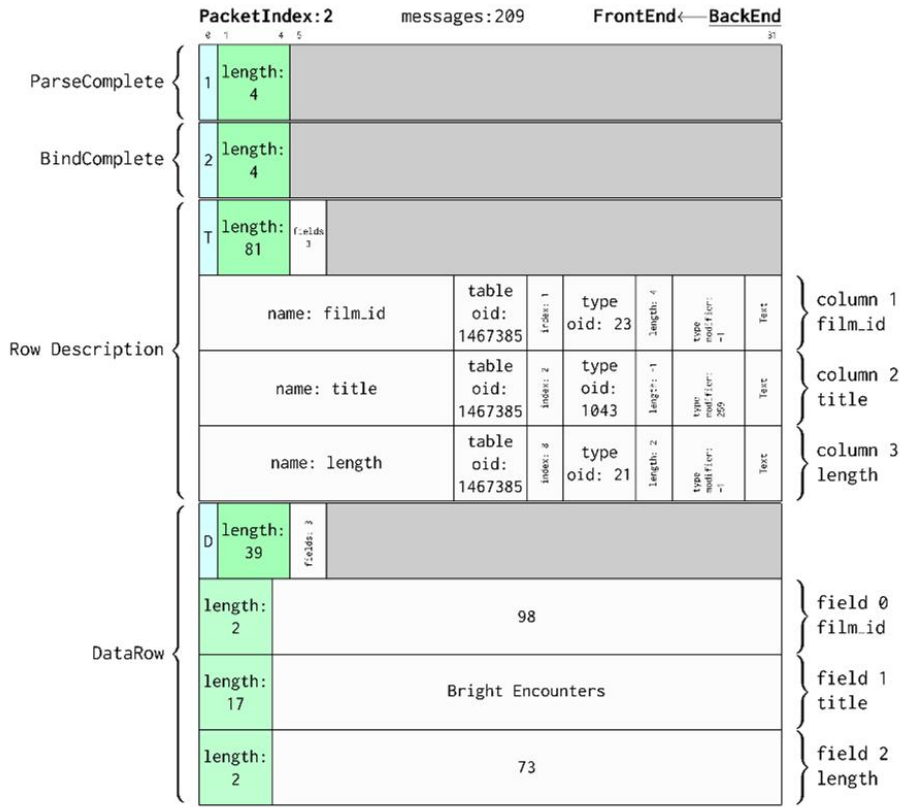


pg_protoexport



LaTeX

Mermaid/PlantUML ASCII




```


1 Packet (F)
2 2025-03-13T06:36:30.1436940Z F 83 Parse statement=, query=select oid, typename, typtype from
pg_type where typtype <> $1 limit 10;, parameter oids=[25]
3 2025-03-13T06:36:30.1436940Z F 19 Bind : formats: [], values len: [1], results format: [1]
4 2025-03-13T06:36:30.1436940Z F 6 Describe Portal: ''
5 2025-03-13T06:36:30.1436940Z F 9 Execute portal='', maxrows=0
6 2025-03-13T06:36:30.1436940Z F 4 Sync SyncMessage
7 Packet (B)
8 2025-03-13T06:36:30.1444240Z B 4 ParseComplete ParseCompleteMessage
9 2025-03-13T06:36:30.1444240Z B 4 BindComplete
10 2025-03-13T06:36:30.1444240Z B 80 RowDescription [oid: 1247, 1, 26, 4, -1, 1, typename: 1247, 2, 19,
64, -1, 1, typtype: 1247, 7, 18, 1, -1, 1]
11 2025-03-13T06:36:30.1444240Z B 30 DataRow [oid:00000047, typename:pg_type, typtype:c]
12 2025-03-13T06:36:30.1444240Z B 35 DataRow [oid:0000004b, typename:pg_attribute, typtype:c]
13 2025-03-13T06:36:30.1444240Z B 30 DataRow [oid:00000051, typename:pg_proc, typtype:c]
14 2025-03-13T06:36:30.1444240Z B 31 DataRow [oid:00000053, typename:pg_class, typtype:c]
15 2025-03-13T06:36:30.1444240Z B 37 DataRow [oid:00000020, typename:pg_ddl_command, typtype:p]
16 2025-03-13T06:36:30.1444240Z B 30 DataRow [oid:000002c1, typename:unknown, typtype:p]
17 2025-03-13T06:36:30.1444240Z B 32 DataRow [oid:00000f40, typename:int4range, typtype:r]
18 2025-03-13T06:36:30.1444240Z B 31 DataRow [oid:00000f42, typename:numrange, typtype:r]
19 2025-03-13T06:36:30.1444240Z B 30 DataRow [oid:00000f44, typename:tsrange, typtype:r]
20 2025-03-13T06:36:30.1444240Z B 32 DataRow [oid:00000f46, typename:tszrange, typtype:r]
21 2025-03-13T06:36:30.1444240Z B 14 CommandComplete SELECT 10
22 2025-03-13T06:36:30.1444240Z B 5 ReadyForQuery Idle
23 Packet (F)
24 2025-03-13T06:36:30.1607910Z F 4 Terminate TerminateMessage
25
  
```

whatever
(like, even PostgreSQL)





Wireshark




capture file



pg_protoexport





psql



whatever

(like, even PostgreSQL or spreadsheet)






Wireshark



pg_protoexport




psql



whatever

PostgreSQL





resha








pg_protoexport

https://github.com/xfischer/pg_protoexport/



Thank you!

Merci vielmal !



xavier.fischer@enterprisedb.com



[@edbpostgres.bsky.social](https://bsky.app/profile/edbpostgres.bsky.social)



[fischerxavier](https://www.linkedin.com/in/fischerxavier)

