

Logical replication Theory and Concepts

by: Ashutosh Bapat

Major contributor, PostgreSQL

Principal SDE, Community Contributors team @Microsoft

Agenda

Physical Replication

Logical Replication

- WAL Sender

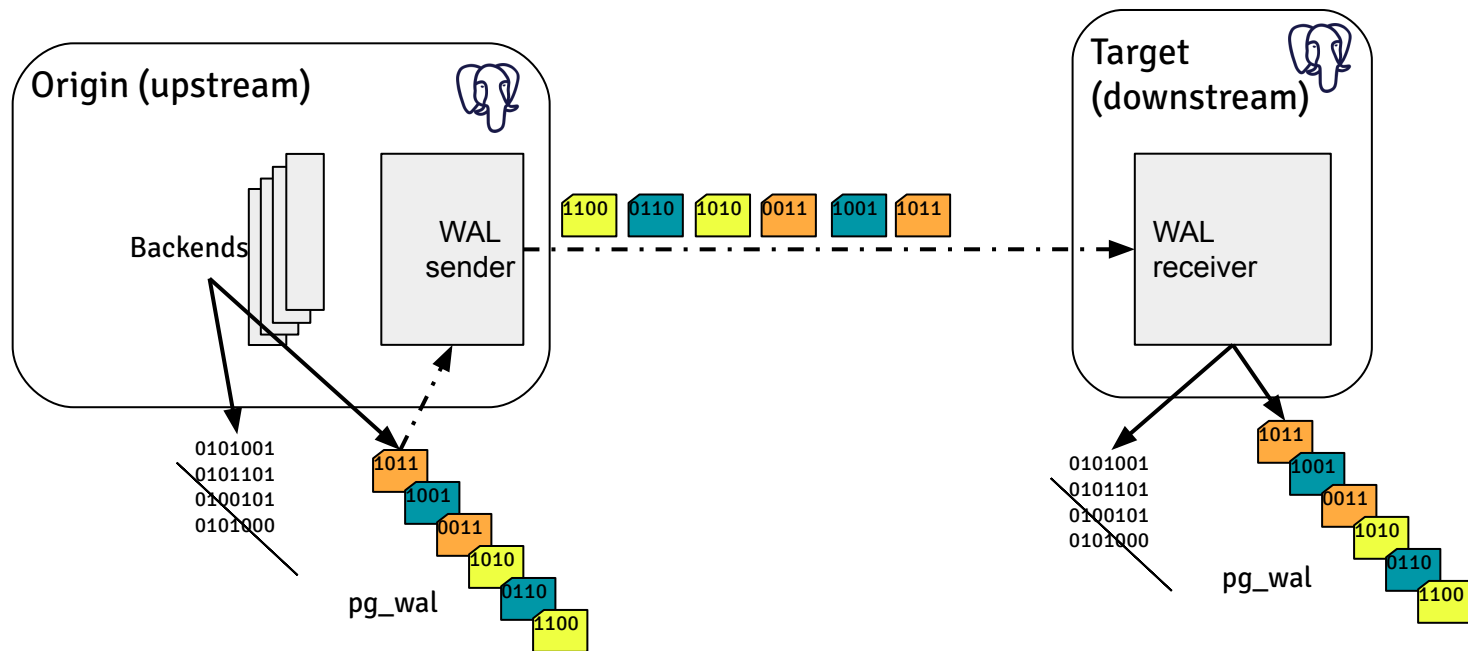
- Logical receiver

- Replication slot

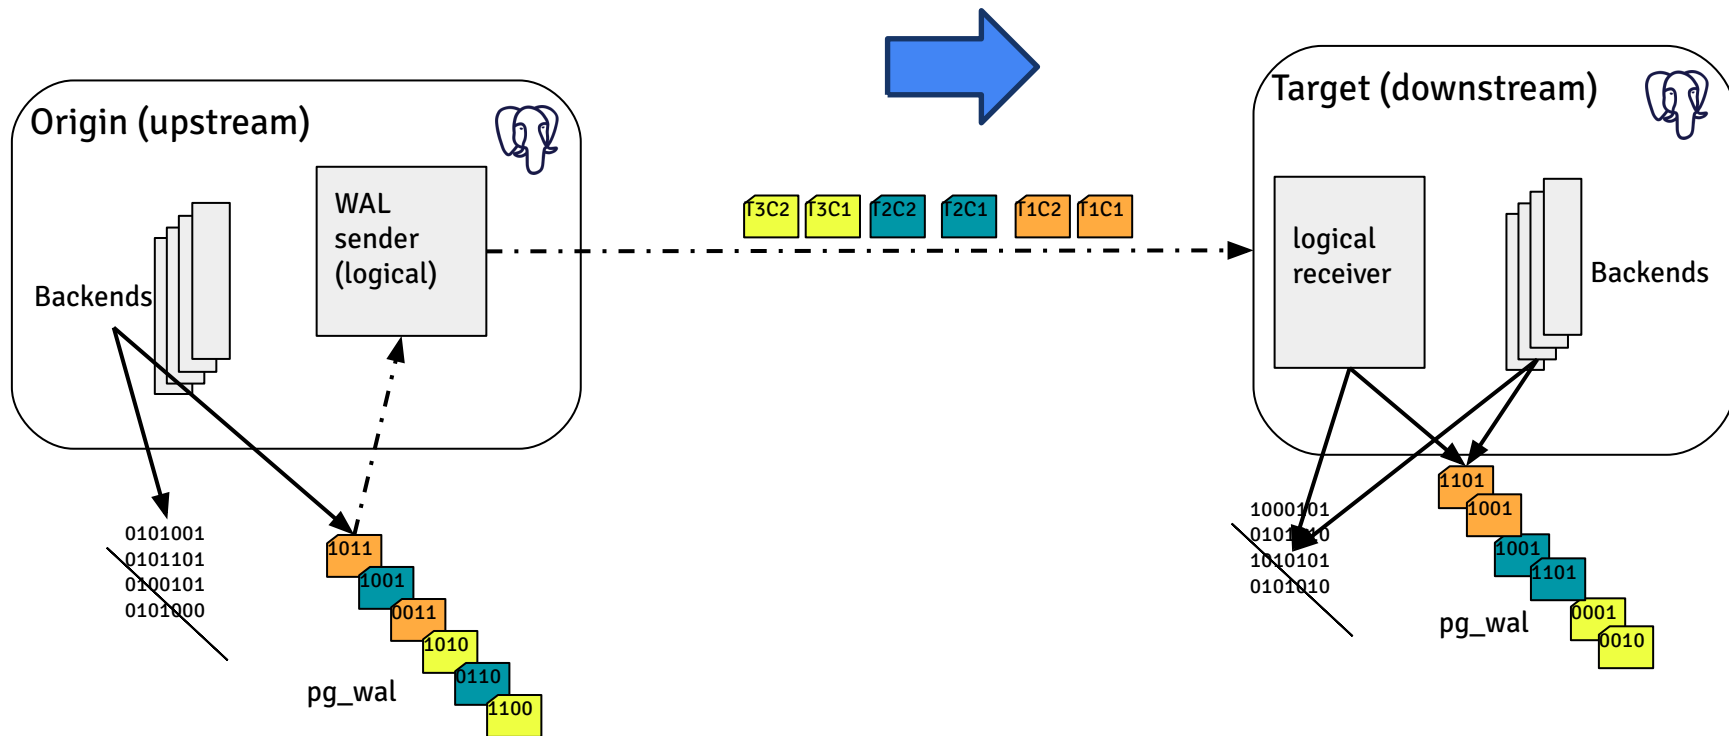
- Error loops

Conclusion

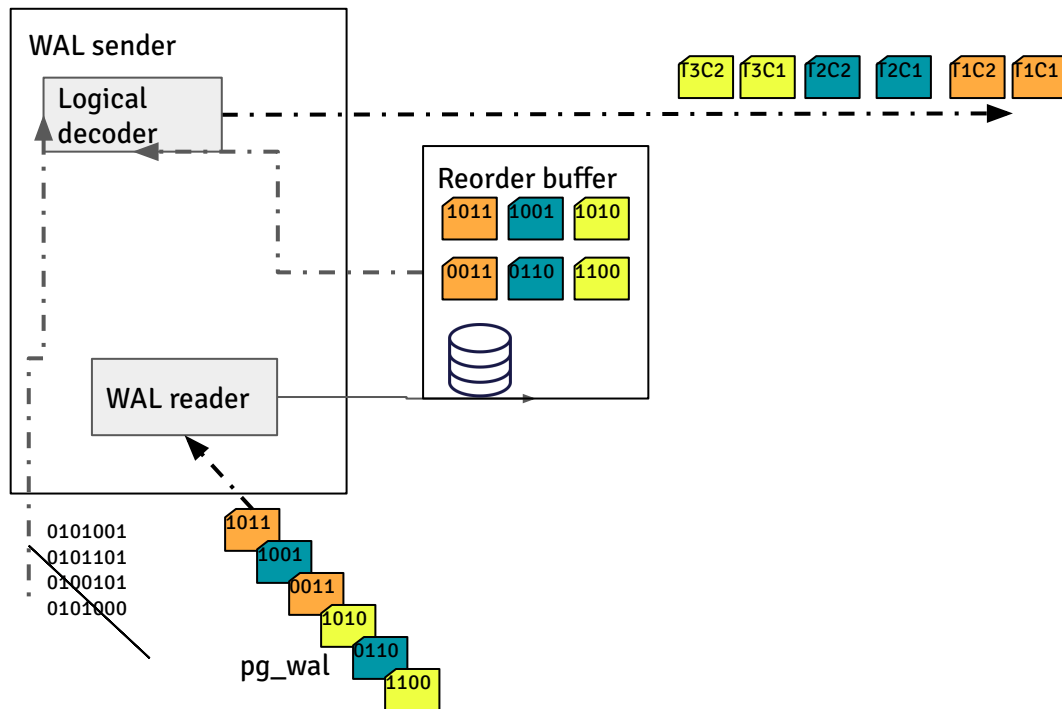
Physical/Streaming replication



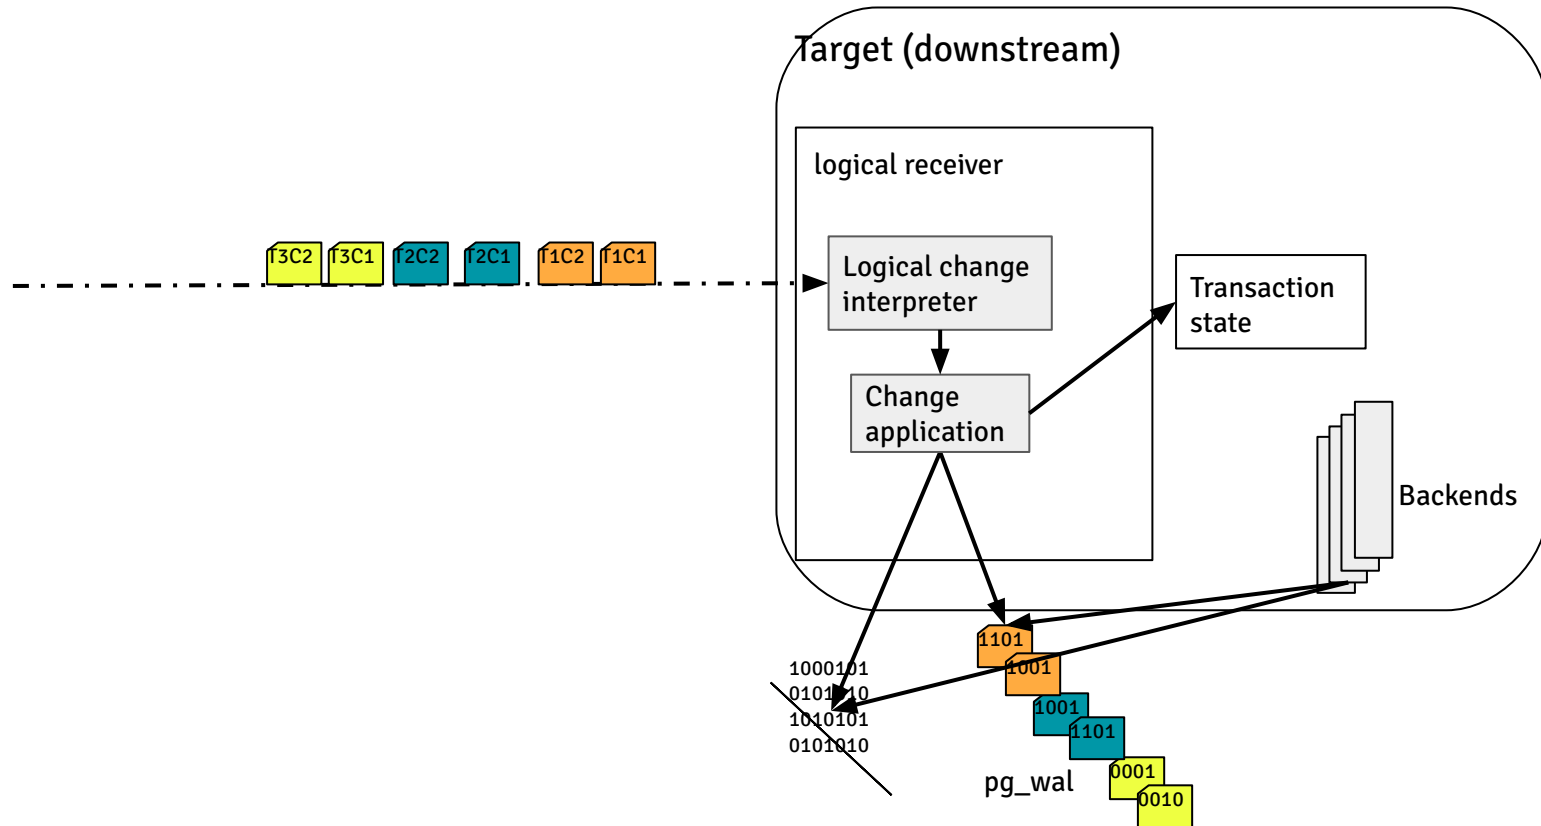
Logical replication



WAL sender: logical change sender



Logical receiver



Replication Slot

Tracks progress of replication

- Transactions sent downstream

- Transactions received and acknowledged by downstream

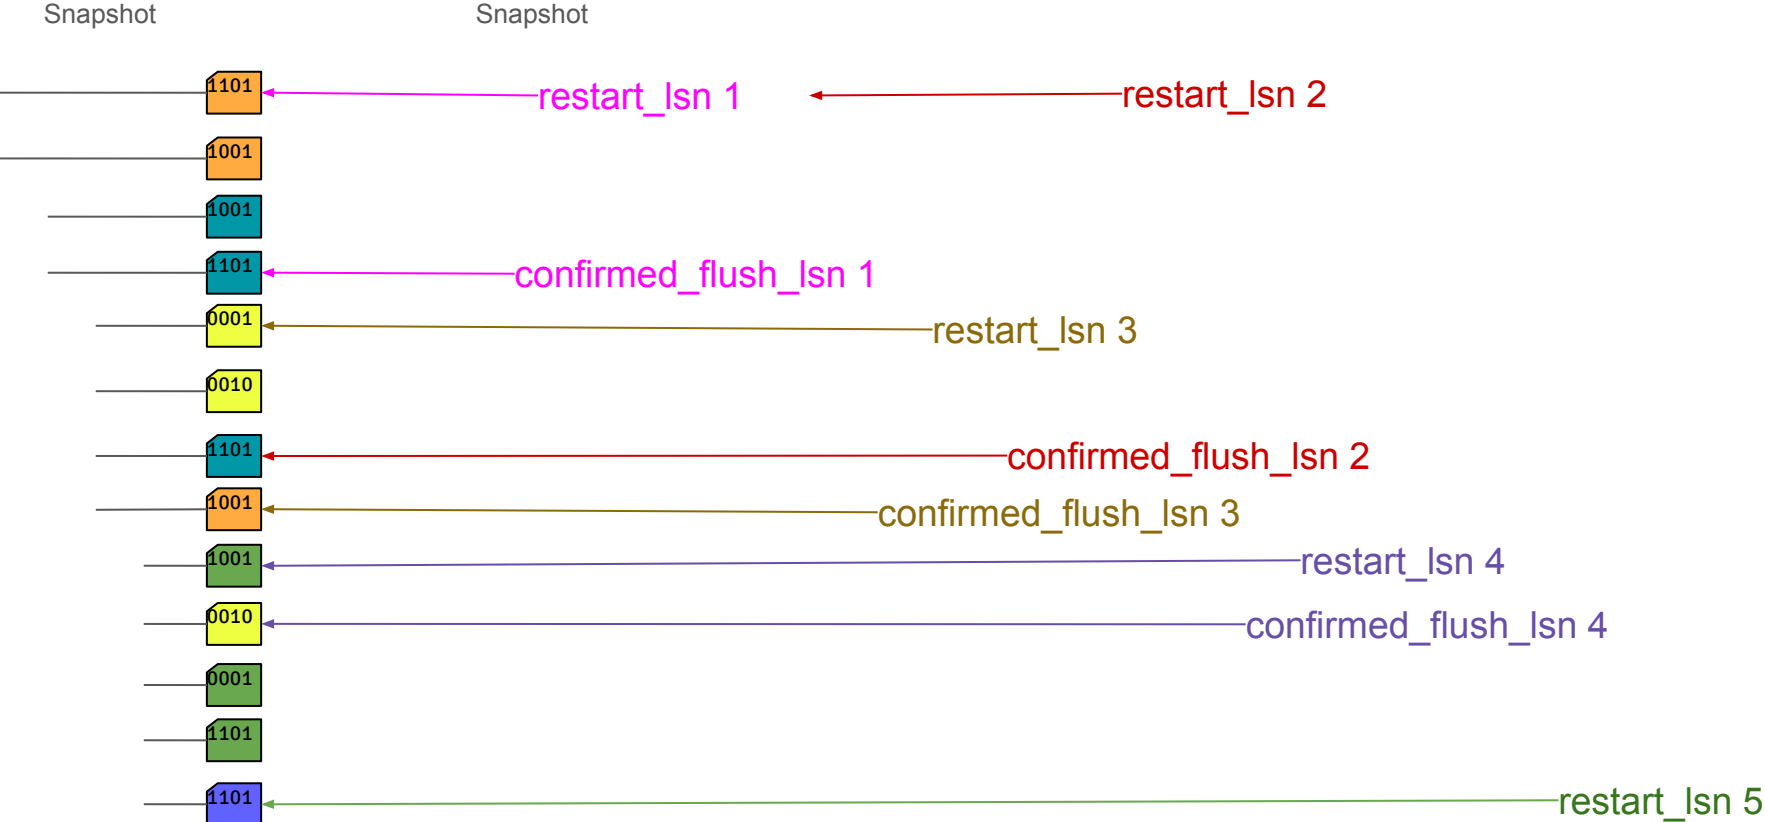
Holds back resources

- WAL

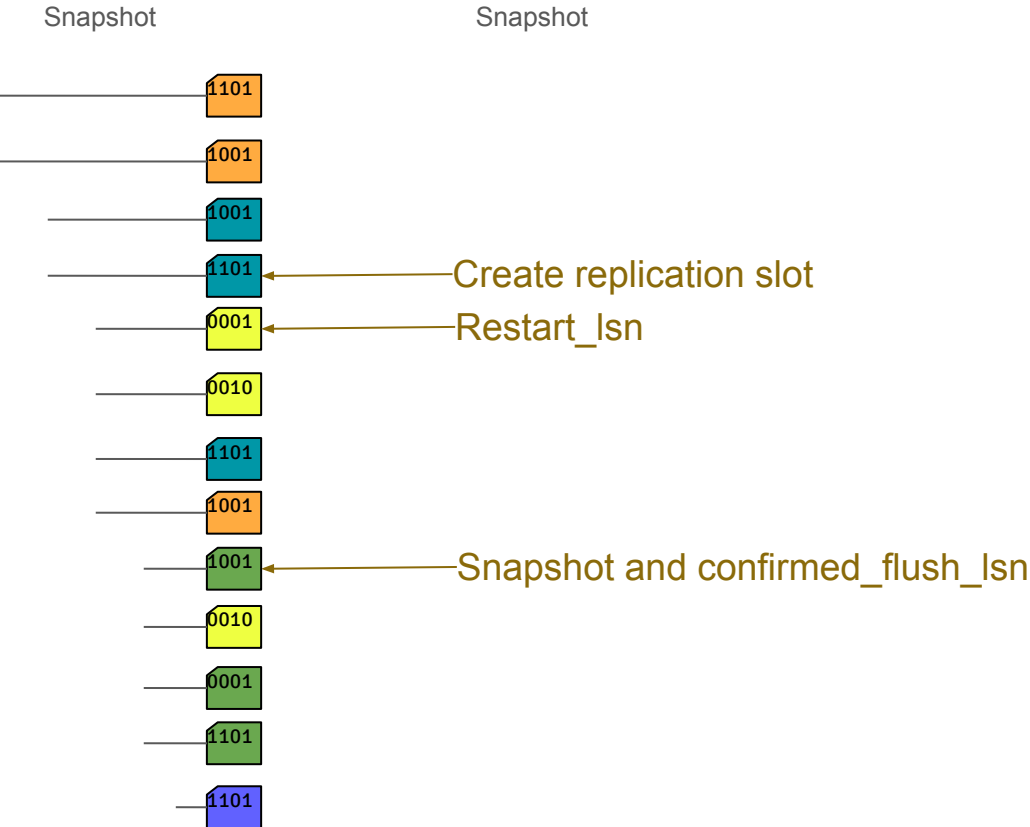
- Catalog tuples

- Min transaction ID

Replication slots: WAL and snapshot

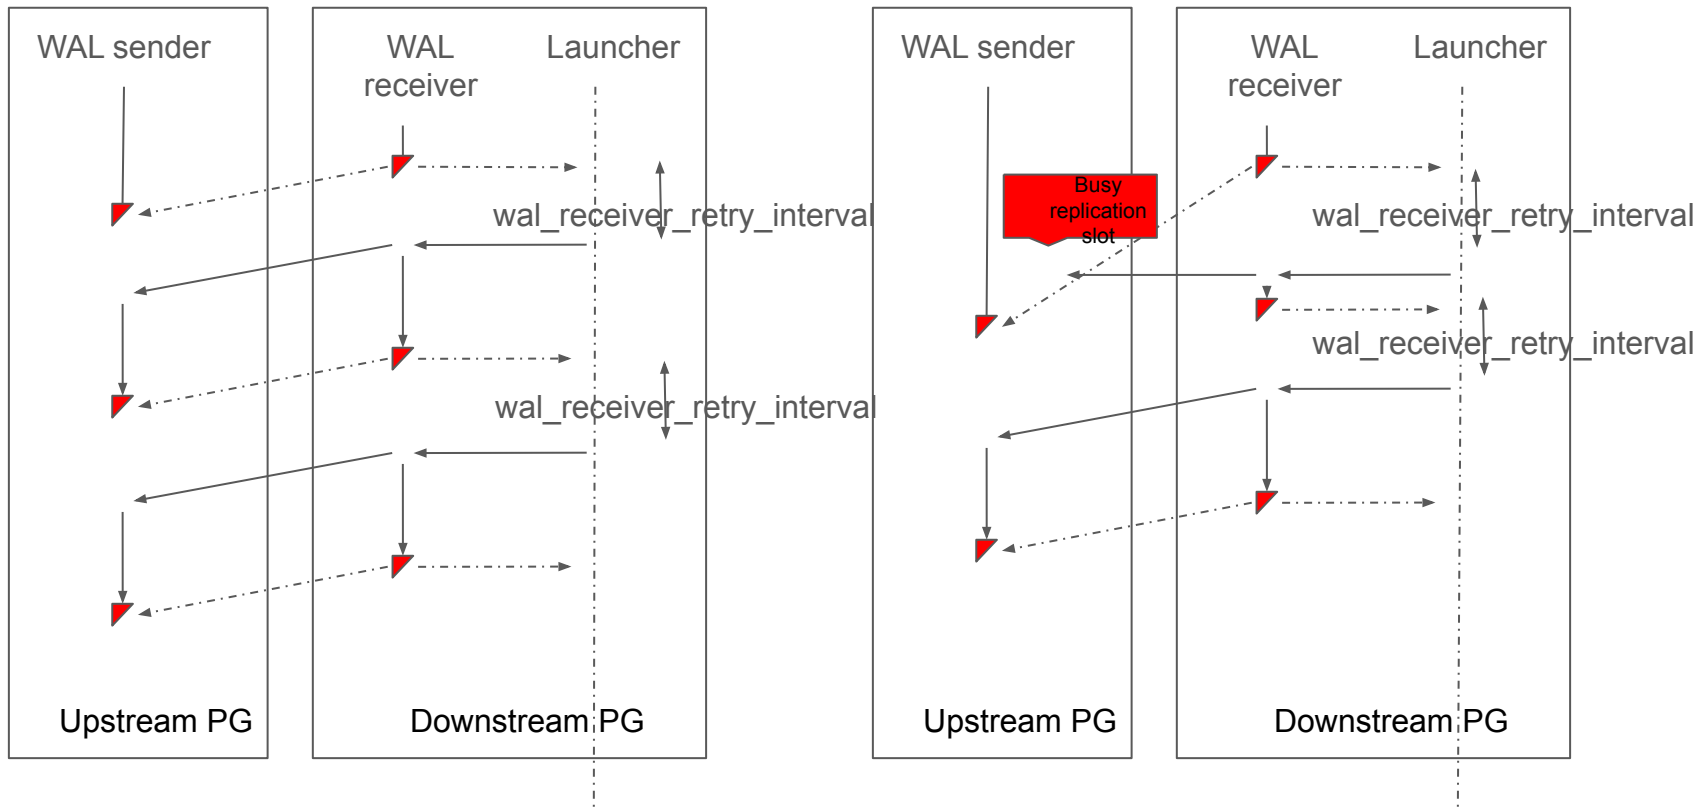


Initial Data synchronization

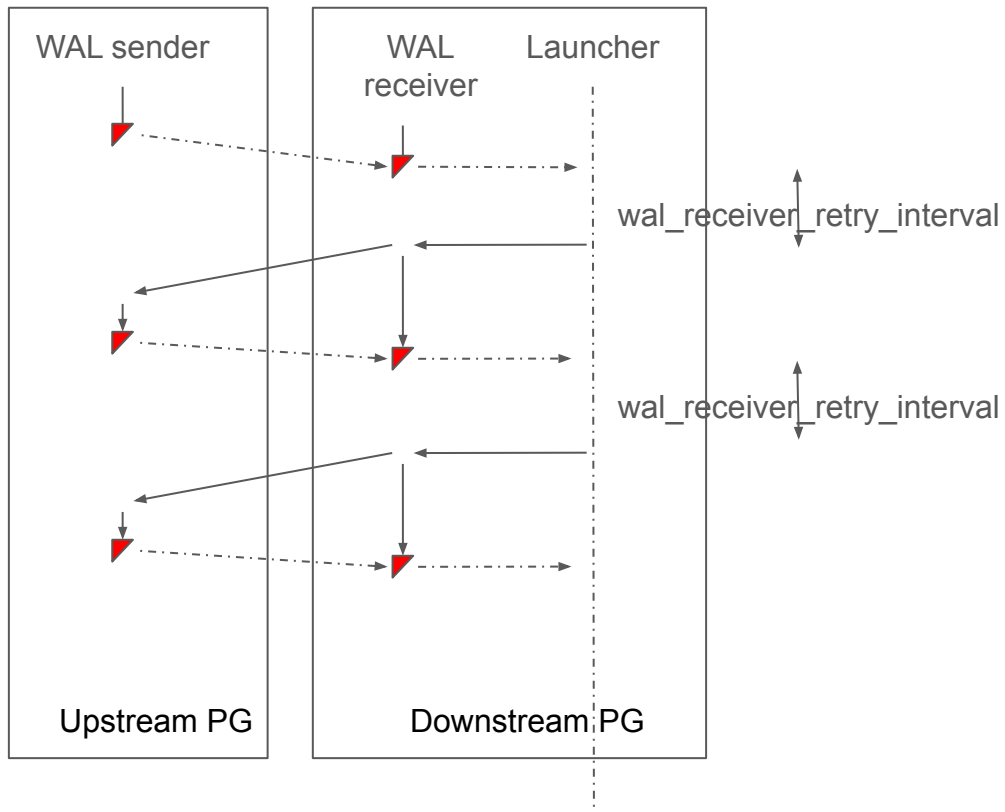


When things go wrong

Receiver error loop



Sender error loop



Conclusion

Use cases

- Cross-version, cross-platform, selective replication

- Heterogenous replication

- Replication to other DBMSes, data lakes

Limitations

- Built-in Bi-directional support

- DDL replication

Things to watch

- Broken replication

- WAL buildup, disk space

- Transaction Wraparound