




# YDB: Open-Source Distributed SQL Database



# YDB — Open-Source Distributed SQL Database

## Relational

- ACID OLTP transactions
- YQL — SQL Dialect

## Consistency

- Strongly consistent
- CAP-theorem, prefer CP
- Serializable transaction isolation level

Horizontal scalability  
across thousands  
of servers

# YDB — Open-Source Distributed SQL Database

## Highly available

- Runs in multiple Availability Zones (AZ)
- Survives AZ plus rack failure w/o human intervention, available for read/write

## Mission critical database

- Works for projects with 24x7 requirements
- Zero Downtime Deployments

## Platform

- Tables, topics, distributed queries, ...

# Open-Source

## Apache 2.0 License

- Platform
- SDK
- Docs







<https://github.com/ydb-platform/ydb>

# Physical Architecture

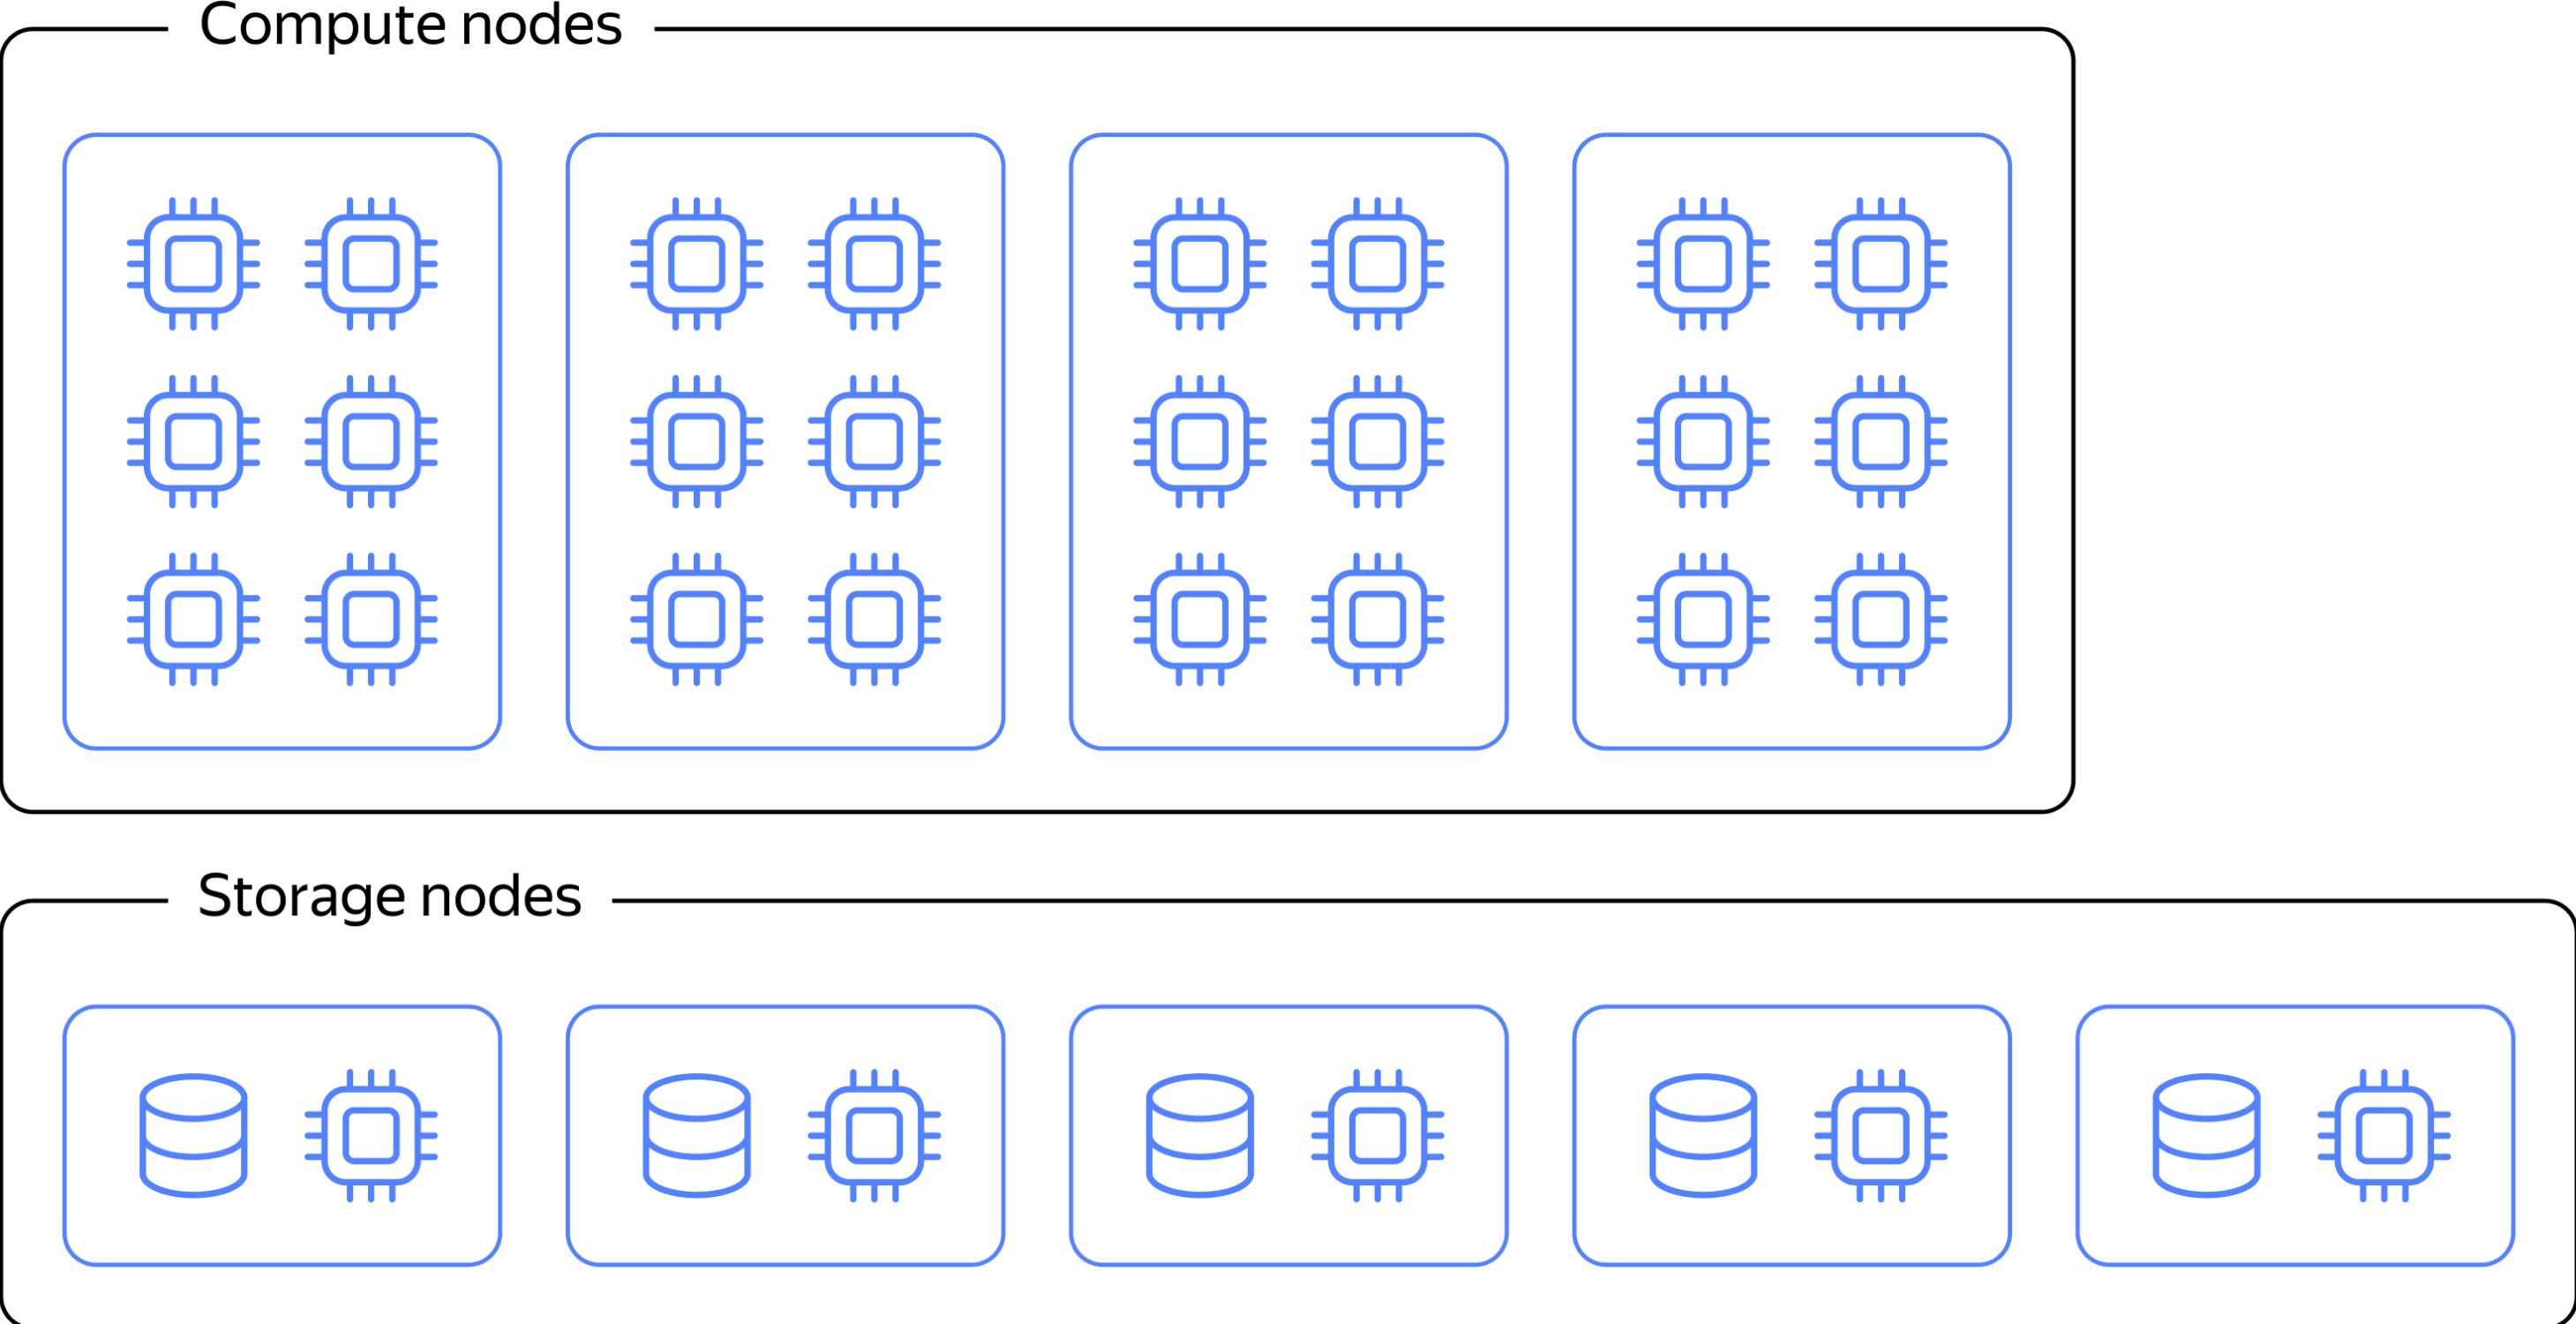
★ Shared nothing architecture

★ Commodity hardware



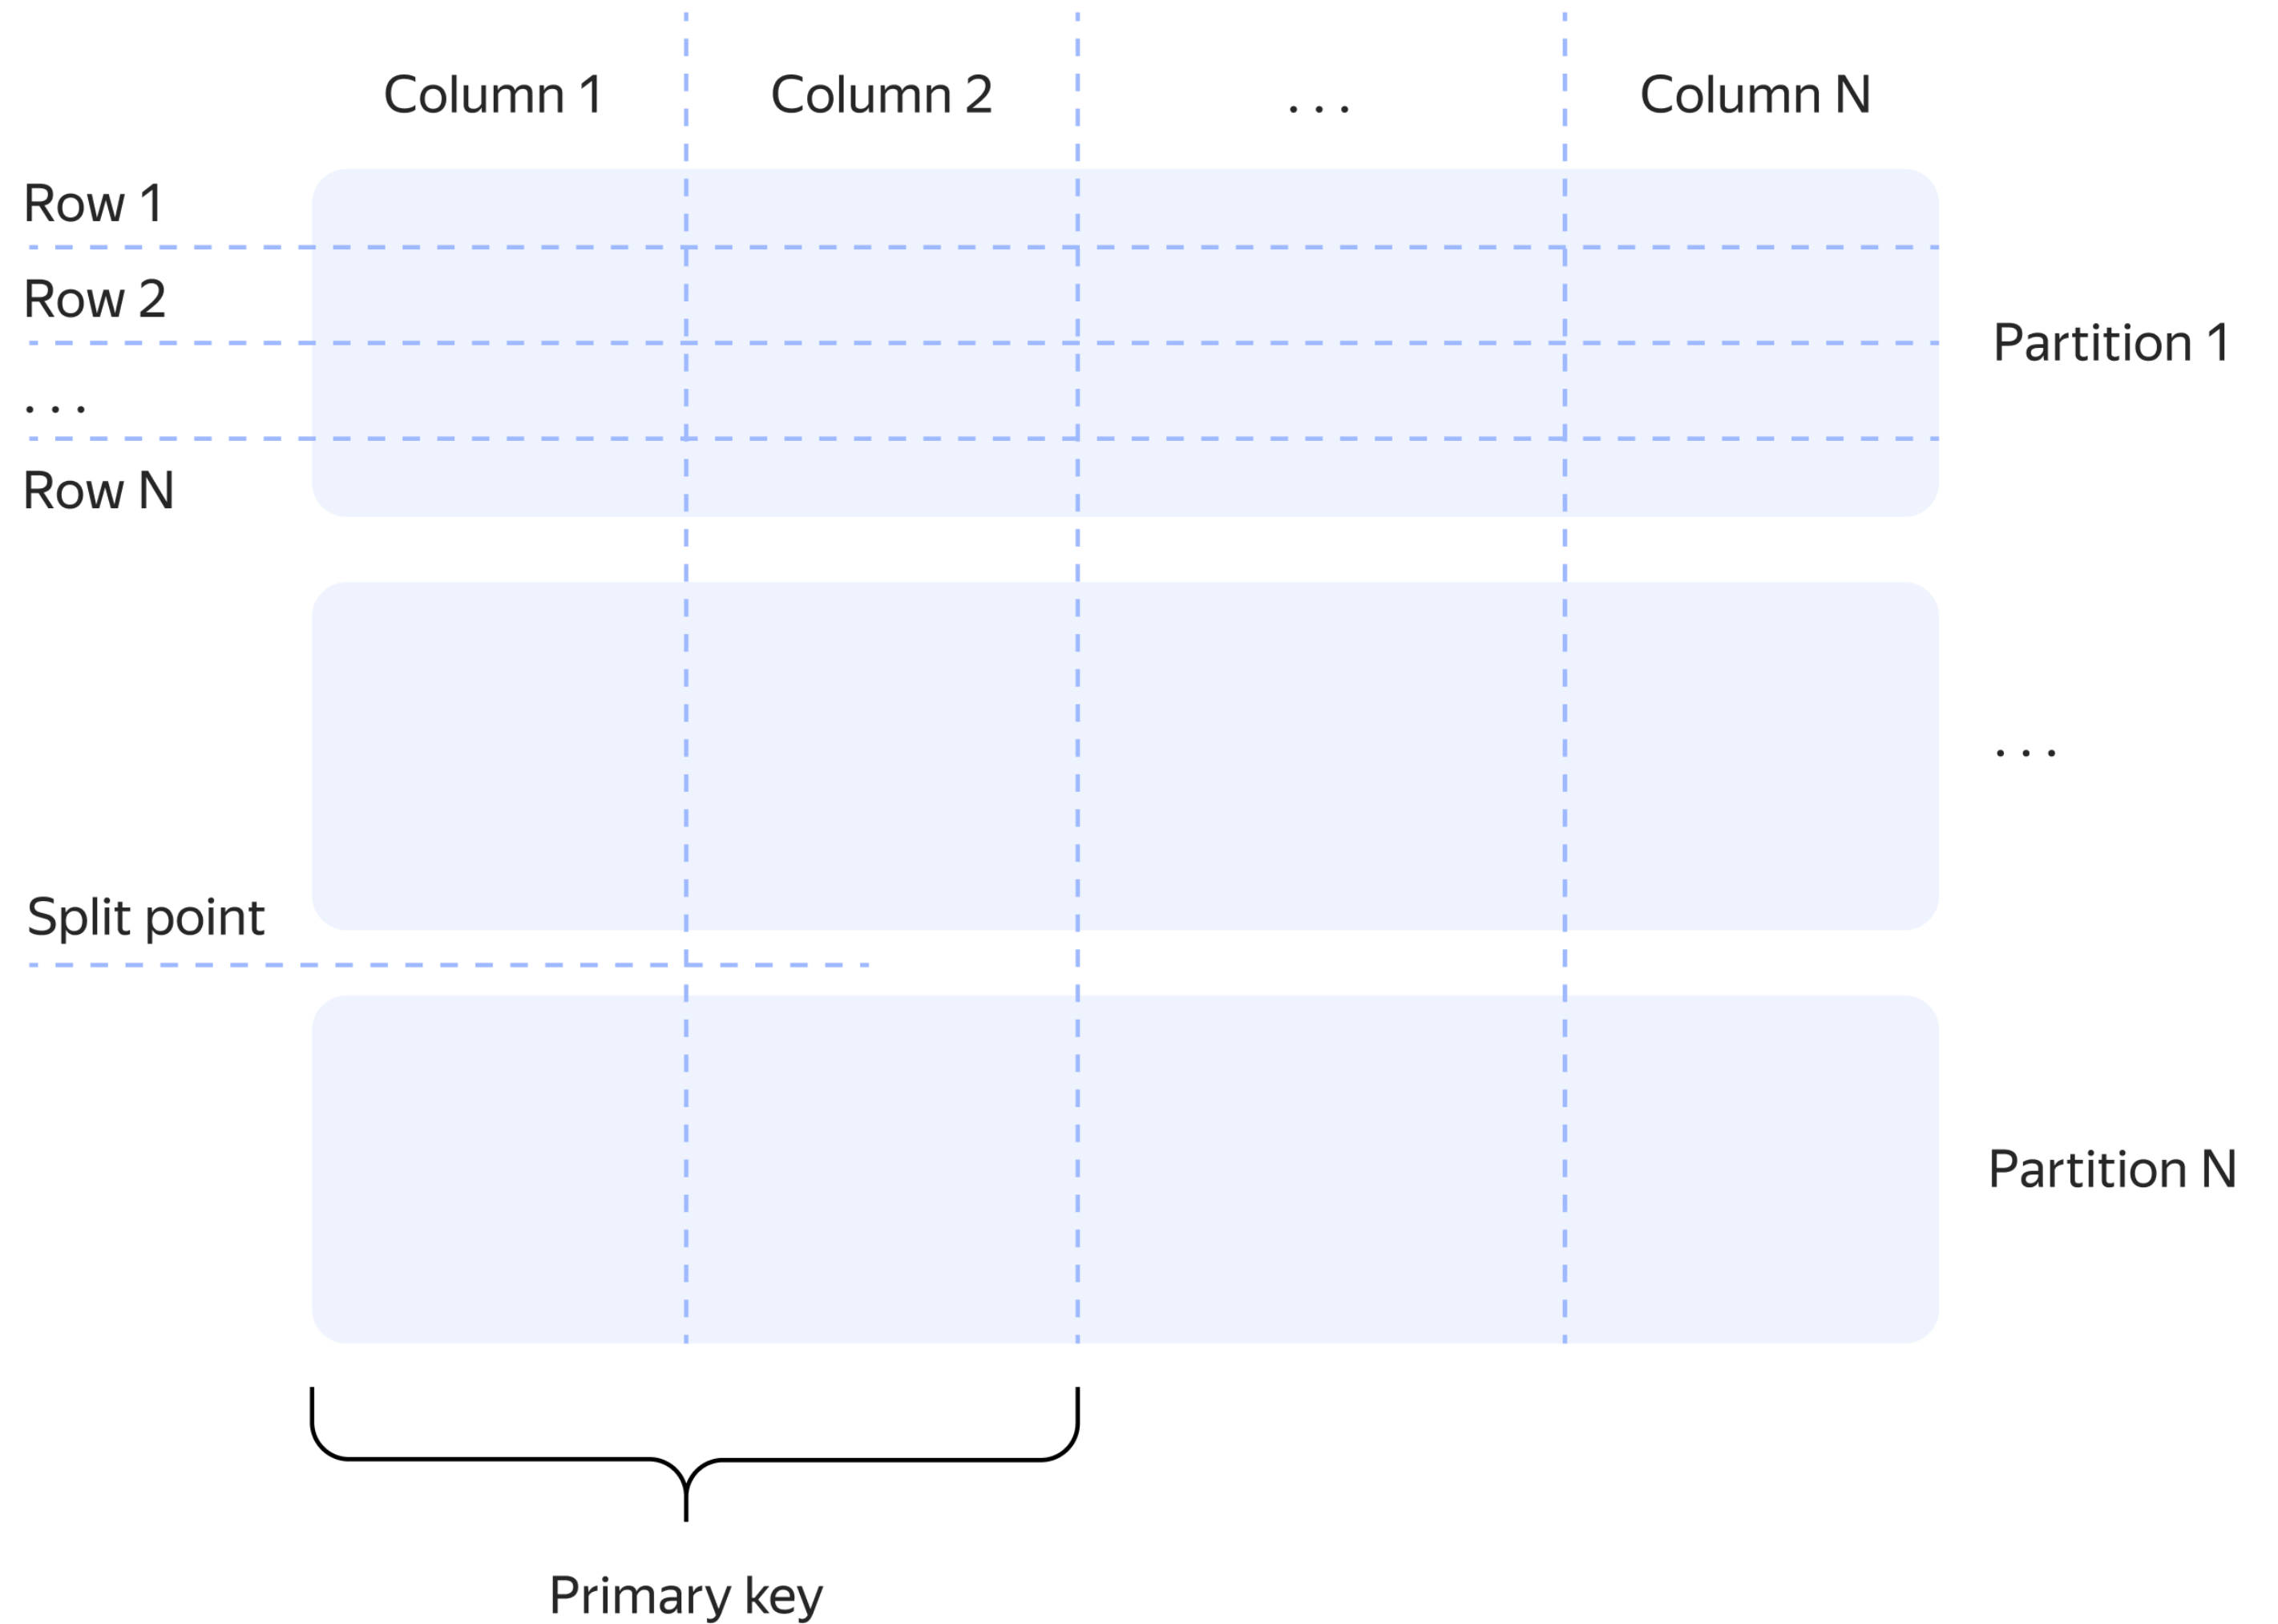
# Separated Compute and Storage

★ Compute and storage are separated, can grow independently




# Table

- Tables have a primary key (PK)
- Tables are sorted by PK
- Tables can grow up to petabytes of data
- Tables are automatically partitioned



# Table Partitions Autosplit and Balancing

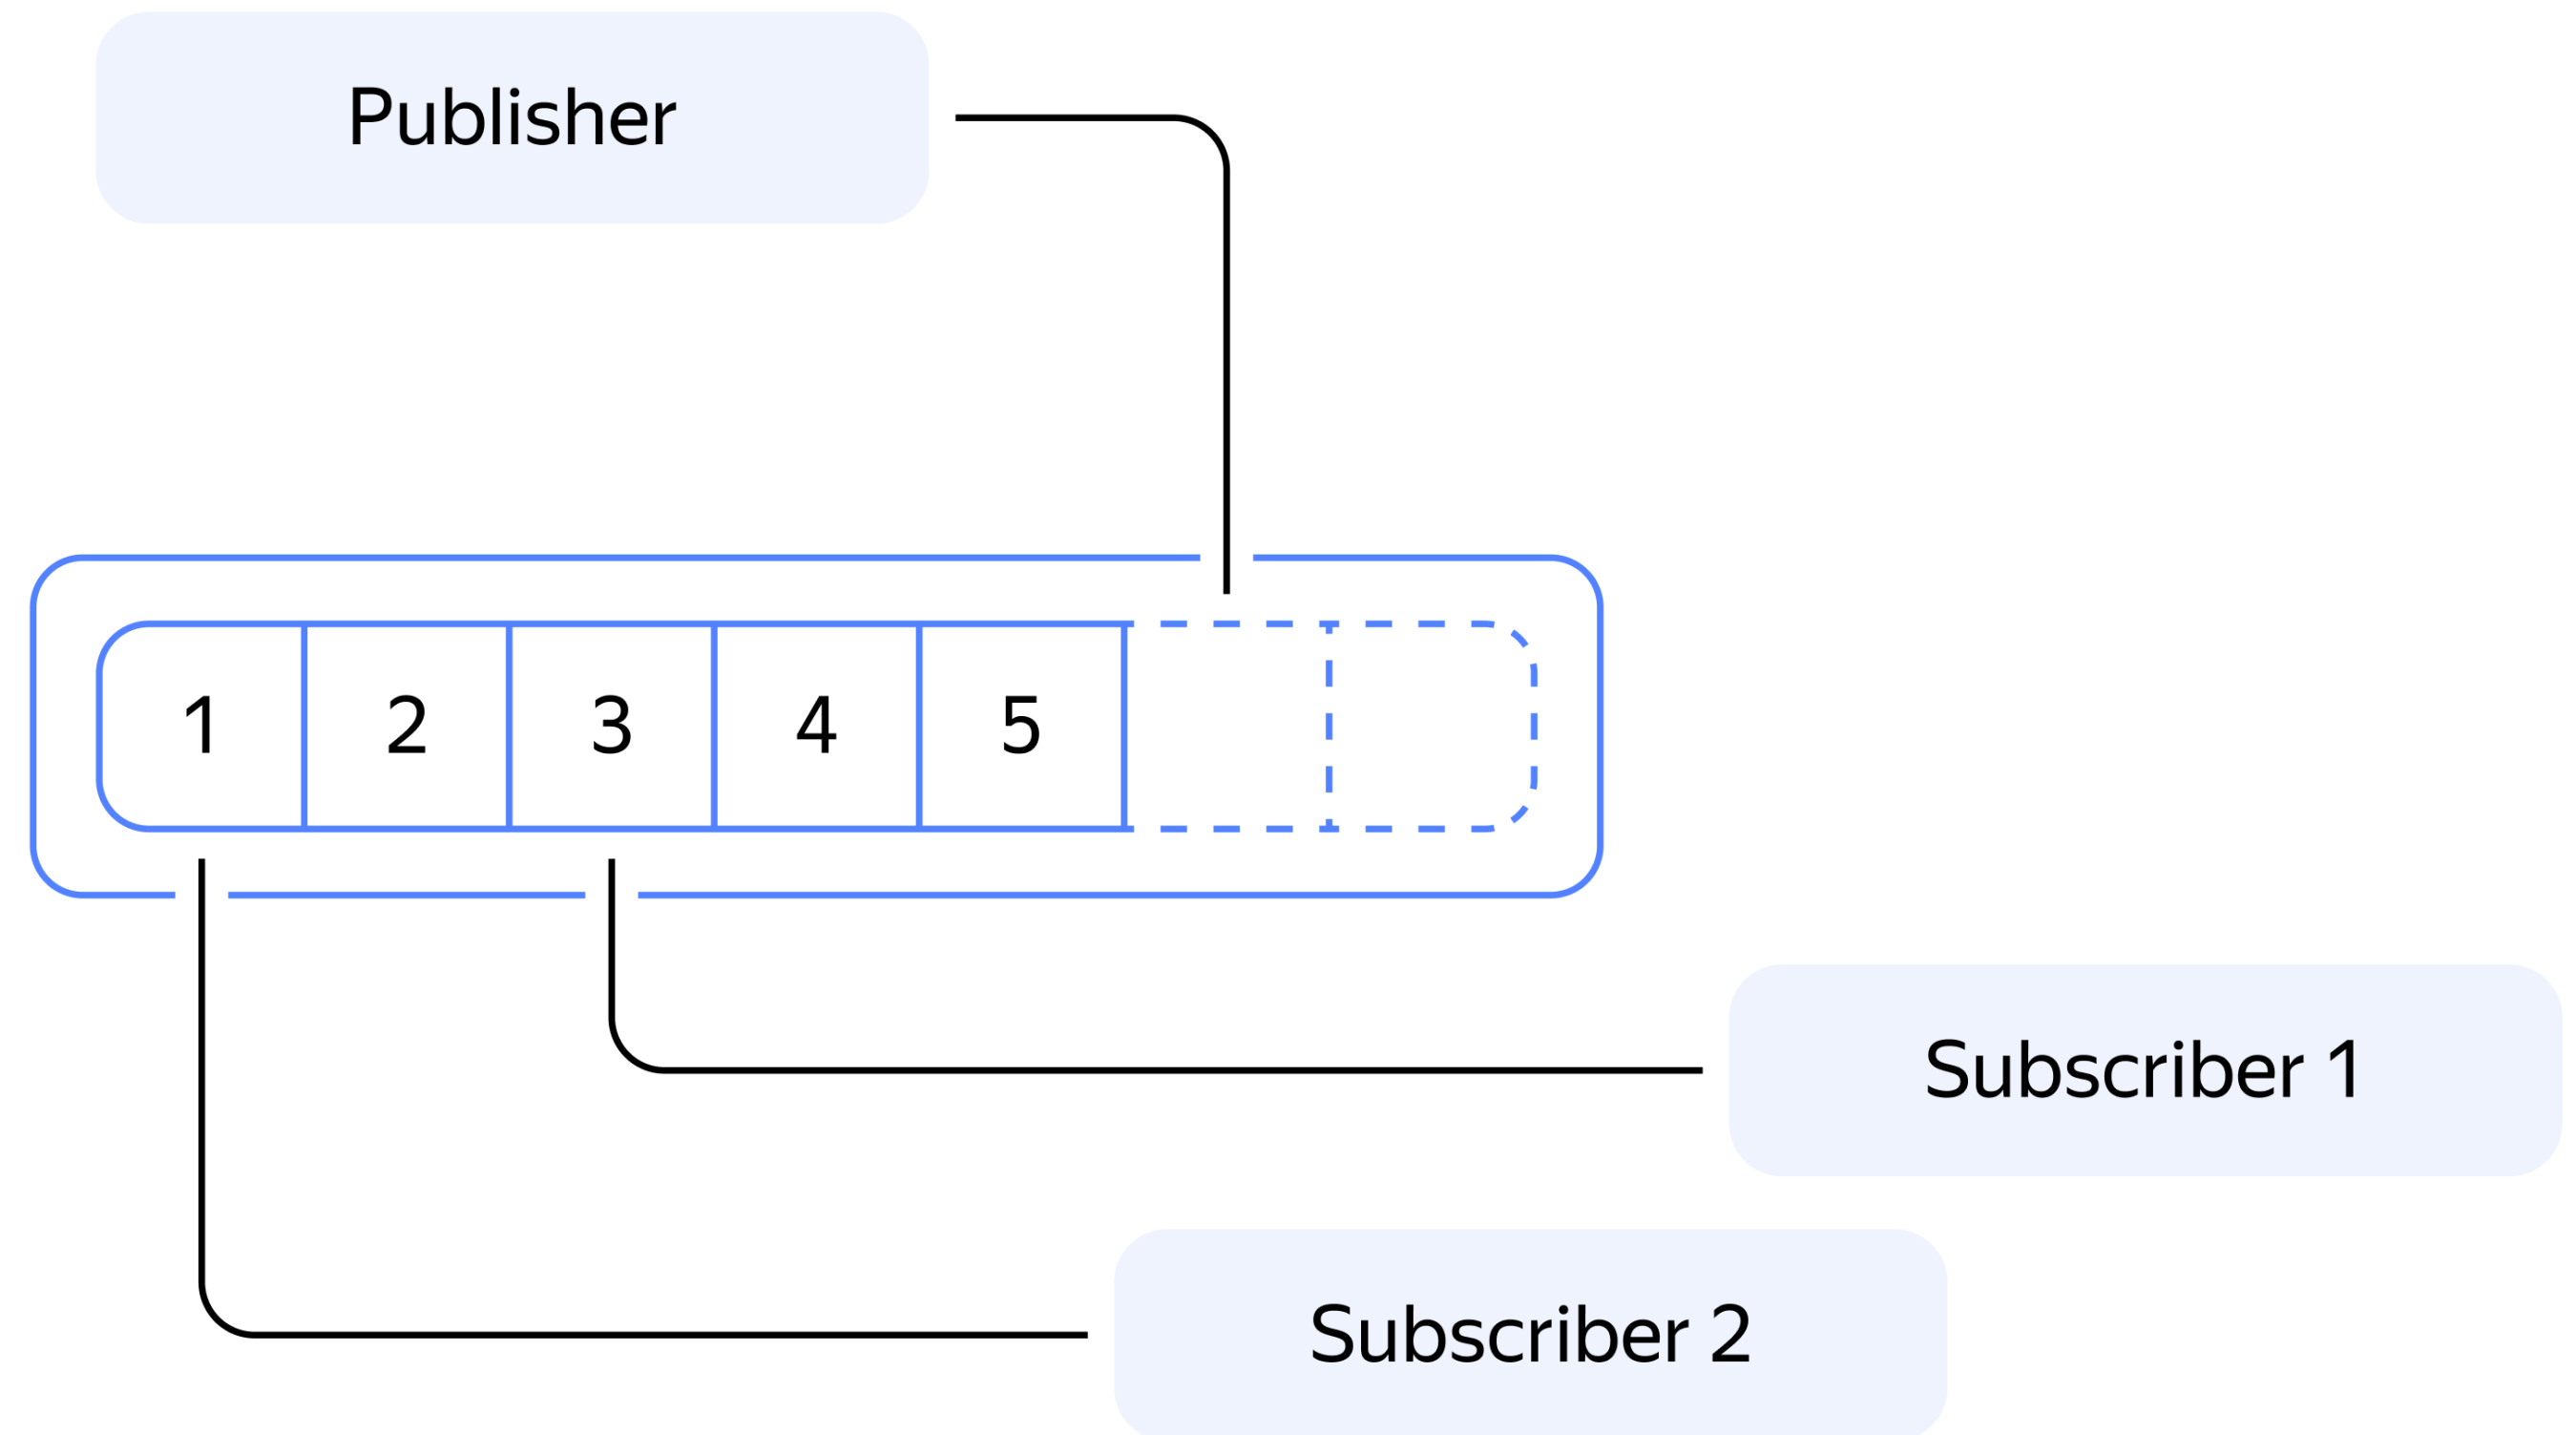


- Split by load
- Split by size
- YDB evenly distributes table's partitions among the nodes



# Topic

- Publish-Subscribe pattern
- Storing unstructured messages
- Delivering messages to multiple subscribers



★  
At-least-once or  
Exactly-once delivery

★  
FIFO message  
processing

★  
Horizontal scaling

# YQL

- SQL dialect
- Strong typing
- Named subqueries
- Built-in functions

DateTime

Regexp

Math

Unicode

JSON

Etc.

- DML

SELECT, INSERT, UPDATE,  
DELETE

JOIN, GROUP BY, ORDER BY

UPSERT / REPLACE

UPDATE ON

DELETE ON

- DDL

CREATE TABLE

DROP TABLE

ALTER TABLE

# YDB Features



Secondary indexes

Synchronous

Asynchronous

Covering

Time to Live (TTL)

Tables And Backup Compression

# YDB Features

Read Replicas

CDC

JSON API

ISO/IEC 9075:2016

Scan (Analytical) Queries



# User Facilities


## SDK

Go, Python, Java, Node.js, PHP, .NET, Rust, C++




## User Interface

DB Navigator, YQL execute, explain, diagnostics, monitoring



## Command Line Interface



# User Facilities

## SDK

Go, Python, Java, Node.js, PHP, .NET, Rust, C++

## User Interface

DB Navigator, YQL execute, explain, diagnostics, monitoring

## Command Line Interface

The screenshot displays a database management interface with a dark theme. At the top, there are tabs for 'Query' and 'Diagnostics', with 'Diagnostics' being the active tab. Below the tabs is a 'Navigation' sidebar on the left, listing various system components like '.sys\_health', 'DeviceGroups', 'DeviceTriggersIndex', 'Devices' (highlighted), 'Experiments', 'ExperimentsUserGroup', and 'External Lear'. Below the navigation is a search bar containing the text 'table'. The main content area is divided into two columns. The left column has tabs for 'Overview', 'ACL', and 'Schema', with 'Schema' selected. It shows a table with columns 'Id', 'Key', 'Name', and 'Type'. The right column has tabs for 'Info', 'Top shards', 'Graph', 'Tablets', 'Hot keys', and 'Describe', with 'Info' selected. It displays various metrics and statistics for the table, including CPU, Memory, Network, Storage, ReadThroughput, WriteThroughput, ReadIops, WriteIops, and Partition Config.

Id	Key	Name	Type
1	⊞	huid	UInt64
3	⊞	id	String
2		user_id	UInt64
4		external_id	String
5		name	String
6		external_name	String
7		type	String

Tablet Metrics		Table Stats	
CPU	12.40	PartCount	64
Memory	729 MB	RowCount	14,946,206
Network	0 B/s	DataSize	22 GB
Storage	36 GB	IndexSize	241 MB
ReadThroughput	13 MB/s	LastAccessTime	Tue, 08 Nov 2022 12:37:26 GMT
WriteThroughput	51 MB/s	LastUpdateTime	Tue, 08 Nov 2022 12:37:26 GMT
ReadIops	167	ImmediateTxCompleted	140,442,389
WriteIops	15,590	PlannedTxCompleted	21,571,803
Partition Config		TxRejectedByOverload	0
FollowerCount	2	TxRejectedBySpace	0
RequireAllDataCenters		TxCompleteLagMsec	1,410
FollowerCountPerDataCenter		InFlightTxCount	697
RowUpdates	33,937,905	RowDeletes	0
RowReads	40,774,176	RangeReads	59,794,458

# User Facilities

## SDK

Go, Python, Java, Node.js, PHP,  
.NET, Rust, C++

## User Interface

DB Navigator, YQL execute,  
explain, diagnostics, monitoring

## Command Line Interface

```
-bash — 147x36
Usage: ydb [global options...] workload [options...] <subcommand>
Description: YDB workload service

Subcommands:
workload      YDB workload service
├── clickbench ClickBench workload (ClickHouse OLAP test)
│   ├── clean  Drop table
│   ├── init   Initialize table (aliases: i)
│   └── run    Perform benchmark (aliases: b)
├── kv         YDB kv workload
│   ├── clean  drop tables created in init phase
│   ├── init   Create and initialize tables for workload
│   └── run    Run YDB KV workload
│       ├── insert  insert random rows into table
│       ├── select  select rows by exactly matching of a
│       └── upsert  upsert random rows into table
├── stock     YDB stock workload
│   ├── clean  drop tables created in init phase
│   ├── init   Create and initialize tables for workload
│   └── run    Run YDB stock workload
│       ├── add-rand-order  Inserts orders with random ID without their processing
│       ├── put-rand-order  Submit random orders with processing
│       ├── put-same-order  Submit orders with same products with processing
│       ├── rand-user-hist  Selects orders of random customer
│       └── user-hist       Selects orders of 10000th customer

Global options:
{-e|--endpoint}, {-d|--database}, {-v|--verbose}, --ca-file, --iam-token-file, --yc-token-file, --use-metadata-credentials, --sa-key-file, --token-file, --user, --password-file, --no-password, --iam-endpoint, {-p|--profile}
To get full description of these options run 'ydb --help'.

Options:
{-?|-h|--help} Print usage

Free args: min: 1, max: unlimited
```

# Diagnostics

Monitoring with any system that supports Prometheus format

- Zabbix
- Grafana

Grafana dashboards

Embedded UI

System tables

- Partitions
- Top queries
- Query details



# YDB deployment

## Kubernetes®

Kubernetes-operator

Kubernetes cluster

Managed Kubernetes (Any cloud)

## Virtual machines

### Bare metal

Docker (development and testing)

# Thank you!

- Open source under Apache 2.0 license
- ACID transactions
- SQL
- Elastic scalability across thousands servers
- Resilience
- Automatic disaster recovery
- Available anywhere



YDB web-site  
[ydb.tech](http://ydb.tech)