

Why Traditional Enterprise favor PG as RDBMS

Bruce Momjian

1. Introduction

2. Proprietary Software vs Open Source

3. How Postgres core team management

4. How to make the enterprise open-source

- 1. Introduction**
- 2. Proprietary Software vs Open Source**
- 3. How Postgres core team management**
- 4. How to make the enterprise open-source**

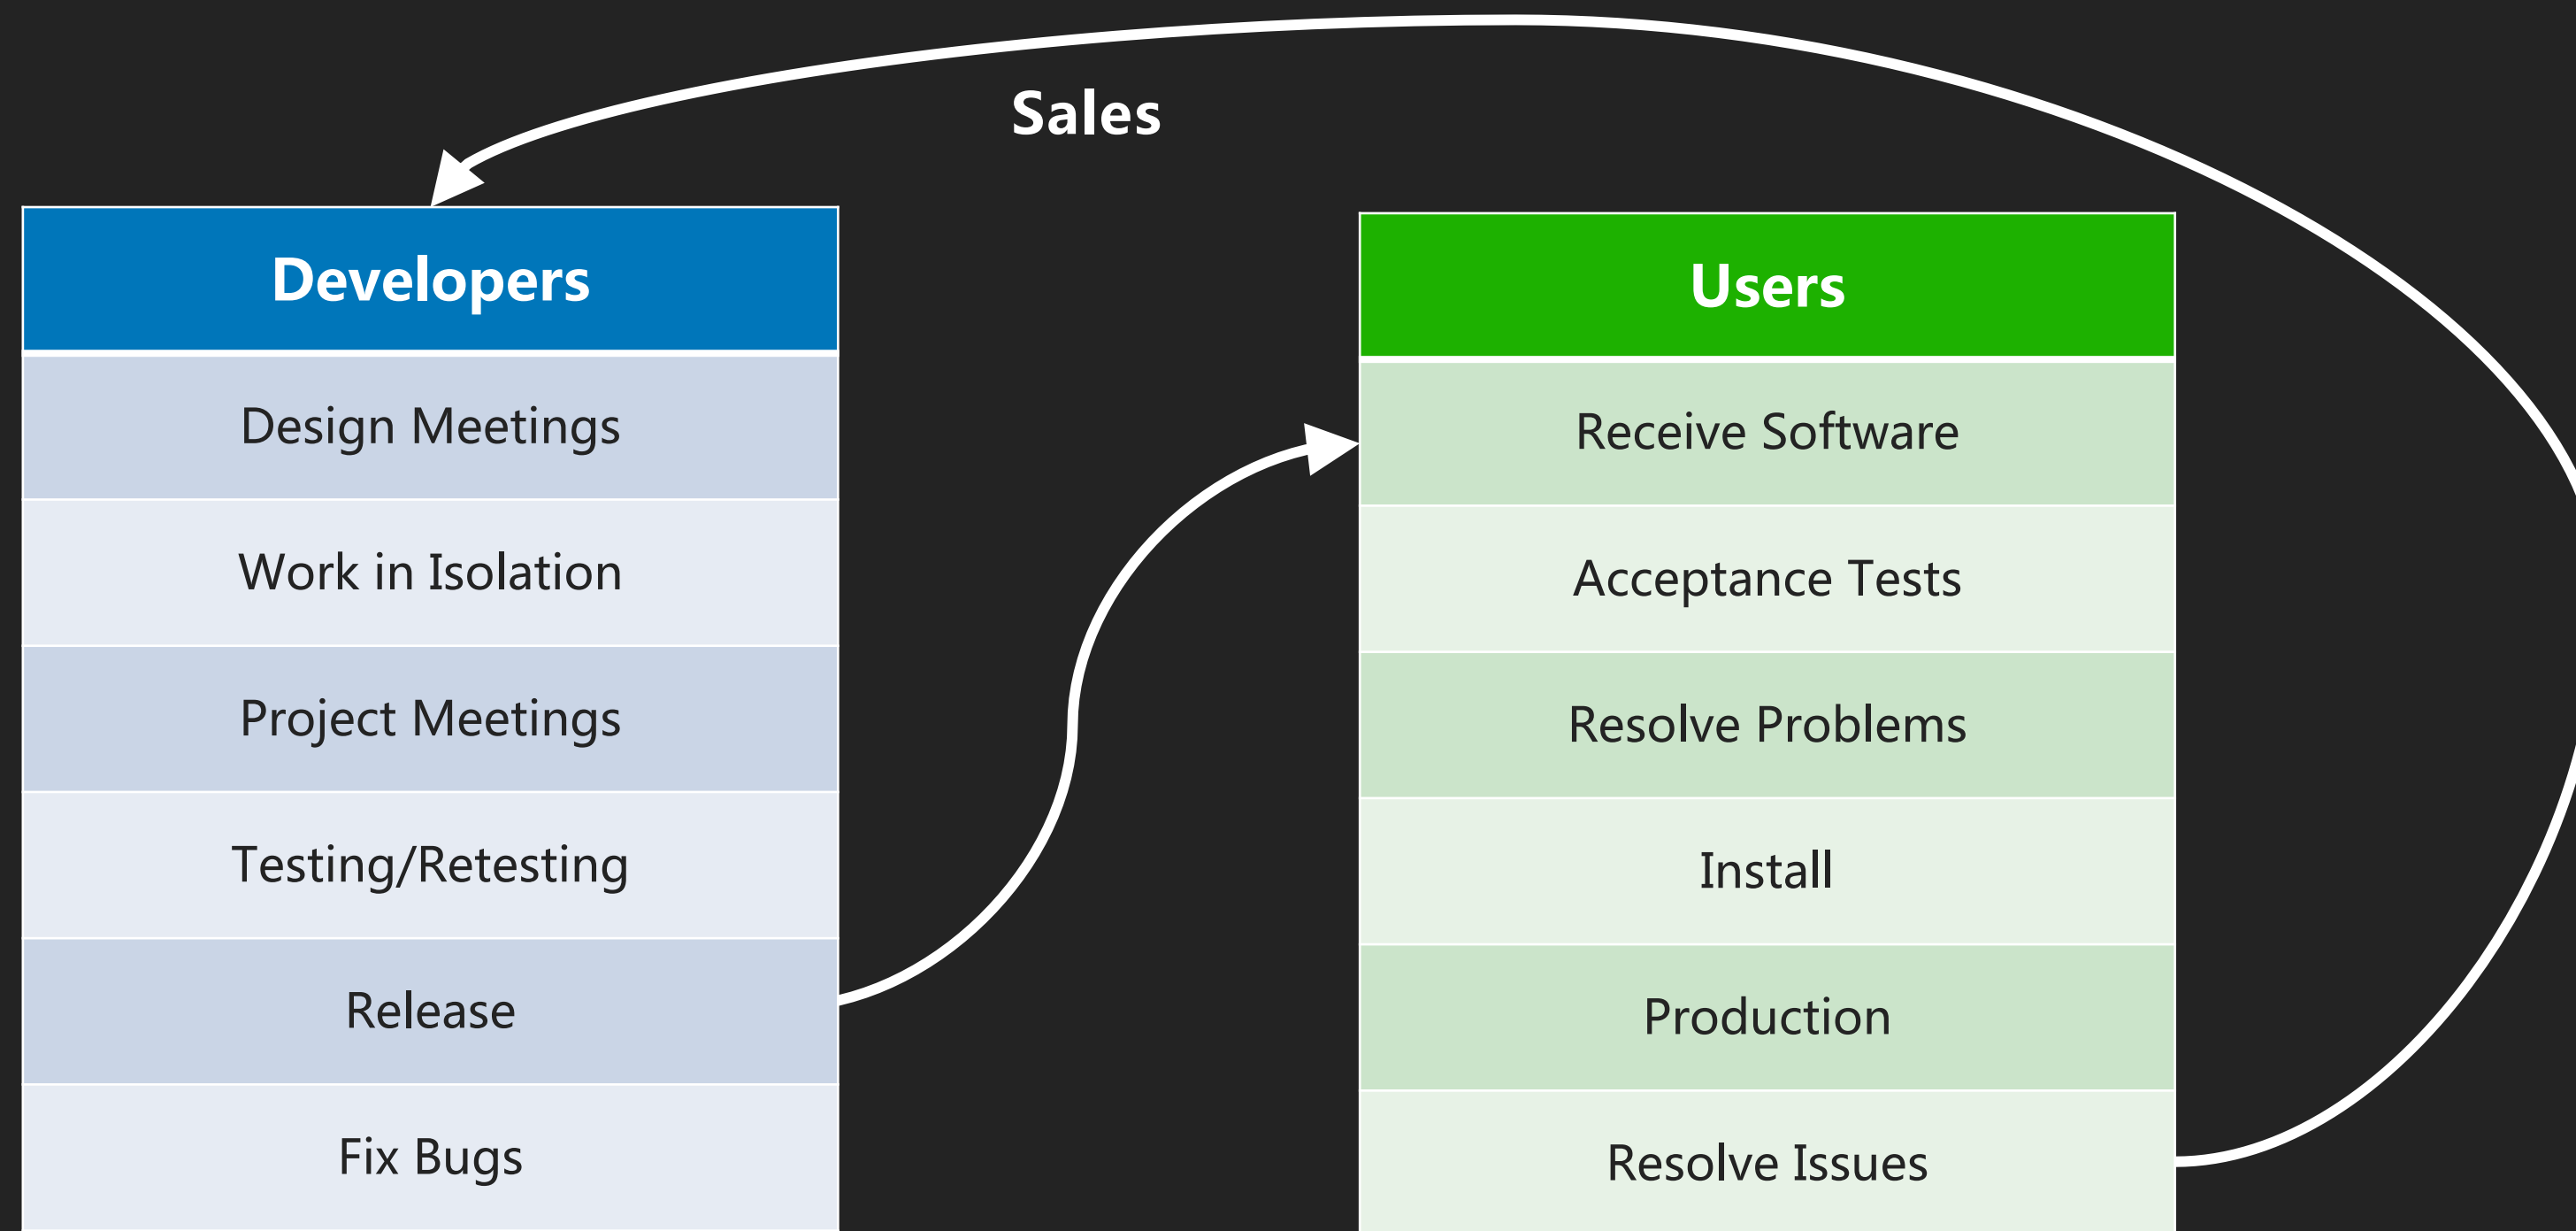
Proprietary Software Life Cycle

1. Innovation
2. Market growth
3. Market saturation
4. *Maximize profit, minimize costs (development, support)*
5. Maintenance mode (no new features, no innovation)
6. End-of-life

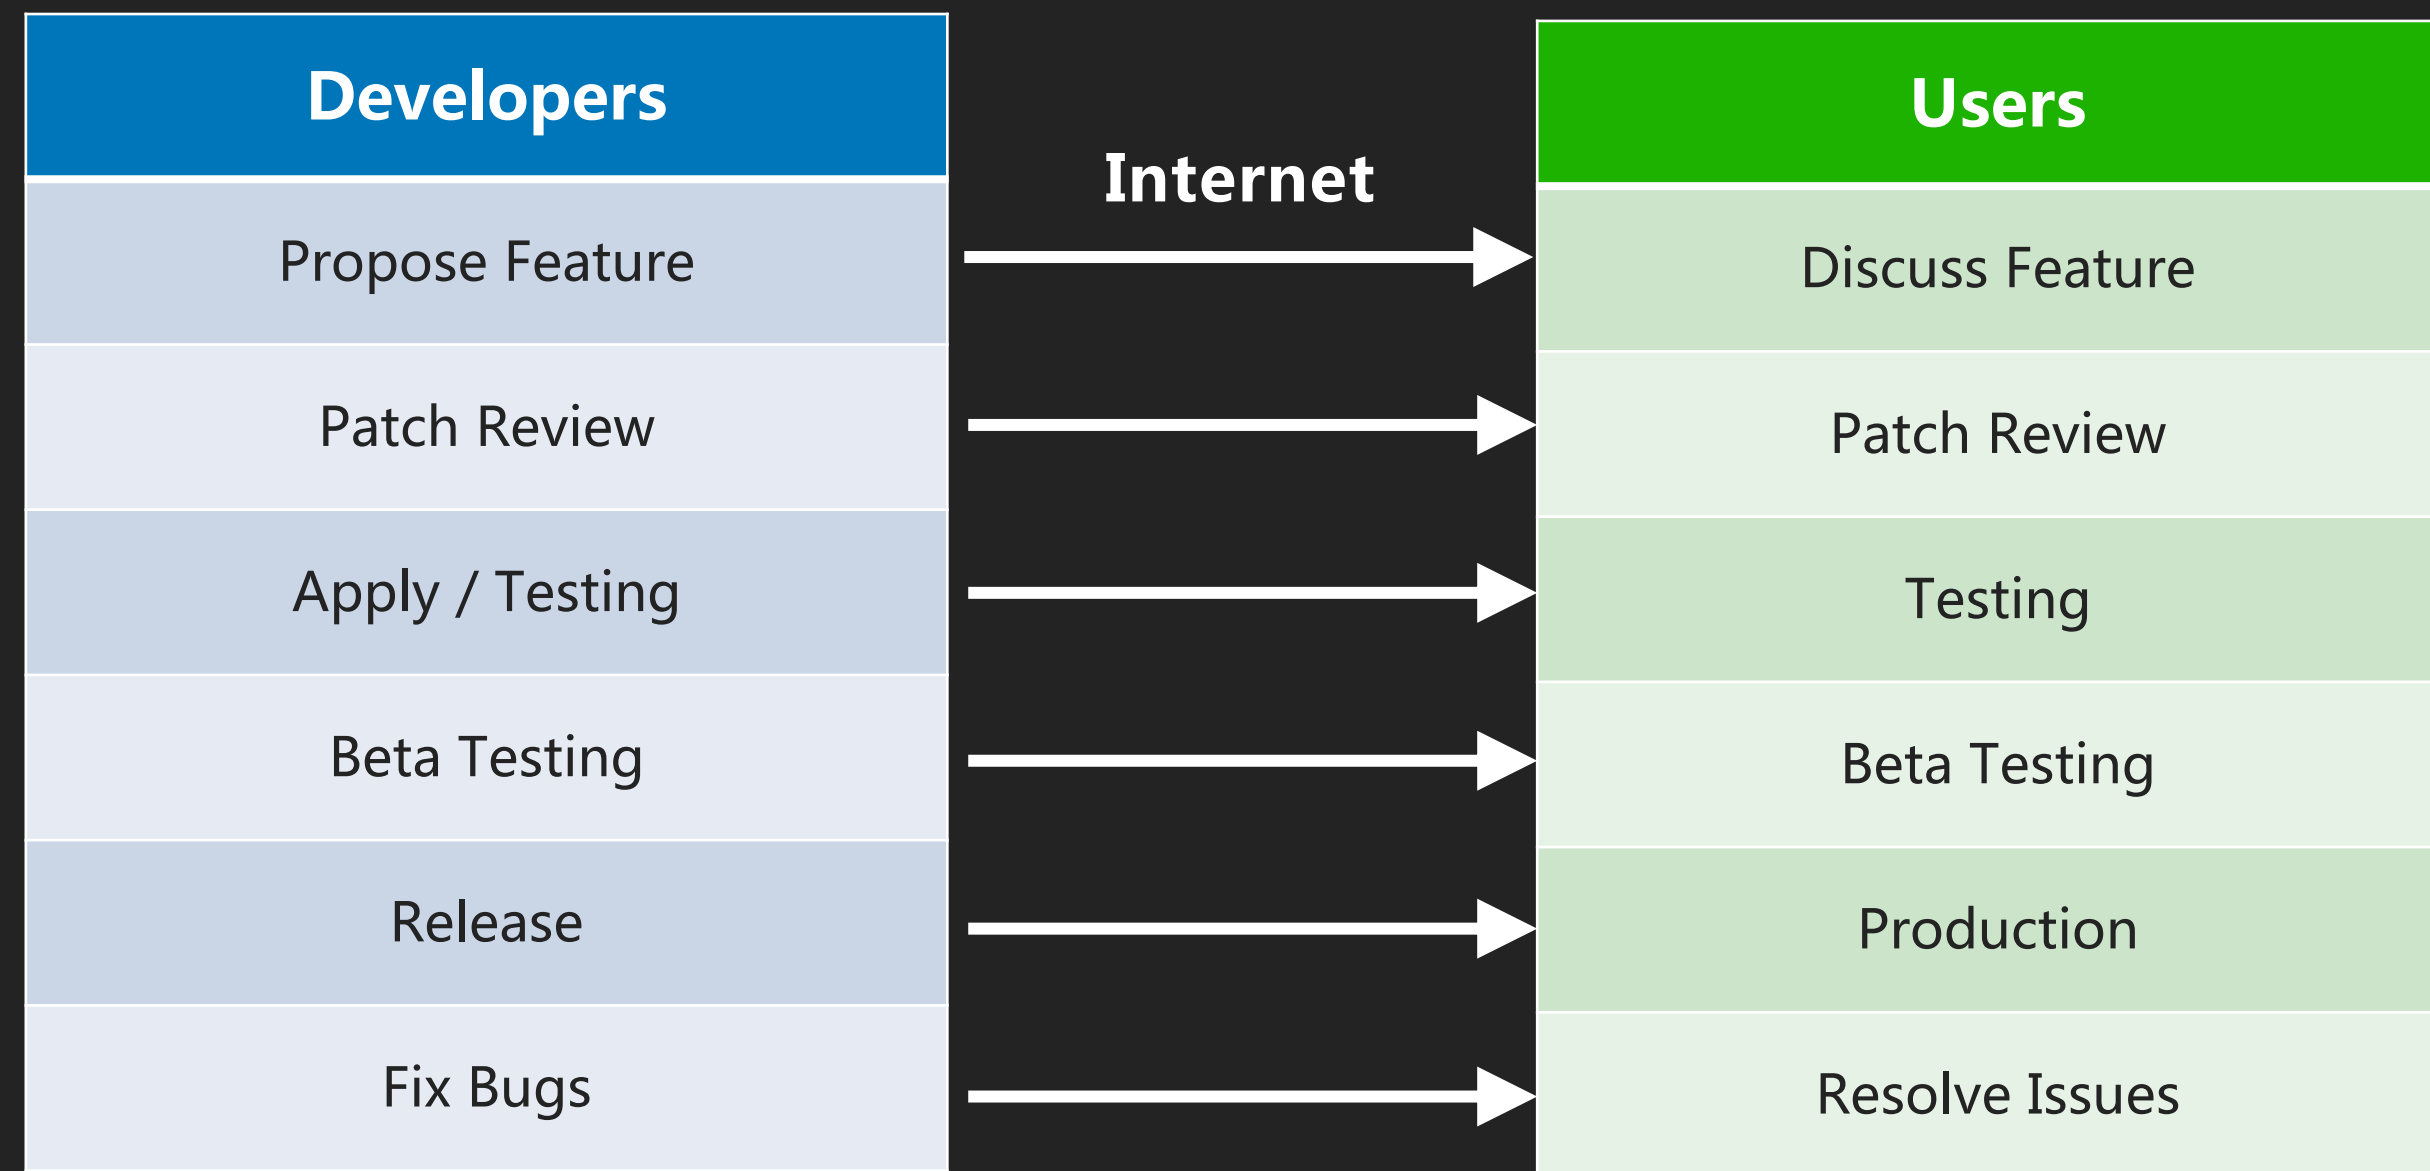
Open Source Software Life Cycle

1. Parity with proprietary software, low cost
2. Market growth
3. *Continue innovation or decline*
4. Source code is always available to continue

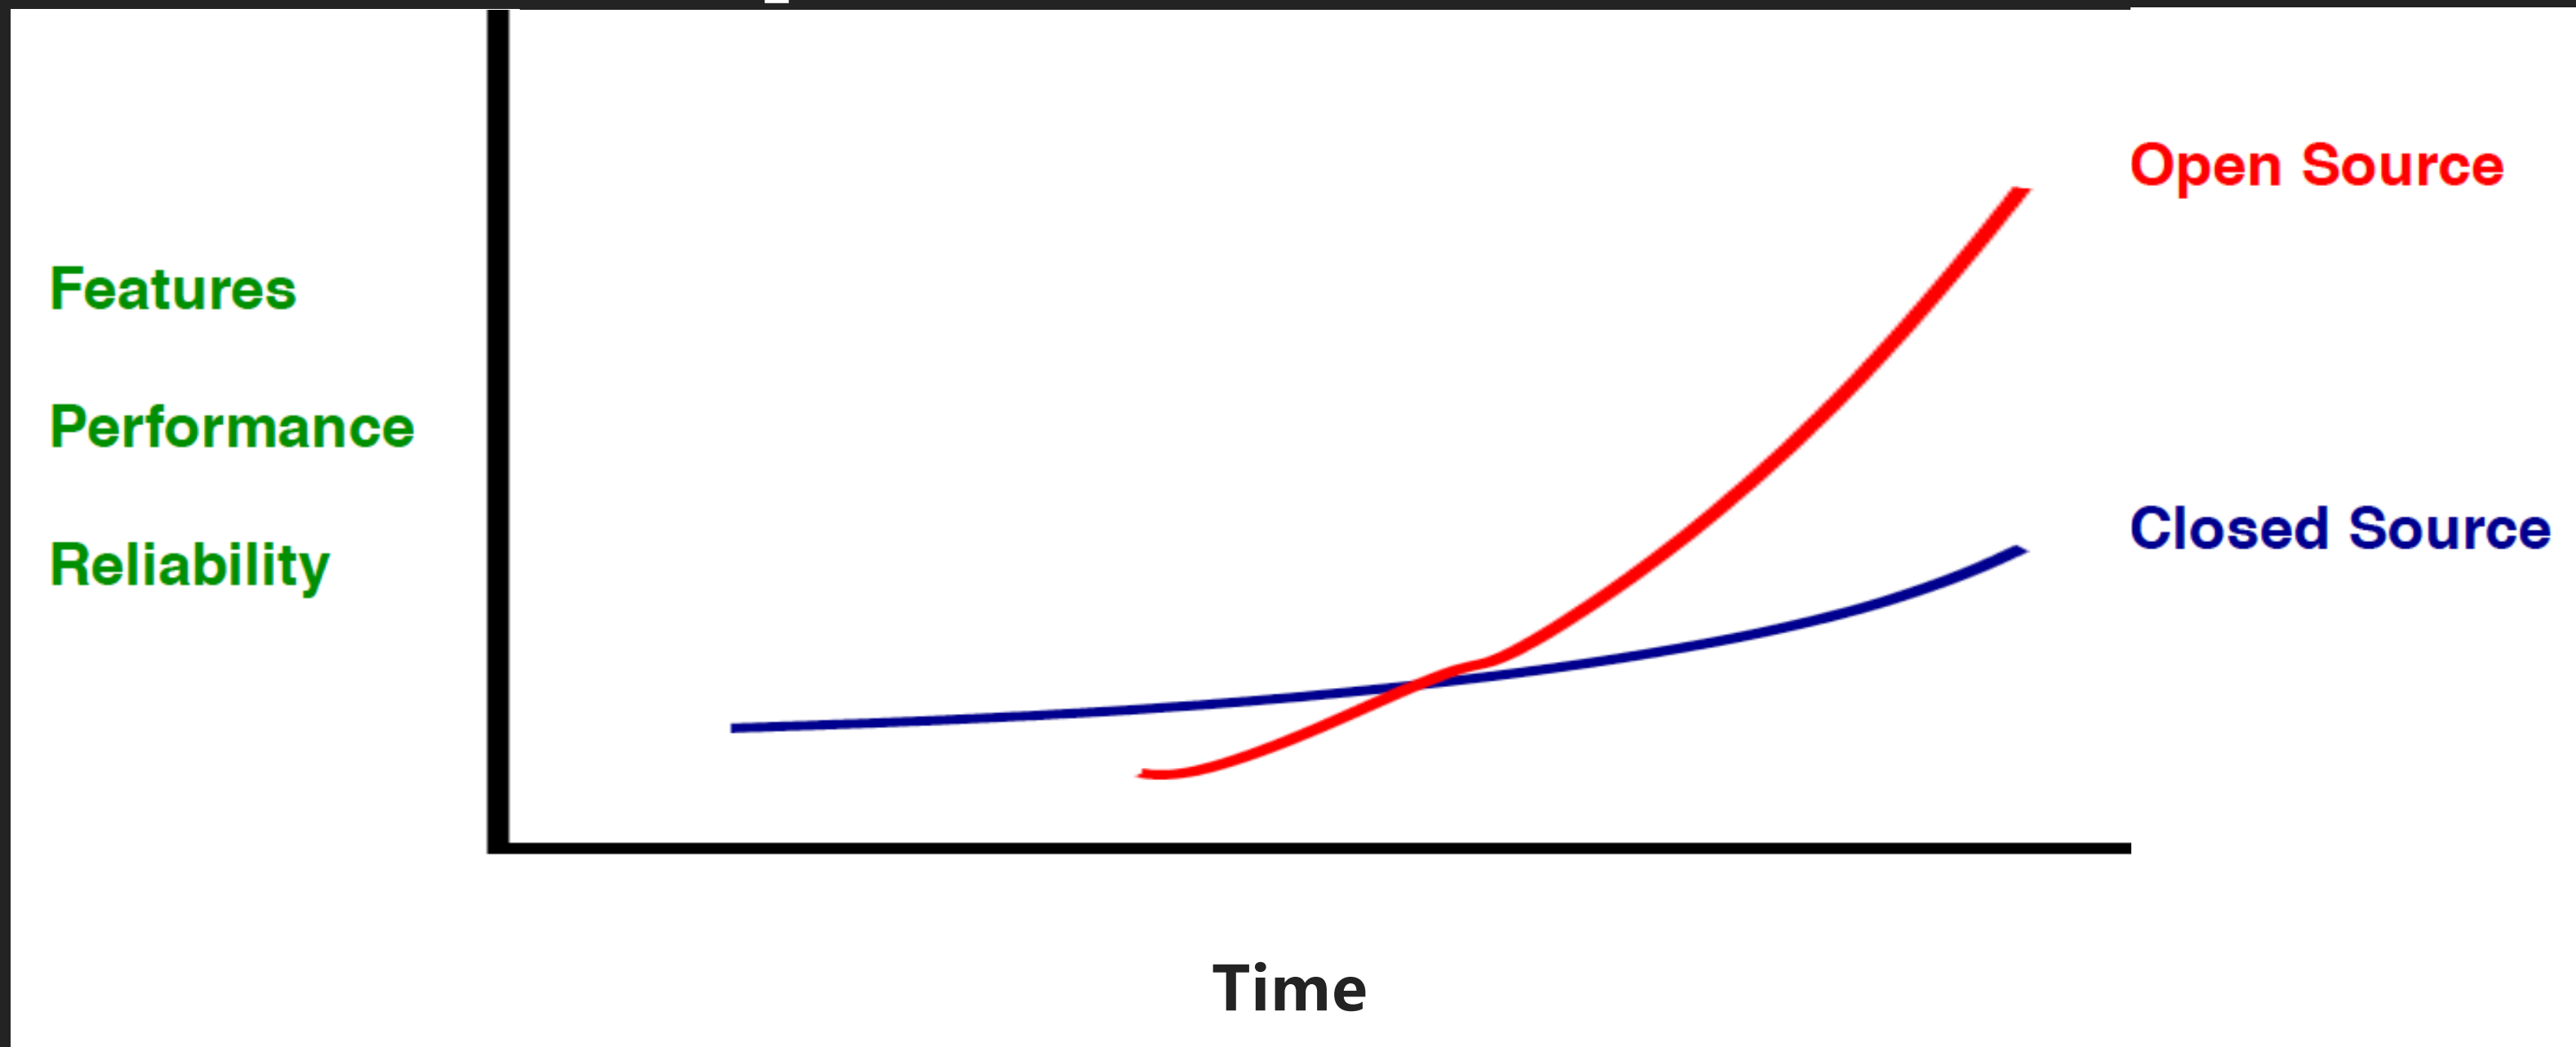
Proprietary Development Flow



Open Source Development Flow



Rise of Open Source



Linux attained feature parity with

1. AIX
2. HP-UX
3. Solaris

and then going on to innovate beyond them.

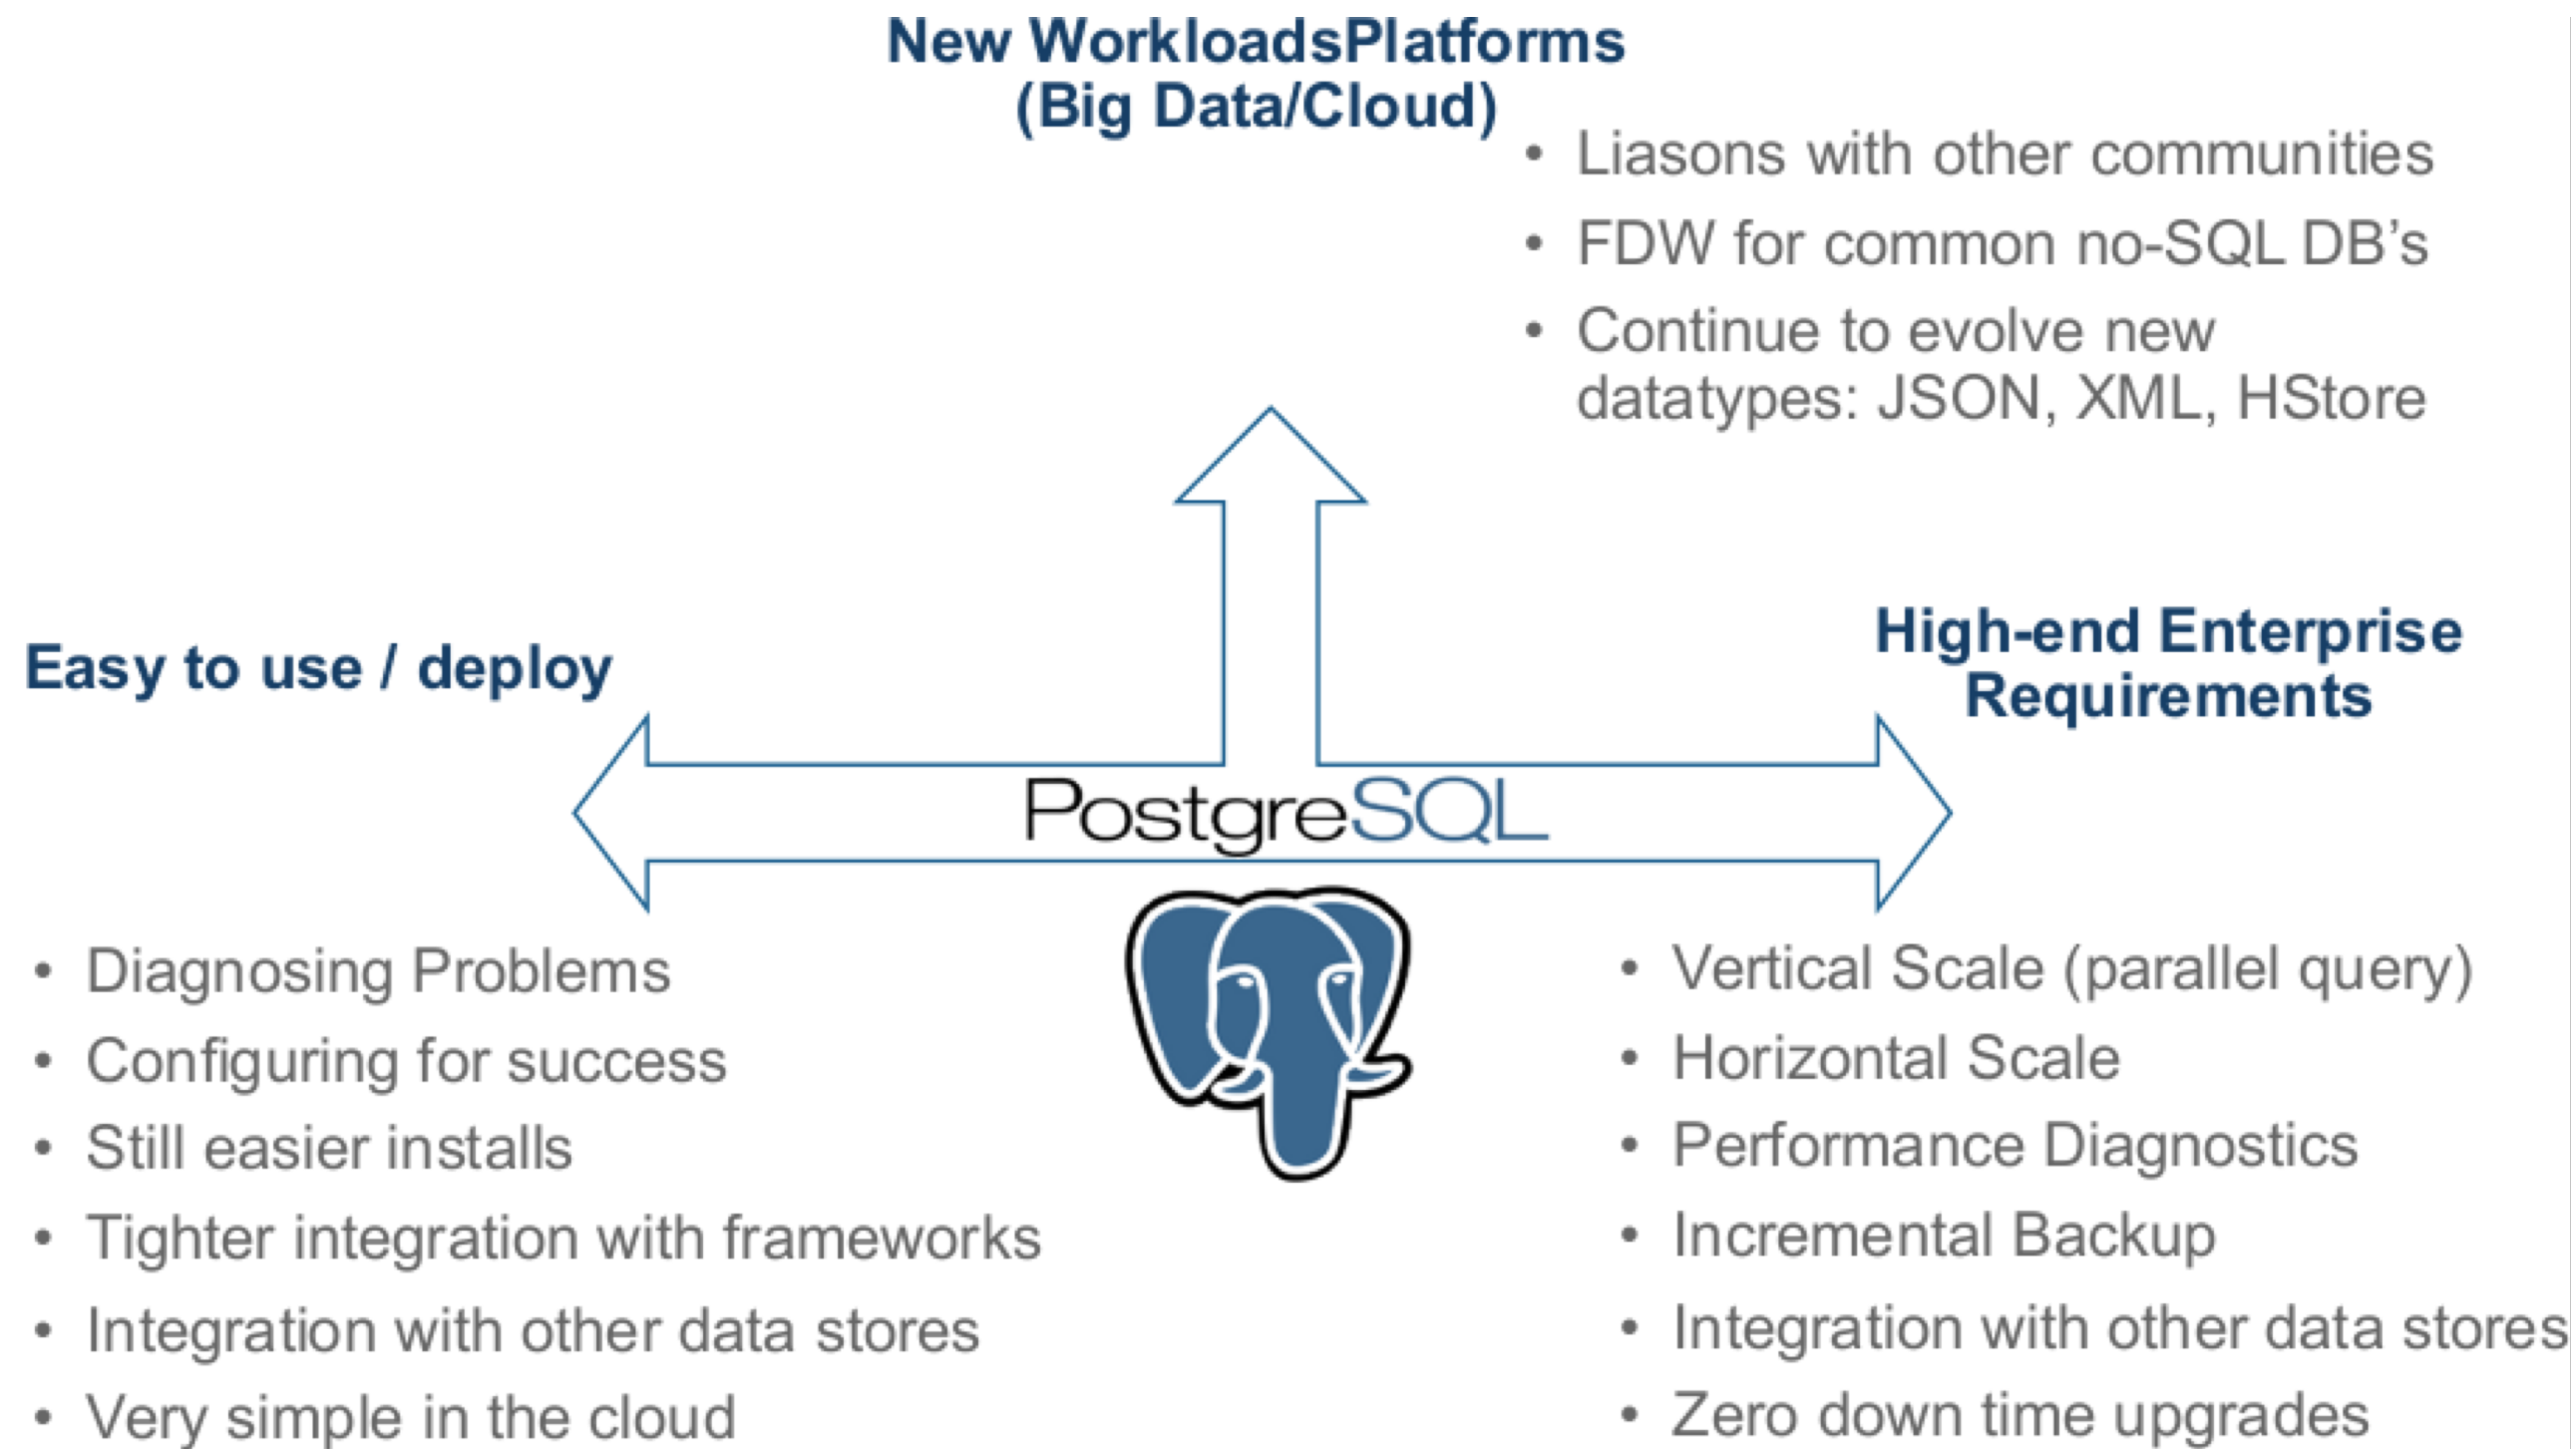


Postgres nearing feature parity with

1. Oracle
2. DB2
3. SQL Server
4. Sybase
5. Informix

and then going on to innovate beyond them.





When Does Software Die?

1. Proprietary software dies when the owner of the source code can no longer profit from it.
2. It declines long before death due to profit maximization.
3. Open source cannot die in the same way.
4. Open source remains active while it serves a purpose.
5. It can always be resurrected if useful.
6. Postgres was given new life in 1996.

Ideas Don't Die

1. Ideas don't die, as long as they are shared.
2. Ideas are shared, as long as they are useful.
3. Postgres will live, as long as it is useful.

Advantages of Open Source

1. Competitive features, innovation
2. Freedom from vendor lock-in
3. Quality of solutions
4. Ability to customize and fix
5. *Cost down*
6. Speed application development
7. Reduce development costs
8. Interoperability
9. Breadth of solutions

- 1. Introduction**
- 2. Proprietary Software vs Open Source**
- 3. How Postgres core team management**
- 4. How to make the enterprise open-source**

Development Process

1. Involve everyone
2. Find each person's motivator
3. Reach out to individuals
4. Harvest the strength of the team
5. There is always someone smarter than you
6. Produce work people can be proud of

Clean Code

1. People can't work on the code if they can't easily understand it
2. You aren't paying people to work, so make it as easy and interesting as possible
3. Produce quality documentation



Manage the Team

1. Lead by example, not from authority
2. See value in other people's opinions
3. Accept failure gracefully
4. Seek consensus
5. Don't be ruled by deadlines

Postgres Extensibility

Extension feature:

1. PostGIS, PG-Strom(GPU)
2. PL/Java, PL/PHP, PL/Python, PL/R (10+)
3. JSON Type, Range Types
4. FDW to MySQL, Oracle, Excel (100+)
5. Full Text Search

.....

Derived database:

1. EnterpriseDB (Oracle compatibility)
2. Greenplum (MPP OLAP)
3. Citus (Sharding)
4. AgenGraph (Graph Model)
5. PipelineDB (Streaming SQL)

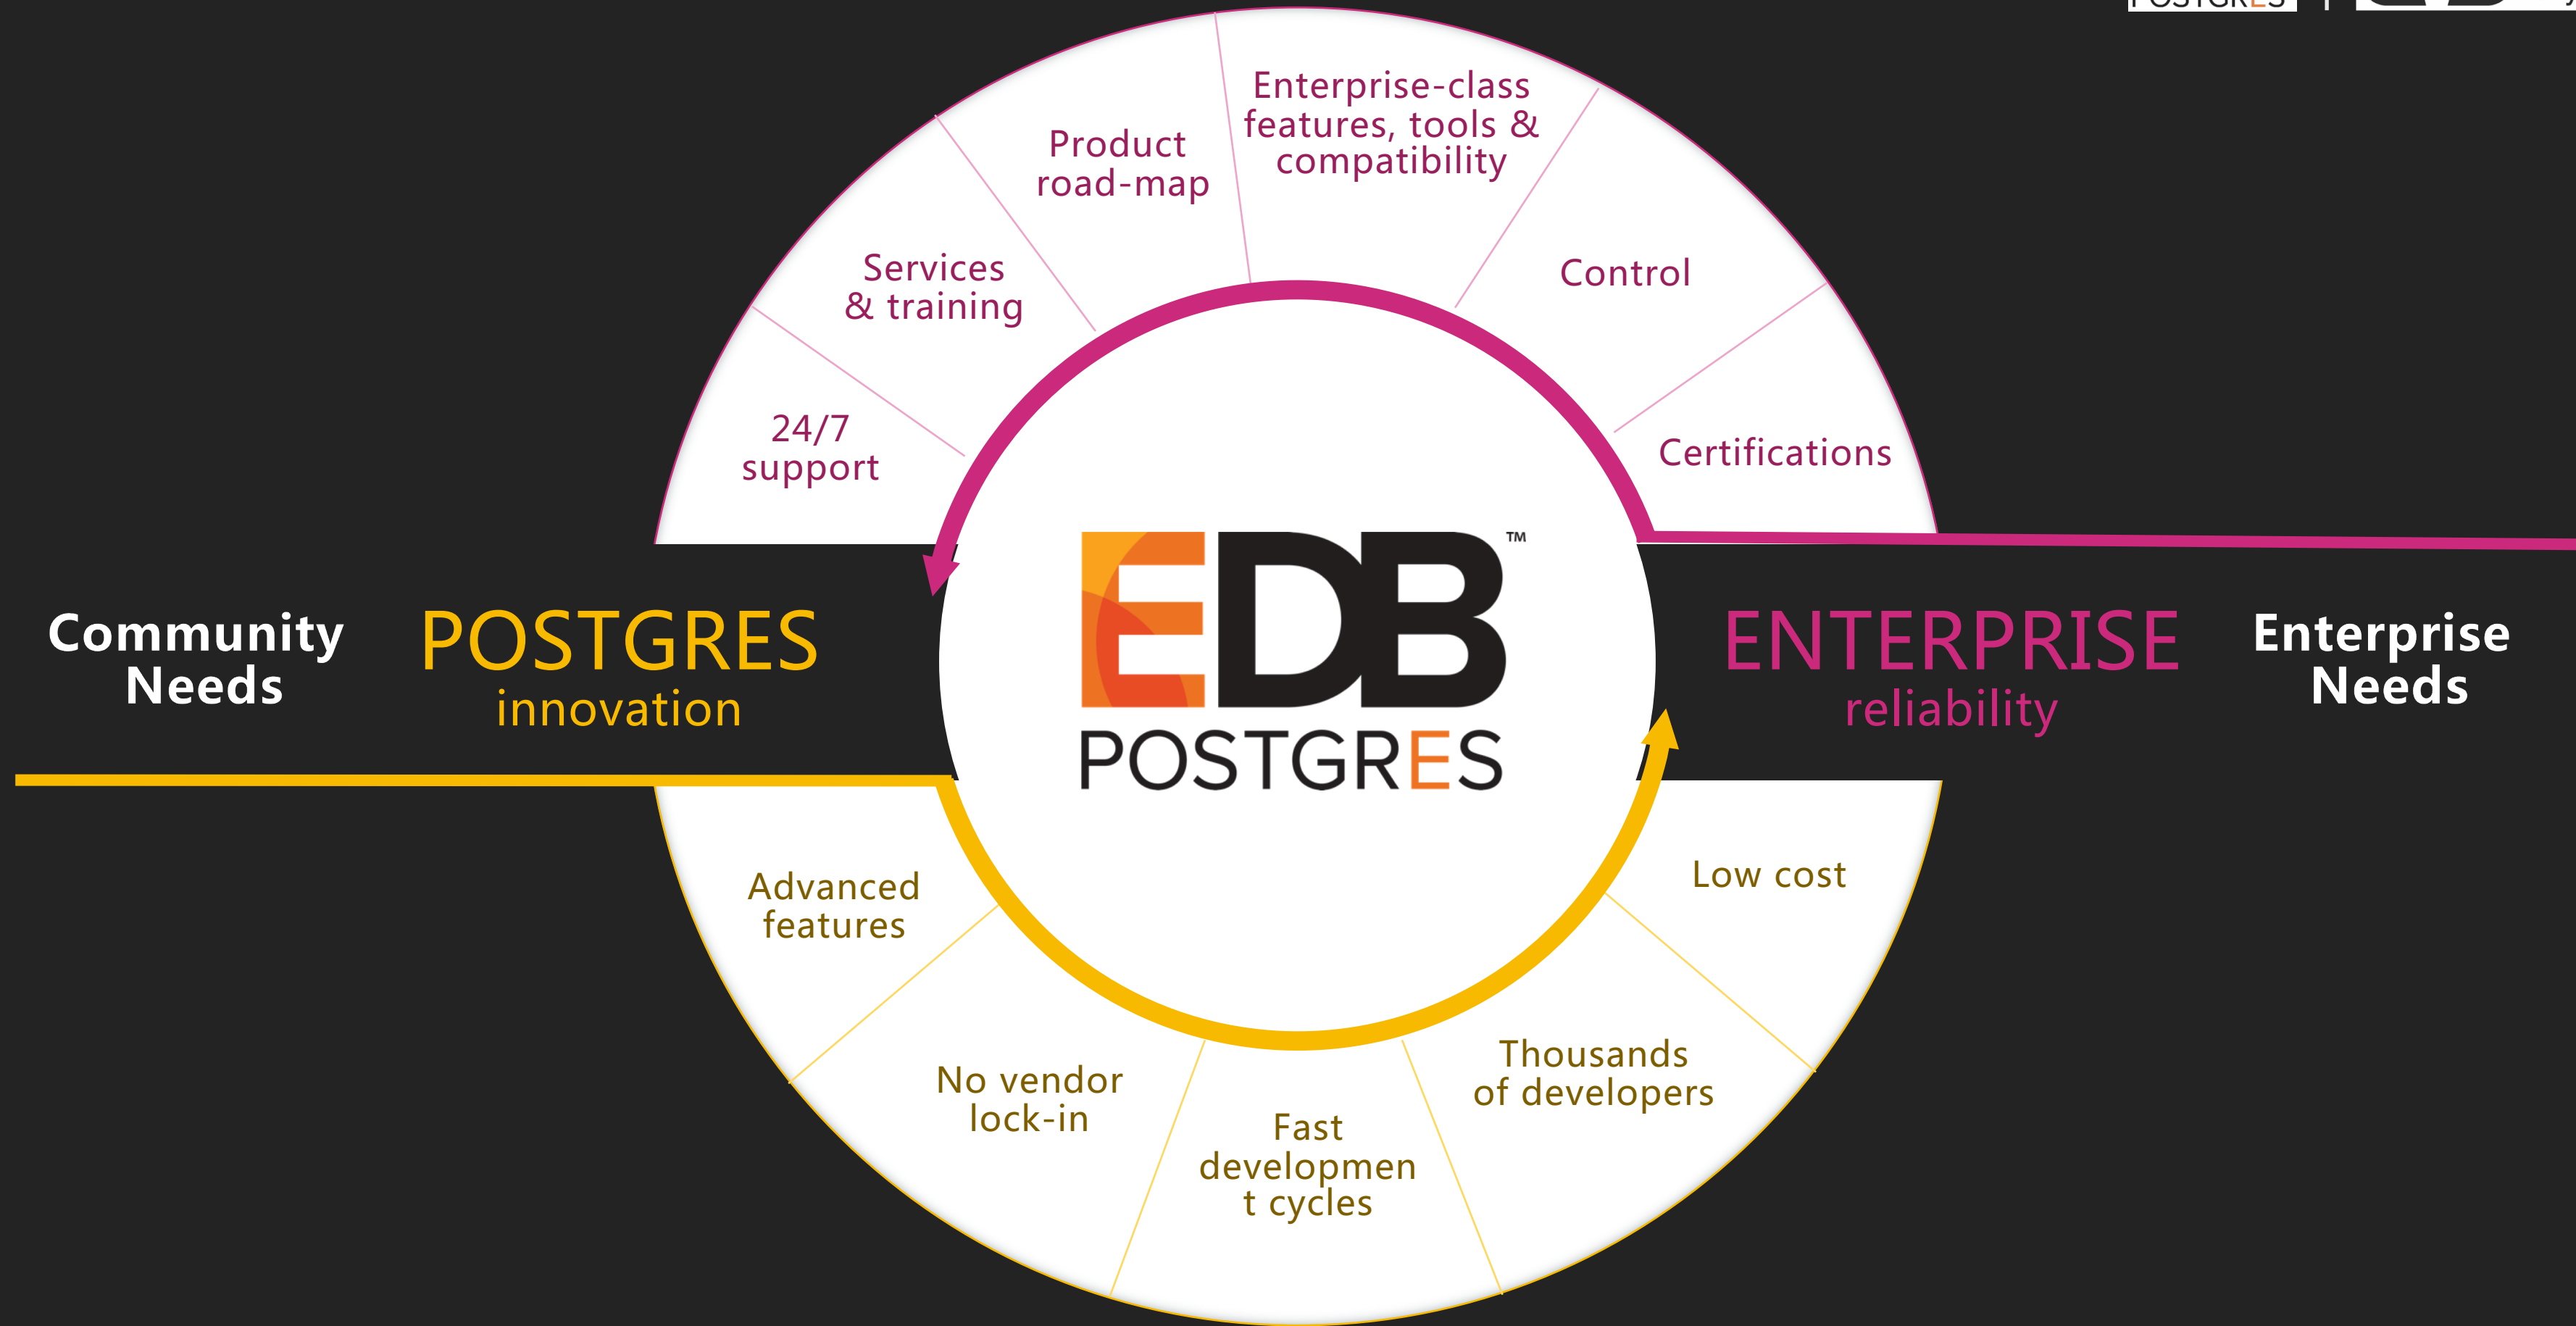
.....

1. Introduction

2. Proprietary Software vs Open Source

3. How Postgres core team management

4. How to make the enterprise open-source



Contributions to community

<i>from PostgreSQL community</i>	EDB contributions to PostgreSQL	<i>from EDB Development</i>
<ul style="list-style-type: none"> Logical Replication Declarative Table Partitioning Quorum Commit for Synchronous Replication 	<p style="text-align: center;">v10</p> <ul style="list-style-type: none"> Parallel Query index, bitmap scans & merge joins Durable Hash indexes postgres_fdw push down of joins and aggregates Trigger transition tables 	<ul style="list-style-type: none"> Re-direct Audit Log Records to syslog EDB Clone Schema Automatic cache pre-warm Customizabl WAL Segment Size
<ul style="list-style-type: none"> Avoid repetitive autovacuum Full-text search for phrases Support for remote joins, 	<p style="text-align: center;">v9.6</p> <ul style="list-style-type: none"> Parallel sequential scans, joins, and aggregates Synchronous replication support for 2+ standbys 	<ul style="list-style-type: none"> Advance queuing Nested sub-procedures Partitioned table performance enhancements EDBLDR enhancements More Oracle compatibility feature
<ul style="list-style-type: none"> Block Range Indexes (BRIN) UPSERT Row Level Security Schema creation for FDW Grouping 	<p style="text-align: center;">v9.5</p> <ul style="list-style-type: none"> Performance: sorting, in-memory hash, concurrency locking Parallelism Infrastructure 	<ul style="list-style-type: none"> Password profiles: User password management Improved performance under high concurrency

<i>from PostgreSQL core</i>	EDB contributions to PostgreSQL core	<i>from EDB Development</i>
<ul style="list-style-type: none"> Logical Decoding for Scalability JSONB Data Type JSONB Indexing Expanded JSON functions Delayed Application of Replication 3x Faster GIN indexes Support for Linux Huge Pages 	<p style="text-align: center;">v9.4</p> <ul style="list-style-type: none"> pg_prewarm ALTER SYSTEM Concurrently updatable Materialized Views Mongo FDW & MySQL FDW 	<ul style="list-style-type: none"> CPU & I/O Resource Management SQL Aggregation with CUBE, ROLLUP and GROUPING SETS Comprehensive UTL_HTTP Package Hash Partitioned Tables Connect By_Root Operator for hierarchical queries SQL/Protect Logging to DB Table EDB*Loader Improved Error handling
<ul style="list-style-type: none"> 64 bit LOBs up to 4TB in size Custom background workers Writable Foreign Data Wrappers 	<p style="text-align: center;">v9.3</p> <ul style="list-style-type: none"> Materialized Views 	<ul style="list-style-type: none"> Partition Read Improvements over 75x Support for 1000s of Partitions Partition write improvements over 400x
<ul style="list-style-type: none"> Cascaded streaming replication JSON support, Range Types 	<p style="text-align: center;">v9.2</p> <ul style="list-style-type: none"> MySQL Foreign Data Wrappers for SQL/MED 	<ul style="list-style-type: none"> Table() function support for nested tables INSERT APPEND hint EDB Postgres Multi-master Replication Expanded Object Type support
<ul style="list-style-type: none"> Synchronous replication Serializable Snapshot Isolation In-memory (unlogged) tables Writeable Common Table Expressions 	<p style="text-align: center;">v9.1</p> <ul style="list-style-type: none"> Index-only scans (covering indexes) Linear read scalability to 64 cores 	<ul style="list-style-type: none"> Row Level Security Declarative Partitioning syntax
<ul style="list-style-type: none"> Deferrable unique constraints and Exclusion constraints Streaming replication Windows 64 bit Support Hot standby 	<p style="text-align: center;">v9.0</p> <ul style="list-style-type: none"> No restore In-place version upgrades 	<ul style="list-style-type: none"> VARRAY support SQL Profiler Index Advisor Parallel Bulk Data Load

Enterprise needs: Compatibility for Oracle®

Your people



Oracle
Developers

Your apps



Oracle
Applications

Your business



Lower Costs and
Increased Agility

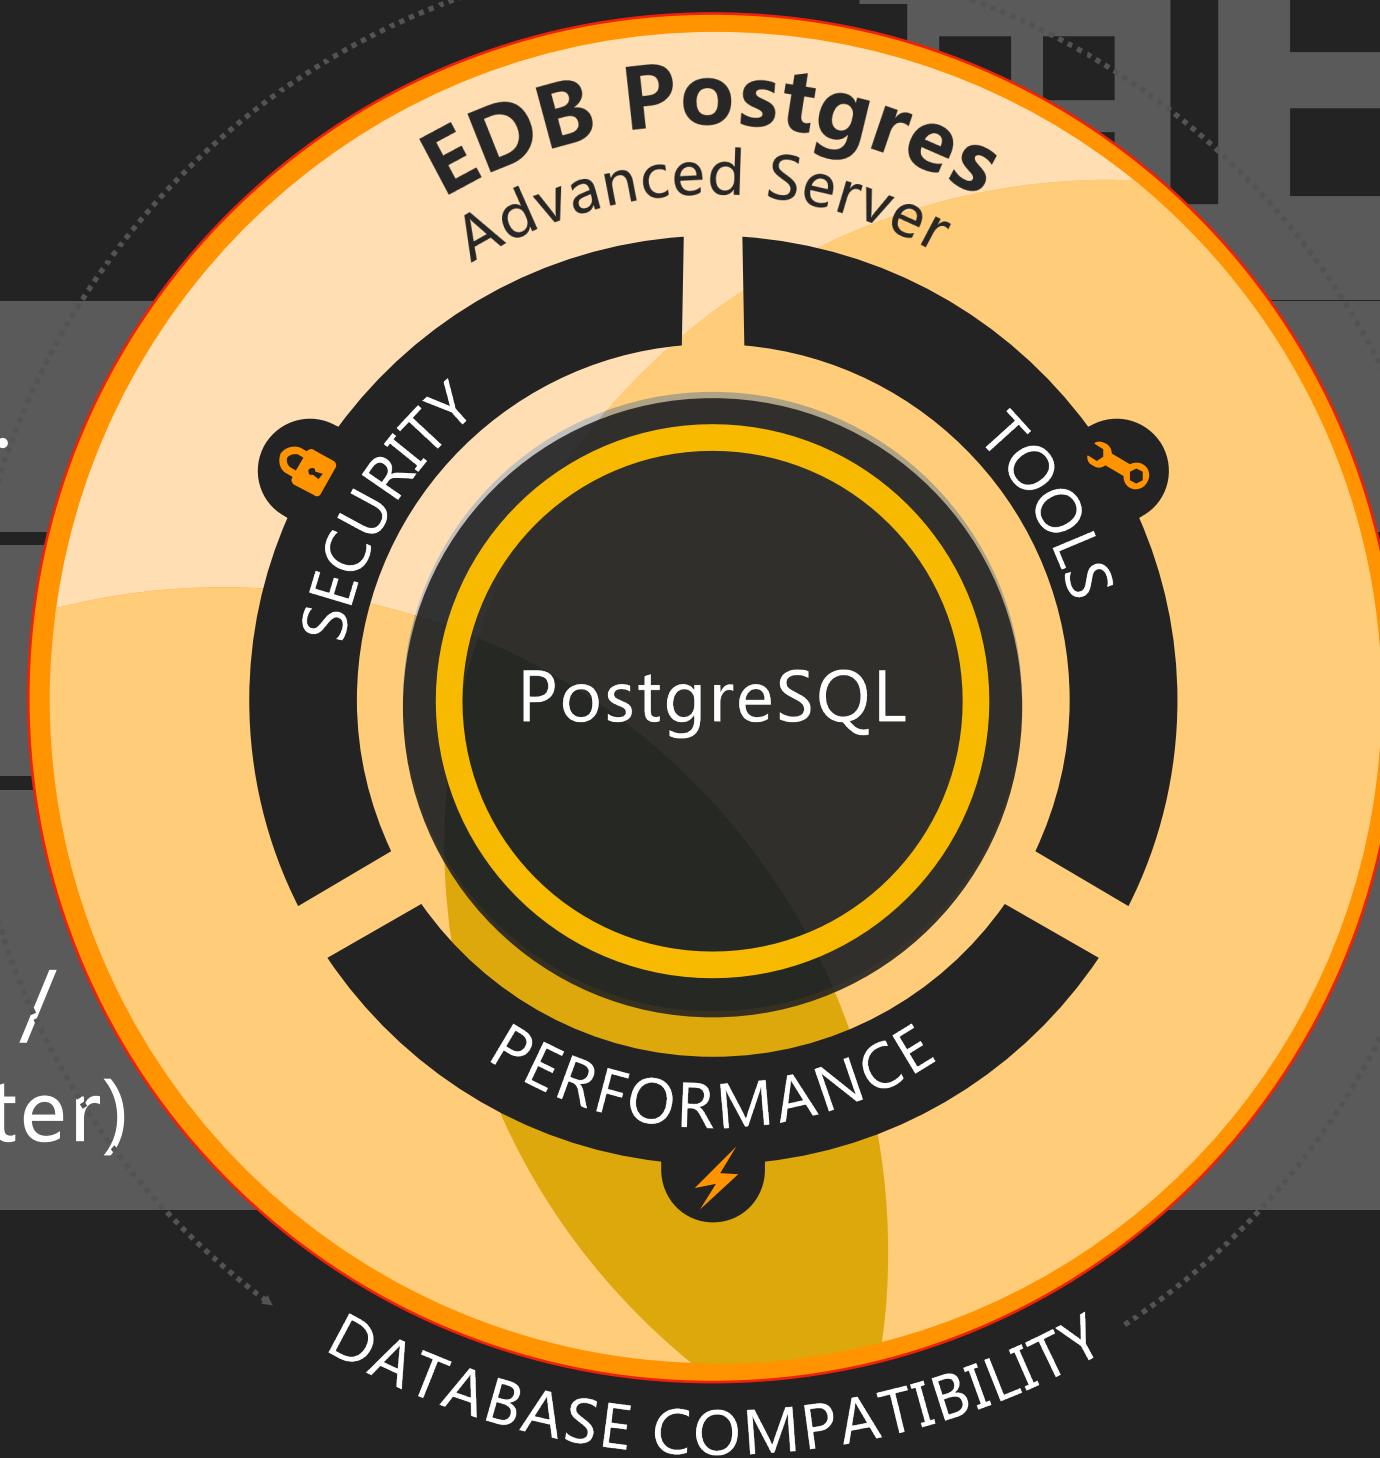
- Deploy new applications on EDB Postgres instead of Oracle
- No need to retrain Oracle DBAs and developers
- Support for PL/SQL language and OCI interoperability
- Replication for easy sharing of data

Enterprise needs: not just compatibility

Security: VPD / SQL Injection protection...

Tools: PEM / Migration / Replication

Performance: Resource Manager / Hints
SQL Profiler / Bulk Data Load (2x faster) /
Partitioning (400x writes & 76x read faster)



驱动数字中国

EMPOWER DIGITAL CHINA