



Self-Managing Databases—If Not Now Then When?

Johann Schleier-Smith
CEO & Founder, CrystalDB

www.crystaldb.cloud

@jssmith

Speaker: Johann Schleier-Smith

CEO & Founder at CrystalDB

- Since 2022

Previously CTO & Founder at if(we)

- 300 million social network members (hi5.com, Tagged.com, etc.)
- Oracle, PostgreSQL, MySQL

Board Member at Sama

- AI supplier to 25% of the Fortune 50
- Leader in computer vision for automotive, retail, and other applications

PhD at UC Berkeley and Stanford



CrystalDB : Operational Excellence for PostgreSQL



Reliability and security

Availability, data protection, trusted as system of record.



Performance and efficiency

Consistently low latency, avoiding wasted resources.



Team effectiveness

People stay focused on projects that move the business forward.

It is 2024



**ROAD
WORK
AHEAD**

JAGUAR
WAYMO

LYON

Lydia & Lyon Market
Liquor Wine Beer ATM
Groceries - 415-775-4480

5WDA993

9GRC786

2811EV2



Over 300 settings in postgres.conf

AREA	NUMBER	EXAMPLES
Cache and Buffer Management	37	shared_buffers, max_stack_depth
Write-Ahead Log	38	wal_buffers, wal_compression
Vacuum	13	autovacuum_vacuum_threshold, autovacuum_naptime
Query Tuning	48	cpu_tuple_cost, effective_cache_size
Connections and Authentication	25	max_connections, password_encryption
Lock Management	5	deadlock_timeout, max_pred_locks_per_page
Error Handling	4	data_sync_retry, restart_after_crash
Replication	24	vacuum_defer_cleanup_age, max_replication_flush_lag
Statistics	12	track_io_timing, log_planner_stats
Others	132	block_size, data_checksums

PostgreSQL 16 Administration Cookbook

Solve real-world Database Administration

SQL PERFORMANCE EXPLAINED

ENGLISH EDITION

A Practical Guide to the Advanced Open Source Database



PostgreSQL

PostgreSQL High Performance Cookbook

Mastering query optimization, data and performance-tuning for Postgr




***** 2024 EDITION *****

POSTGRES DBA v16 HANDBOOK

Fixing Data Consistency Issues in Amazon RDS

PostgreSQL Query Optimization

The Ultimate Guide to Building Efficient Queries
—
Second Edition



Contents selected according to industry requirements

Explanation with diagrams

Downloadable VM image for tests

Postgresql Administration

mehmet devcisoölu | dataerca

RELIABLE PARADIGM®



High Performance Database Tuning & Optimization

UNLOCK DATABASE PERFORMANCE MASTERY
A COMPREHENSIVE GUIDE CLEARLY EXPLAINED



PostgreSQL Configuration


Best Practices for Performance and Security

— Baji Shaik

The Pragmatic Programmers

High Performance PostgreSQL for Rails

Reliable, Scalable, Maintainable Database Applications



What can go wrong?

Lock contention

Connections exhausted

Slow transactions

Out of memory

Missing indexes

Slow recovery

Xid wraparound

Table bloat

Inefficient partitioning

Insufficient CPU

Insufficient indexes

Load spikes

Insufficient IO

Excessive indexes

Wasted capacity

Lost revenue

Lost data

Lost sleep

Consistency errors




Project delays















2024




Self-Managing DB vs Managed Service

-  Retained responsibility
-  Shared responsibility
-  Delegated responsibility

	Self-Managed	Managed Service	Self-Managing
Application SQL tuning			
Workload-specific DB configuration and optimization			
Workload-agnostic DB installation, provisioning, backups, failover			
Cloud Infrastructure			

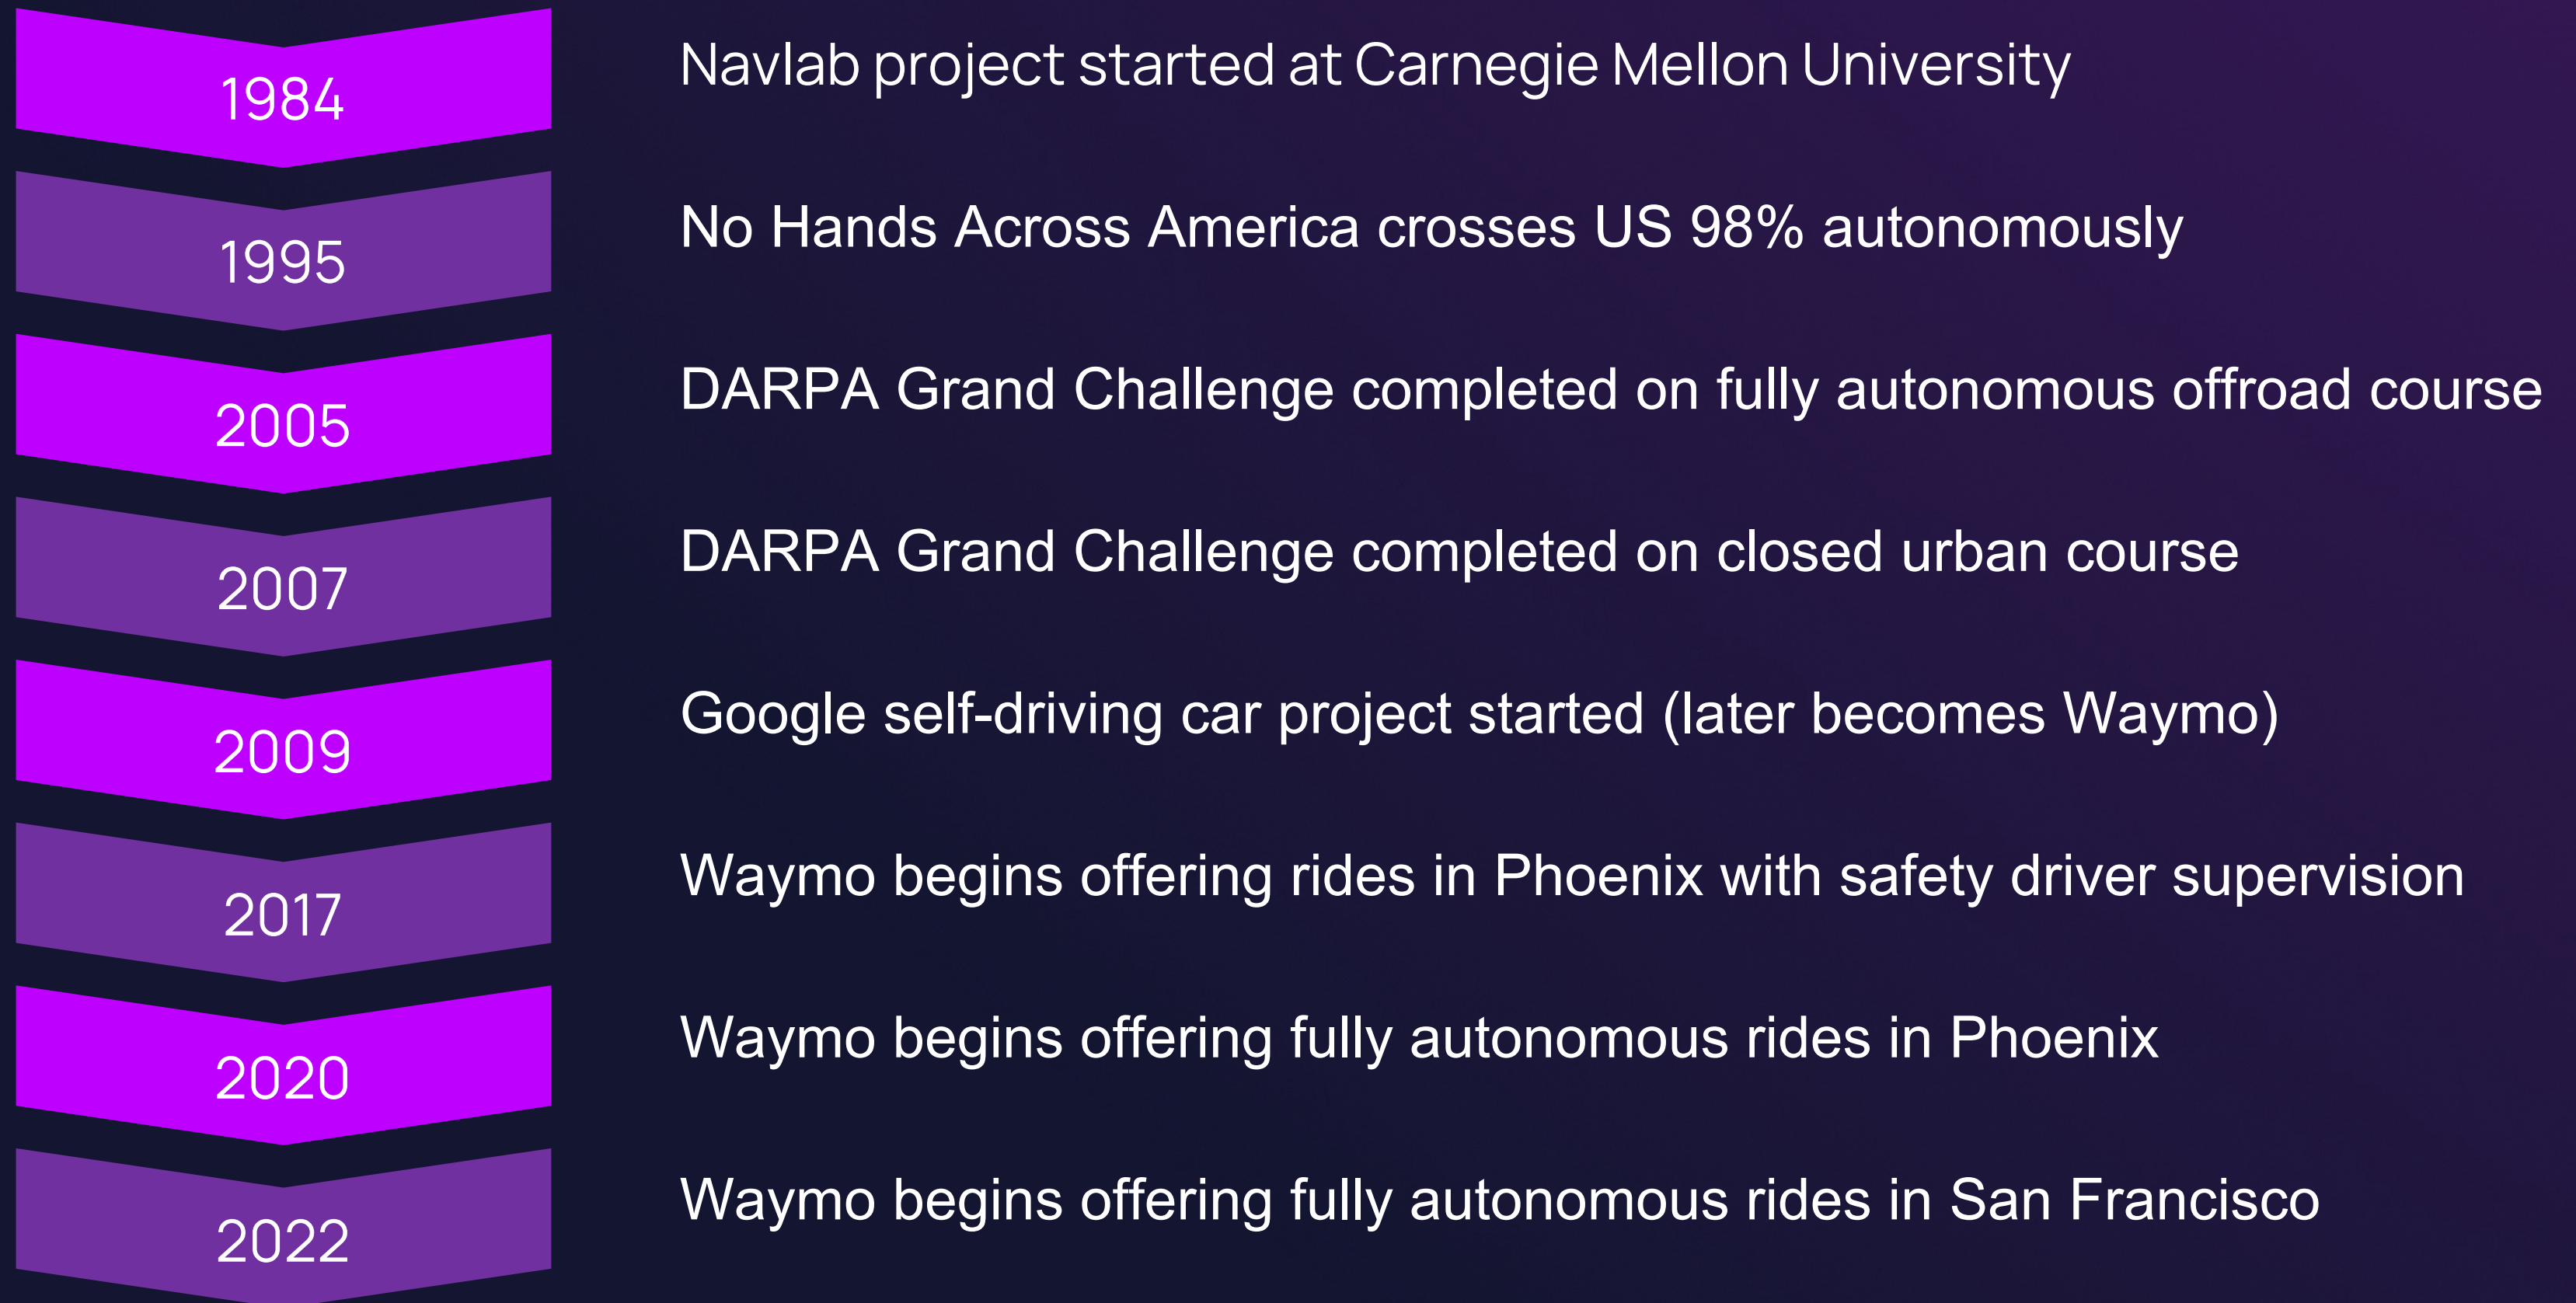
Overview

- Self-driving cars primer
- History of database automation
- Self-managing PostgreSQL at CrystalDB



ELECTRICITY MAY BE THE DRIVER. One day your car may speed along an electric super-highway, its speed and steering automatically controlled by electronic devices embedded in the road. Travel will be more enjoyable. Highways will be made safe—by electricity! No traffic jams . . . no collisions . . . no driver fatigue.

Selected milestones for autonomous vehicles






The DARPA 2005 Grand Challenge Winner: Stanley.

Typical Self-Driving Architecture Has Three Parts

Perception



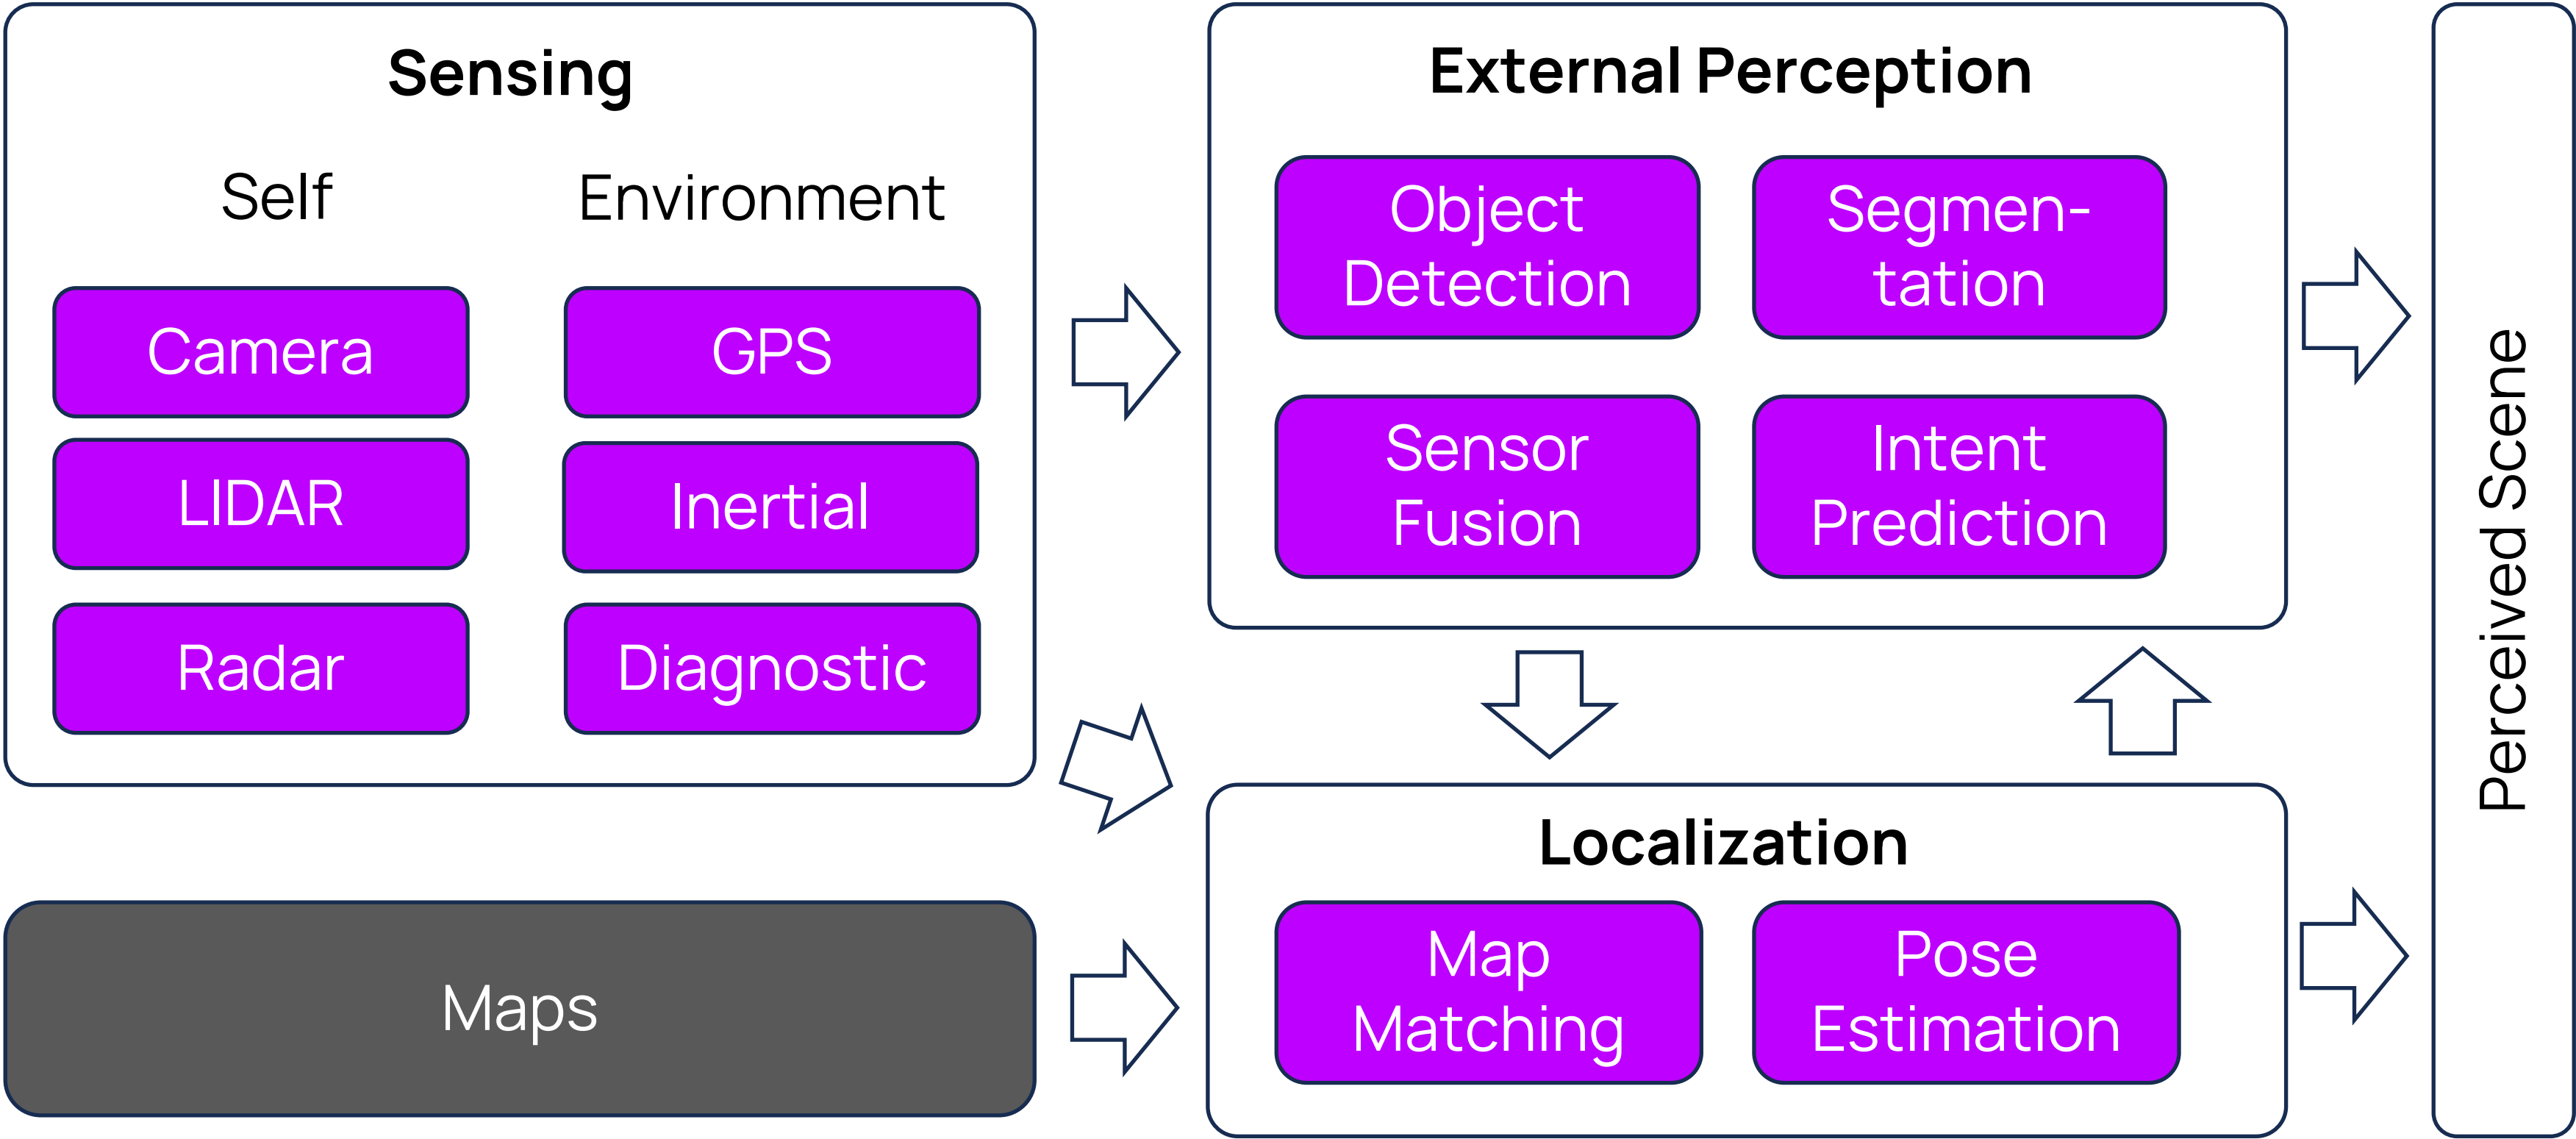
Planning



Control



Perception



weather: rainy
timeofday: daytime
scene: highway

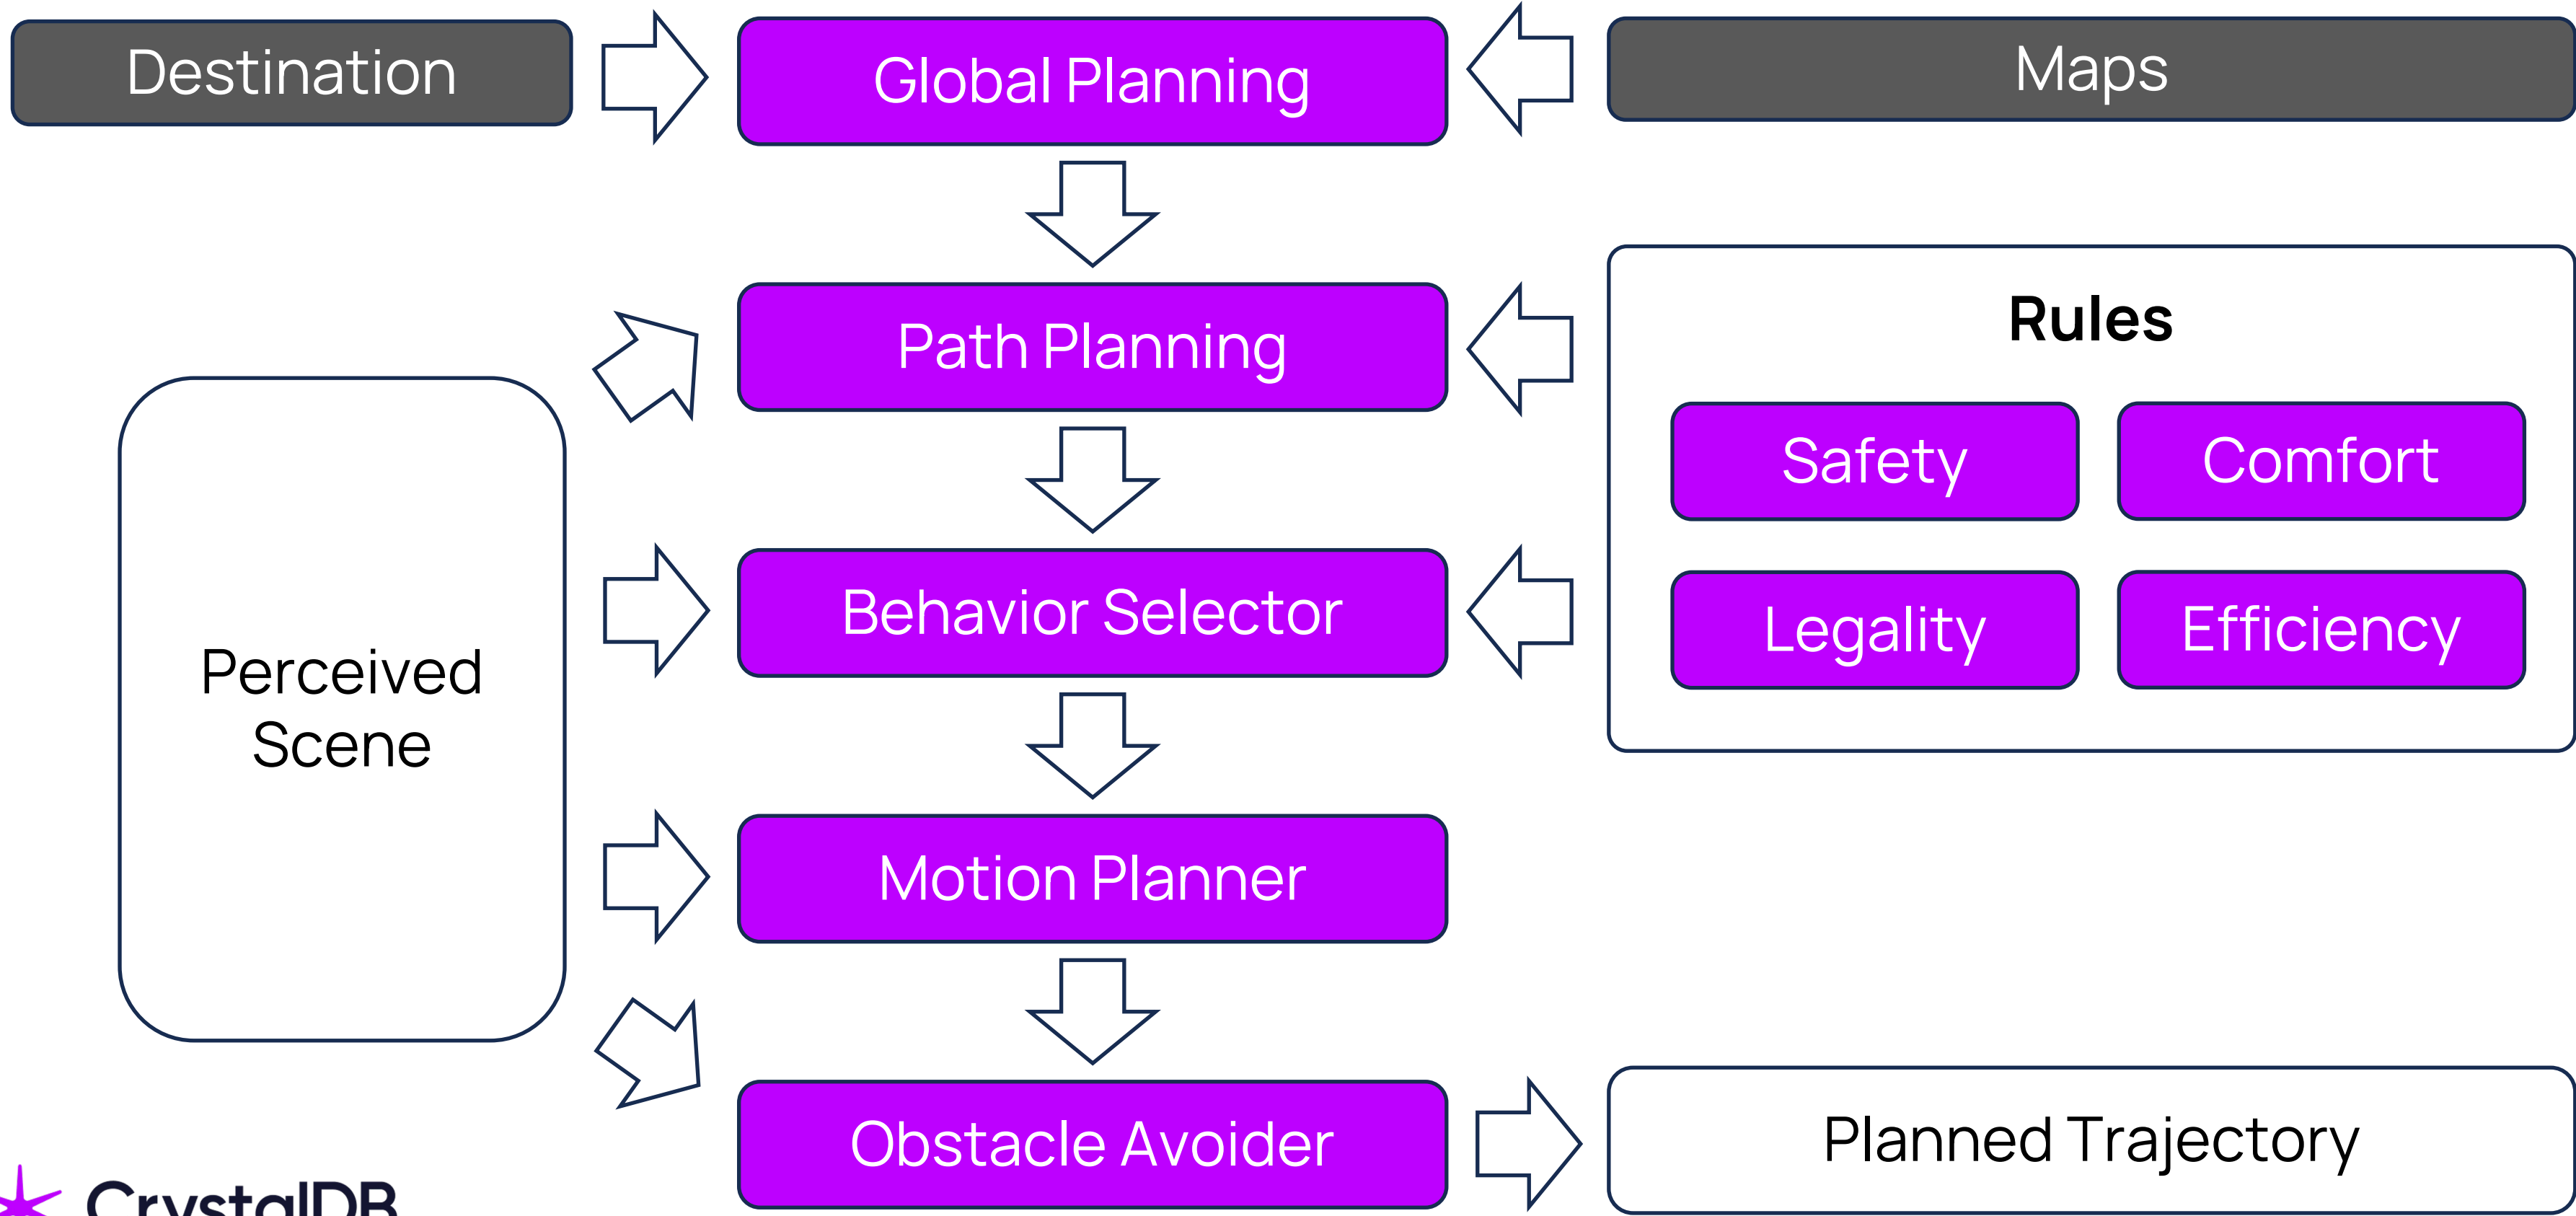




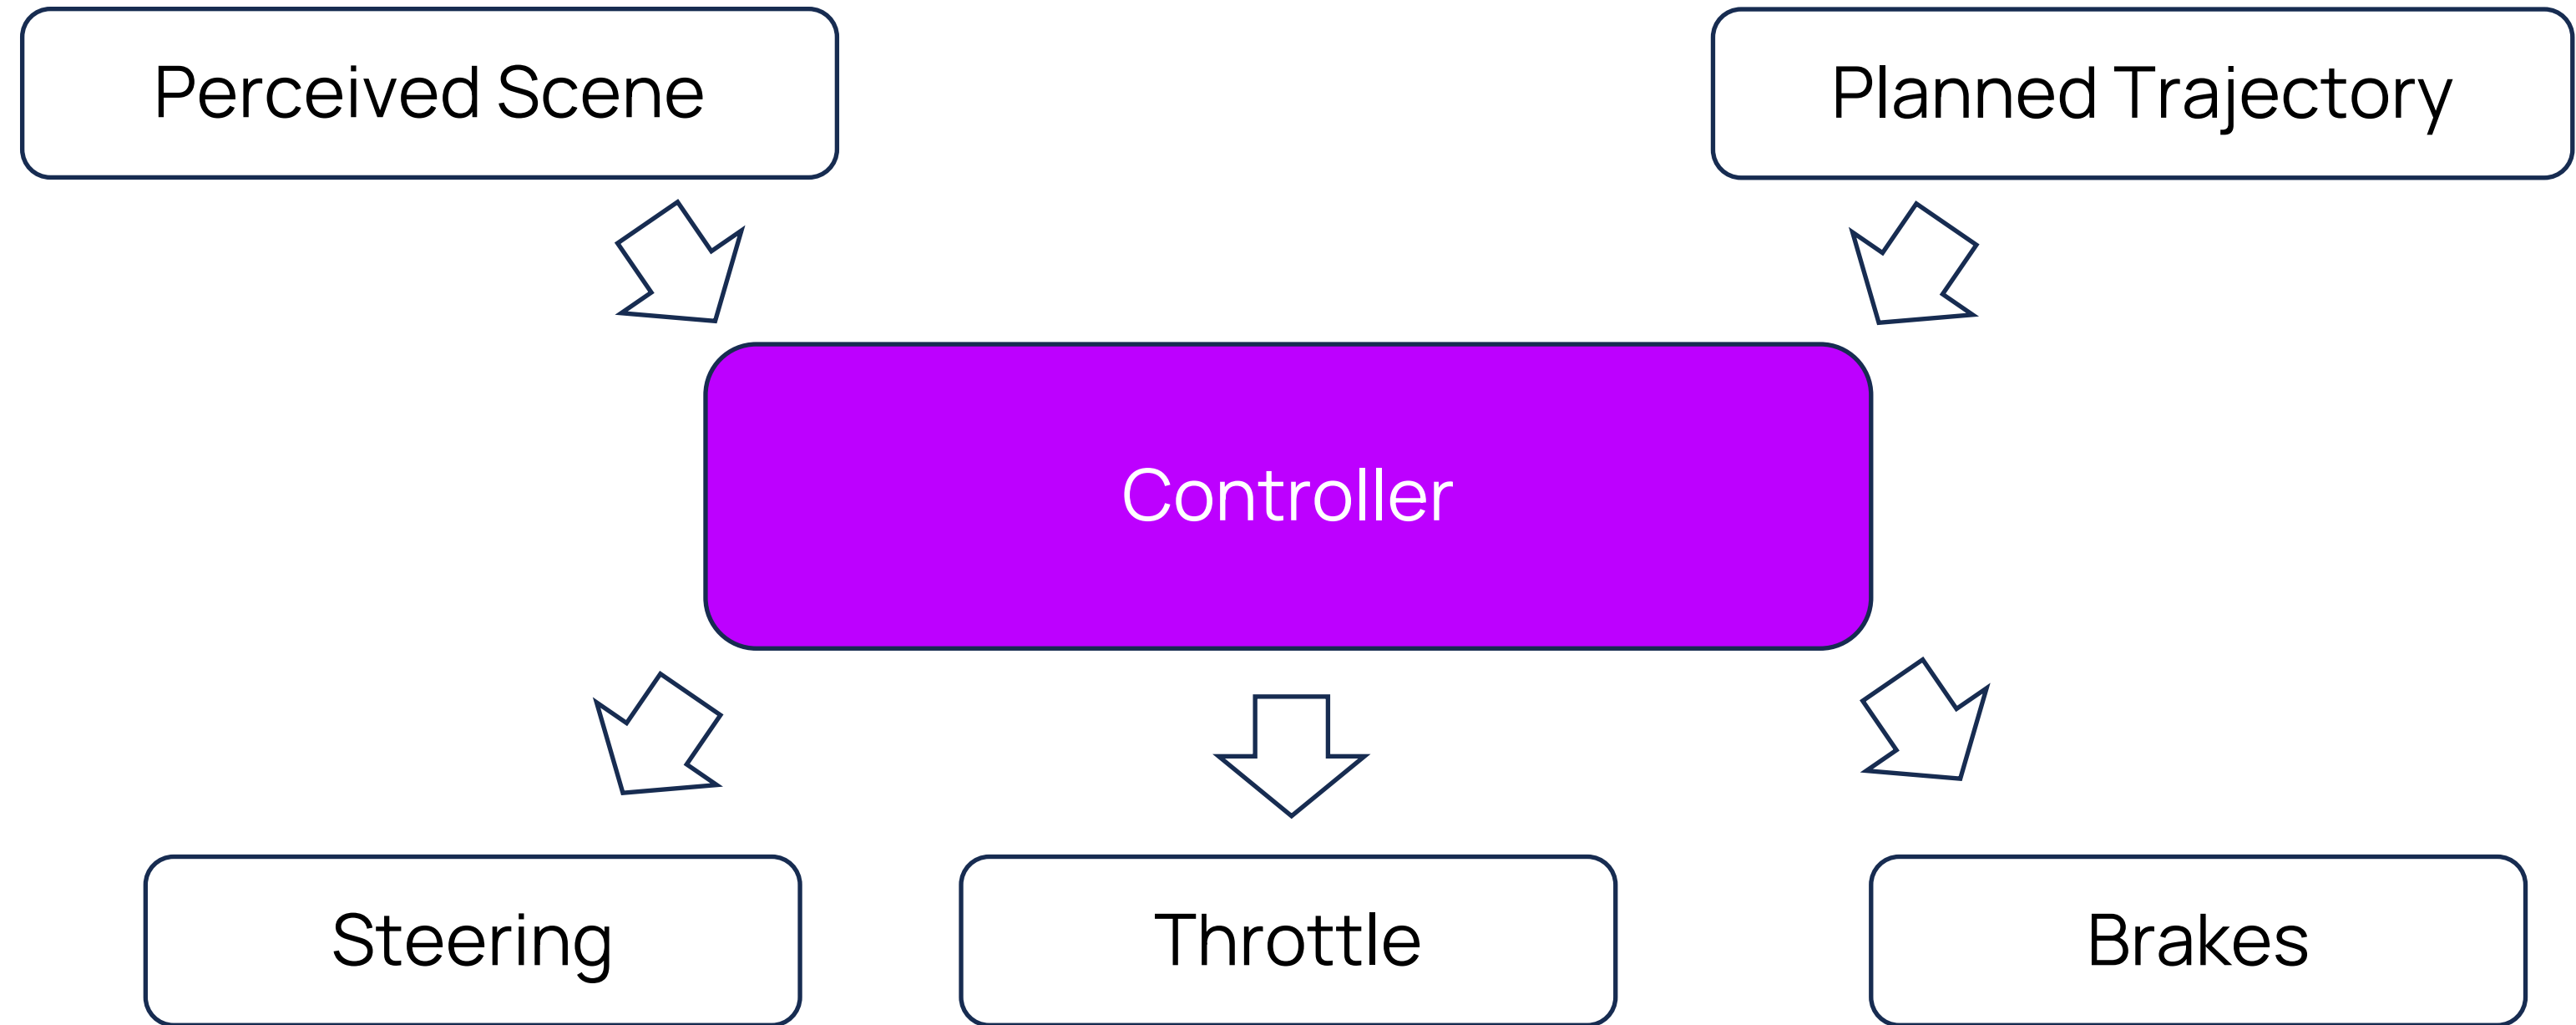




Planning



Control





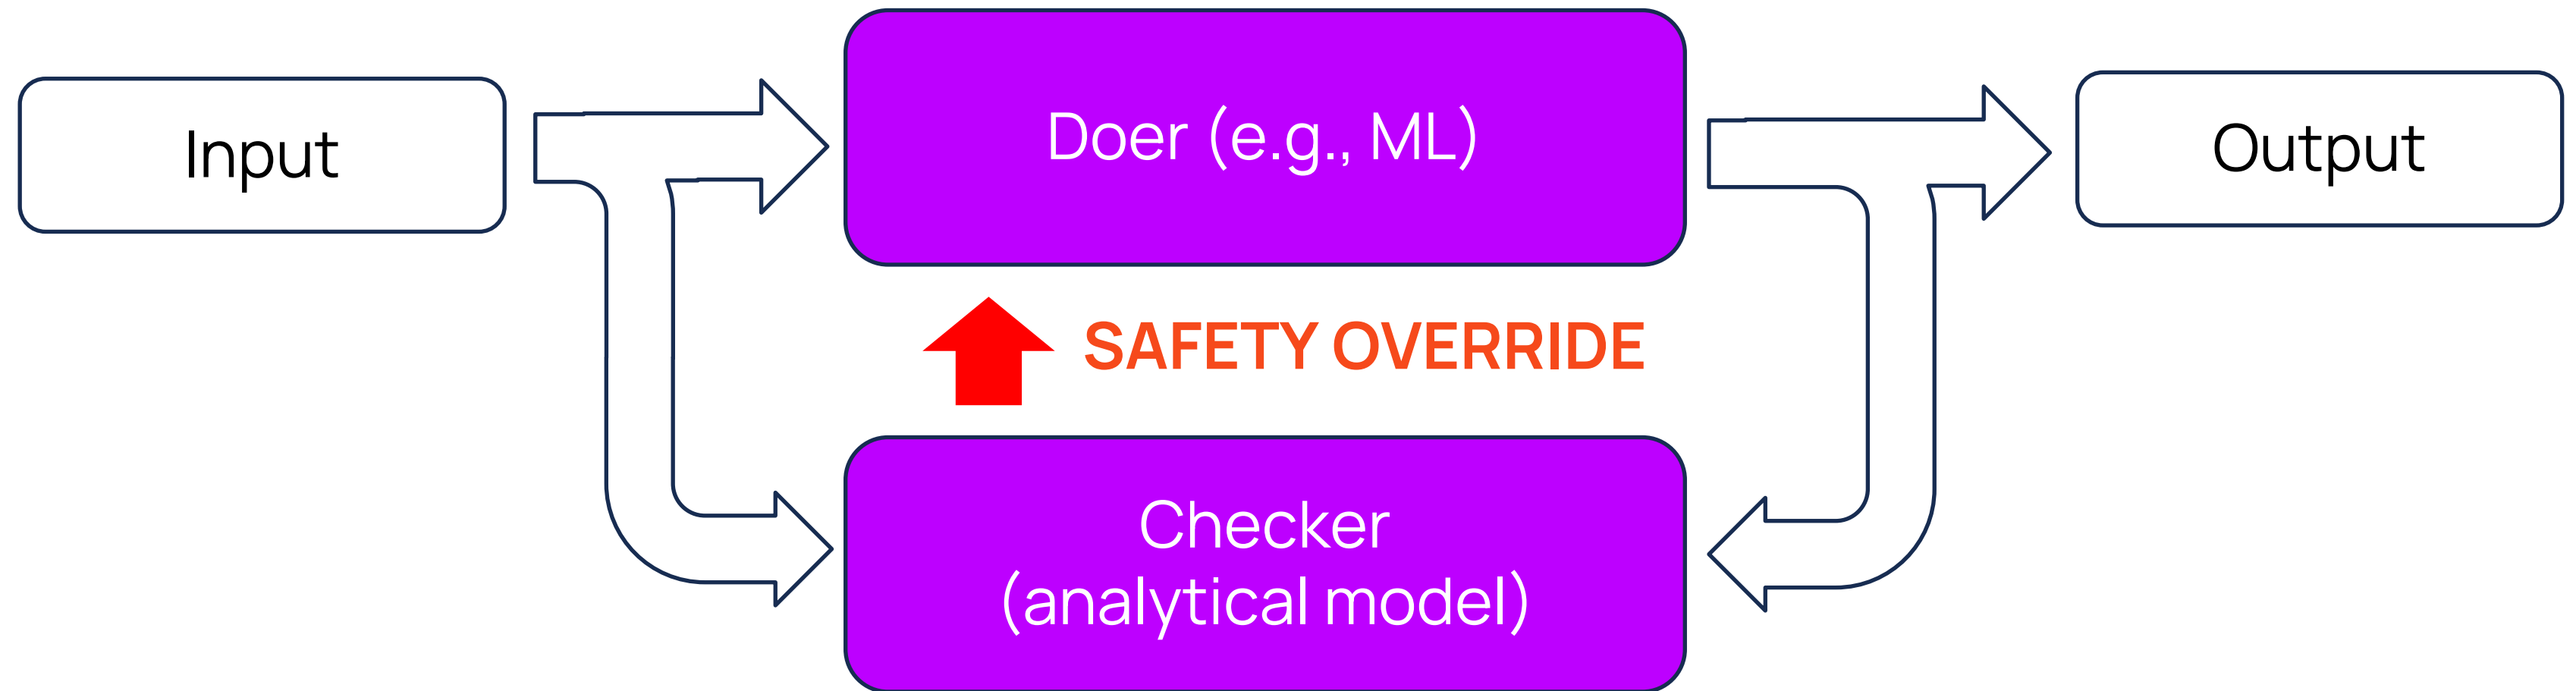
Safety

Safe

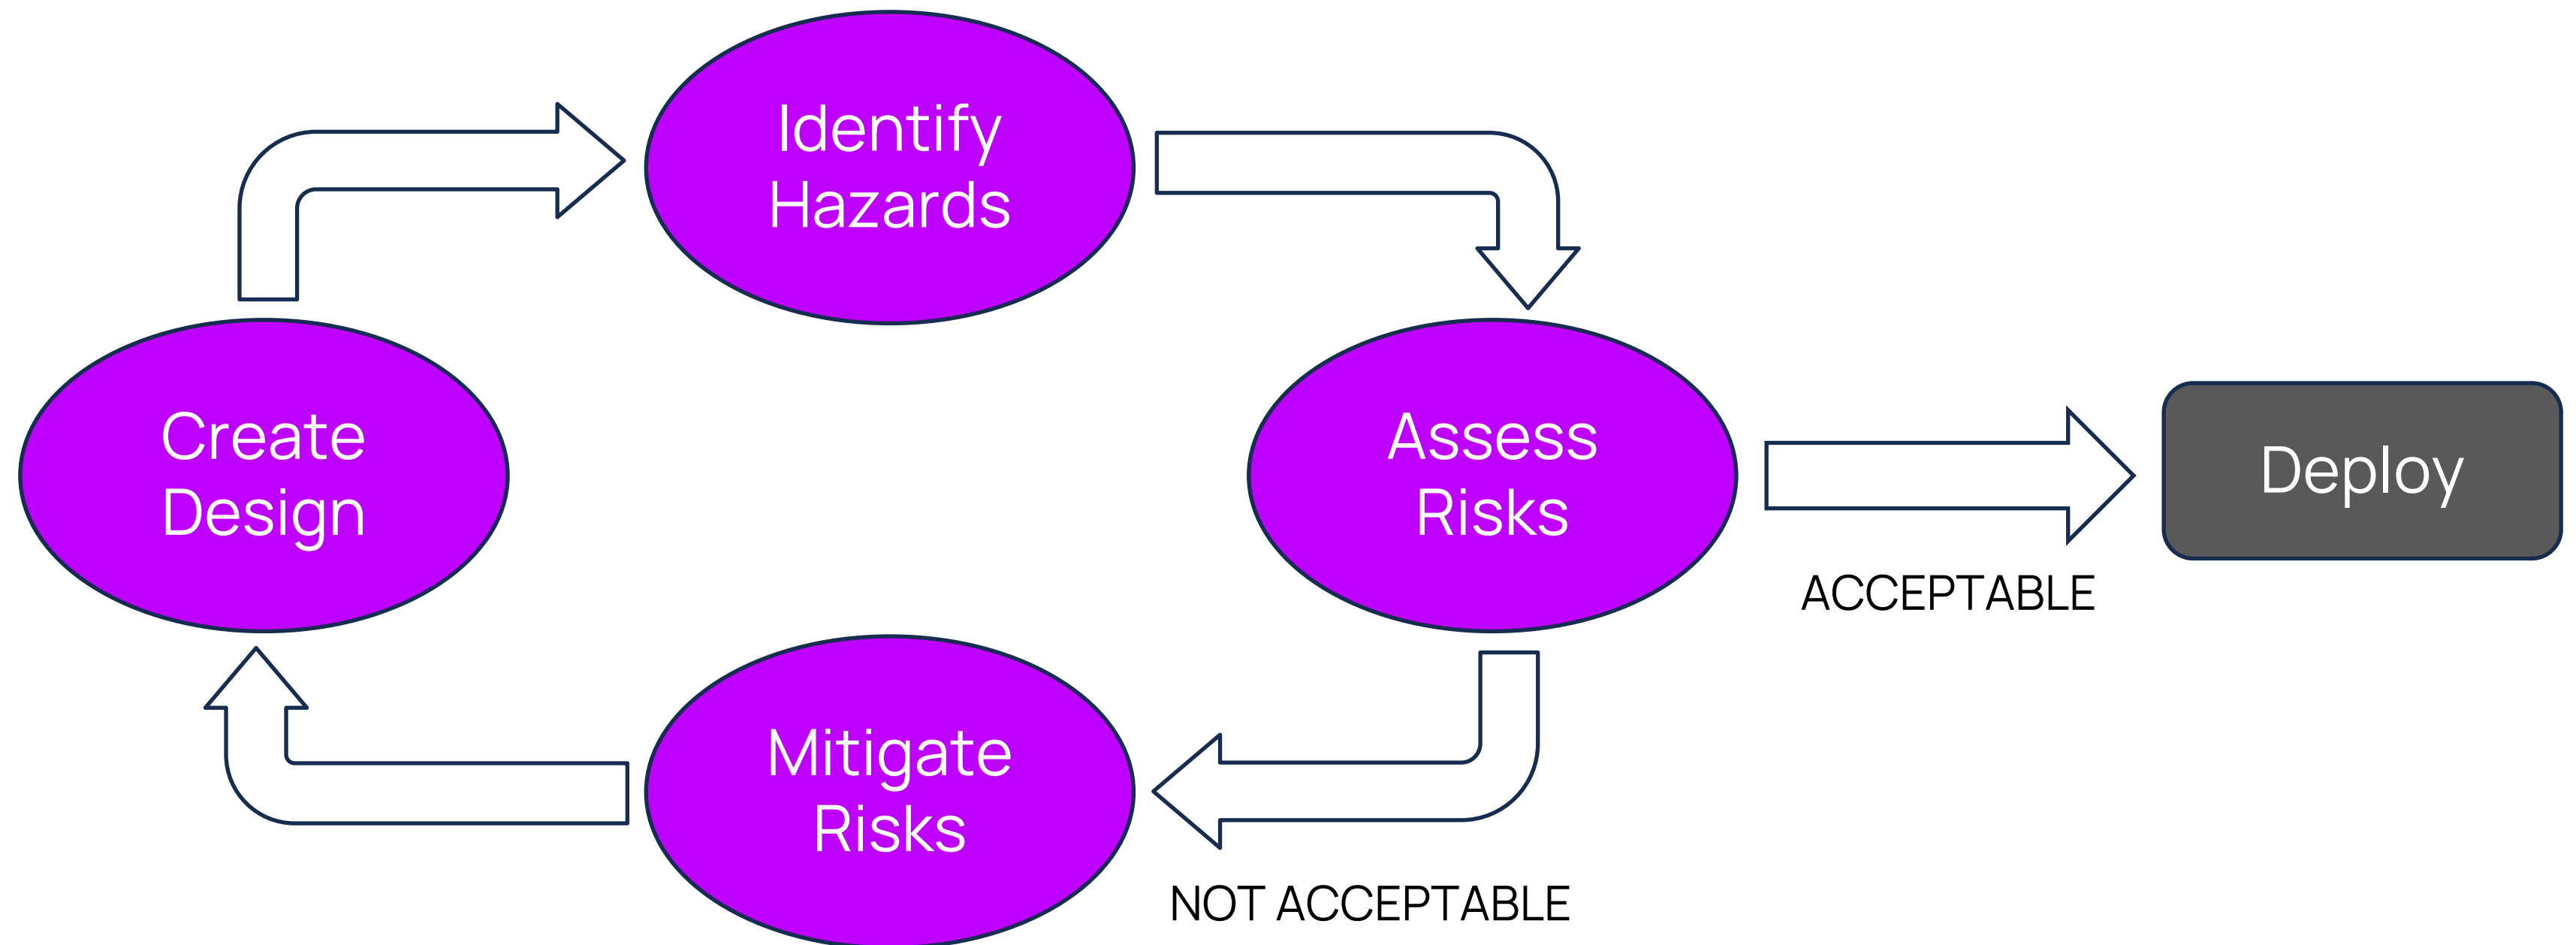
Provably
Safe

Unsafe

Ensuring a safety envelope



Hazard and Risk Analysis



Failure Mode and Effects Analysis

- Potential failure mode
- Potential failure effect
- Severity
- Potential causes
- Expected frequency
- Current controls
- Likelihood of detection
- Computed risk priority
- Recommended action

Process Step	Potential Failure Mode	Potential Failure Effect	SEV ¹	Potential Causes	OCC ²	Current Process Controls	DET ³	RPN ⁴	Action Recommended
What is the step?	In what ways can the step go wrong?	What is the impact on the customer if the failure mode is not prevented or corrected?	How severe is the effect on the customer?	What causes the step to go wrong (i.e., how could the failure mode occur)?	How frequently is the cause likely to occur?	What are the existing controls that either prevent the failure mode from occurring or detect it should it occur?	How probable is detection of the failure mode or its cause?	Risk priority number calculated as SEV x OCC x DET	What are the actions for reducing the occurrence of the cause or for improving its detection? Provide actions on all high RPNs and on severity ratings of 9 or 10.
ATM Pin Authentication	Unauthorized access	• Unauthorized cash withdrawal • Very dissatisfied customer	8	Lost or stolen ATM card	3	Block ATM card after three failed authentication attempts	3	72	
	Authentication failure	Annoyed customer	3	Network failure	5	Install load balancer to distribute work-load across network links	5	75	
Dispense Cash	Cash not disbursed	Dissatisfied customer	7	ATM out of cash	7	Internal alert of low cash in ATM	4	196	Increase minimum cash threshold limit of heavily used ATMs to prevent out-of-cash instances
	Account debited but no cash disbursed	Very dissatisfied customer	8	• Transaction failure • Network issue	3	Install load balancer to distribute work-load across network links	4	96	
	Extra cash dispensed	Bank loses money	8	• Bills stuck to each other • Bills stacked incorrectly	2	Verification while loading cash in ATM	3	48	

Source: isixsigma.com

Safety Cases


Claim: Vehicle stops at stop signs

Argument: Recognizes stop signs and plans accordingly


Evidence: Simulation testing, road testing

Sub-Claims:



















- Recognizes stop signs
- Recognizes limit lines
- Adjusts speed to ensure it can stop



Simulation

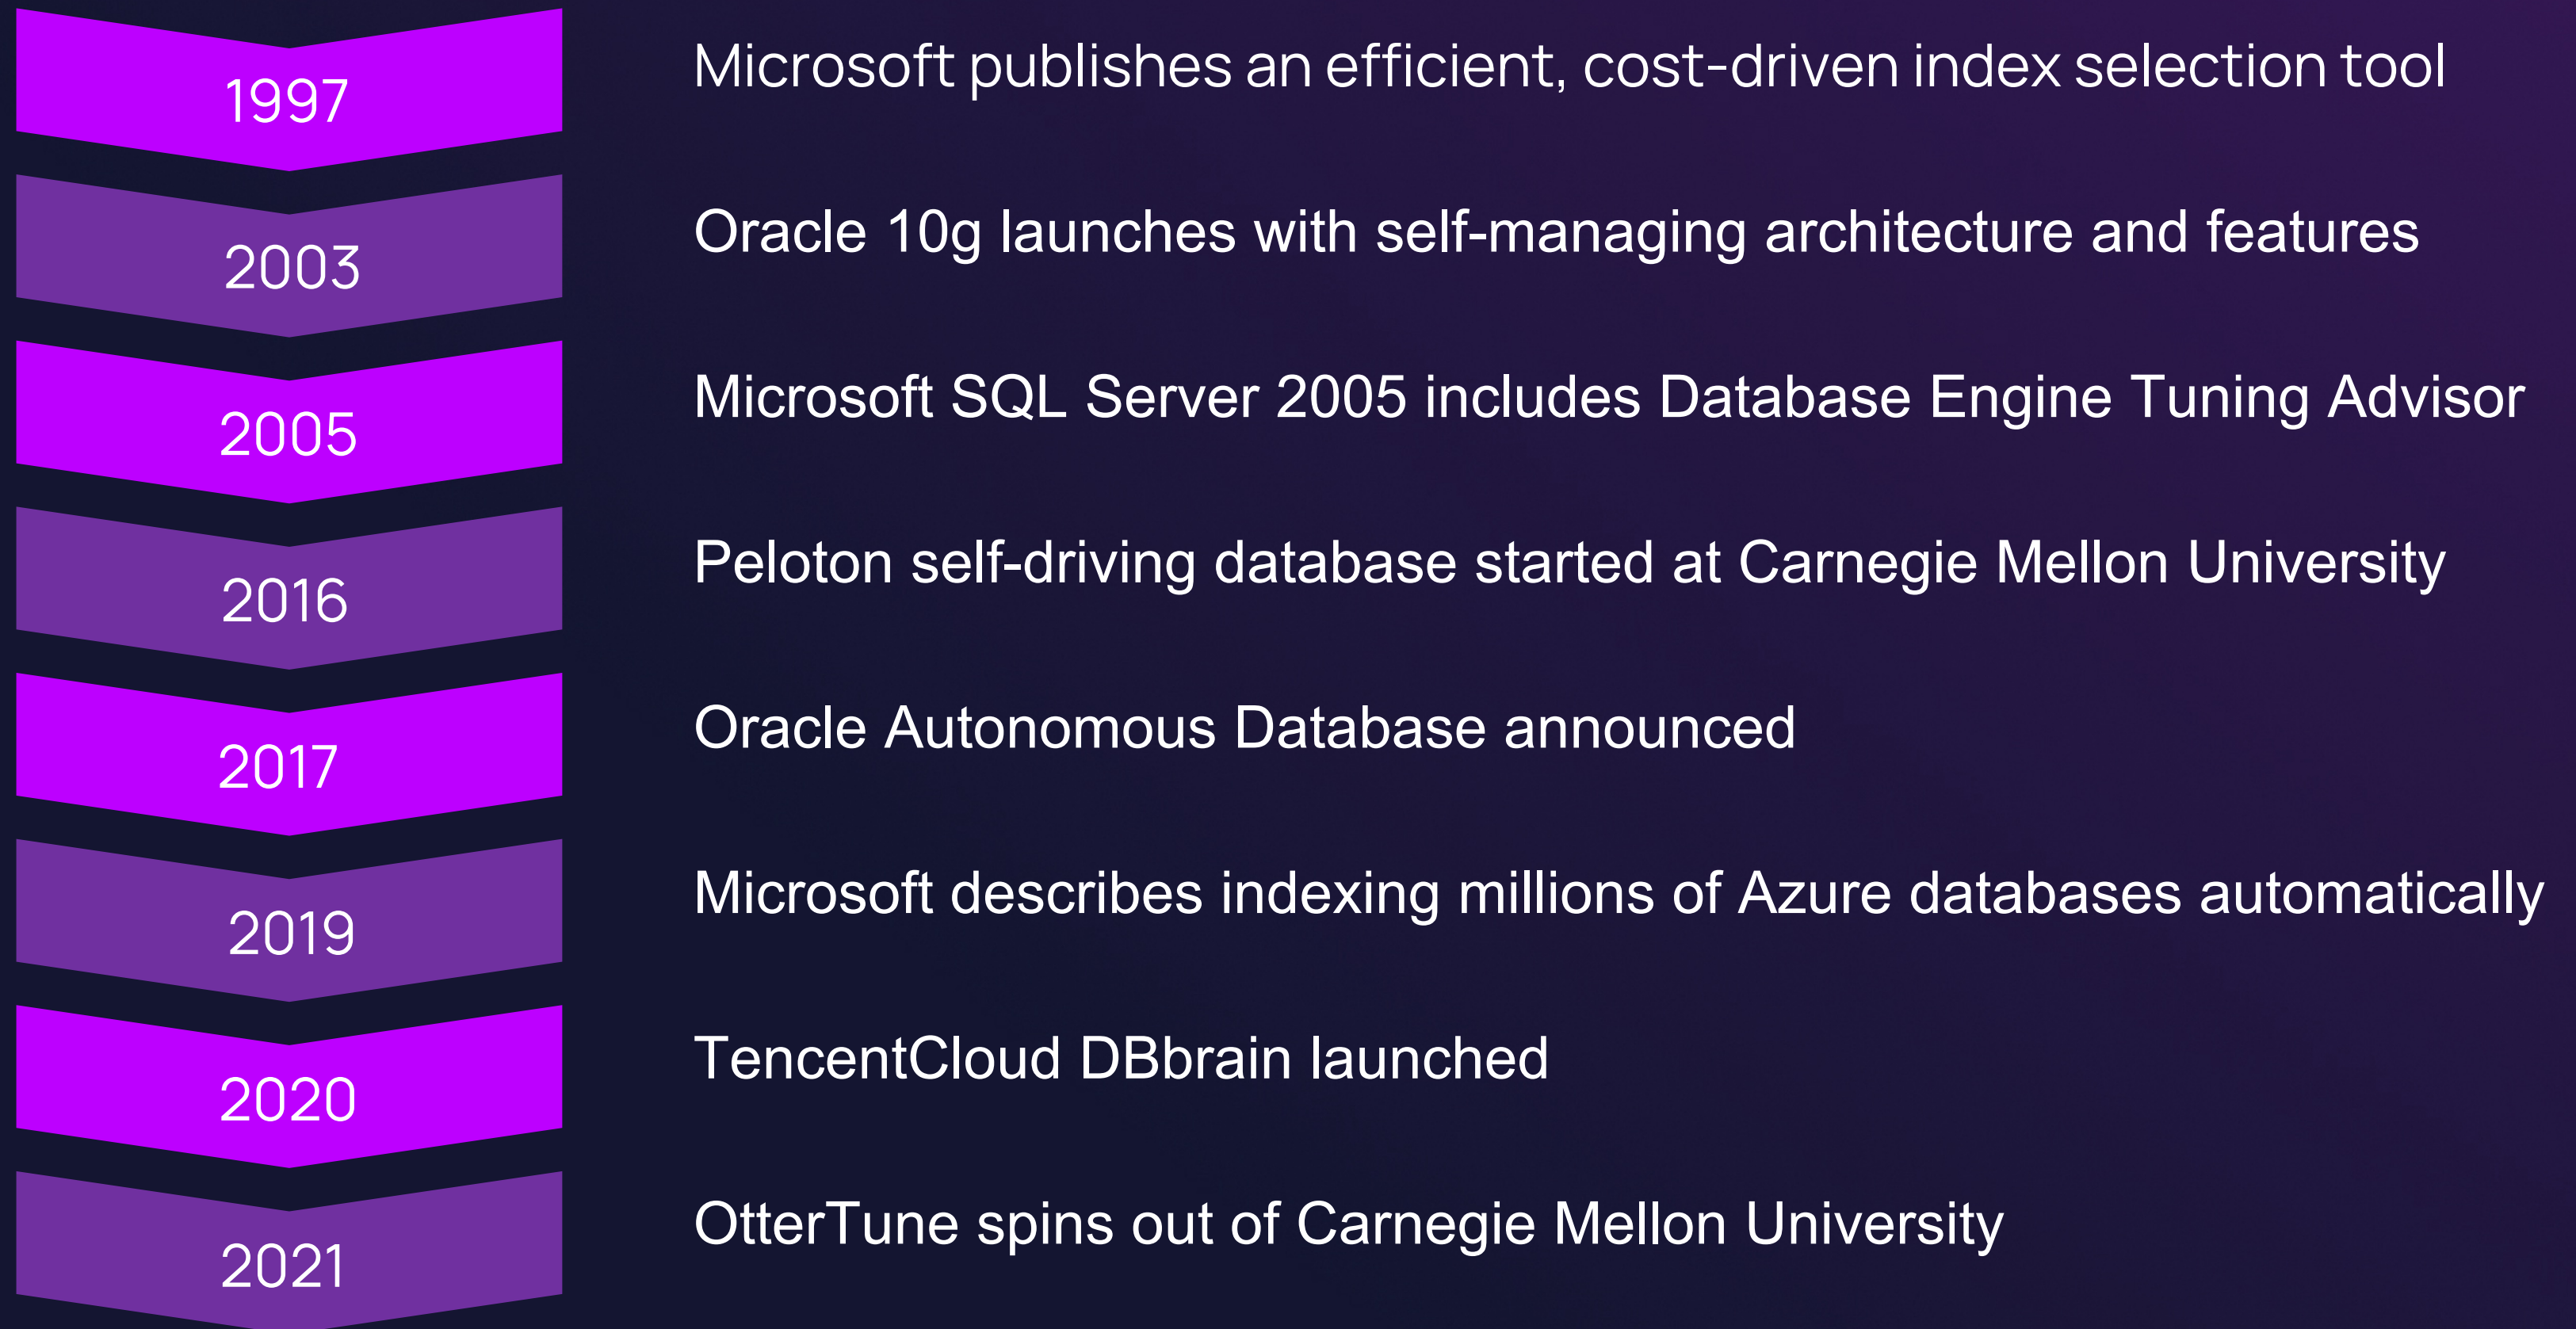


Levels of Automation (SAE J3016)

			Steering and acceleration/ deceleration	Monitoring of driving environment	Fall-back when automation fails (DDT fall-back)	Operational Design Domain
Human driver monitors the road	0 NO AUTOMATION				LIMITED	
	1 DRIVER ASSISTANCE				LIMITED	
	2 PARTIAL AUTOMATION				LIMITED	
Automated driving system monitors the road	3 CONDITIONAL AUTOMATION				LIMITED	
	4 HIGH AUTOMATION				LIMITED	
	5 FULL AUTOMATION				UNLIMITED	



Selected milestones for autonomous databases



Aims, Areas, and Approaches

Areas to automate

- Physical design (indexing)
- Query tuning
- Knob settings
- Hardware selection
- Troubleshooting
- High availability
- Security


Aim of automation

- Assist DBAs
- Replace DBAs

DB System support

- Metric collectors
- Query optimizer integration

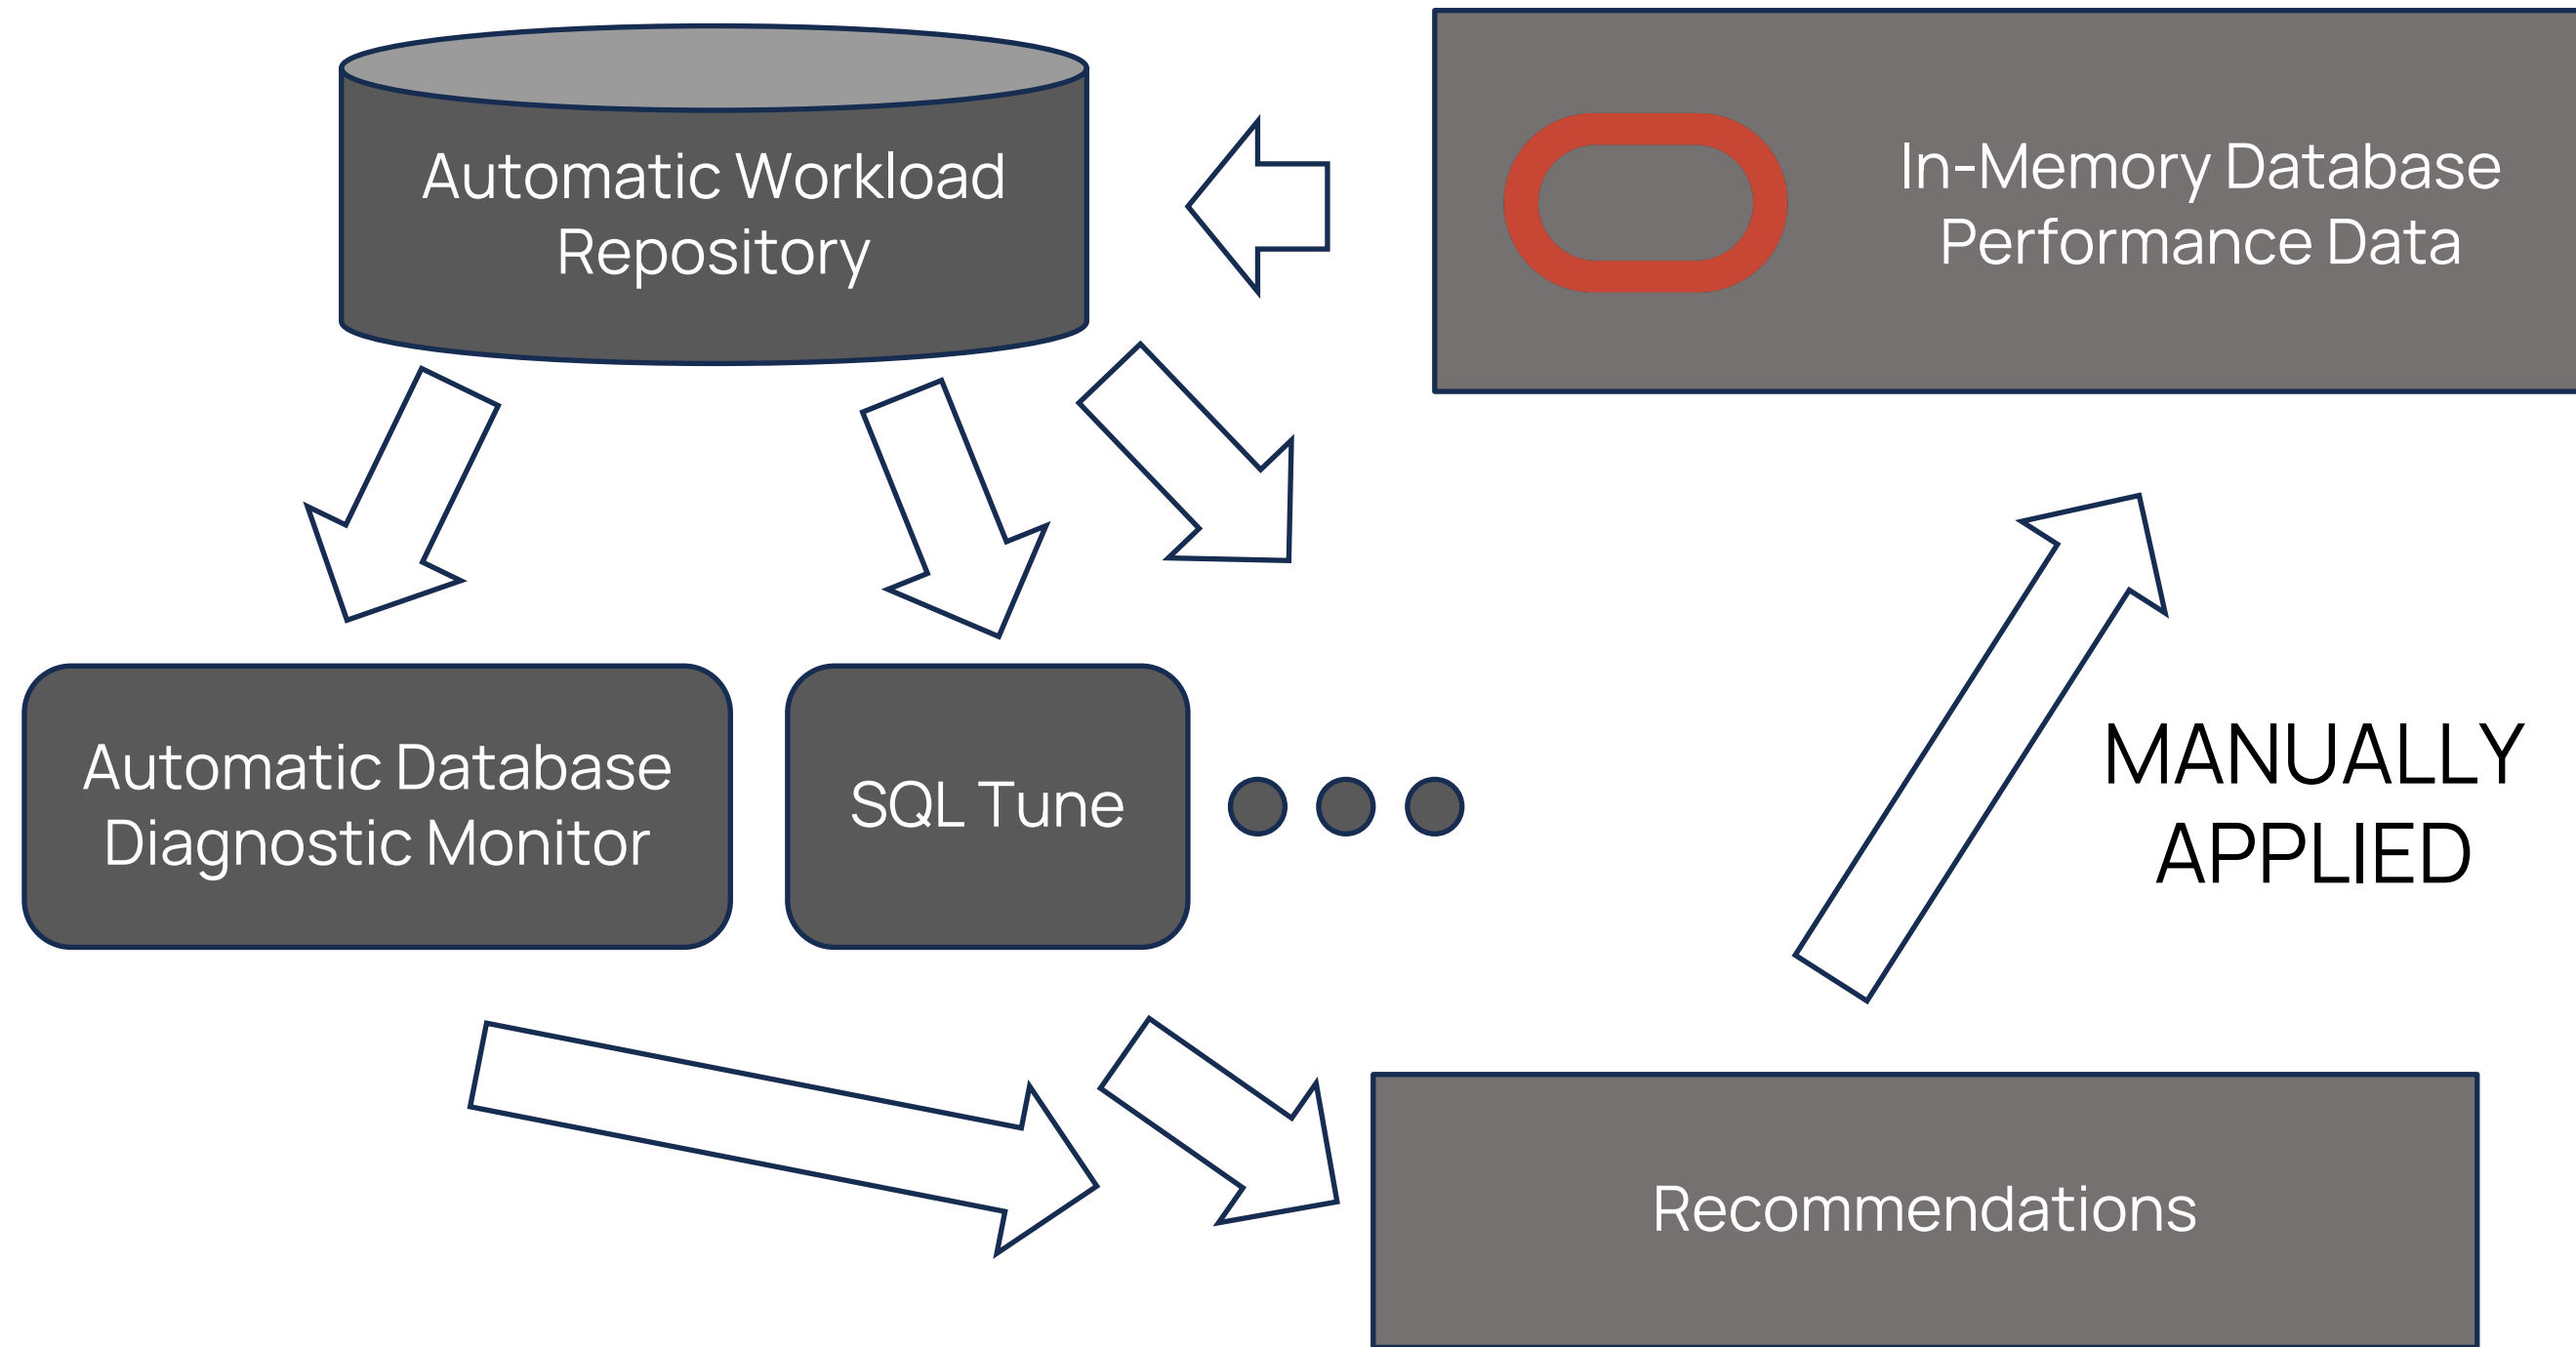
Oracle Autonomous Database



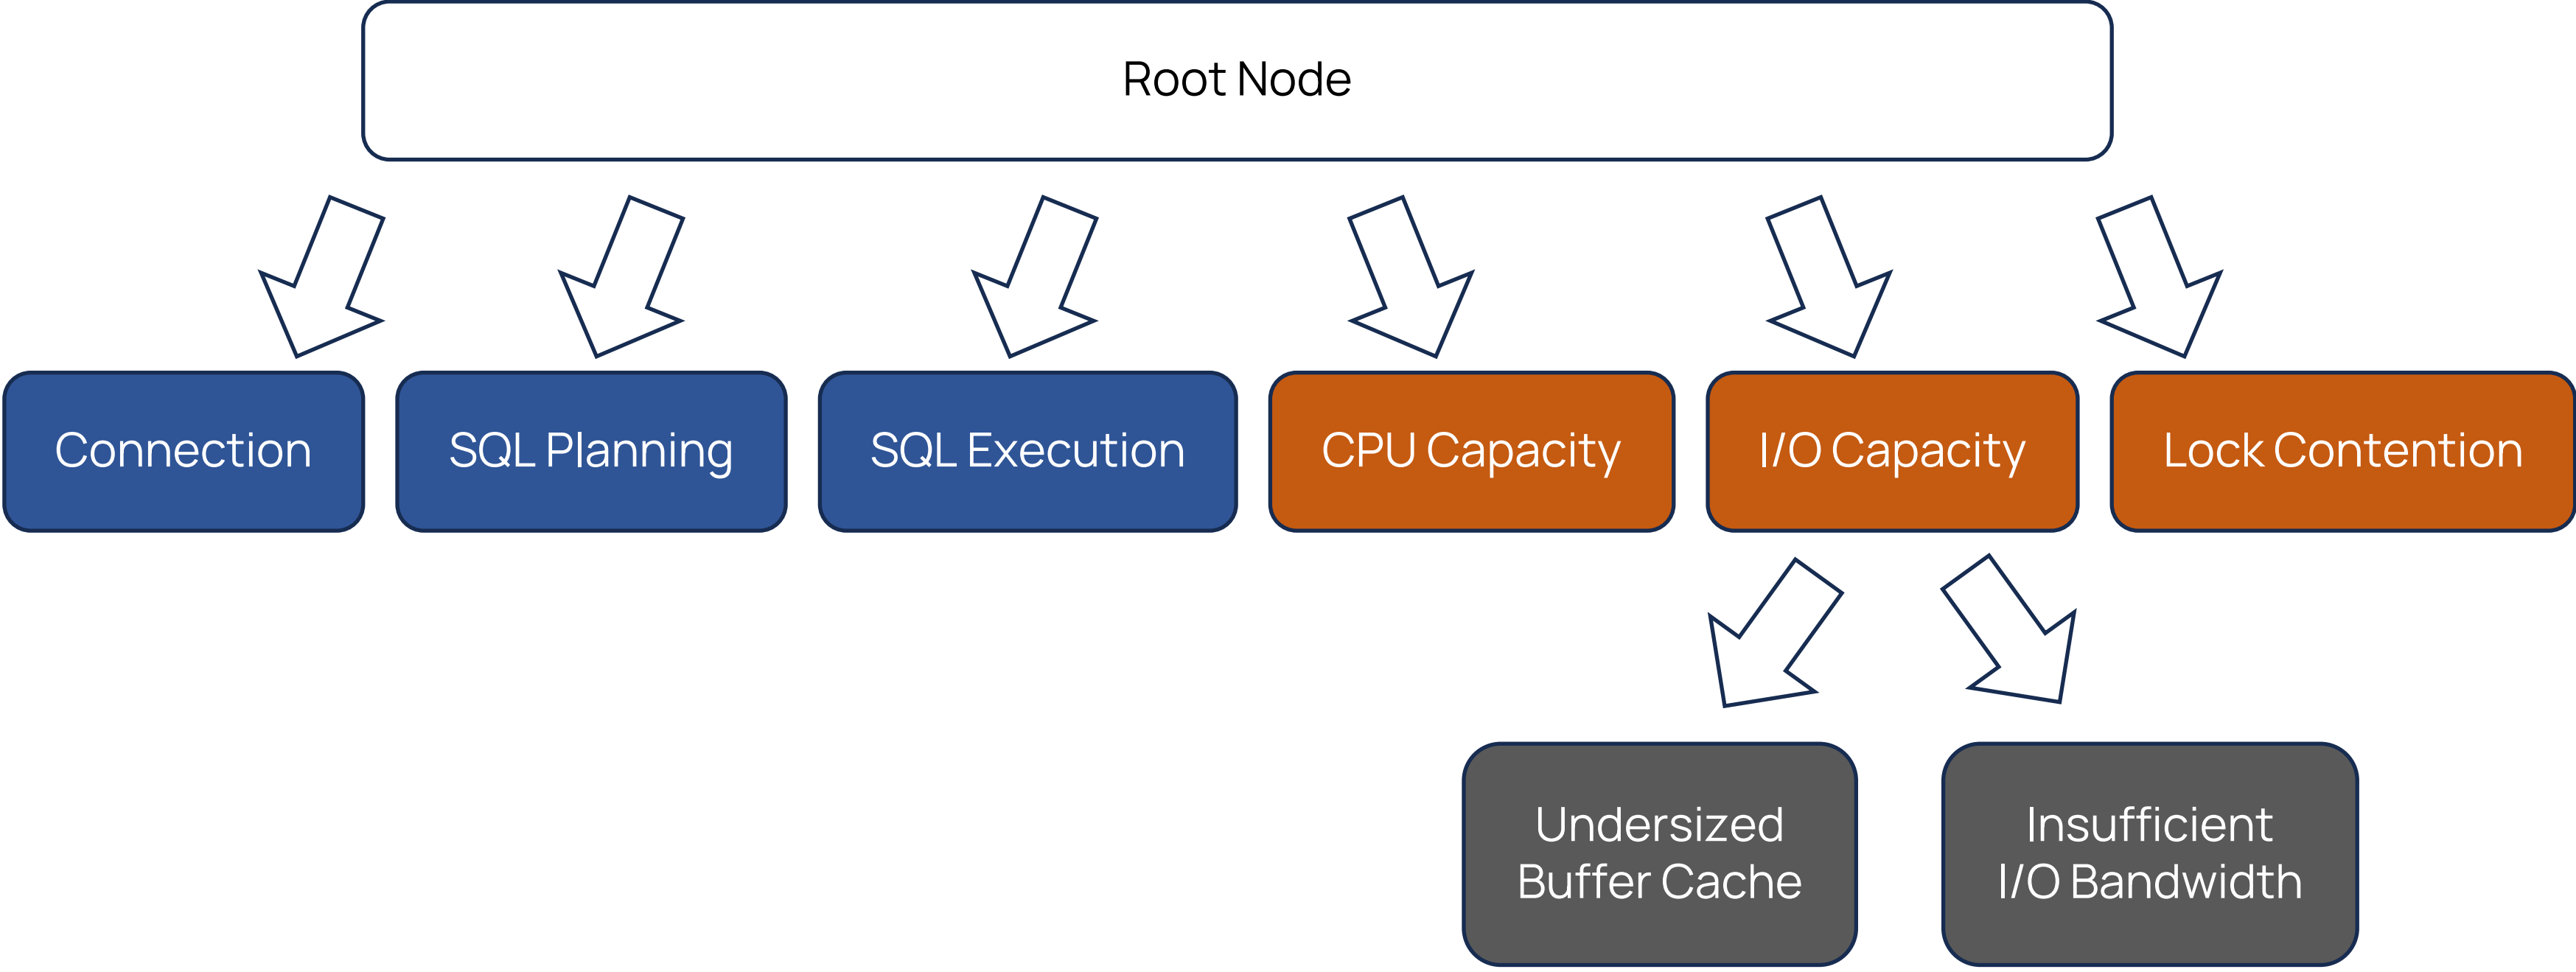
Release 10g in 2003

- Automatic storage management
- Automatic workload repository
- Automatic database diagnostic monitor
- SQL tuning advisor
- SQL access advisor
- Automatic memory management

Oracle 10g Architecture



Database time statistics support bottleneck analysis

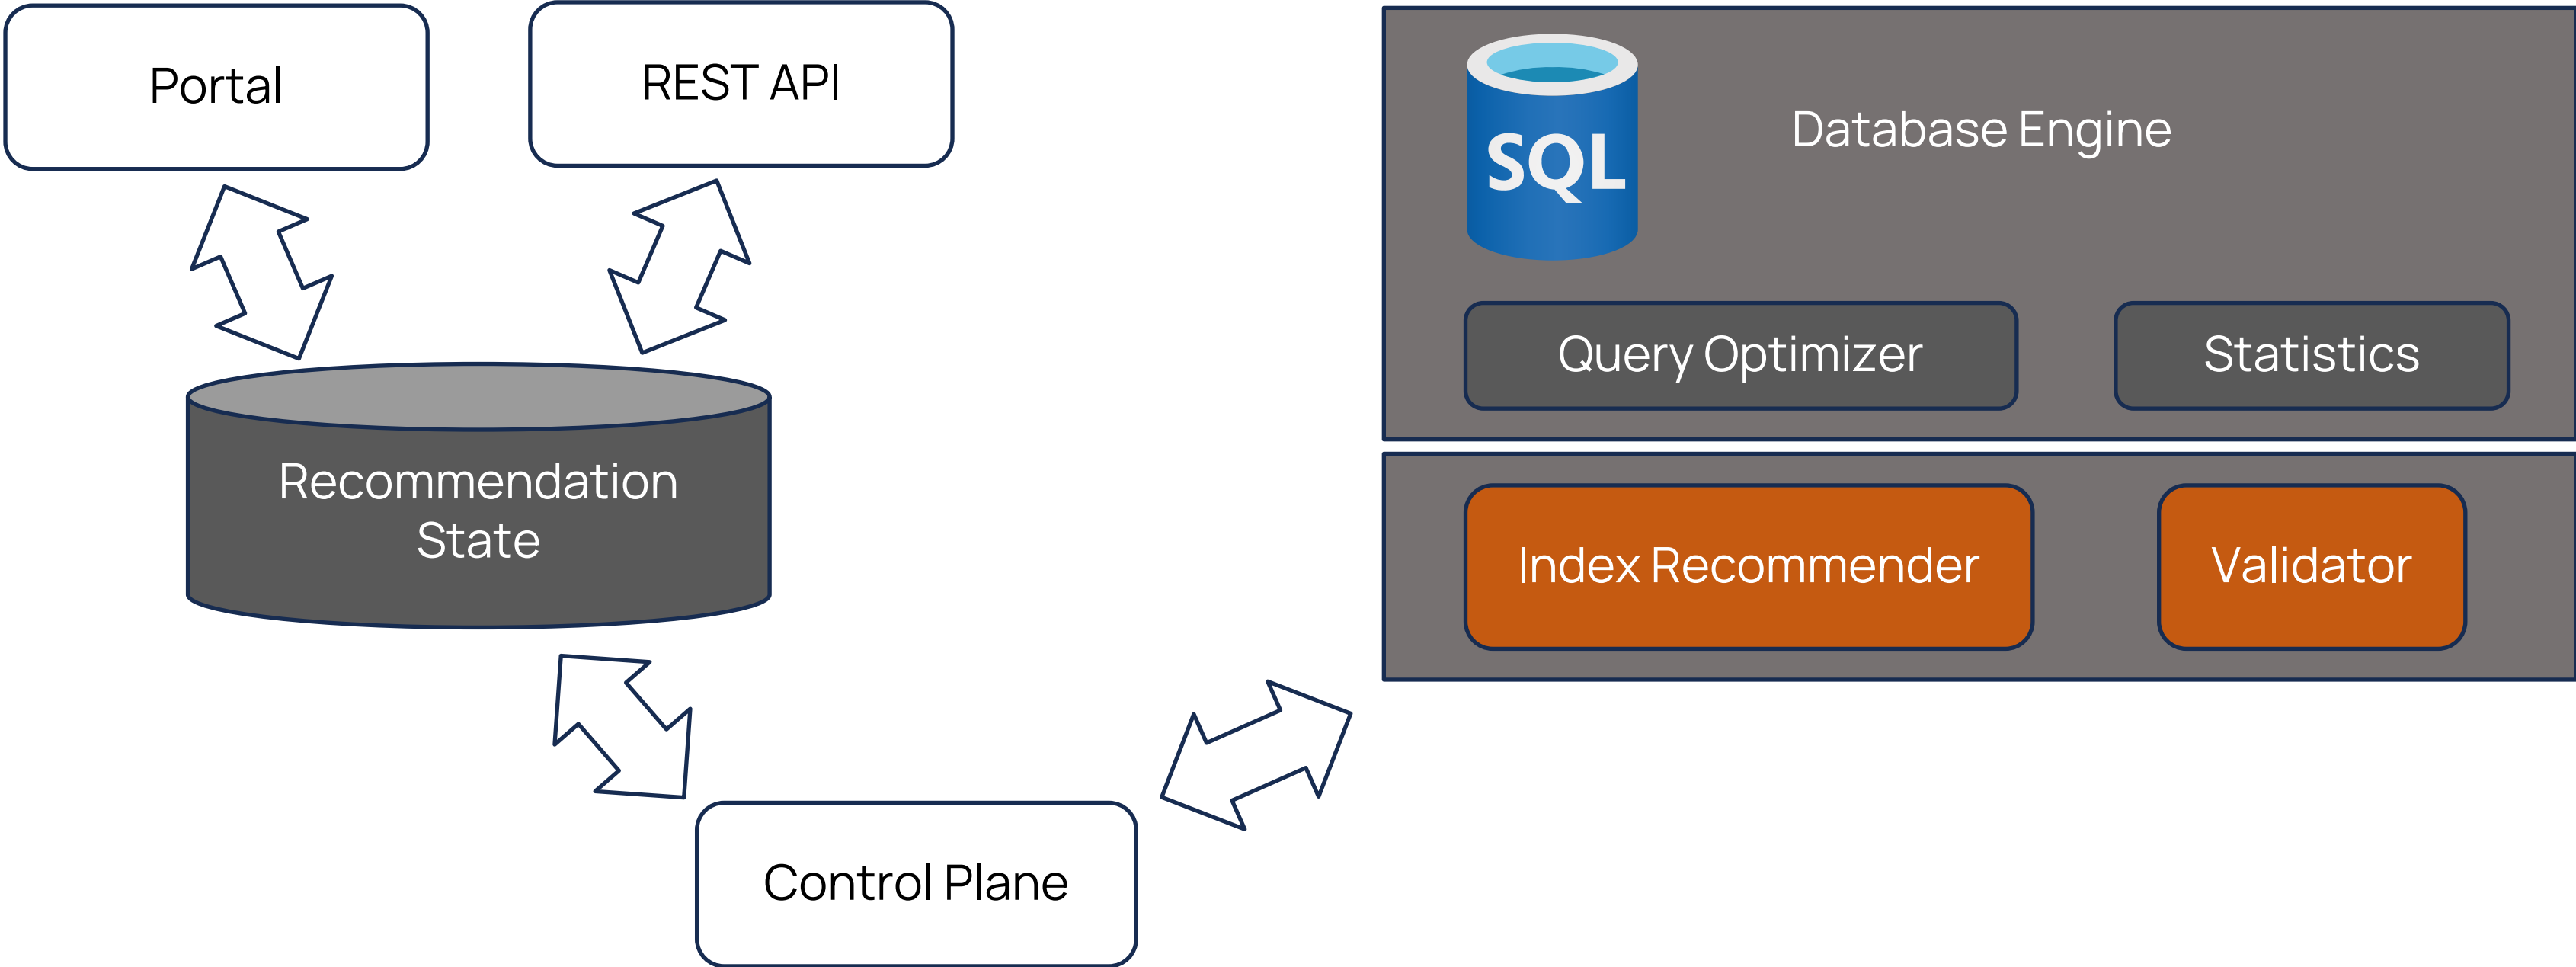


Automatic Tuning in Azure SQL Database

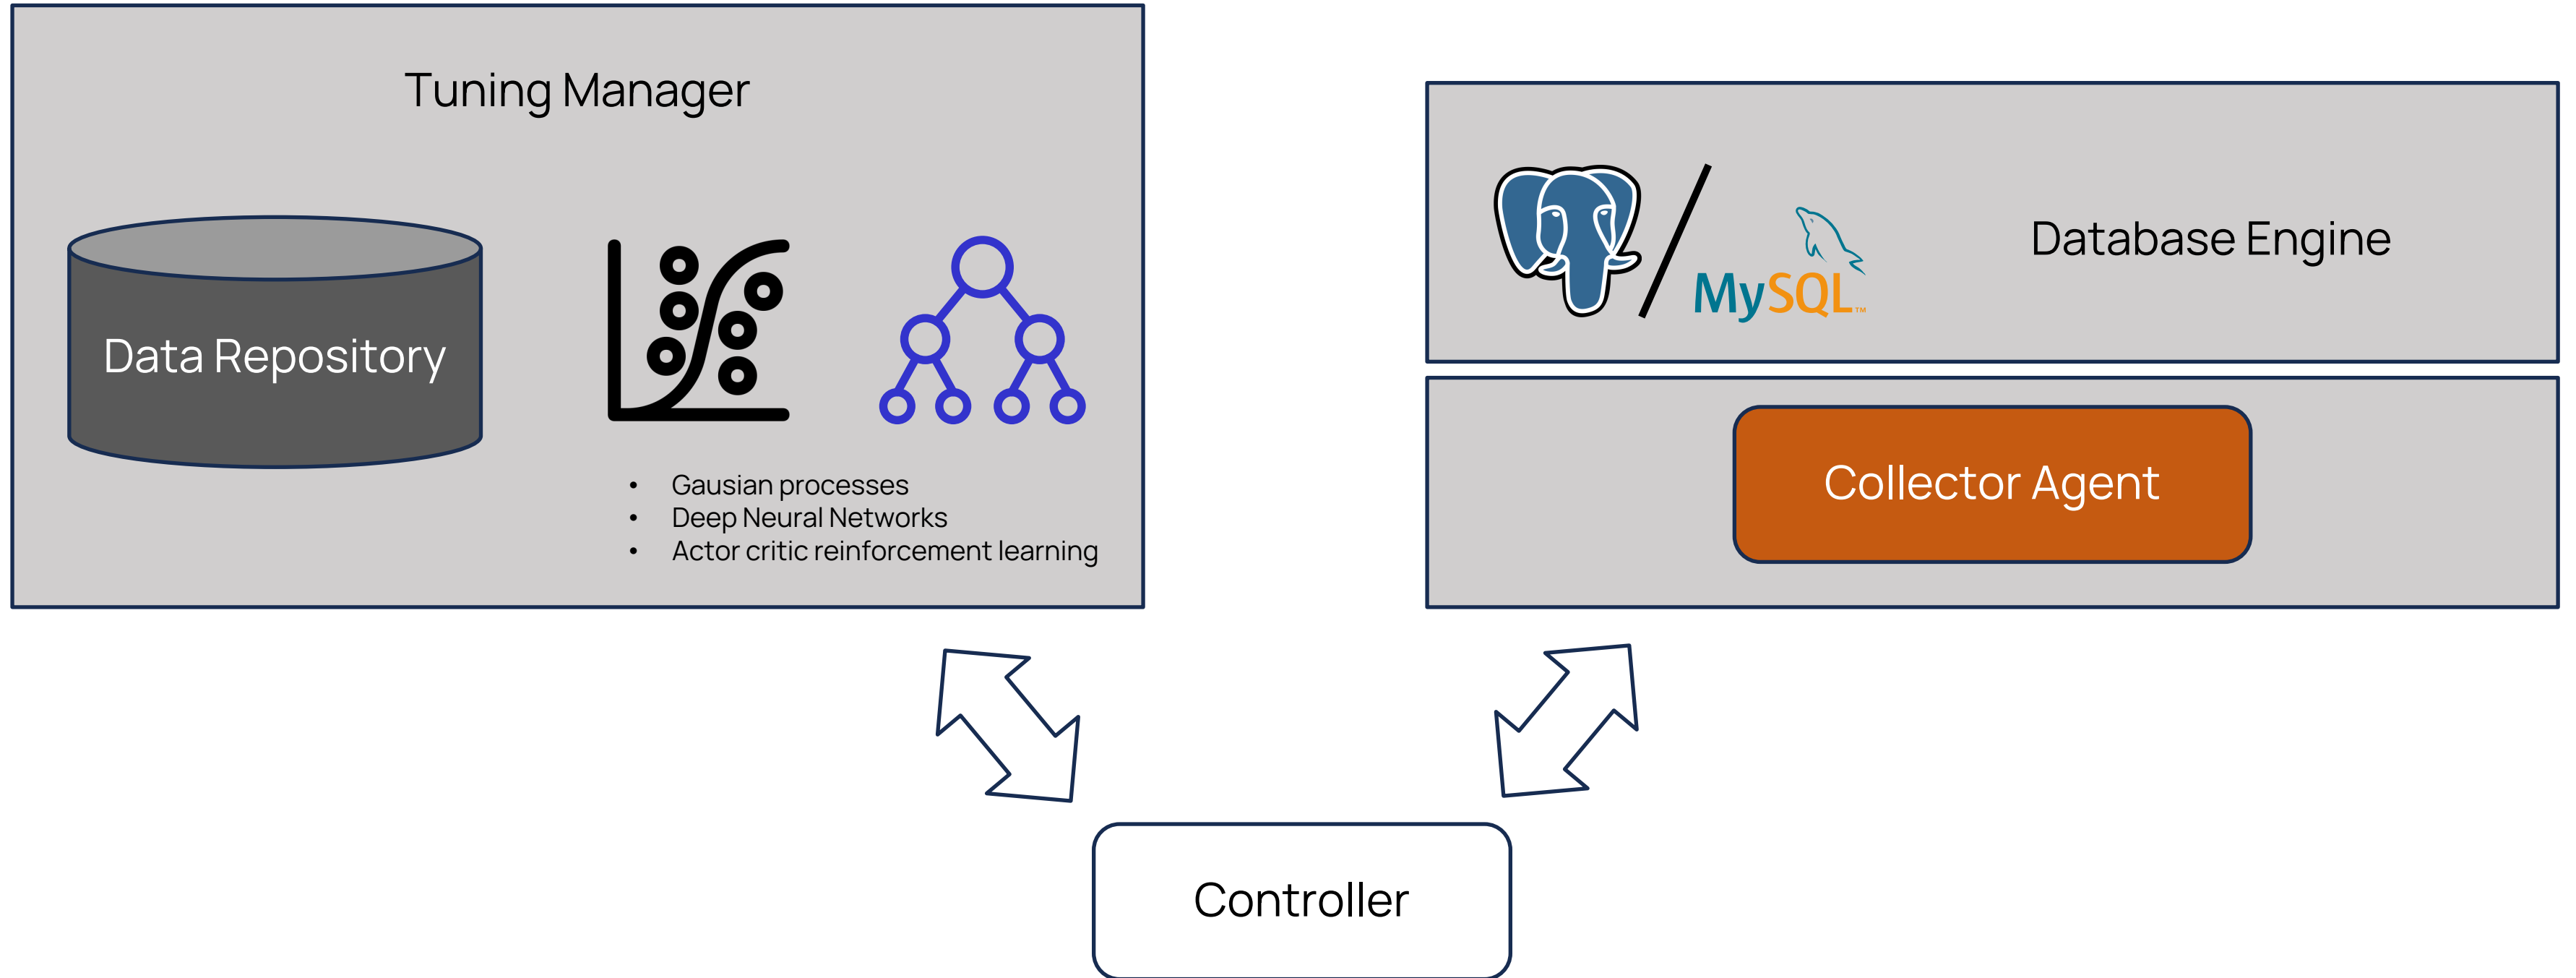


- AutoAdmin research (1997) created technology
- SQL server 2005 – Database Engine Tuning Advisor
- Today in Azure SQL Database: Create index, drop index, last good plan

Microsoft Azure SQL Database 2019 Architecture



OtterTune Architecture



Claim: “SQLite just works”

“Other database engines may run great once you get them going. But doing the initial installation and configuration can often be intimidating.”




- No installation
- No setup
- No server process
- No access permissions
- No configuration files
- No manual recovery

How Does SQLite Do It?

	SQLite	PostgreSQL
Connections	In-processes	Network service
Caching	OS file system cache	Database buffer cache
Concurrency	Single-writer	Multi-version concurrency control
High-availability	Storage replication	Database replication, many variants

PostgreSQL Tuning Extensions and Tools



pgTune: Configuration calculator

HypoPG: Evaluate query plans with new indexes

Dexter: Automatically explore HypoPG plans

pgAnalyze: Indexing engine (commercial)

Percona PMM: monitoring and management

CrystalDB



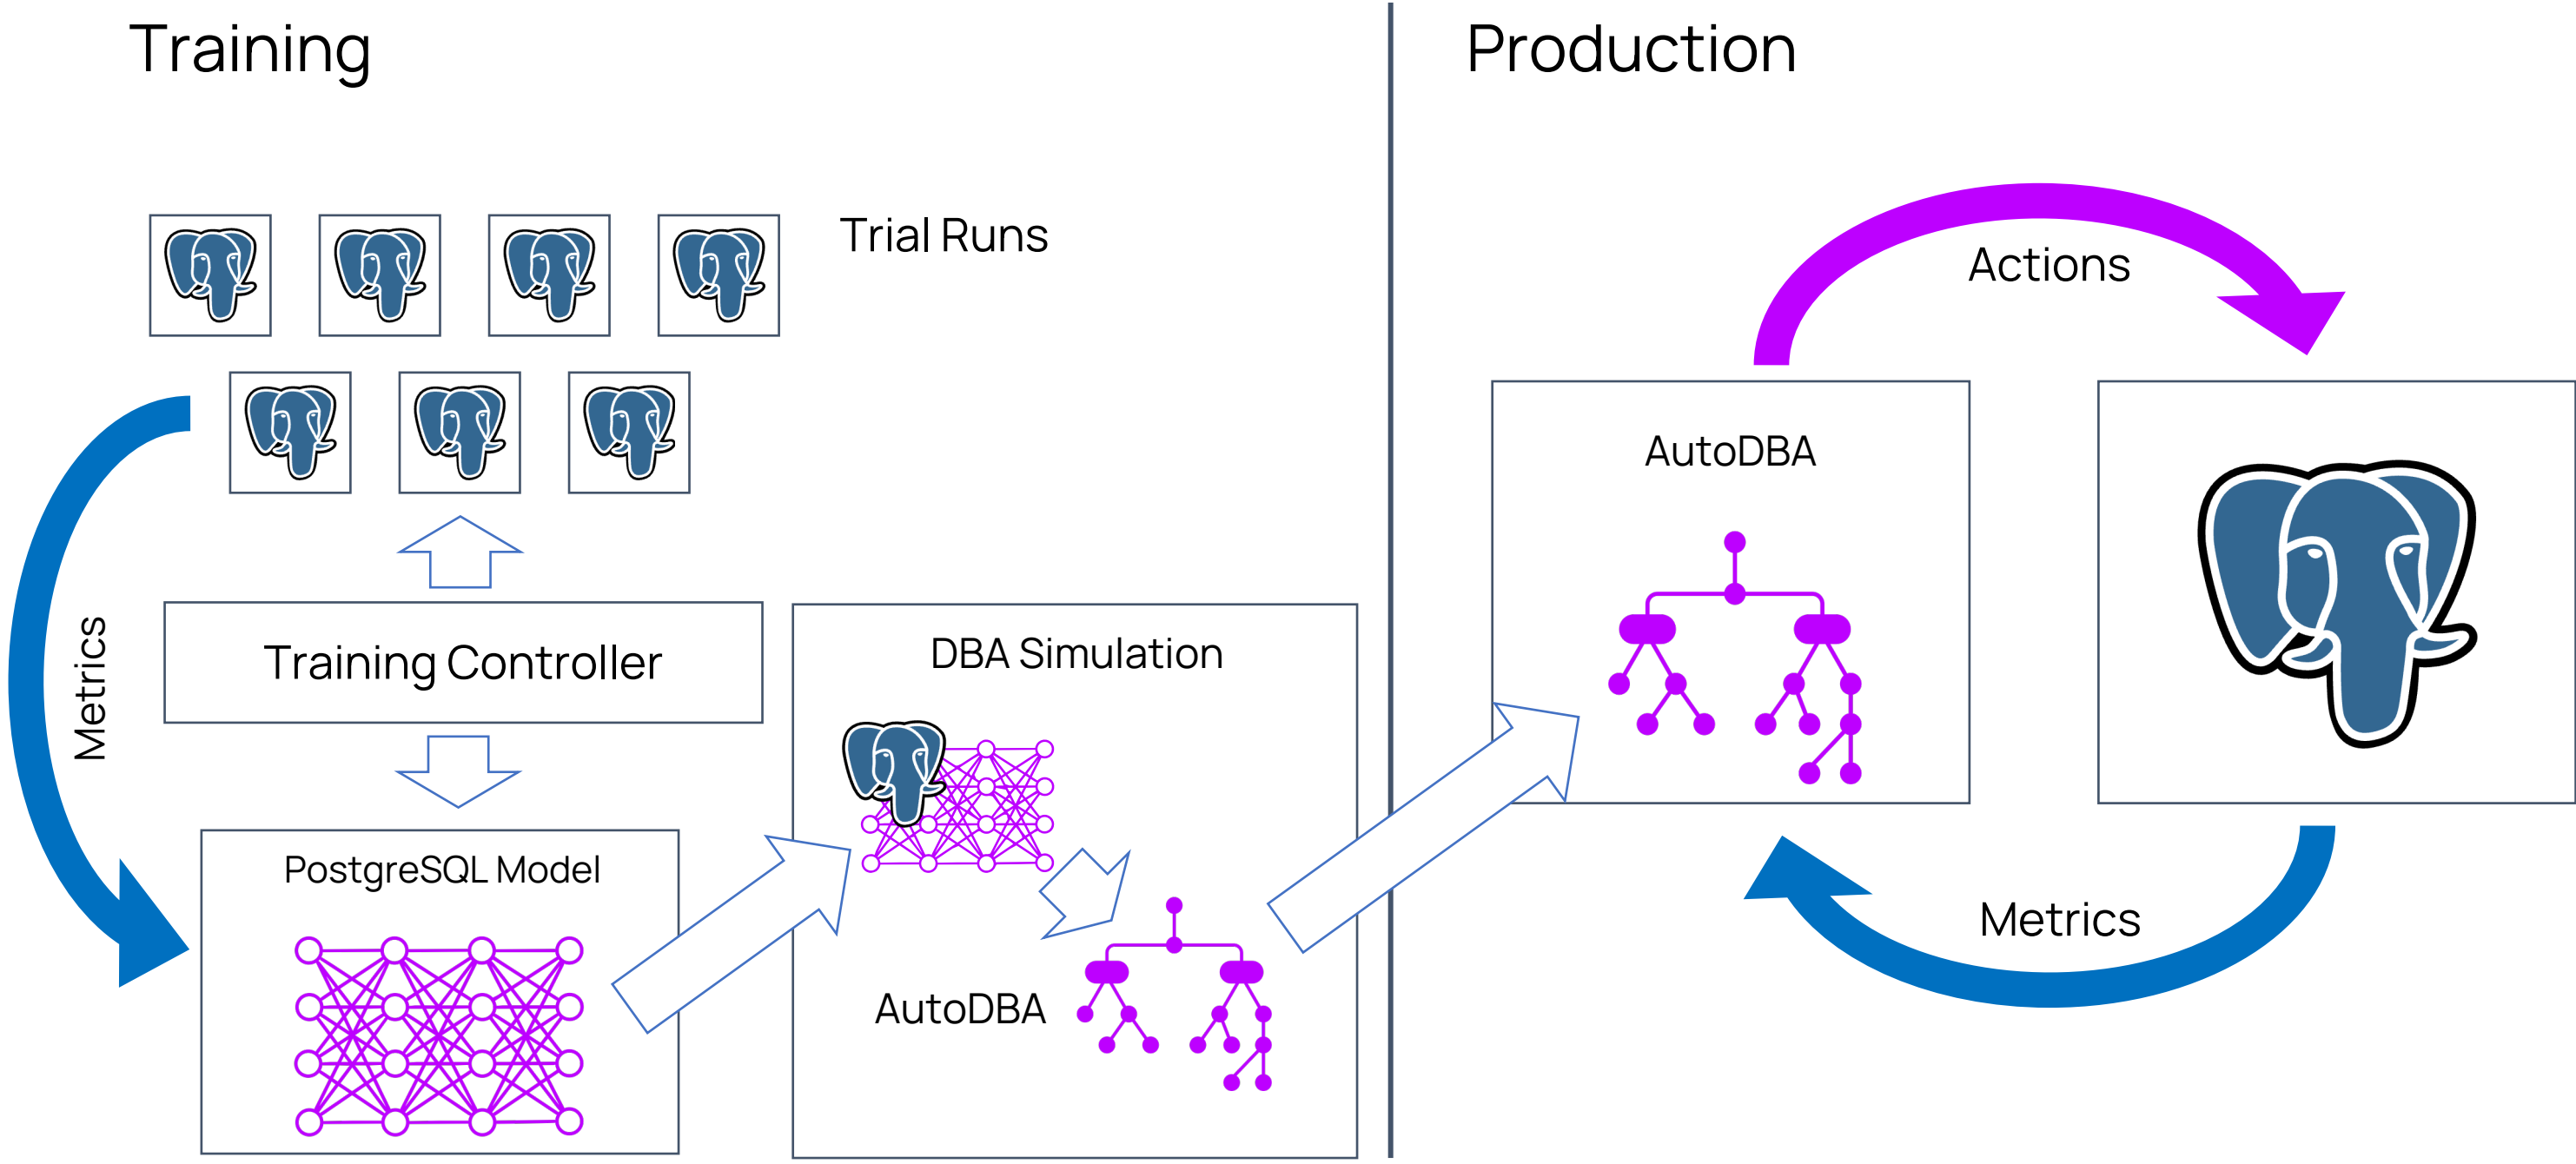
Approach

- Engineering for **full automation**, not a DBA tool
- Self-driving vehicles hold **practical lessons**, not only not metaphorical ones
- First objective is **safety**, next is performance

Deployment Models


1. **Database service:** fully automated application-specific optimization so anyone can run production database well
2. **Add-on for managed services:** enhancing RDS and other PostgreSQL installations with partial automation

CrystalDB Vacuum Controller Architecture




+ Old active transactions, which might go unnoticed, will increase the: 1) database bloat, 2) transactions latency, and 3) probability of XID wraparound.

Suggestions		
1454	Housekeeping	Automatic vacuum done to improve performance.
657	Housekeeping	Automatic vacuum done to improve performance.
400	Housekeeping	Automatic vacuum done to improve performance.




Bloat Level

Lower bloat Level is better





Latency

Lower latency is better




XID Horizon

Lower XID horizon is better



Vacuum control with deep reinforcement learning



Thank You

- Self-managing PostgreSQL is coming, will be super-human this decade
- Safety engineering makes it possible
- Benefits are better applications, better efficiency, better use of people's time