


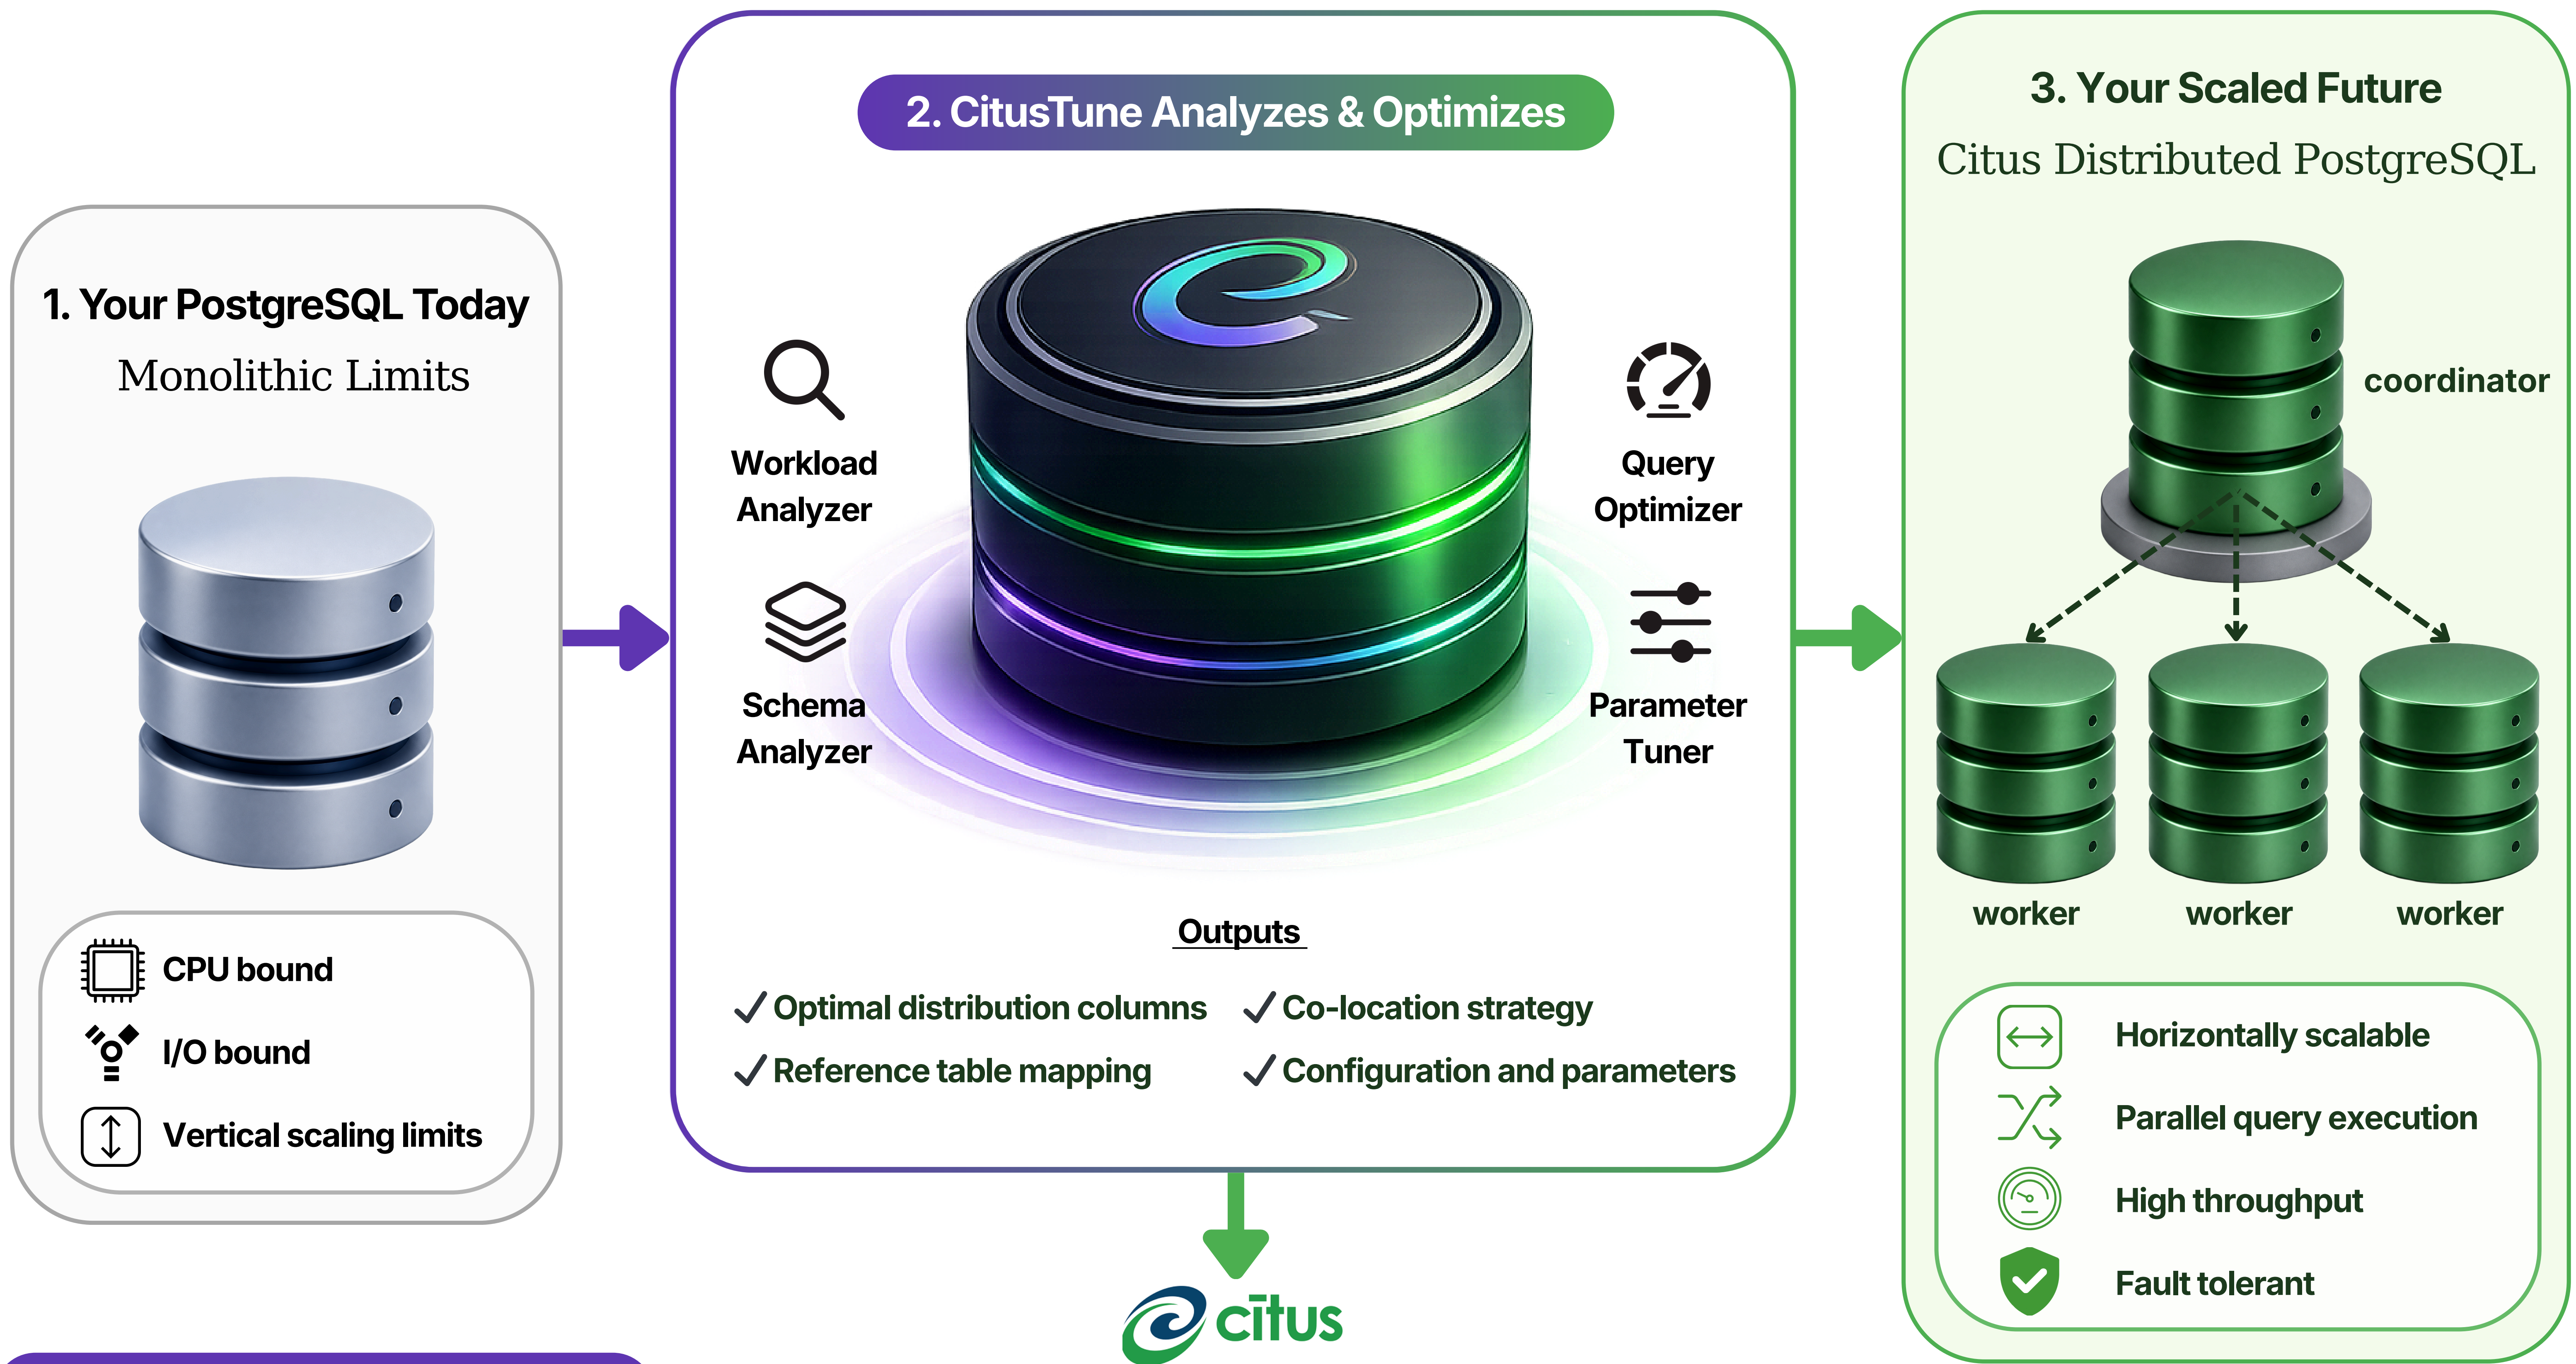
# Citustune

Intelligent Optimization Engine for Citus



## We analyze. Citus distributes. You scale.

Citustune analyzes your workload and schema, then automatically configures Citus to unlock massive scale on distributed PostgreSQL.



### Citustune Capabilities


- Smart Distribution**: Automatically selects the optimal columns to maximize data locality and minimize network communication.
- Reference Table Mapping**: Intelligently decides which tables to replicate across workers for fast joins and reduced network overhead.
- Co-location Optimization**: Groups related tables on the same workers to enable local joins and eliminate data shuffling.
- Query & Workload Analysis**: Understands your query patterns and workload to make smarter distribution and tuning decisions.
- Automated Tuning**: Auto-tunes Citus and PostgreSQL parameters for your hardware and workload.
- LLM-Powered Explanations**: Explains configurations after applying them. Users can understand the reasoning behind decisions.

### Benchmark Results

Up to **7x**

**Faster Queries**  
Real-world workloads on large scale datasets

**Query Performance Improvement**  
(Higher for better, some examples)



**Evaluation Environment**

**Coordinator Node:** 1 instance (4 vCPUs, 16 GB RAM)  
**Worker Nodes:** 3 instances hosted on 12 vCPUs, 48 GB RAM  
(Results may vary based on infrastructure and workload)

Scan the QR code to view detailed benchmark results, including example workloads and test datasets.



**10B+** rows tested

**1000+** queries analyzed

**1.5x - 7x** performance gain


### Tech Stack




Rust



Python



React



PostgreSQL



Citus

Aneliya Abdimalik

Alper Çelebi

Atahan Çelik

Efekan Korkut

Kemal Ünlü

Advisor: Ebru Aydın Göl