Leveraging open-source data platforms to build cost-effective and reliable Supply Chain Systems

Tim Steward

Principal Data Enterprise Architect

Leveraging open-source data platforms to build cost-effective and reliable Supply Chain Systems

FUJITSU

The evolution of supply chain management systems (SCM) increasingly requires cost effectiveness, reliability and scalability, while operating in dynamic data environments.

In this session we will outline how PostgreSQL can deliver a foundation for creating reliable SCM systems with all the advantages of an open-source database from improved innovation to reduced costs.

The Crude Oil we need

What petroleum products are made from crude oil?

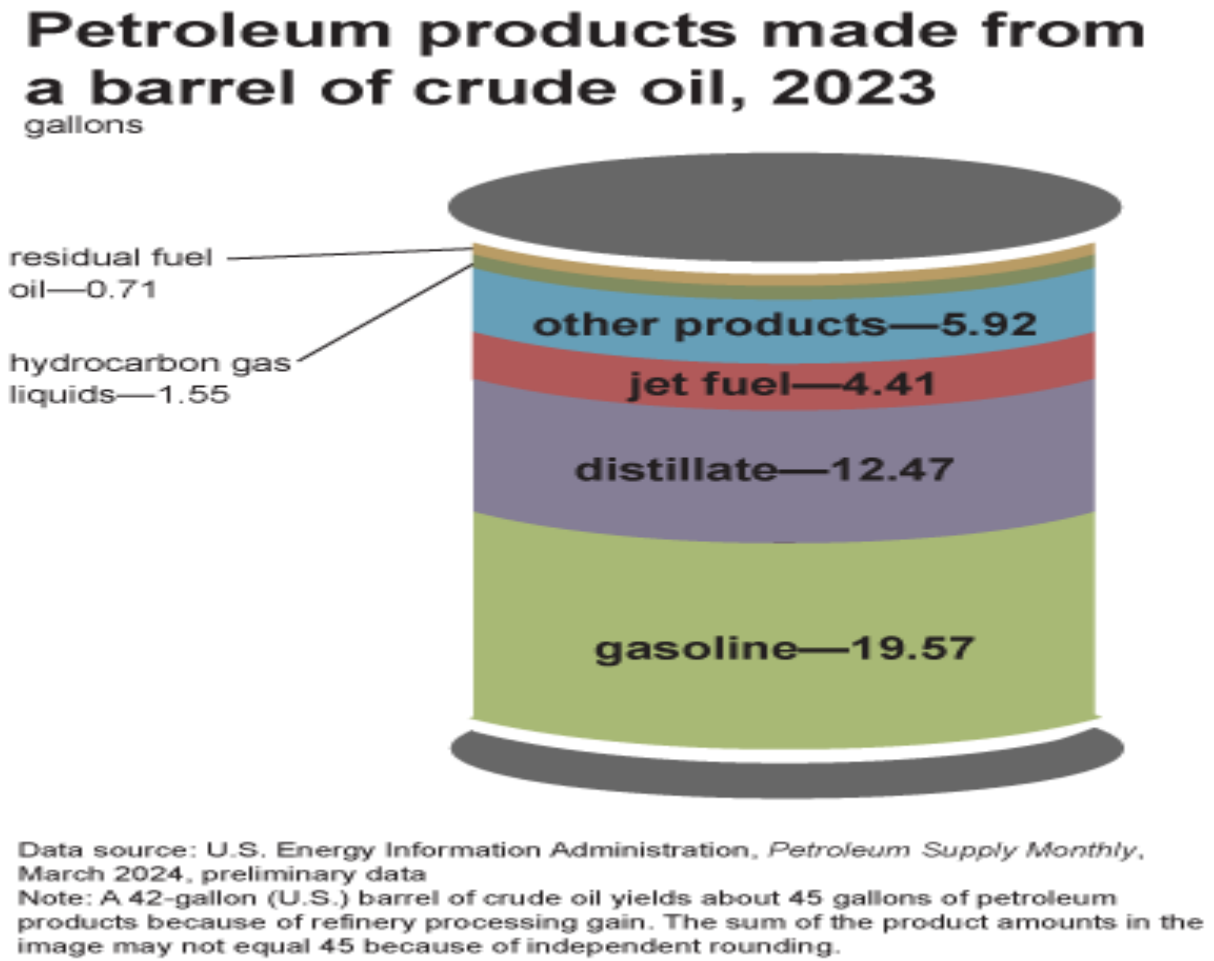

Data source: U.S. Energy Information Administration, *Petroleum Supply Monthly*, March 2024, preliminary data
Note: A 42-gallon (U.S.) barrel of crude oil yields about 45 gallons of petroleum products because of refinery processing gain. The sum of the product amounts in the image may not equal 45 because of independent rounding.

Where does the oil go?

Petroleum products produced from one 42-gallon barrel of oil input at U.S. refineries, 2023

Product	Gallons
Finished motor gasoline	19.57
Distillate fuel oil	12.47
Kerosene-type jet fuel	4.41
Petroleum coke	2.06
Still gas	1.68
Hydrocarbon gas liquids	1.55
Asphalt and road oil	0.84
Residual fuel oil	0.71
Lubricants	0.38
Naptha for feedstocks	0.34
Other oils for feedstocks	0.21
Miscellaneous products	0.21
Special napthas	0.08
Finished aviation gasoline	0.04
Kerosene	0.04
Waxes	< 0.01
Total	**44.65**
Processing gain	**2.65**

Data source: U.S. Energy Information Administration, *Petroleum Supply Monthly*, March 2024; preliminary data

The current state of Supply Chain Tech

The burden of legacy systems	Lack of flexibility and change
The problem	Many supply chains run on rigid, monolithic ERPs or proprietary middleware
Pain points	• High licensing fees (vendor lock-in) • Inability to scale dynamically during peak seasons (e.g., holidays) • Data silos between warehouses, suppliers, and last-mile delivery.

The open-source advantage

FUJITSU

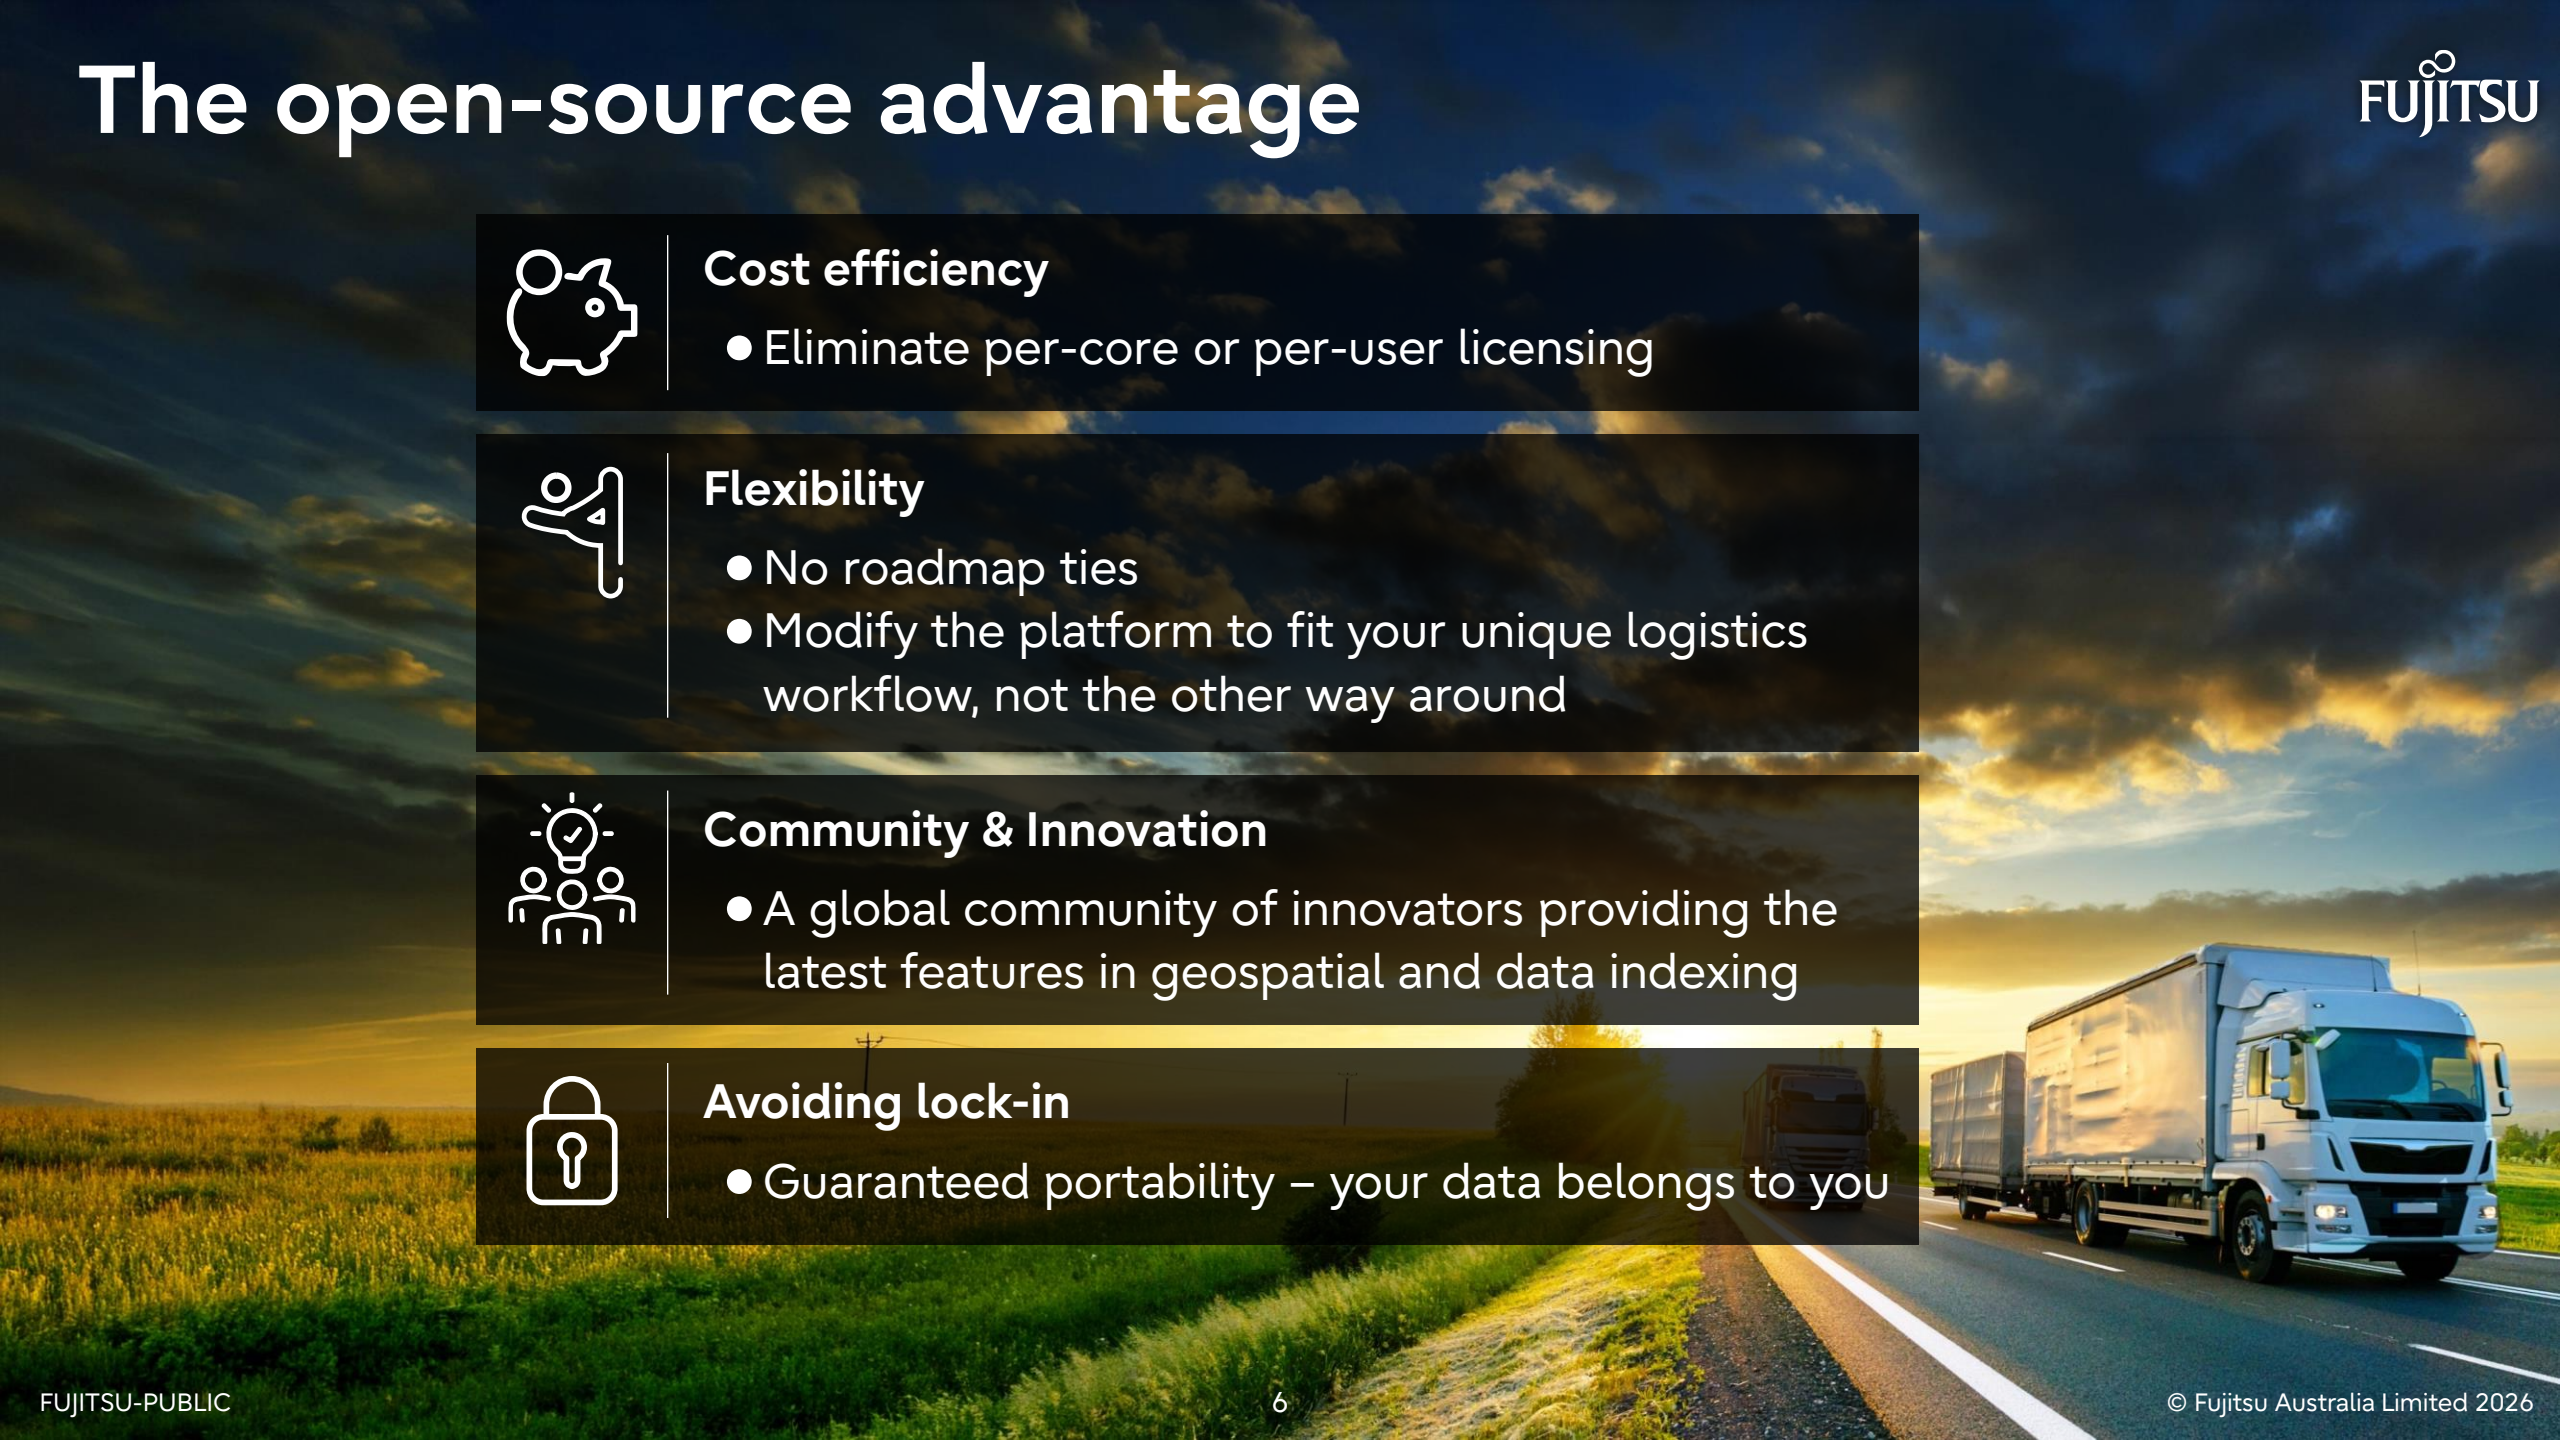

Cost efficiency

- Eliminate per-core or per-user licensing

Flexibility

- No roadmap ties
- Modify the platform to fit your unique logistics workflow, not the other way around

Community & Innovation

- A global community of innovators providing the latest features in geospatial and data indexing

Avoiding lock-in

- Guaranteed portability – your data belongs to you

The architect's choice: PostgreSQL

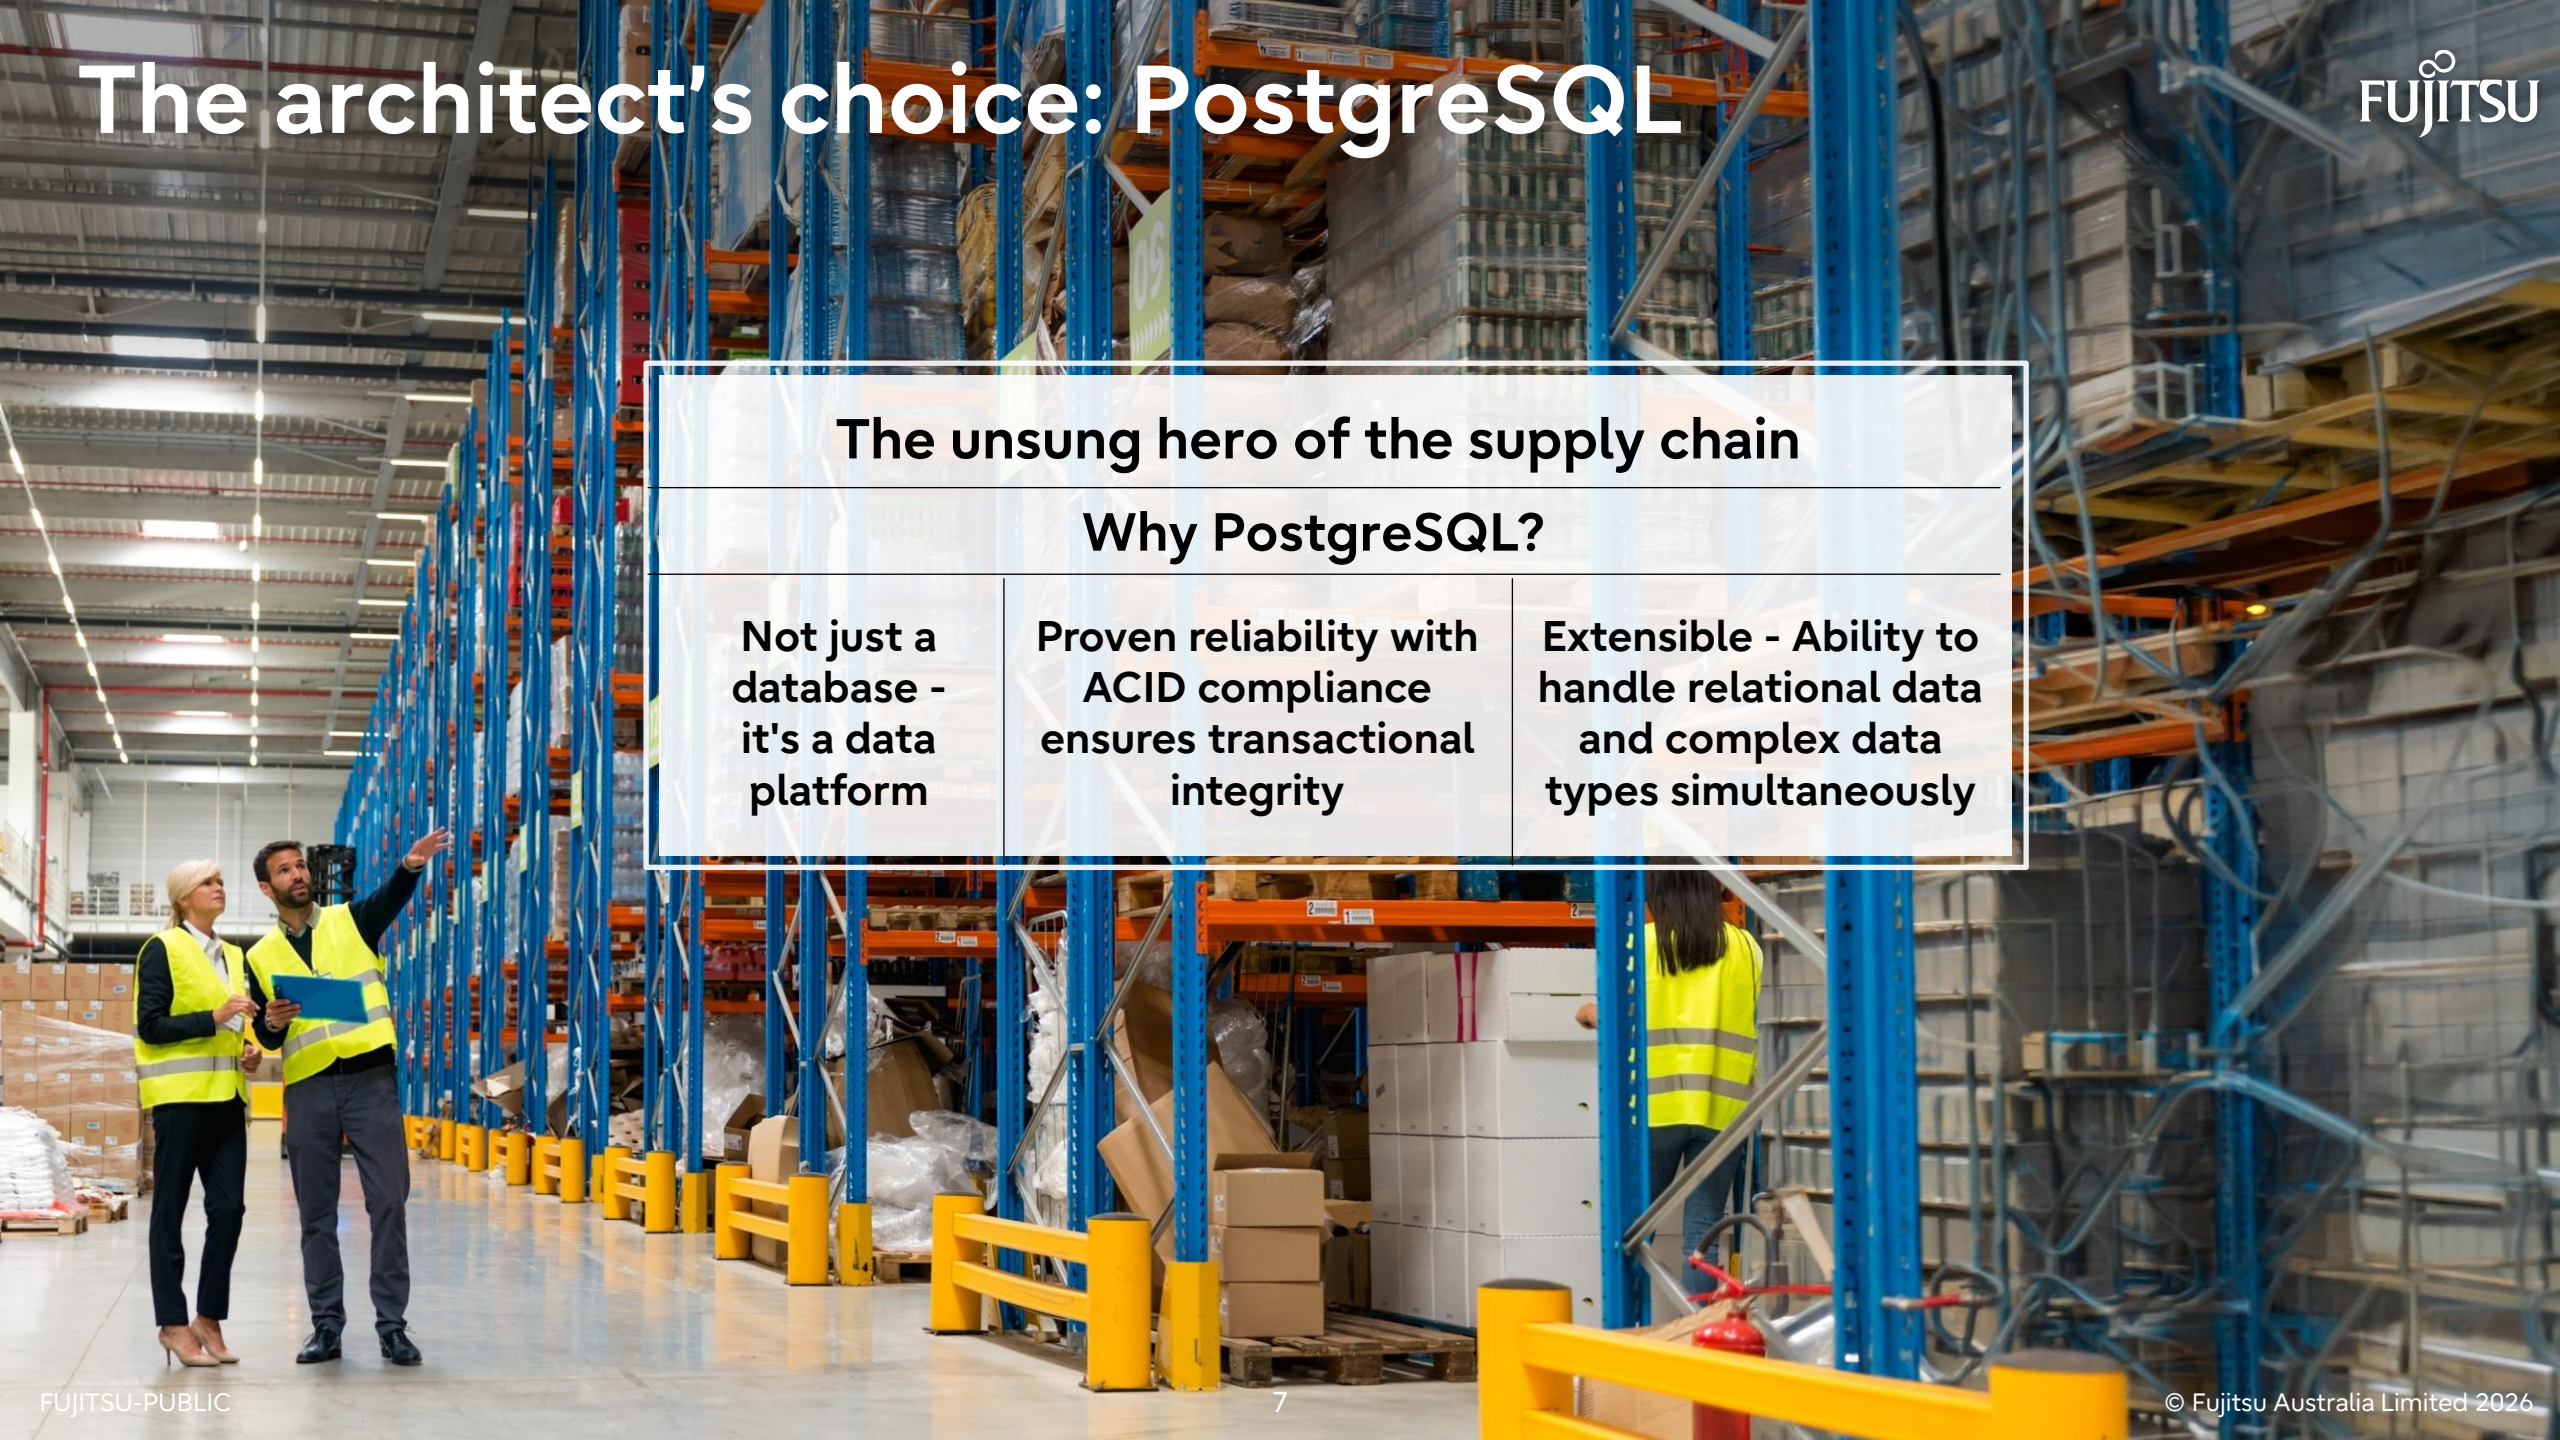

The unsung hero of the supply chain

Why PostgreSQL?

Not just a database - it's a data platform	Proven reliability with ACID compliance ensures transactional integrity	Extensible - Ability to handle relational data and complex data types simultaneously

Solving supply chain puzzles with PostgreSQL

FUJITSU

Four pillars of a cost-effective system

Inventory optimization

Using complex window functions to calculate safety stock levels and reorder points in real-time

Real-time tracking

Leveraging JSONB data types to ingest semi-structured IoT data from GPS trackers without having rigid schemas

Geospatial logistics

Using PostGIS to calculate distances, optimize routes, and find nearest warehouses and distributors

Time-based analytics

Analyzing lead times and carrier performance using sophisticated date/time manipulation

Why we need route optimization

FUJITSU

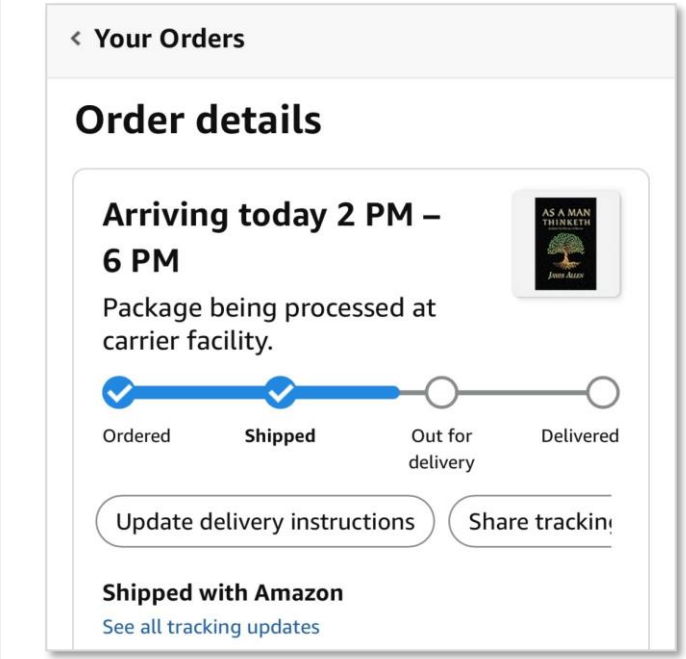

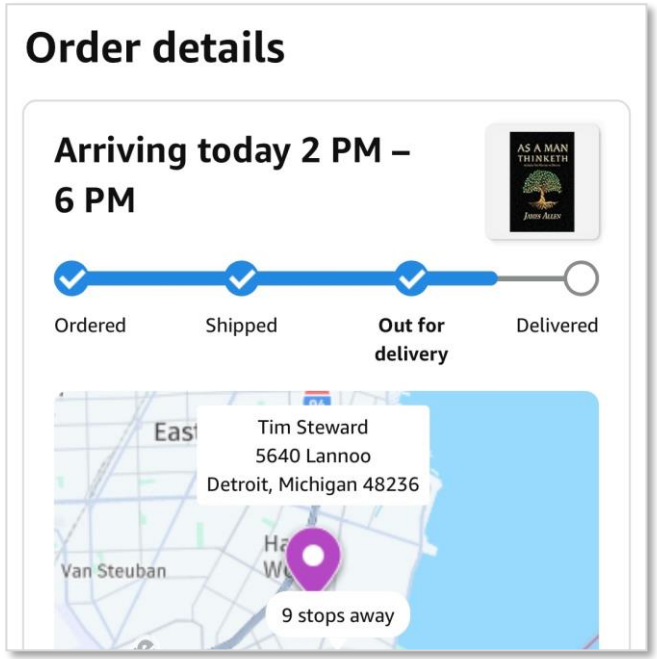

Route optimization with PostGIS

Case study

The last mile

FUJITSU

The PostgreSQL way

- Store geographic points (lat/long) as geometry or geography types
- Use PostGIS functions like ST_Distance and ST_DWithin to filter nearby drivers
- Calculate convex hulls (ST_ConvexHull) to define delivery zones

Alessandro Nava, Luciano Greco,
The Impact of Last-Mile Logistics: a Case Study on the Optimisation of Commercial Fleets through the European Union,
IFAC-PapersOnLine,
Volume 56, Issue 2, 2023,
Pages 2371-2376,
ISSN 2405-8963,
https://doi.org/10.1016/j.ifacol.2023.10.1209.
(https://www.sciencedirect.com/science/article/pii/S2405896323016130)

Predictive Analytics - from reactive to proactive

FUJITSU

The setup

Instead of exporting data to a separate warehouse, run ML inside the database

The tool

PostgresML

Application

Run regression models directly on historical shipping data to predict:

- Future demand spikes
- Potential supplier delays based on weather data joined with route data

Benefit

No expensive ETL (Extract, Transform, Load) processes; insights are generated where the data lives

AI and advanced analytics integration

FUJITSU

Modern supply chains use predictive analytics and AI

Generative AI Support

Secure pgvector extension for vector similarity search in demand forecasting

RAG Application Development

Integration with LangChain for LLM-powered supply chain insights

Embedded Python

Run data science algorithms directly inside the database

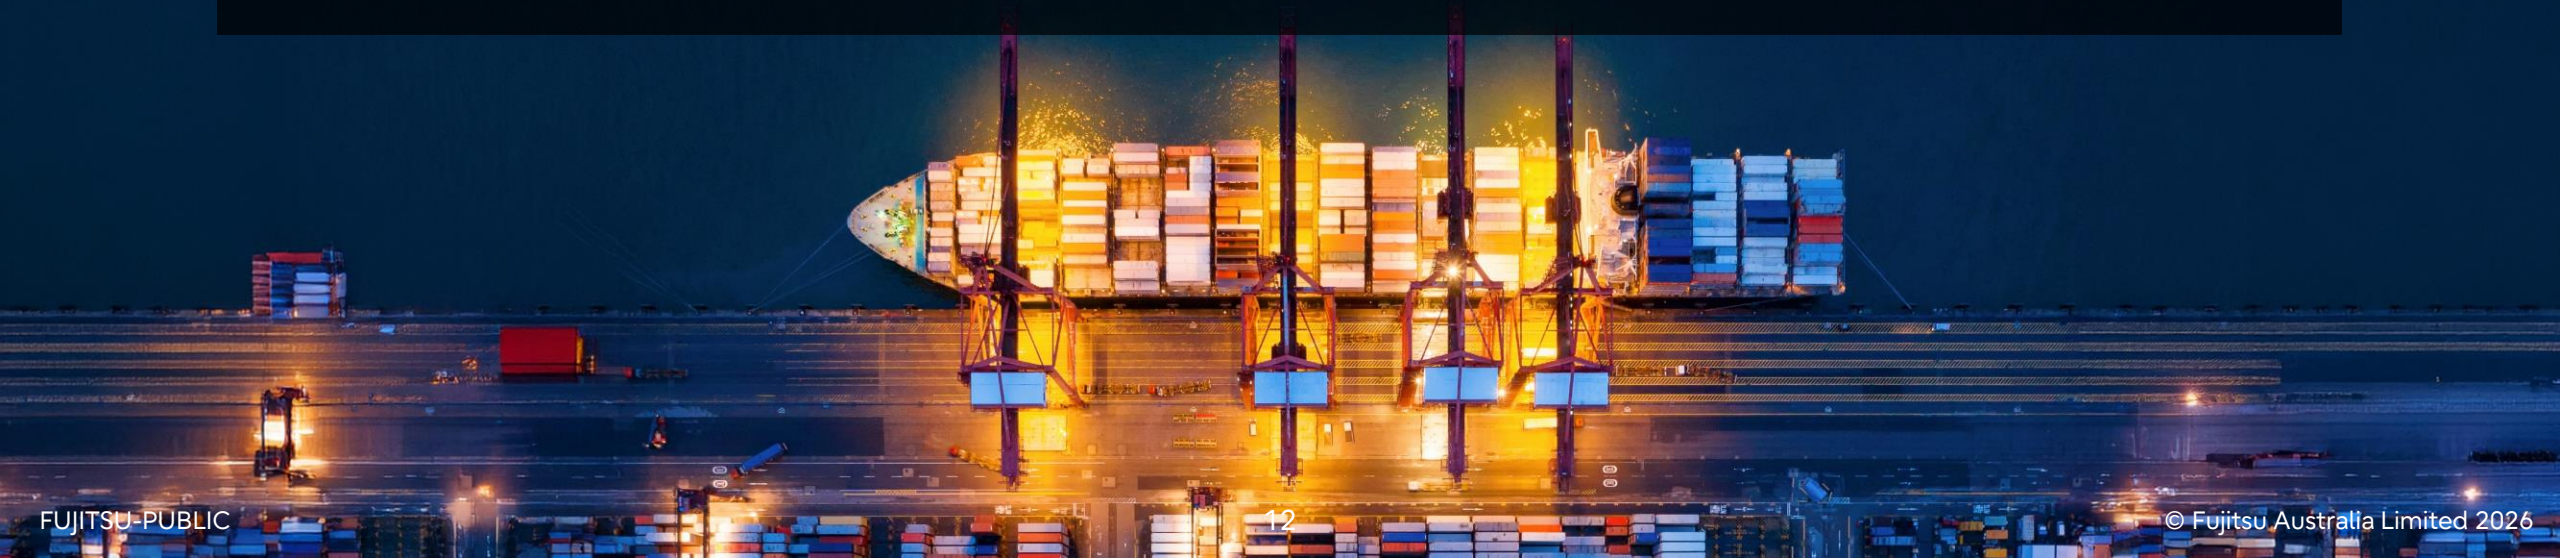

The cost breakdown

FUJITSU

Proprietary stack

- Expensive DBMS licenses
- Separate ETL tool costs
- Specialized GIS software costs

Open-source stack

- Software cost: $0 (Licensing)
- Extension cost: $0 (PostGIS, PostgresML)

Cost center

- Only infrastructure (cloud VMs or bare metal) and engineering talent

Ensuring reliability

FUJITSU

Open source doesn't always mean *unsupported*

High Availability

Tools like Patroni and Repmgr manage automatic failover. If a primary server dies, a replica takes over in seconds

Backup & Recovery

Point-In-Time Recovery (PITR) allows you to restore data to the exact moment before a mistake happened (e.g., a bulk price update error)

Replication

Use logical replication to stream data in real-time to reporting instances without slowing down the transaction system

Building the Modern Stack
A reference architecture

FUJITSU

PostgreSQL as the central source of truth

Edge/IoT	Ingestion	Processing	Storage & Analysis	Output
Trucks/warehouses send JSON/ geolocation data	Data lands in a queue (Kafka/RabbitMQ)	Stream processor (Apache Flink) cleans data	PostgreSQL (with Citus for sharding)	Dashboards (Grafana), Mobile Apps (API), and AI models

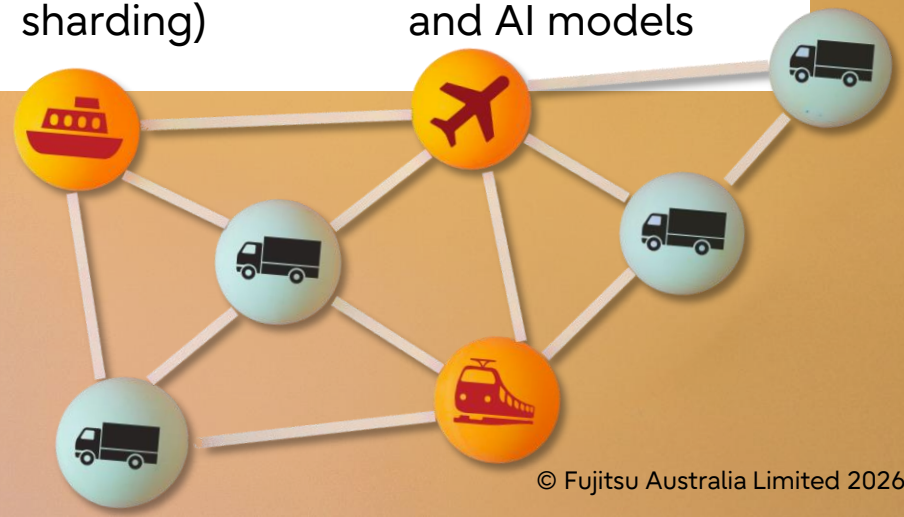

Potential challenges & solutions

FUJITSU

Keeping it real

Challenge:
We need enterprise support

Solution
Companies like Fujitsu, PgEdge and EDB offer enterprise-grade support for PostgreSQL

Challenge:
Our team doesn't know open source

Solution
PostgreSQL syntax is standard; most SQL Server/Oracle developers adapt within weeks - the community is massive for troubleshooting

Challenge:
Scaling writes

Solution
Use extensions like Citus to distribute PostgreSQL across clusters horizontally

The Real Deal

FUJITSU

Who is using PostgreSQL?

Magaya : A leading logistics provider used PostgreSQL to overcome severe performance bottlenecks.
Win: Achieved a 10x performance boost, reduced upgrade costs by 83%, and maintained 99.99% availability.

Infios : Global software provider serving over 5,000 customers.
Win: Migrated from Oracle to PostgreSQL (Amazon Aurora) to save $2.2 million annually. Reduced data footprint from 128 TB to 60 TB and cut database timeouts by 45% .

PL Developments : A major manufacturer of OTC healthcare goods.
Win: Replaced legacy Oracle system with a modern stack using Postgres. Eliminated heavy ETL jobs that consumed nearly 30% of ERP resources, enabling real-time analytics .

Bayer Crop Science : Manages geospatial data and billions of observations.
Win: Moved to a PostgreSQL-compatible database to handle massive growth. Reduced response times by 50% and increased throughput by 5x during critical harvest seasons .

The Real Deal

FUJITSU

How are they using PostgreSQL to solve complex challenges?

Hybrid Transactional/Analytical Processing : PL Developments uses PostgreSQL for fast-changing operational data, combined with other engines for heavy analytics—keeping each system in its "performance zone"

Real-Time EDI Integration: PostgreSQL can sync EDI documents (like purchase orders) directly into database tables. This replaces slow batch files and allows real-time SQL queries on transactions, reducing compliance penalties

Master Data Management : Building MDM solutions directly within PostgreSQL using stored procedures. They use PL/pgSQL for data cleansing, fuzzy matching to find duplicates, and "survivorship rules" to create a single "golden record" for customers or products

Analytics & Optimization: Reference examples like the TPC-H benchmark (which models a wholesale supplier business) show how PostgreSQL can run complex analytical queries to answer critical questions, such as identifying the lowest-cost supplier or forecasting revenue impact .

The future is open

Summary

Open source eliminates licensing waste, free up budgets for actual logistics innovation. PostgreSQL, with its extensions, handles the unique complexity of supply chains better than any monolithic ERP.

Call to Action

- Start Small Migrate one non-critical tracking service to PostgreSQL
- Measure as you go comparing the operational cost vs. your legacy system cost
- Scale for the future - Build your next-gen supply chain data platform on open ground

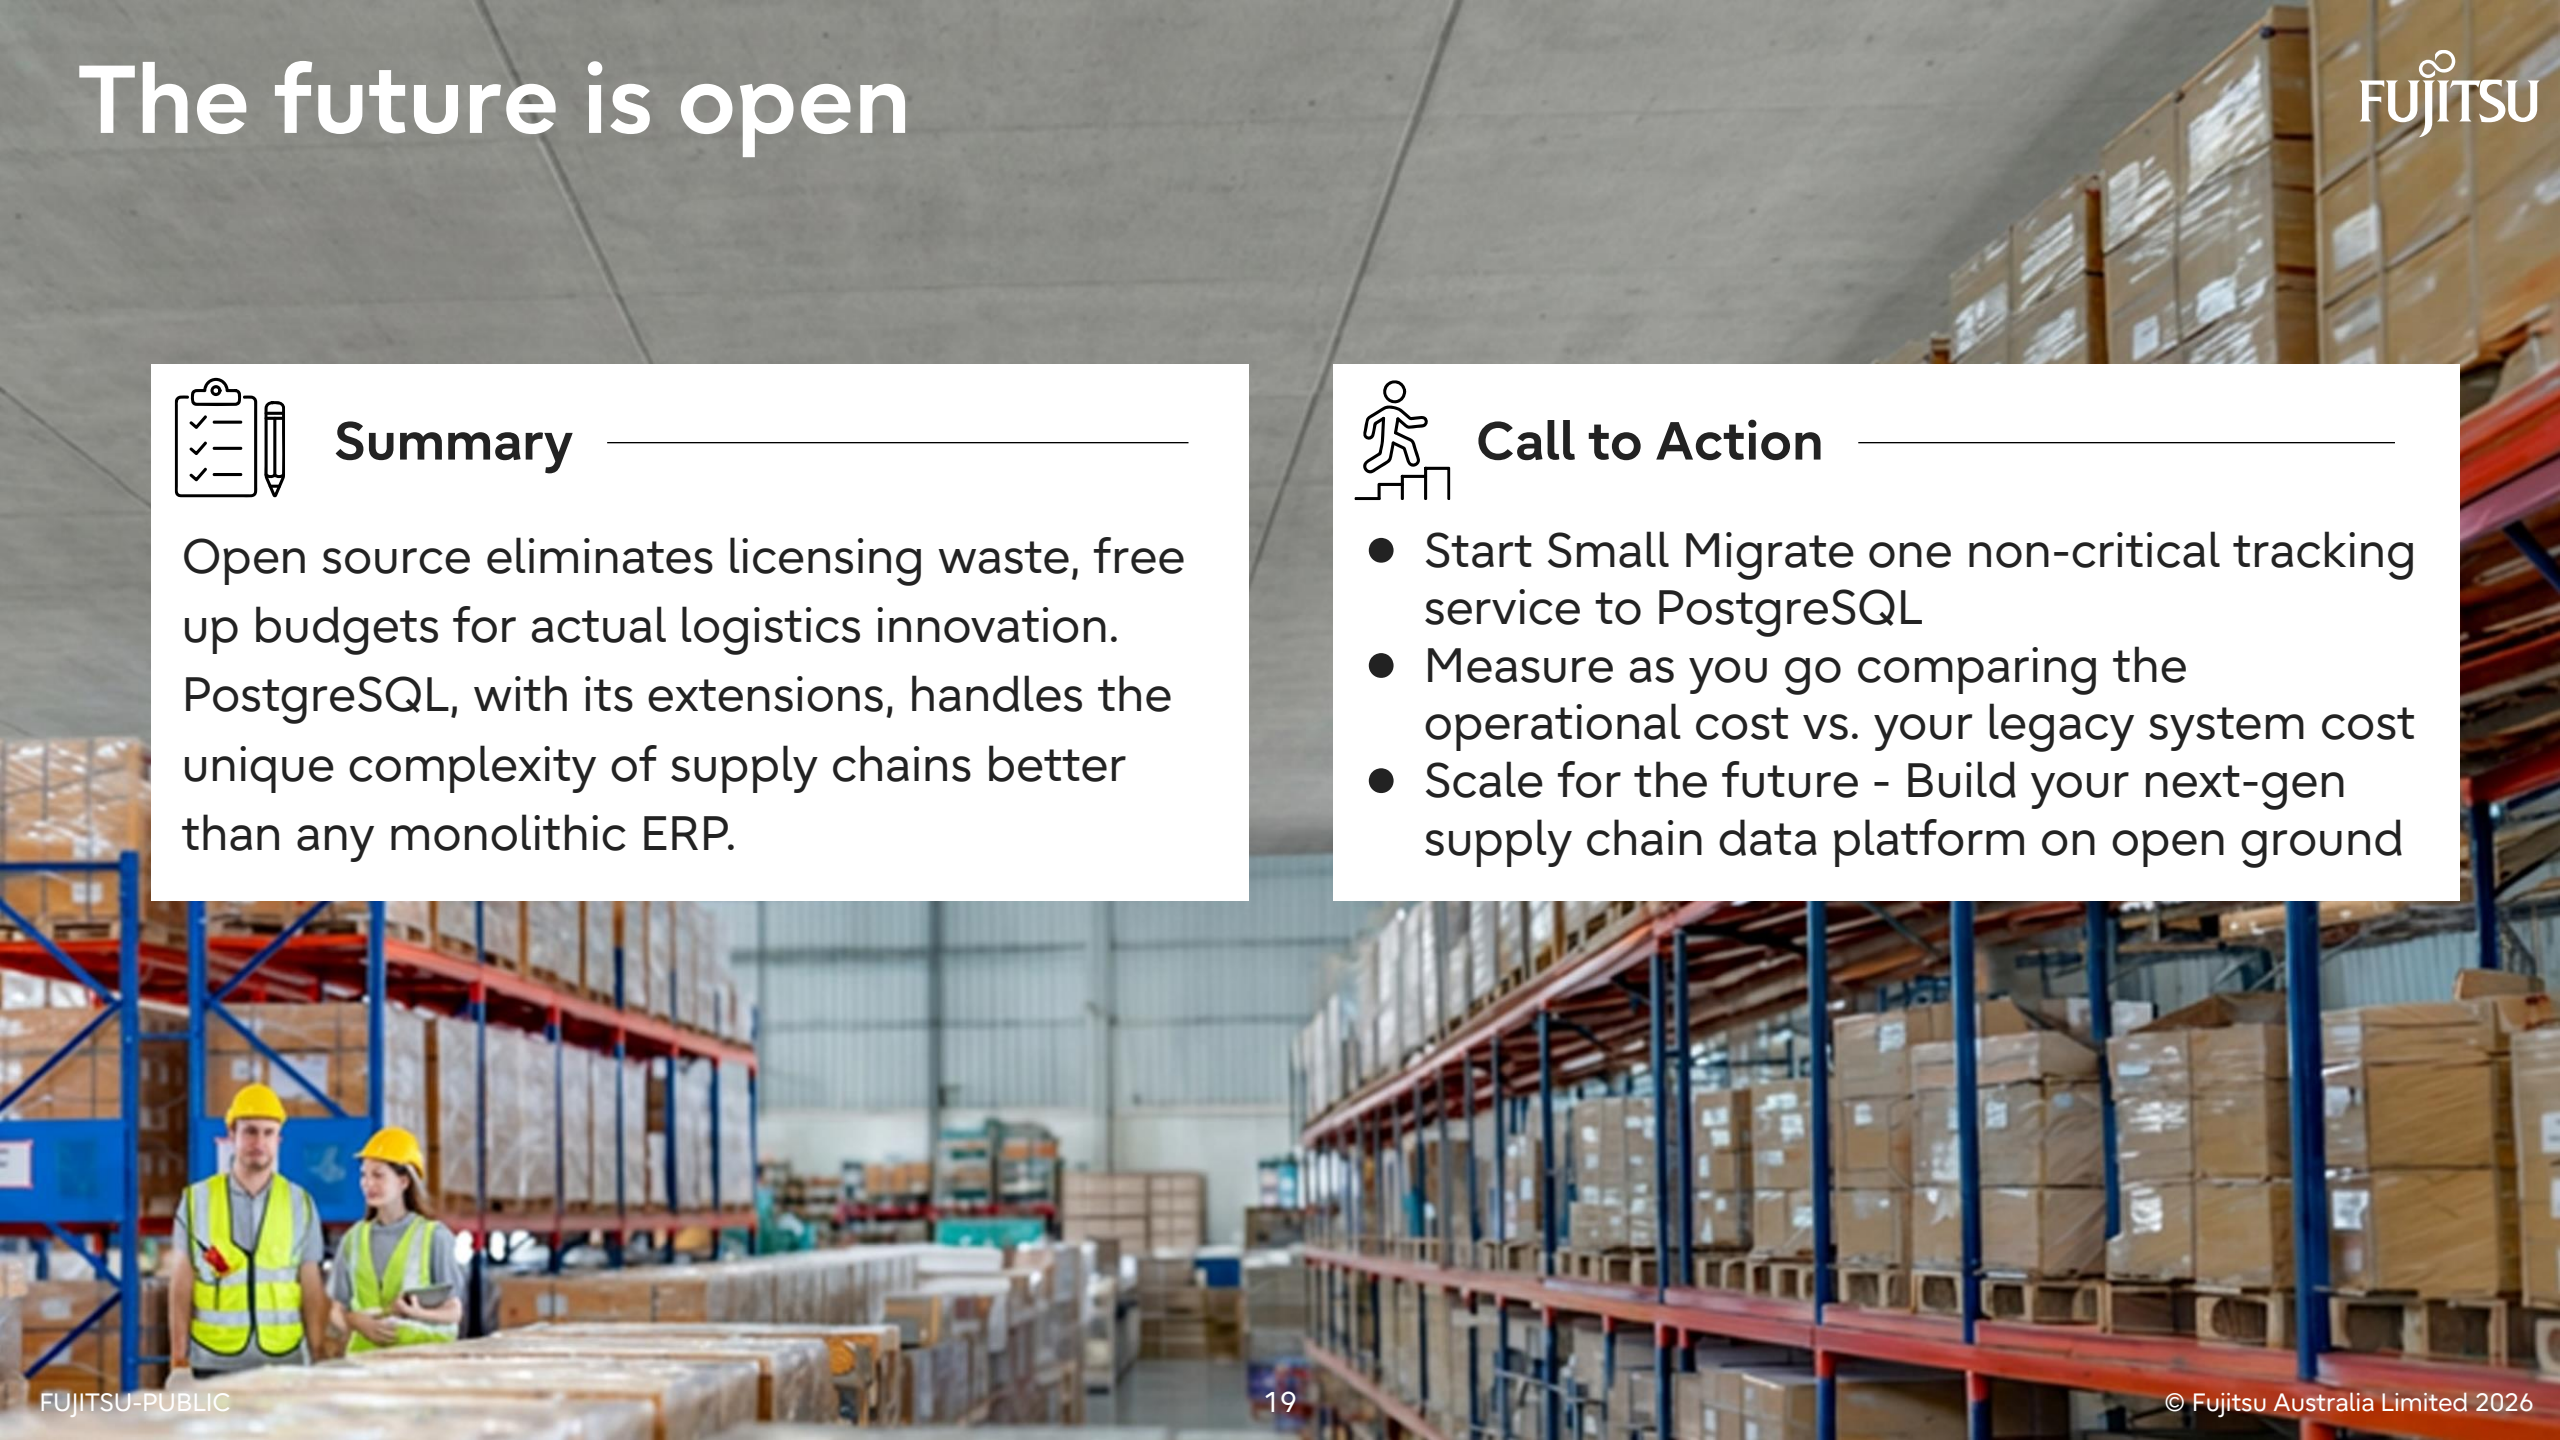

Thank you

FUJITSU

Tim Steward

Principal Data
Enterprise Architect