



AECC PROOF OF CONCEPT

Optimal Edge Selection for Realizing Digital Twins and Green Energy Utilization

By Oracle, KDDI, Equinix and Toyota

Abstract

As connected car applications evolve, service providers must intelligently route traffic to the most suitable edge servers to ensure efficient and flexible service delivery. This proof of concept (PoC) demonstrates how AECC’s solutions enable traffic digital twins and optimize green energy utilization by selecting the best path within the mobile network based on key performance metrics.

Building on previous work in [Traffic Load Balancing of Edge Server Access](#), which confirmed the effectiveness of dynamic traffic distribution, this PoC takes the next step by enhancing vehicle-oriented applications. By leveraging Telco APIs and real-time network data, it optimizes traffic distribution for applications such as digital twins and green energy-aware server selection.

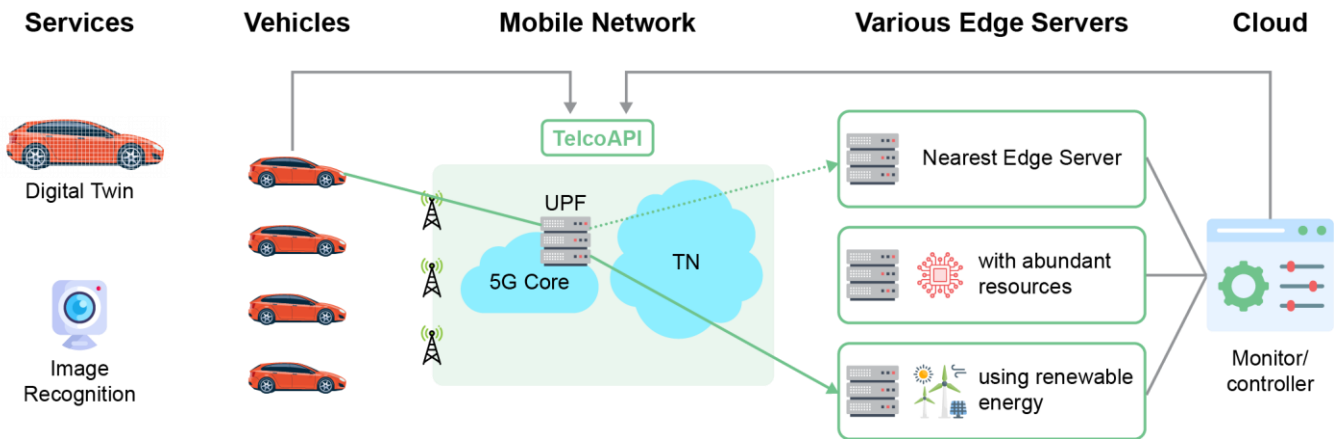


Figure 1: An overview of the network architecture and workflow.

The PoC incorporates distributed computing architecture, edge data offloading, access network selection, and opportunistic data transfer as defined by AECC’s Working Group 2. Through functional validation, it confirms the feasibility of dynamically steering traffic to the optimal edge server, offering a practical approach to implementing AECC’s proposed architecture.

Business Strategy

This PoC demonstrates the value of AECC’s approach for a number of players in the connected vehicle services ecosystem:

- **For mobility service providers:** This PoC offers a glimpse into the future of mobility services, demonstrating how edge infrastructure and intelligent breakout point selection can enhance service delivery in next-generation mobile networks.
- **For mobile network operators (MNOs):** By validating a dynamic approach to user protocol data unit (PDU) session breakout point selection, this PoC highlights how MNOs can optimize network efficiency and deliver seamless connectivity tailored to real-time demand.
- **For edge providers:** This PoC showcases real-world applications of edge infrastructure in the automotive industry, illustrating the design and deployment of AI workloads that power advanced mobility solutions. These insights are essential for developing high-value, future-ready edge infrastructure that meets the evolving needs of connected vehicles.

Proof of Concept Objective

In [phase 1](#), the engineering team for the PoC successfully implemented load balancing across edge servers using TelcoAPI/LISP, ensuring consistent service quality.

In phase 2, the PoC team expanded on this foundation by demonstrating how flexible edge selection, guided by multiple metrics, enhances application performance. This phase focused on two key use cases:

1. **Traffic digital twin:** Dynamically switching edge servers based on vehicle movement to create a real-time digital replica of traffic conditions.
2. **Green energy optimization and load balancing:** Balancing network load while prioritizing edge servers powered by sustainable energy sources.

These advancements highlight the potential of intelligent edge selection to drive efficiency and innovation in connected vehicle applications.

Proof of Concept Scenario

There were two scenarios in this PoC, one for each use case. But for both use cases, the engineering team used five different containerized edge servers based in different locations across Japan.

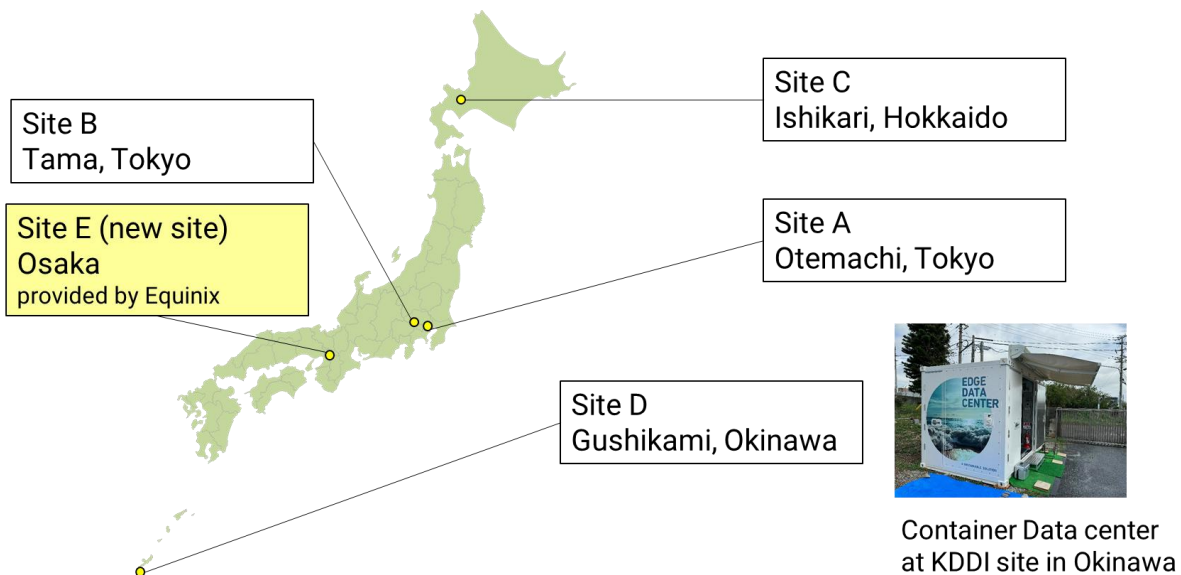


Figure 2: Locations of the five edge servers used in this PoC, plus a photo of a KDDI container data center.

Digital Twins Overview

Creating digital twins for connected vehicles requires low latency and efficient data transmission to maintain real-time accuracy. While processing data on local edge servers can improve service delivery, vehicle movement across regions presents a challenge — edge servers must adapt dynamically to transfer the location information from server to server and ensure seamless performance.

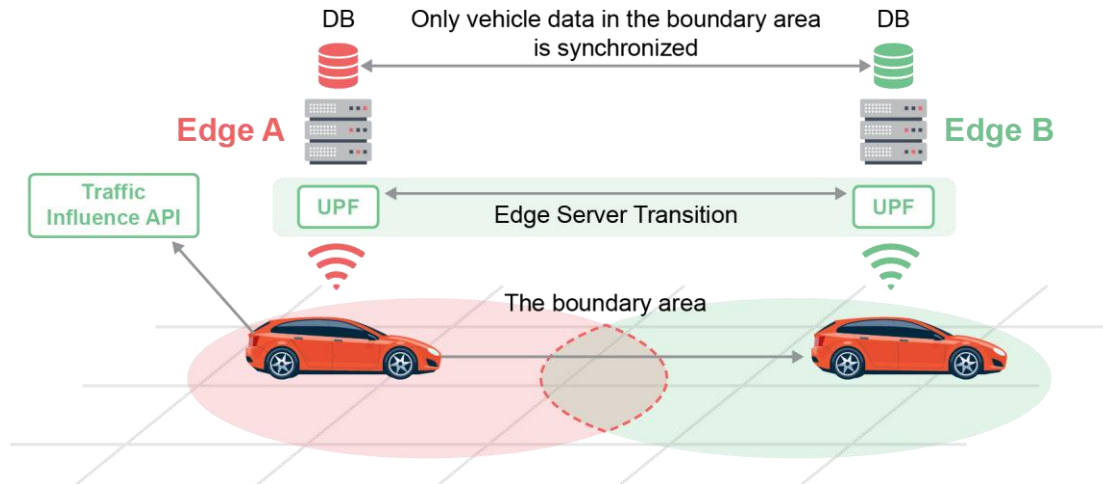


Figure 3: Digital twins functional architecture overview.

In this scenario, the engineering team for this PoC demonstrated two key capabilities:

1. **Intelligent edge switching with the Traffic Influence API:** As vehicles move between regions, the Traffic Influence API dynamically shifts their connection to the most suitable edge server, optimizing performance and responsiveness.
2. **Seamless data synchronization across boundaries:** To prevent service disruptions, only the necessary vehicle data in boundary areas is synchronized between edge servers, enabling smooth transitions without unnecessary data transfers.

Green Energy Overview

The engineering team set out to develop a system that intelligently selects edge servers based on both sustainability and performance. The goal was to prioritize solar-powered edge servers whenever feasible, without compromising the low-latency requirements essential for many connected vehicle services.

To achieve this, the system dynamically adjusted distribution metrics in real time. When solar energy was abundant and CPU load was low, the system routed traffic to solar-powered edge servers to maximize energy efficiency. However, during nighttime hours or periods of heavy cloud cover (i.e. when solar generation was limited or CPU load becomes high), the system automatically redirected traffic to ensure responsiveness remained the top priority.

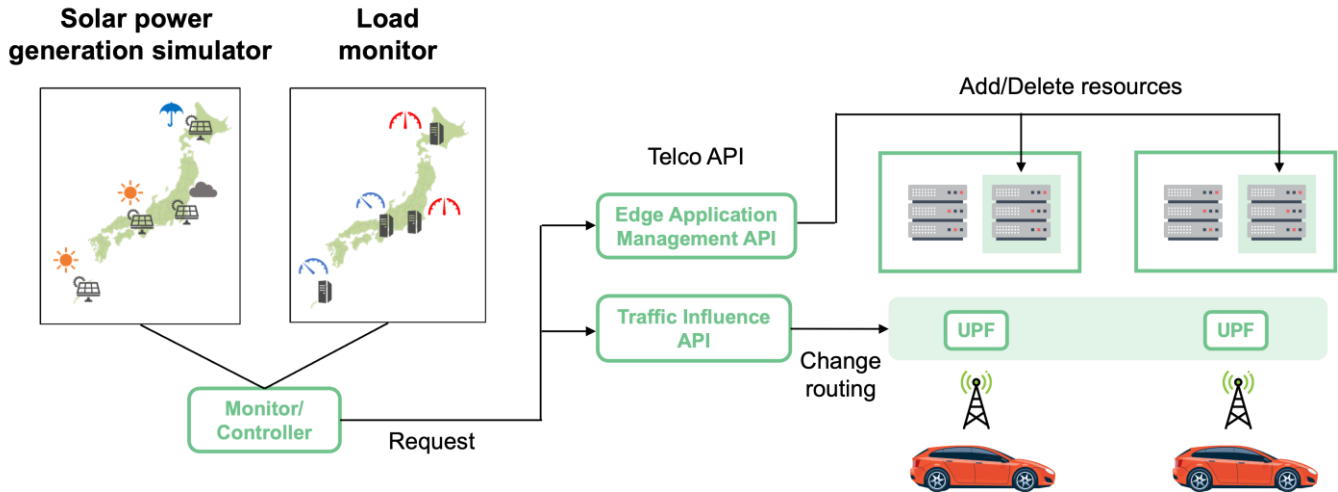


Figure 4: Green energy functional architecture overview.

Process

Traffic Digital Twin Use Case Process

This use case demonstrated that a system could dynamically switch edge servers based on vehicle position. The system ensures low-latency processing and seamless data synchronization between geographically distributed edge nodes.

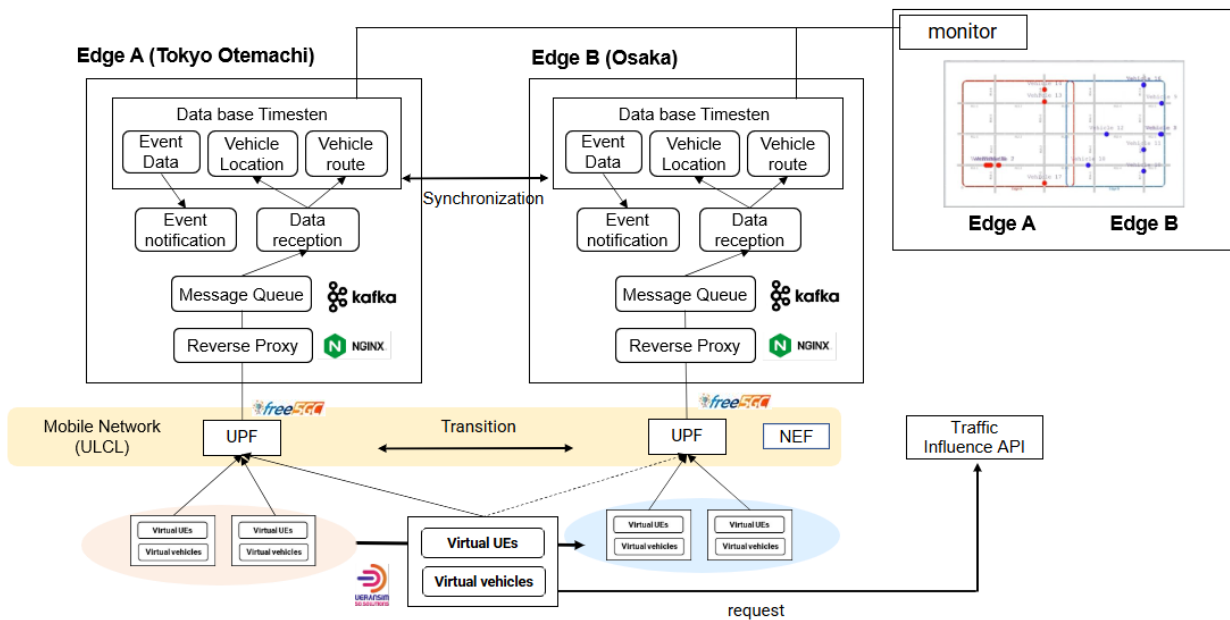


Figure 5: Digital twins system configuration.

- Vehicle data collection and processing at edge nodes:** Vehicles and virtual user equipment (UEs) generate real-time data, including event data, vehicle location, and route information. This data is stored in a time-series database (Timesten) at the edge servers.

2. **Event notification and data reception:** The edge nodes use Kafka-based message queues to handle event notifications and data reception efficiently. An NGINX reverse proxy manages secure and efficient data flow between system components.
3. **Edge server synchronization:** As vehicles move between regions, their data needs to transition seamlessly from Edge A (Tokyo Otemachi) to Edge B (Osaka) and vice versa. A synchronization mechanism ensures that only relevant vehicle data near the boundary is shared between edge nodes, preventing unnecessary data transfer.
4. **Mobile network and traffic steering:** The user plane function (UPF) in the mobile network uplink classifier (ULCL) directs traffic to the most appropriate edge server. When a vehicle moves across regions, the network transitions its traffic session to the corresponding UPF at the nearest edge location.
5. **Traffic influence API and dynamic edge selection:** The Traffic Influence API plays a crucial role in determining the optimal edge server for processing, based on vehicle movement and network conditions. When a virtual vehicle moves across regions, it sends a request to the Traffic Influence API. As a result, traffic from the virtual vehicle is dynamically routed to the appropriate edge server, ensuring efficiency and low latency.
6. **Real-time monitoring and visualization:** A monitoring system provides a visual representation of vehicle movements and the switching process between Edge A and Edge B, ensuring smooth operations and validating the PoC's effectiveness.

Green Energy Use Case Process

This use case shows how connected vehicle services can dynamically balance computational loads across solar-powered and traditional edge servers, optimizing both energy efficiency and service responsiveness.

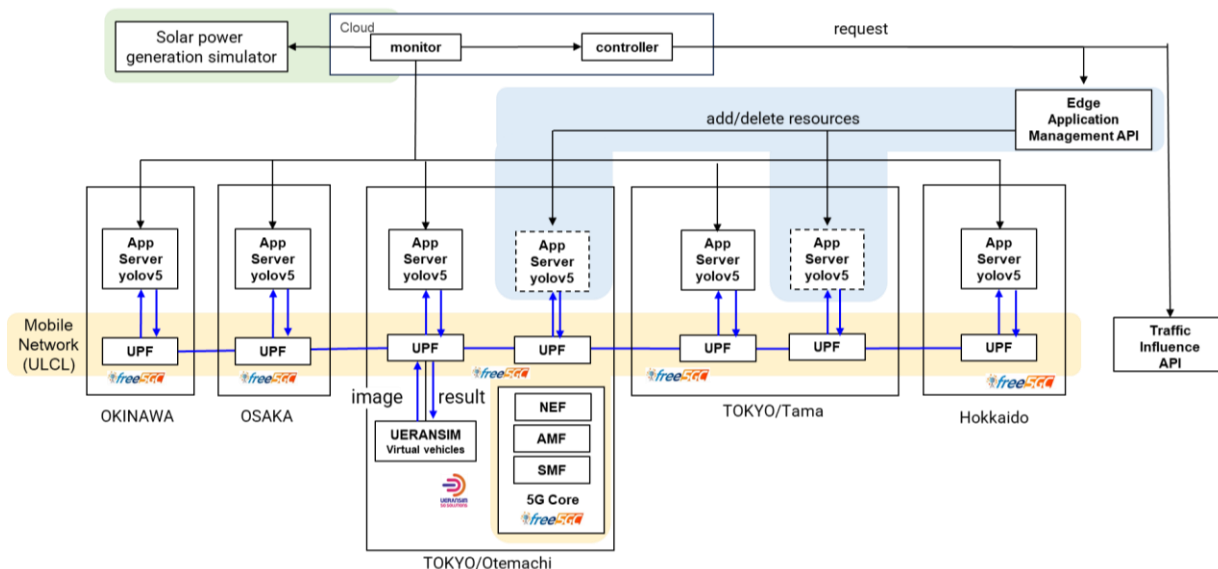


Figure 6: Green Energy PoC system configuration.

1. **Energy monitoring and traffic steering initialization:** A solar power generation simulator continuously estimates the availability of solar energy at different locations. The estimated solar power generation, predicted based on weather forecasts and expected sunlight levels at each site, is sent to a controller that makes traffic distribution decisions. The system aims to prioritize solar-powered edge servers while ensuring that service latency remains low.

2. **Dynamic traffic distribution based on energy metrics:** Traffic Influence API gathers vehicle movement and network conditions. The controller assesses real-time solar power generation and computational loads across multiple edge locations (Okinawa, Osaka, Tokyo/Tama, Hokkaido). Traffic steering is executed based on a combination of solar power availability, current CPU utilization at each edge location, and network latency considerations.
3. **Image recognition processing on edge servers:** Vehicles generate image recognition tasks, processed using YOLOv5-based application servers deployed at edge nodes. Each edge node is connected to the mobile network through user plane functions (UPFs) to ensure efficient data routing. The 5G Core at Tokyo/Otemachi facilitates communication between the virtual vehicle simulator (UERANSIM) and edge nodes.
4. **Load balancing with edge application management API:** If solar-powered edge servers have sufficient energy and low CPU load, they are prioritized. If energy availability decreases (e.g., nighttime or cloud cover or CPU load becomes high), workloads are dynamically shifted to other edge locations using the edge application management API. This API coordinates with edge nodes to distribute processing tasks efficiently, preventing bottlenecks.
5. **Real-time processing and results delivery:** Images from virtual vehicles are processed at the selected edge node. The results are then returned to the originating vehicle via the mobile network and 5G core.

Proof of Concept Results

Digital Twins Results and Takeaways

This digital twin demonstration highlights how adaptive edge selection enhances digital twin technology, paving the way for more efficient and reliable connected vehicle applications. It successfully demonstrates a scalable approach to traffic steering and edge computing for connected vehicles, ensuring:

- **Optimized data transmission** for real-time applications.
- **Seamless edge transitions** as vehicles move across regions.
- **Efficient resource utilization** through intelligent workload distribution.

By leveraging mobile network capabilities and edge-based processing, this system paves the way for next-generation intelligent transportation solutions.

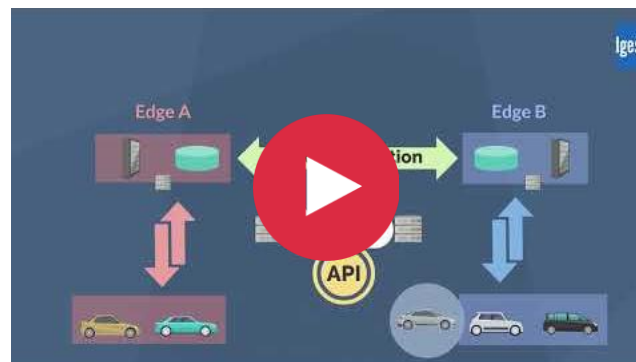


Figure 7: AECC PoC video: Optimal Edge Selection for Realizing Digital Twins (click to watch on YouTube).

Green Energy Results and Takeaways

The green energy use case demonstration showcased how adaptive edge server selection can balance green energy utilization with the performance demands of connected vehicle applications, paving the way for more sustainable and efficient mobility solutions.

Key takeaways include:

- **Optimized resource allocation:** The system ensures efficient use of solar power while maintaining low-latency processing.
- **Intelligent workload balancing:** Dynamic traffic steering prevents overload at any single edge location.
- **Scalable for real-world applications:** This use case lays additional groundwork for sustainable, energy-aware edge computing in connected vehicle ecosystems.

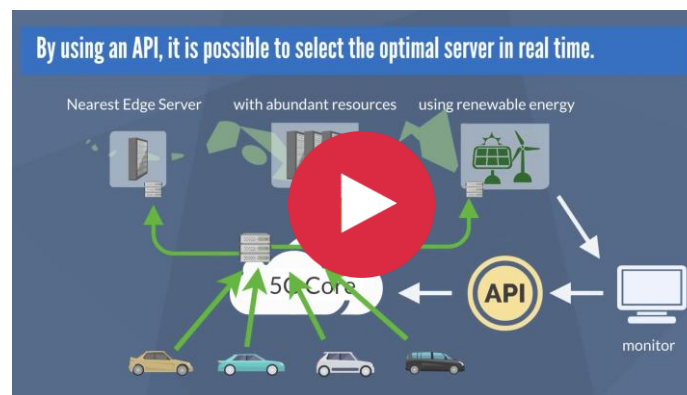


Figure 8: AECC PoC video: *Optimal Edge Selection for Realizing Green Energy Utilization* (click to watch on YouTube).

Next Steps

As the next phase, the engineering team is planning to apply this system to more advanced AI services (e.g., AI agents or digital twins) and enhance the system using new APIs.