

Top 10 Automotive Manufacturer Makes the Business Case for OpenStack

OPENSTACK WHITE PAPER

Contributors:

SOLINEA:

Francesco Paola, *CEO*

Seth Fox, *Vice President Operations*

Brad Vaughan, *Vice President Service Delivery*

OPENSTACK FOUNDATION:

Kathy Cacciatore, *Consulting Marketing Manager*



This work is licensed under the **Creative Commons Attribution-NoDerivatives 4.0 International License**.
To view a copy of this license, visit <http://creativecommons.org/licenses/by-nd/4.0/>



Executive Summary

The open source cloud operating system that has gained an enormous following by web-based organizations such as eBay, PayPal and Cisco WebEx , and tech businesses such as Intel, AT&T and Comcast, is quickly making itself at home in traditional industry vertical enterprises.

Large enterprise companies such as Disney, Bloomberg and Wells Fargo are adopting OpenStack to enable the agility needed to achieve competitive advantage, at a fraction of the cost of proprietary solutions. Every day OpenStack makes its way further into the enterprise because of its openness and flexibility to give companies the freedom to achieve goals on their terms.

OpenStack is the only solution that supports mixed hypervisor and bare metal server environments and allows users to choose from a number of competing solutions with the same underlying infrastructure. This competitive market helps customers lower costs, increase operational efficiency, and achieve true vendor independence on a stable platform supported by hundreds of IT vendors around the world.

This white paper describes how a top 10 car manufacturer used OpenStack to create an enterprise-wide private cloud to support a big data initiative. The adoption of OpenStack is serving as a catalyst to drive change across the technology organization, from chipping away at development silos to reducing dependence on closed-system appliances.



A HUGE NEED

The company needed a solution to collect, store & analyze petabytes of data globally from millions of sensors embedded in vehicles.

A FIRST ATTEMPT

Commercial legacy appliance costing in excess of \$1.5 Million could not handle the massive data processing workload.

THE RIGHT SOLUTION

\$125K OpenStack setup scaled from POC to 10-rack system in 6 months, easily handling complex data analysis.

Customer Need

Opportunity to Embrace Agility

The automotive industry is quickly becoming a software-driven business as the number of onboard computers and sensors take the wheel of the total driving experience. The automotive industry is projected to be the second largest generator of data by 2015, greatly influencing the way cars are manufactured and driven.

Telematics data produced by cars and drivers provide manufacturers with new insights to the way their products perform and how they're used. This car manufacturer saw the opportunity to leverage massive amounts of data being generated to create new features based on real-time customer insights and gain competitive advantages in research and development, product development, and customer service.

To start the process of cloud adoption, the company first had to architect a solution that could scale with the compute, storage and analytics demands associated with big data.

The company's IT deployment process was based on a very traditional enterprise environment, as well as being in lock step with the traditional design, engineering and deployment of cars. This led to duplication of efforts, and silos being created with no sharing of data.

For example, the quality assurance group gathered data from test cars to replicate infield failures while research and development gathered similar data from alternate sources for new car features and designs. Both saved and evaluated data from their own specialized appliances and vendor solutions from multiple geographic locations across the world.

Solution Requirements

Data as a Service Around the World

The company needed a cloud platform that enabled teams across the company and divisional lines to share data. They needed to provide data as a service across the organization and prepare its infrastructure for the petabytes of data they would eventually collect and analyze from multiple sources. The company had to be ready to analyze data from its product, as well as non-product sources, such as the servicing departments to drive quality. Data was also included from customer online forums and social media like Twitter and Facebook, which would yield insights to customer sentiments to feed into product prioritization as well as brand marketing.

The company had to store and analyze structured and unstructured data across multiple dimensions to gain a competitive advantage and develop product and customer insights.

They looked at OpenStack as a catalyst for sharing data and embracing open standards during the development lifecycle. In addition they were able to reduce costs and enable efficiency to cope with the explosion of data and compute resources on a global scale. To be successful, the OpenStack-based cloud had to excel in six categories.

Open Innovation

To provide a wide range of applications from diverse business units that could be managed centrally to reduce costs.

Data Analysis

For data scientists across business units to gain insights and competitive advantages in areas such as quality assurance and research and development.

Shared Data

To enable data as a service to be accessed globally by different divisions and subsidiaries across the organization.

Agility

To deploy new services quickly and respond to market demands.

Real-time

To quickly respond to vehicle quality issues and customers. This included the ability to respond to real-time data to make accurate decisions for consumers and vehicles.


Costs

To create a cost-effective infrastructure and software solutions that could also reduce the expense associated with managing enormous amounts of data.

Proof of Concept

Phased Rollout, Quantifiable Results

Based on these criteria, the company teamed with Solinea, a leading OpenStack software and services company, to architect a car cloud for internal lines of businesses to share information with each other as well as with consumers and third-party partners. This cloud would also connect with multiple groups to leverage data from legacy data warehouse applications.



They began with a small proof of concept (POC) that would validate the car cloud on OpenStack comparing use cases against the company’s legacy appliance. The POC started by defining metrics and KPIs to set quantifiable objectives that would help make the case to senior executives.

In three short weeks the team built the cloud and tested it with several practical use cases, including in-car mapping applications and back-end big data analytics with telematics data loaded into Hadoop.

OpenStack was deployed on a half-rack of commodity equipment, leveraging the following software components:

- Compute (code name Nova)
- Object Store (Swift)
- Image Service (Glance)
- Block Storage (Cinder)
- Identity (Keystone)
- Dashboard (Horizon)
- Orchestration (Heat)

The half-rack POC was completed at a cost of \$125,000, compared to the \$1.5 million legacy appliance.

During the test, the team loaded and updated mapping information in real time from vehicles to the cloud infrastructure. They also completed comparison tests to demonstrate the ability to scale and meet fluctuating workloads in a distributed environment. The automotive company preloaded vehicle telematics data into Hadoop, then transformed the data via map reduce jobs, and ran several analytics use cases to validate the accuracy and performance.

The comparison between the OpenStack car cloud and the company's legacy appliance was dramatic. In the use case with telematics data and Hadoop, the performance on the OpenStack car cloud clocked in at 40 minutes (before optimization), while the same job on the legacy appliance was unable to compute the function.

Beyond the performance gains, the test team was able to fine-tune the OpenStack architecture on specific workloads, which wasn't possible on the legacy solution. For example, they could reconfigure the cloud with more nodes in the Hadoop cluster to match the workload demands. Or they could run larger nodes to gauge performance gains.

With such stark quantifiable results, specific metrics and flexibility, they were able to justify investments quickly to gain funding for the OpenStack Hadoop car cloud, scaling from a half-rack POC to a 10-rack deployment in six months.

The final deployment was planned out over three phases:

Phase 1: Building out a production infrastructure with big data analytics for existing applications on a single rack consisting of 10 compute nodes, 48 TBs of block storage and 72 TBs of object storage.

Phase 2: Expand solution to focus on big data analytics for production vehicle telematics and in-car services in limited regions, upgrading the architecture and expanding the car cloud with seven racks consisting of 64 compute nodes, 181 TBs of block storage and 859 TBs of object storage.

Phase 3: Expanding the car cloud globally with a focus on in-car services with 26 racks consisting of 192 compute nodes, 1 PB of block storage and 5 PBs of object storage.



PoC: Half-rack with OpenStack

Cost: \$125K + Services

Timeframe: 3 weeks

Performance: 40 minutes

Legacy Appliance

Cost: \$1.5mm

Timeframe: 2 weeks

Performance: Did not compute

Insights to Achieving Adoption

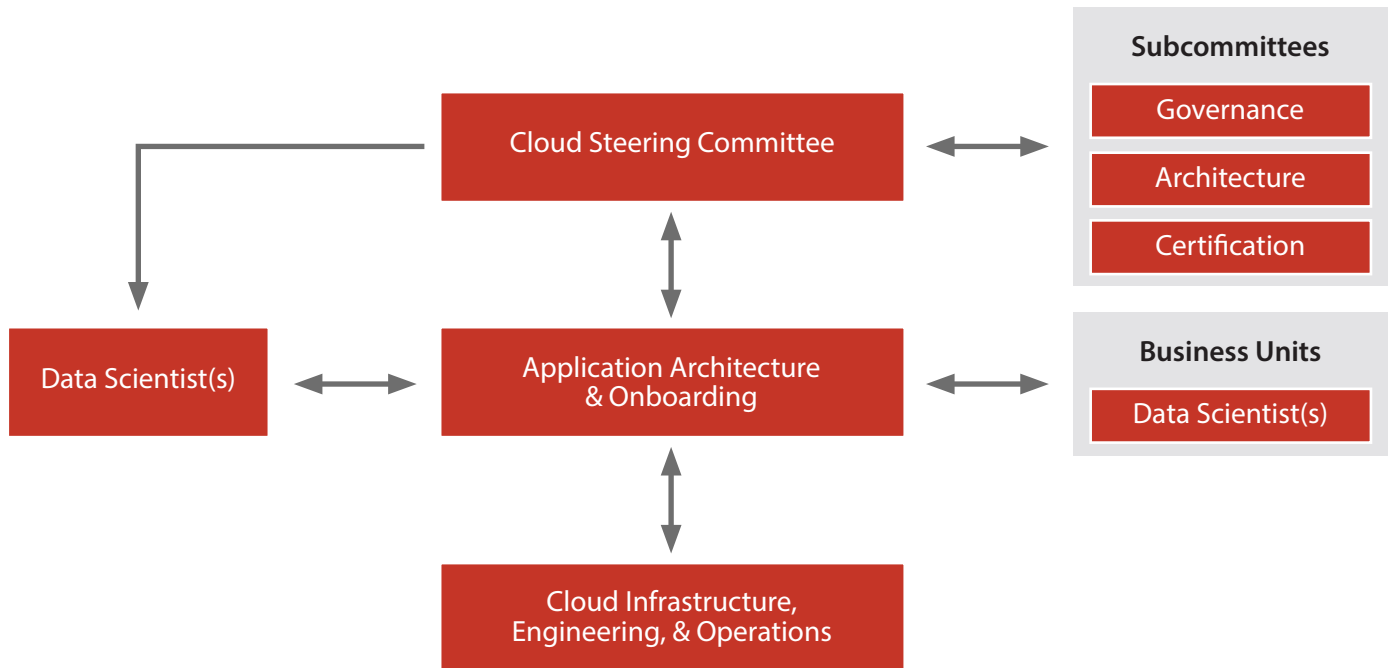
Organizational Change

A key component to the car company's adoption of OpenStack was the creation of an incubated Cloud Competency Group outside of IT. This group drove adoption as well as governance of the car cloud deployment across the entire company.

The Cloud Competency Group included a cloud steering committee that consisted of IT executives, business unit executives and specific engineers within the organization to drive the strategy and flow of the applications deployed on the cloud.

It included an Application Architecture and Onboarding team that set standards for architecting solutions to be optimally deployed within the OpenStack environment. Whether it was big data or mobile applications, having the centralized group to help guide development teams across the organization was critical to ensuring applications ran effectively. For this deployment, big data analytics required the team to build a pool of cross-LOB data scientists that could guide the application performance and query requirements.

The core of the Cloud Competency Group was comprised of cloud infrastructure engineering, DevOps and operations resources, critical to the successful management, monitoring and operations of the car cloud.



Information Exchange among the Cloud Competency Group

Cultural Change

The adoption of an open source based cloud solution to support a new data driven business model across the organization required a change from a traditional waterfall solution design to a more agile approach. Rapidly progressing from user stories to implemented solutions and gradually building more features and capabilities into the solution became a critical success factor.

This change impacted the:

People

New skills in engineering and design transitioning from vendor products to open source solutions

Process

Focus on data collection, metrics, engineering and self-support

Technology

Best of breed toolsets instead of large single vendor frameworks.

In addition to forming a Cloud Competency Group and facilitating cultural change, Solinea recommended these steps to ensure the success of an OpenStack cloud implementation:

- Maintain executive sponsorship during the entire deployment by staging the rollout in small phases and demonstrating quantifiable results along the way.
- Choose carefully how to position and communicate your rollout in the organization. Rather than positioning it as a product/project rollout, promote it as a long-term deployment of shared services across the entire organization.
- Make sure business units, including IT and operations are aligned on the cloud objectives.
- Standardize processes and offerings to minimize duplication and overlap.
- Review internal policies and procedures to prepare for a hybrid public cloud.

Conclusion

With an agile approach to the rollout, this top 10 car manufacturer was able to develop a car cloud that propelled it into big data and a software-based automotive arena. It enabled them to increase the velocity of their development while decreasing costs, and breaking down development and communication silos to share information as a service across multiple business units.

Agility in this case was about being able to deploy services faster with insights from big data analytics to improve products and services. Other enterprise companies can help replicate similar results with the same agility through small yet productive steps. To ensure success, the company validated the concept, progressively upgraded its skills, transformed the governance, and moved to a service-based approach.