NAVIGANT
RESEARCH

**WHITE PAPER**

# Connected Vehicles and the Cloud

## Enabling the Future of the Mobility User Experience


**Published 2Q 2018**

**Commissioned by Amazon**



**Sam Abuelsamid**
Senior Research Analyst

**John Gartner**
Director

# INTRODUCTION

In just over a century, the automobile has done more to enable personal mobility than almost any other device, allowing people to travel almost anywhere, whenever they like. Over the past 2 decades, the mobile phone revolution and ubiquitous high speed data connectivity has transformed the way people see the world and has provided a glimpse into aspects of life in faraway lands. It has changed how people interact, creating connections between those whom have never met in person and enabling them to do more in less time. The transformations from in-person conversations and written letters to wired networks and then to a vast trove of data constantly floating through the air would have been hard to conceive in the early 1900s.

However, as life gets busier, the personal mobility machine people know as the car and the ubiquitous yet utterly intangible communications network will fuse. By 2020, virtually every newly built vehicle in North America—and most vehicles globally—will be connected. Harnessing that communications power requires managing where the messages go, making sure the wrong messages do not end up at the wrong destination, and keeping it all running. Not so long ago, that meant every company would have to build its own unique infrastructure, something most are ill-equipped to do, particularly as the number of connections and complexity of the interactions increase. Today, fast, reliable, and robust cloud computing platforms provide the underpinnings to build new businesses and services that cross manufacturer platform boundaries seamlessly.

Anyone with a great idea can now execute it faster and reach a wider addressable market than ever before. New business models can be deployed on cloud platforms at a fraction of the cost of building out dedicated network infrastructure and scaled up easily to meet customer demand. Cloud providers also have the expertise to keep the platform more secure and resilient. Synergy is a vastly overused and abused word, but the connected car and the cloud definitely create the opportunity for products that are more than just a sum of parts.

# SUPPORTING THE CONNECTED CAR FROM THE CLOUD

Two decades after General Motors' (GM's) launch of its OnStar services in the US, the concept of the connected car is hardly new. However, a confluence of factors is driving a rapid expansion of deployment. These include accelerating wireless speeds, declining service costs, and the ability to more easily build services using readily available cloud infrastructure resources.

## Why Is the Car Connected?

Modern life is increasingly connected, with most people around the world now carrying at least a mobile phone most of the time. With more than 2 billion active Android devices and 1.3 billion iOS devices, nearly half of the world's population now has at least one device with them that outperforms the super computers of a generation ago. These mobile connected computers share and collect data from almost every part of the world.

In many cases, these portable devices have faster wireless data connections than people's home internet service—analytics company StatCounter reported that mobile internet use surpassed desktop use in 2016. With all that ubiquitous connectivity, it makes sense that the most mobile device in most people's lives, the car, should be part of the same ecosystem. Through the course of the 20th century, the automobile changed the way humans live, work, and play, adding a degree of mobility never previously imagined.

***Figure 1*** ***Stock Traffic Congestion Image***

*(Source: Osvaldo Gago)*

In the 21st century, changing lifestyles, busier schedules, and increasing urbanization mean that many people spend more time in the car commuting than ever before. The Texas Transportation Institute report, *2015 Urban Mobility Scorecard*, estimated that Americans wasted an average of 42 hours annually due to congestion, and that commuters in the worst congested cities lost double that amount. Staying connected gives commuters a means to be productive, entertained, or educated during those hours of playing hurry up and wait.

Connectivity also provides a means to take friction out of life and add convenience by reducing the need to remember mundane things like, “when did my car last get serviced,” “where did I park,” or “where can I grab a quick lunch on the way to my next meeting.” The next generation of assistants will add contextual awareness of where a person is and what they intend to do, combined with their preferences to make the vehicle more personalized and proactive, rather than reactive. As cars become ever more automated, connectivity will be a key enabler with continuously updated maps, communication with other vehicles, people, and devices nearby, and the ability to be fully productive or thoroughly entertained while heading to a destination.

## The Internet of Everything

The internet began as a means of connecting disparate computer systems in far-flung locations to provide a means of data sharing, application distribution, and improved fail operational capability for networks. Over the past decade, processing power and wireless communications capabilities have become exponentially faster, and smaller and cheaper computing capabilities have increasingly been added to almost any imaginable device and spawned the Internet of Things.

***Chart 1 LDV Population with Built-In Telematics by Region, World Markets: 2016-2025***

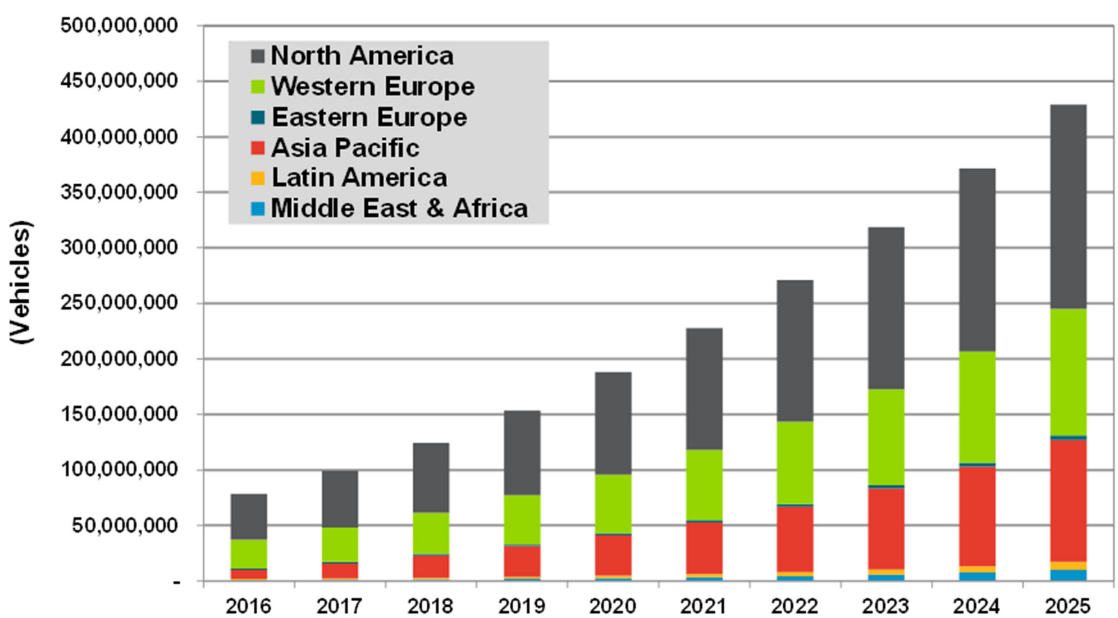


*(Source: Navigant Research)*

Today, objects as mundane as light bulbs and as advanced as self-driving cars are nearly all interconnected in some way. As people commute home from work, approaching a grocery store can trigger a notification to stop and get a few supplies for dinner. Walking back to the car can trigger a signal to the thermostat at home to bring the temperature to a more comfortable level. Being stuck in traffic can trigger a delay in pre-heating the oven, and pulling up to the driveway can automatically turn on the lights in the doorway. Everything can be connected to anything, but it all runs through servers in the cloud that route signals, determine which vehicle to send the air conditioning on signal to, and perhaps pre-order the pizza you are about to pick up for dinner. The cloud powers it all.

## Why Is the Cloud a Key Enabler?

The emergence of high speed connectivity in the 21st century has enabled a fundamental rethinking of how computing resources are managed. At the birth of the worldwide web, almost everything was built locally at the edge of the network and then connected to the internet backbone to provide communications. As the web expanded and became commercialized, every site needed to individually scale-up its own infrastructure, building out for its expected peak loads. This approach dramatically increased costs and technical support needs and contributed to the initial boom and bust of the internet at the turn of the century as companies spent vast sums to build out both local and backbone infrastructure and then ran out of cash.

That investment was not a total loss, however. The data capacity provided by all the fiber laid in the previous years was combined with new, faster, and more efficient servers. By eliminating the bottleneck that limited the applications that could be executed using centralized data centers, this upgraded network paved the way for the modern cloud computing revolution.

Cloud computing is the on-demand delivery of compute power, database storage, applications, and other IT resources through a cloud services platform via the internet. With cloud computing, developers do not need to make large upfront investments in hardware and spend a lot of time managing their infrastructure. Instead, they can provision exactly the right type and size of computing resources needed to power the newest bright idea or move their entire IT department to the cloud. The cloud economics model is based on pay-as-you-go pricing, with no upfront investment; effectively providing unlimited computing resources on-demand and automatically scaling up and down as needed.

***Figure 2 The Connected Car and the Cloud***

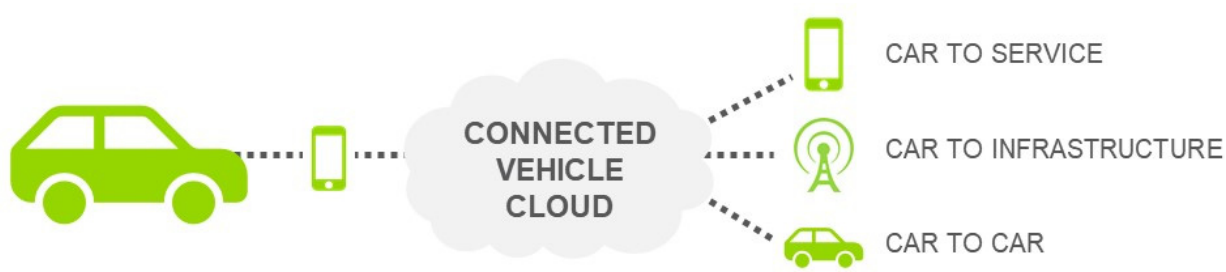


*(Source: Navigant Research)*

Utilizing computing resources in the cloud provides crucial advantages for everyone from independent startups to vast multinational manufacturers:

- Trade CAPEX for variable expense
- Increase speed and agility
- Benefit from massive economies of scale
- Eliminate the need to guess capacity requirements
- Stop spending money on running and maintaining data centers
- Scale globally in minutes

### Ubiquitous Connectivity Makes It Possible to Use the Cloud Anywhere

The emergence of wired (and later wireless) broadband connections and improved virtualization technologies enabled the creation of the modern cloud computing platform. A developer with an idea is now able to focus on what they had that is unique and rapidly deploy it on a cloud platform. Instead of building infrastructure from scratch for every new application, the platform is now a service available on-demand at a fraction of the cost it would have been a decade earlier. Hugely popular services like Netflix and Snapchat were created with relatively modest initial investments, but have been able to scale-up to massive businesses because they were born on such cloud platforms.

As vehicles become ever more connected and generate more data that can be utilized to build new service offerings, cloud platforms become a crucial component of the ecosystem. With more than 270 million registered vehicles in the US alone and more than 1 billion globally, the number of potential moving network nodes is simply too large to provide direct peer-to-peer communications. Having every service provider trying to directly communicate with each vehicle is not scalable, especially since much of the available data may be utilized by multiple services.

The cloud enables the deployment of abstraction layers that can aggregate data from millions of vehicles, anonymize it, manage it, and make it available to service providers through application programming interfaces. This mechanism enables developers to target a much larger addressable market across multiple brands while dealing with fewer potential proprietary interfaces.

Without the need to focus on infrastructure deployment, small teams with big ideas can launch new services with minimal upfront investments. Concepts that would have required millions of dollars in venture funding in the 1990s can now be executed much more rapidly by smaller teams with only tens or low hundreds of thousands of dollars in seed money. Even if an idea does not become the next WhatsApp or Skype, the low cost of operating on a cloud platform can make a niche product financially viable. The cloud enables innovation to be put back in the hands of the creators, and not in an IT or finance department.

## Computing in the Middle, at the Edge, or Somewhere in Between

In 1970, cars were essentially entirely analog and generated no usable data. By 1980, a small trickle of sensor data was being fed into early engine management computers and anti-lock brake systems. However, it was used only for real-time controls with no way to get it out of the car. At the turn of the century, the onboard diagnostics port had become standard and GM and other manufacturers were installing first-generation cellular telematics systems to enable features like automatic emergency response, remote locking/unlocking, and other safety and convenience capabilities.

Today's cars are generating upwards of 200 GB/hour and automated vehicle prototypes are producing 4 TB/hour of raw data. The vast majority of these bits and bytes are consumed in the process of real-time control for everything from the brakes to steering to blindspot monitoring. However, there are significant portions of this that have usefulness elsewhere in the value chain.

Even with the most optimistic projections for wireless connectivity advances, transmitting unprocessed data from the car will always outpace the capabilities of the network. The onboard compute platforms in the vehicle will have to do most of the parsing, filtering, and compressing of data that has value beyond real-time control.

***Figure 3 Cloud Computing and Edge Computing***

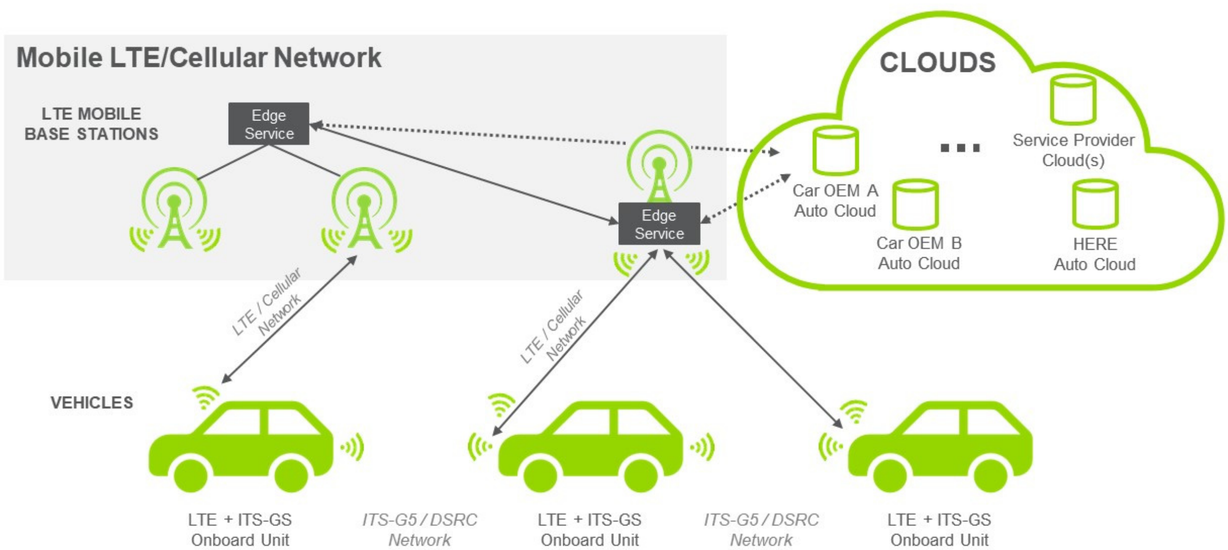


*(Source: Navigant Research)*

Curated and compressed data can be further processed at the edge of the network where a wide range of near real-time information relevant to traffic management systems and other vehicles can be put to work. This includes data about traffic and road conditions that can be used in signal management systems and shared back out to the rest of the fleet to be factored into navigation routing and vehicle control systems. For example, information about slippery road conditions can be shared back to vehicles approaching that area to alert drivers or automation systems to slow down.

The data center provides a hub to aggregate and process vast quantities of data from road users. Utilizing big data analysis tools, meaningful insights can be extracted from the data that can provide value to road users, manufacturers, suppliers, insurers, transportation authorities, and a vast array of other service providers.

### Bridging the Connectivity Gaps

Applications that apply relevant data as soon as it becomes available can be used to trigger follow-on events that inform or launch other services. In the emerging mobility as a service (MaaS) ecosystem, interlocking services will need to run on cloud platforms while also communicating with the smart nodes at the edge of the network that are connected and shared automated vehicles (AVs). Most of the companies developing AVs today are also working on new business models that can take advantage of the flexibility these technologies can enable while providing a steady stream of recurring revenue.

MaaS models go beyond just transporting people to delivering food, groceries, and other cargo. The ability for developers to create applications that can live on the cloud and interoperate will be crucial to optimize this ecosystem. These new services are just beginning to come to life and will grow in coming years. Having the ability to create these apps in a way that can scale from supporting a few dozen prototype vehicles to tens of millions of production models in the future without requiring a total rebuild of the platform will be crucial to the economics of this new ecosystem. Cloud services that can support this sort of deployment without requiring an in-house server management team will be essential for both startups and century-old automakers that want to be a part of this space.

However, as good as modern wireless connectivity has gotten, it is still far from perfect. There will always be times when a cloud connection just is not possible and the ability to run some of those applications locally for a time may be necessary. Thanks to virtualization, it is possible to run a subset of the cloud platform locally with the same applications and then sync relevant data back to the cloud once a connection is resumed. All of this is possible without even changing any of the application code. It can be as seamless as if there were no gaps in the connection.

### Cost Reduction

Developing a great consumer- or driver-facing product or service is a huge challenge. It is also a vastly different challenge from managing network infrastructure. If every automaker or service provider has to do that independently, significant additional staffing and hardware are required, much of which may remain vastly underutilized much of the time.

A serverless computing infrastructure running on a cloud platform can be utilized to execute on-demand. When an application is needed by user, provider, or manufacturer, a virtual server can spin up almost instantly, handle the required tasks, then shut itself down, making the resources available to other customers as needed.

The utilization of shared computing and management services in the cloud can slash information technology expenses by eliminating much of the upfront investment in infrastructure and staffing. The cost of computing for the user shifts CAPEX to OPEX. When no operations are needed, there is little or no expense. Cloud providers can locate data centers to take advantage of lower power costs and greater efficiency. At the same time, they can provide a degree of redundancy that would not be practical for many companies running their own infrastructure. Since a digital business can only make money when it is available, reliability is paramount and cloud providers can be hugely beneficial.

## The True Power of Distributed Computing

Once the cloud has been harnessed to process vast quantities of data, the real innovation begins. Machine learning algorithms are trained by data aggregated from a large population and then triggered by each user's data, which can personalize results in tune with that user's context.

Unlike the traditional automobile that retains largely the same feature set throughout its useful life after leaving the factory floor, the connected car has the potential to gain new functionality built in the cloud until the day it is retired. However, this continuous improvement must be tempered with a focus on maintaining security.

# AWARENESS AND UNDERSTANDING

In this section, a deeper look is given to some of the services that become possible when merging computing resources locally in the vehicle, at the edge of the network, and in the cloud.

Until the emergence of the connected car, every vehicle was an isolated node. Messaging was limited to whatever drivers could transmit via nonverbal communications such as a nod of the head, a hand gesture, or a flash of the headlights. The cars of today and the future know where they are, who is in the vehicle, where they are headed, and a plethora of information about the current status of the vehicle itself. By feeding this into data platforms and algorithms trained by past data, a wealth of new and relevant services can be offered up to travelers.

## Context Awareness

In any sort of interaction, whether between human and human or between human and machine, context is always a key factor. Understanding what the other side of the interaction knows, is capable of doing, and wants determines success or failure. While humans are flawed, they have evolved to have a surprisingly capable understanding of nuanced communications, especially when they already understand the context. Memory is crucial to that context awareness—memory of past interactions and desires. In the realm of digital services and human-machine interactions, cloud platforms that connect disparate devices and services can help to enhance that memory and corresponding context awareness.

### Local Services, Targeting the Connected Consumer

Until the age of the internet, advertising and attracting customers was largely done by brute force. Often little knowledge was available about who was seeing any given message, where they were, and what they wanted. In the digital age, it is possible to generate a much deeper understanding of who an individual is and what information, products, and services they may be interested in by correlating data across multiple sources. With such individual profiles, services can now be targeted at those most likely to be interested.

With the combined knowledge of a driver's location, the service history of their car in the cloud, and current vehicle data like tire pressures and fuel consumption, platforms have the potential to detect when a tire might be in need of repair and provide an ad for an instant discount at a nearby service facility. Knowledge of a driver's calendar and culinary preferences could provide a prompt to make a dinner reservation at a popular new restaurant that just had a cancellation. Some people might find such awareness of their preferences unsettling, but increasingly, many find it compelling and valuable.

Getting that value requires compiling data and analyzing it for patterns. This is much more effective on a cloud platform where petabytes of data can easily and economically be aggregated with little or no friction, and can be processed using the virtually unlimited computing capacity in the cloud.

## Voice Recognition

Over the millennia of human civilization, voice has been the most enduring means of communication. As vehicles get ever smarter, voice is also becoming a primary means of letting cars know what people want.

### *In the Car*

Until AVs become widely accessible—which is unlikely until at least the mid-2020s—the primary task of the vehicle operator must be driving. Eyes on the road, hands on the wheel. Any task that takes attention away from controlling the vehicle is likely to contribute to accidents, injuries, and fatalities. As the traffic fatality rate has reversed its long-term decline in recent years, distraction is seen as one of the key culprits.

***Figure 4 Talking to the Car***

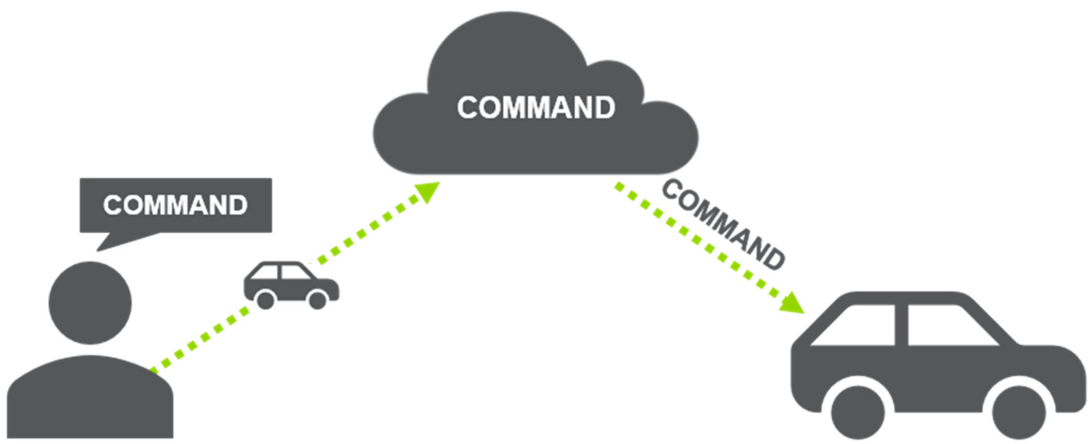


*(Source: Navigant Research)*

Adding access to features and services in the vehicle must be done in a way that does not add to the driver's cognitive load. Searching for targets on a touchscreen or buttons on a console is problematic, but it is an issue that can be addressed by reliable voice control.

Voice recognition is a computing-intensive task, and until recently it was executed entirely by systems embedded in the vehicle. Unfortunately, the long-lead times, challenging acoustic environment, and need for long-term reliability have made in-vehicle voice control troublesome at best and often more frustrating than useful for drivers. As a result, many automakers have taken significant downgrades in quality surveys and customer satisfaction resulting from frustration with this important in-vehicle technology.

The recent addition of smartphone projection systems with cloud-based voice recognition to the vehicle infotainment system has taught drivers that more is possible. The ever-increasing ubiquity of embedded connectivity is now enabling the addition of hybrid voice recognition in the vehicle that utilizes both an embedded component and a cloud component.

More powerful in-vehicle compute platforms are improving the capability of the embedded voice control. However, the connected component is where the real leap forward is happening. The connected vehicle can harness vastly more computing resources to provide nearly natural language voice recognition and instant search. Rather than the limited vocabulary of traditional embedded recognition, users can ask for a service, play media, or set a destination without using specific command syntax.

Cloud-based digital assistants can further build on the information people share vocally by aggregating data from all the devices that are used on a daily basis, including phones, tablets, and home voice assistants. The powerful machine learning algorithms that can only realistically run on cloud platforms can take clues from human responses to the millions of interactions a day to improve both word recognition and semantic understanding. A better grasp of the meaning and context of people’s word choice will enable assistants to use the bots that are increasingly available from network services to continue sanding down the rough spots in daily life.

#### *From the Outside*

The advent of the modern plug-in vehicle coincided with increasing wireless data speeds which in turn have enabled new means for drivers to take control of their vehicles remotely. Smartphone apps are now becoming voice enabled, with apps that allow EV owners to monitor battery charging, remotely pre-condition the climate control while the car is plugged in, and perform other functions. Instead of having to dig out phones to tell the car to warm itself up on a cold winter morning, people can now simply ask a device on the kitchen counter to do it while they pour their coffee and eat their cereal.

Enterprising developers and tinkerers are building applications that connect every device in a seamless way. For instance, if a user is listening to an audiobook over breakfast, they can pick up from the same location when they sit down in the pre-warmed car with their navigation preset to the destination of their first appointment. Voice is the ultimate context-aware communications system.

### Advanced Mapping and Navigation

In-vehicle navigation is a technology that has been spreading for more than 2 decades. Until recently, it was a blunt tool, mostly useful for automating the process of the routing people used to do with fold-up maps and atlases. For a time, it appeared that mapping apps from smartphones might supplant the pricey, embedded in-car navigation system. Digital maps have come a long way in the past decade from flat two-dimensional representations of streets and some landmarks to vastly more detailed three-dimensional representations.

While those smartphone maps helped people find their destinations in unfamiliar places, the embedded navigation system is making a comeback as it is utilized for more than turn-by-turn directions. A number of cars on the market today use elevation data in the maps as a signal to help intelligently control hybrid and electric powertrains for better efficiency. Real-time traffic layers are used to generate routes that can save time and energy. Semi-automated cars use maps as long-range sensors to adjust speed when approaching curves before visual sensors can detect the road geometry. Traffic signal data can be used to inform the driver of the state of upcoming lights on the route and adjust their speed accordingly.

Individual vehicles are beginning to contribute the data collected by on board sensors back to cloud-based mapping services. This is being used to generate the high definition maps that will be required for AVs to more accurately localize themselves as they navigate through the world. As road conditions and configurations change, maps can be updated in near real-time in the cloud and rapidly fed back to the vehicles in the field—potentially within minutes rather than annually as traditional navigation systems did.

### Predictive Maintenance

Getting stranded at the side of a busy highway because of a mechanical breakdown can be both inconvenient and dangerous. The sensors on modern cars feed signals to electronic control systems that manage functions in real-time, automatically adapting to changing conditions. However, just as humans have patterns of behavior, so do machines.

Data scientists have devised ways to detect anomalies in the behavior of systems from the sensors that monitor them. By collecting these signals and aggregating the data in the cloud across a large enough population, patterns that signal conditions like impending failure can be characterized. Comparison of individual vehicles to the data patterns in the cloud can be used to provide alerts to drivers that they should probably get their vehicles serviced soon to avoid getting stranded.

Dealer service departments and independent shops can use APIs from cloud-based data brokers to target special offers at drivers based on those notifications and locations. Even more mundane information such as recommended service intervals and odometer data can be used to target potential customers. A driver receiving an offer that they cannot take advantage of immediately may be able to immediately call a service provider to make an appointment, later receiving a reminder from their calendar.

For fleet operators, ensuring the availability of their vehicles is even more important to ensuring profitability. Being able to predict when service will be required can minimize unscheduled downtime that adds to costs.

## Solving the Hard Problems

The combination of vast quantities of data and previously unimaginable compute power spread across thousands or millions of servers provides the potential to tackle problems in ways never before imaginable.

### Machine Learning

Since the dawn of electronic computing, most tasks have been handled by deterministic algorithms that follow known sequences of steps for any given set of input signals. This approach has worked well when trying to solve a constrained problem set with a known scenario. However, as computer scientists and engineers tackle ever more complex artificial intelligence (AI) problems like automated driving, a purely deterministic approach is inadequate.

This is where machine learning algorithms come into play. The real world of driving, especially in urban environments, has a nearly infinite number of variables including the participants, weather, and road conditions. Machine learning requires feeding in vast quantities of data to help train AI software.

In the transportation realm this includes data collected from vehicles. The data must be parsed, annotated, and labeled to train the algorithms, which is where cloud computing can be hugely important. Highly automated vehicles can generate up to 4 TB of data per hour and there are hundreds of test vehicles already on the roads accumulating millions of miles a year. In the process of teaching cars to drive themselves, data centers are a critical tool that can process this information and then feed it into simulations.

However, machine learning extends well beyond that to applications in use today. For several years now, insurance companies have been collecting data from connected adapters provided to customers in exchange for premium discounts. This sort of data can be fed into risk models to provide more accurate pricing and better returns. Real-time data about the activation of driver assistance features, such as anti-lock brakes and stability control, can feed into weather forecast and traffic models. By recording and collecting data from connected vehicles, a wide variety of predictive models can be powered by machine learning in the cloud to estimate the life of key components like batteries, tires, and brake pads with high accuracy.

### Bots to Reduce Friction

When most people think of robots, they are likely to visualize a machine first seen in films—like *Star Wars* or *The Terminator*. Closer to reality are devices like Honda's cute and diminutive bi-pedal Asimo or the industrial robots that have populated factories for decades. Industrial machines have been developed to handle tasks that may be too difficult or dangerous for humans to handle, thus reducing friction in production processes.

In recent years, a new breed of bot emerged. Existing purely in the digital realm, it smooths out the interactions between applications running across servers. The general public has begun to experience some of these bots through a variety of messaging applications that can answer questions or perform a variety of tasks ranging from finding a restaurant and making a dinner reservation to ordering a car from a carsharing or ride-hailing service.

***Figure 5 Cloud Bots to Provide Automated Services***

*(Source: Navigant Research)*

The real power of software bots lies behind the scenes. These are essentially autonomous applications that talk to each other, sharing a broad range of data and providing numerous functions. Bots can be used to manage carsharing services, tracking where vehicles are, where they are likely to be requested at various times of the day, what the weather and traffic conditions are, and then dispatching staff to reposition them or bring them in for service.

For end users requesting a carshare, a bot can automatically fetch and load their preferences upon arrival. When they unlock the car using a temporary code received from the cloud and transmitted by Bluetooth LE, the radio could be preprogrammed with their preferred stations and the climate control set to a comfortable temperature. If the car is equipped with memory seats, even the position last used by the driver with the same type of car could be restored.

Today, commuters use connectivity and the cloud to request rides from transportation network companies on their smartphones. In coming years, the bot will play an increasingly important role as the transition occurs from humans driving vehicles on ride-hailing platforms to highly automated vehicles. Instead of the user pulling out a phone and opening an app, bots will likely access the online calendar, current location, the user's mobility services subscription, along with weather and traffic information.

A user's personal bot can be aware of their preferences to help improve the experience of commuting. It can talk to the bots of other providers and have a ride from the nearest available service that they have an account with ready and waiting to get them to the next appointment on the calendar with no direct interaction. When they sit down in the car, the audiobook they were listening to over breakfast can automatically resume, or their favorite playlist can fire up. All payments can also be handled automatically and the vehicle may even make additional stops to pick up other people going to the same destination by checking their calendars. This nearly frictionless user experience is only possible thanks to the cloud.

### Making It Easier to Build Solutions

A diverse ecosystem is always more robust and resilient than a monoculture. This is just as true in business as it is in nature. At the same time, there are significant advantages to scale and be able to target a broad audience. The mobile application and services business has grown and thrived over the past decade in large part because developers have had powerful platforms to target. These platforms aggregate audiences so that developers and entrepreneurs do not need to create custom applications for every device.

In the automotive landscape, manufacturers have long focused on proprietary solutions designed to provide product differentiation. This differentiation often extends to individual models within a manufacturer's lineup. With dozens of major automakers around the world, this approach has traditionally made it difficult to build services aimed even at a single company's entire lineup—much less the broader transportation market.

The cloud enables the creation of abstraction layers where data pulled from many vehicle models and manufacturers can be normalized into common formats and then made accessible to service providers. This enables developers to target a more manageable set of platforms while still leveraging the scale of the larger marketplace. A provider building a system to manage urban parking can aggregate availability from many lots and garages and access multiple vehicle data platforms used by any number of automakers. Drivers can then access apps from the vehicle that share their location and estimated time of arrival to find, reserve, and even pay for a parking spot. All of this would be impractical if providers had to individually target every vehicle platform.

The converse is also true for automakers. They understand that having the most popular or innovative apps and services available in the vehicle can help drive sales. For years, they have struggled to get developers to target their proprietary platforms because they were difficult to work with, incompatible with those from competing brands, and had limited target audiences. From the user's perspective, a proprietary platform adds little or nothing to the experience except speed bumps to getting features and services they want to use.

Feeding data into cloud-based data platforms can enable manufacturers to focus on providing unique features and services to customers that set them apart, many of which can provide new revenue streams. Thanks to the cloud they can do this without investing in dedicated infrastructure that does not add value.

### Building Powerful Services without Infrastructure

When the commercial internet initially began to flourish in the mid-1990s, countless entrepreneurs conceived of new ways of conducting business online that were more efficient, and entirely new lines of businesses that could not be executed before. However, when an online business became popular, it had to scale its computing infrastructure rapidly to meet peak load demands.

Unfortunately, those peak loads were often well above the average demand. In an era when every business basically needed its own server infrastructure this meant it was frequently underutilized. Building and managing that infrastructure was generally not a core competency for a new business just getting started. This often led to unreliable service and magnified costs.

The development of modern, more efficient data centers, virtualization technologies, and rapid provisioning technologies has changed all of that. Cloud platforms specifically designed to run an infinitely variable array of services now make it possible for a business to sign up for an account and deploy applications in a matter of hours—or days in many cases. They never need to worry about the server hardware, power supplies, or even scaling. In many cases, a business can add dozens or thousands of additional servers on-demand without ever setting foot in the data center.

Experts in managing large-scale infrastructure can handle those tasks for thousands of businesses at a time while the developers focus on creating great products for customers and democratizing advanced user functions like neurolinguistic programming.

### Continuous Improvement for Life

The traditional model of the auto industry has focused on building a great product that appeals to customers at the time of purchase. For the best automakers, continuous improvement is focused on the manufacturing process and the development of the next generation of product. Connectivity and the cloud now provides the opportunity for continuous improvement from product cradle to grave.

# FUTURE-PROOFING THE CAR

For more than a century, the automotive business model has revolved around designing, building, and selling products to consumers and then moving on to the next product. Aside from service parts and some aftermarket upgrades, compliance recalls were often the only reason for an automaker to continue development on an existing product. Effectively, the car purchased on day one was the product the buyer had for the life of the vehicle. In the 21st century that is no longer a viable option.

## Over-the-Air Updates

The advent of electronically controlled systems in automobiles began in the 1970s, and has led to cars that are cleaner, more efficient, safer, and better performing than at any time in the history of these machines. Today's cars often contain up to 75 or more discreet electronic control units (ECUs) with a combined total of more than 100 million lines of code. Software drives everything from power windows and heated seats to advanced driver assist systems.

Until the launch of the Tesla Model S in 2012, if an automaker needed to update any of that software except for some of the telematics control units it required a visit to a dealer to reprogram an ECU. Tesla changed the ball game by using built-in cellular and Wi-Fi connectivity to deliver software updates that fixed bugs and added new capabilities over-the-air (OTA) without ever going to the service facility. Customers are increasingly accustomed to getting such OTA updates on smartphones, tablets, and smart Internet of Things devices at home such as thermostats and mesh Wi-Fi routers. Those same customers now question why their non-Tesla cars cannot do the same.

***Figure 6 Over-the-Air Updates for Functional Improvement, Safety, and Security***


*(Source: Navigant Research)*

As the car becomes ever more complex and defined by software, automakers are working to implement the same sort of capabilities for all new models going forward. Incremental updates to improve performance, capability, and security will become the new norm in the auto industry. Unlike Tesla with its more limited product line and configurations, most other manufacturers offer more variety, which adds to the challenge of what to update on any individual vehicle. Managing OTA updates for hundreds of millions of vehicles will be incredibly challenging; robust connectivity and a secure cloud platform are essential components for success.

A cloud-based deployment platform can be used to query individual vehicles to determine update status and configuration and then deliver customized updates. As the world moves forward into an era of automated driving, this will become essential. No competent software engineer will be under any illusions that completely bug-free software can be delivered every time, regardless of how rigorous testing is.

Early AVs will inherently have limited capabilities as manufacturers learn how to make the technology work in the real world where AVs must coexist with human driven vehicles and variable weather. Just as Tesla delivered a stream of software updates to customers with its AutoPilot semi-automated system, every AV will require performance and bug updates provided from the cloud.

### Cybersecurity and Privacy

With ubiquitous connectivity comes new attack surfaces for bad actors hoping to exploit security vulnerabilities for financial or potentially political gain. Today’s vehicles are more often equipped with physical ports, such as USB and OBD-II, and wireless connections such as cellular, Bluetooth, and Wi-Fi. The transition from purely mechanical to electronically controlled actuation for steering, braking, and propulsion creates the potential array of bad outcomes that can be relatively benign such as data theft or a ransomware attack that disables a vehicle.

However, unlike the purely online world where the consequences of a successful cyber attack are rarely life threatening, the same cannot be said of the connected car. If perpetrators were to penetrate a control system that managed anything from OTA updates to dispatch of AVs, the impact could be far more serious. A bad actor that managed to find a way to command millions of cars simultaneously could wreak havoc in many locations.

It is crucial to build the entire value chain around the connected vehicle with security in mind. Since it is virtually impossible to guarantee absolute security and reliability in a complex system, both attack detection and resistance and functional resiliency need to be designed in. A secure chain of trust is required in the design, validation, manufacturing, and service process.

With network infrastructure being essential to delivering OTA updates and other data-driven services, securing the network is essential. Just as operating a network probably is not a core competency of most businesses, securing one generally is not a core competency of most automakers and service providers. With security being one of the most challenging aspects of managing a network, it is definitely a task best assigned to those that know it best. A cloud platform provider’s business relies on ensuring security and resiliency for its service so it has a vested interest in doing it right.

## CONCLUSIONS

As it has for more than a century, the car can exist in relative isolation. Aside from fuel, it does not need other vehicles in order to transport people or goods from one location to another. However, as more vehicles are put on the road, congestion inevitably increases and more time is wasted going nowhere. Time is a finite resource for people and it is not renewable.

Adding connectivity to the car and harnessing that connectivity to make life more convenient by providing a broad array of services can help to relieve the time crunch. Smarter navigation powered by data collected from connected cars can help release some of the lost time, while context-aware services can enable more effective use of travel time, reducing the number of trips required to complete a series of tasks.

In the future, the advent of the AV has the potential to provide huge societal benefits by reducing the number of vehicles needed to move more people and eliminating most of the crashes that kill more than 1 million people around the world every year. Energy and resource use can be slashed and time can be returned to travelers.

Achieving these laudable goals requires immense computing power and coordination. Cloud computing platforms are the connective tissue that ties all the disparate moving nodes and service providers together in a secure, scalable, and cost-effective manner.

Published 2Q 2018


1375 Walnut Street, Suite 100
Boulder, CO  80302 USA
Tel: +1.303.997.7609
http://www.navigantresearch.com



*Note: Editing of this report was closed on April 24, 2018.*