Introducing Ghost

The ideal AI software partner for L4 autonomy.

Self-Driving for Everyone

Ghost is an automotive technology company on a mission to enable attention-free self-driving for everyone. The Ghost Autonomy Engine, a software-based autonomy platform, is making it possible for leading automakers to deliver Level 4 autonomous driving in their next generation of consumer vehicles.

Built for safety, the Ghost Autonomy Engine leverages a new form of artificial intelligence capable of handling attention-free driving in its operating domain, so safe that it will no longer require driver monitoring or intervention. Built for scale, the Autonomy Engine runs on industry-standard sensors and chips, designed to be flexibly integrated into OEM vehicles that can reach hundreds of millions of drivers across the world in the next decade. Only by solving both fundamental safety *and* delivering scalable economics can self-driving technology gain the mass-market adoption necessary to make driving truly safer for the next generation.

Ghost is the ideal AI software partner to help make the promise of attention-free self-driving a reality for everyone and help re-imagine the driving experience in tomorrow's vehicles.

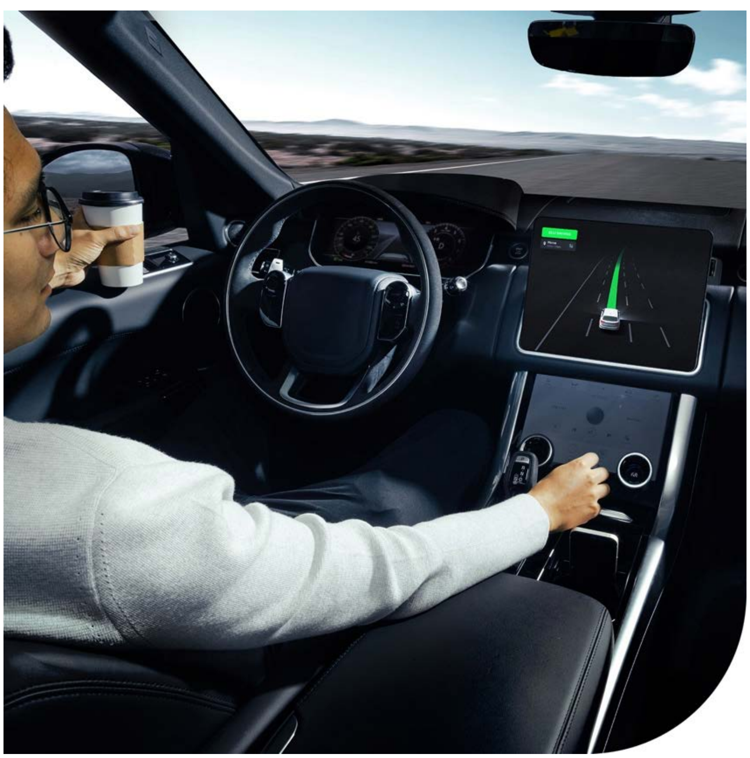

Driver Assistance Is Not Self-Driving

Despite 15 years and $100B+ of investment in self-driving and driver assistance technologies, drivers are still driving. No one has delivered technology safe enough for consumer vehicles to fully take responsibility for driving on the road. Even the most promising prototypes are far from scaling to lots of cars in lots of places.

Existing systems have long struggled to safely navigate the infinite long tail of objects, scenes, and driving events, behaving unpredictably, disengaging, or outright crashing when encountering new scenarios or unrecognized objects. As these "rare" events can happen on the road at any time, drivers must still pay attention, ready to intervene at a moment's notice. Ironically, this widespread reliance on human attention is destroying the actual consumer benefit that automating driving promised to deliver.

The Challenge: The Infinite Long Tail

Although many companies are building self-driving technology, most are pursuing minor variations of a similar robotics-inspired blueprint. Image-based object recognition married with HD map localization is at the crux of this approach – training neural networks with billions of images to detect and recognize the road actors, markers, and hazards in a scene to reason about their position and predict their behavior, and then matching that in real-time with a HD map of the world to drive from.

Unfortunately, this approach hasn't shown to deliver fundamental safety. While it is fairly straightforward to train a network on 99.999…% of what a car will encounter on the road, it is an inherently brittle approach since it is impossible to try and train a network on *everything* it might encounter on the road, in every rotation and orientation, in every color, in every lighting and weather condition. The objects that a vehicle might encounter are unbounded, no matter how large the training set, and rare objects are infrequently encountered in real-world training, forcing new attempts at simulation to try and fill training gaps. The number of training images or simulated miles, now often numbering in the billions, is often used as a proxy for completeness of the training set, but this is a weak proxy, as training on billions of common events doesn't make up for missing just a few rare ones. As it is simply not feasible to train on everything – the infinite "long tail" problem of current autonomous driving – image-based object recognition approaches have been incapable of safely taking full responsibility for driving without human supervision.

In response to these safety challenges, the complexity and cost of AI training and deployment has increased exponentially. To chase the infinite long tail of possible objects and scenarios, many vendors have turned to a strategy based upon more and more sensors, sometimes outfitting a vehicle with 30+ cameras, radar, and LiDAR sensing devices. While more sensory input can be good, it can also be challenging to fuse sensors when they don't agree, as neural networks output confidence probabilities of perception, not absolute perception. To determine positioning, these sensors must agree and then localize themselves with cm-accurate HD maps in real-time, which is a major source of errors and disengagements (or even crashes in extreme situations as in a recent well-publicized example), when algorithms, sensors and maps fail to align. In light of these multiple layers of challenges, some operators are turning to remote monitoring centers to both monitor and guide self-driving vehicles out of perception or logic jams. Even if this recipe eventually achieves an acceptable level of safety, it does not scale down to make autonomy possible for the mass-market passenger vehicle. The mainstream architectures of today are either incomplete systems that rely on humans to rapidly take over, or highly complex solutions with supercomputers in the trunk, only suitable for robo-taxis and luxury cars. Generalizing across the industry, neither is truly reliable and rely on some form of human intervention for when the neural networks encounter the long tail event – neither solves the problem of attention-free driving at scale. It's time for a different approach.

The Solution: Breakthrough AI For Safety and Scale

Ghost was founded in 2017 with the express idea providing a universal solution to the challenge of the infinite long tail in a completely different way, using new forms of artificial intelligence. Ghost's technology strategy revolves around:

1. **Universal detection with physics-based, motion neural networks.** Relying on object recognition makes safety probabilistic, and despite exponentially increasing data and compute resources for neural network training, it's impossible to train on everything that might exist in the real world. Ghost pioneered a universal set of physics-based video neural networks that can detect (without explicit identification) any obstacle, big and small, in any shape, trajectory, rotation, lighting condition, without prior experience/training of that particular object. These networks are based on the rules of physics, can be reliably trained with a fraction of the training data, can be mathematically proven to have a complete training set, do not require any human data labeling, and run on low power in-car compute.

2. **Sensors built for scale & AI: camera + radar.** Sensor stack design should optimize for complete near, far, and surround perception, performance across a multitude of lighting and weather conditions, redundancy, and economic viability for mass-market application. When considering all these factors, Ghost believes that camera + radar offers the best combination of capabilities. Both cameras and radar offer incredibly rich raw data streams that are ideal for interpretation by new AI methods. They are also highly complementary sensors, able to confirm one another's inputs (distance and velocity measurement, for example), and add signal in areas of strength where the other is limited (radar provides good perception in challenging weather and light, while cameras provide rich scene context, color information, and higher spatial resolution).

3. **Attention-free driving starts on the highway.** Drivers are typically most interested in reclaiming long stretches of driving, say their daily commute or leisure trips, but that benefit only comes with attention-free self-driving, often called "hands-off, eyes-off" autonomy. By initially focusing on the highway commute domain, Ghost can achieve attention-free driving first for the driving domain that has the most value to consumers and auto OEMs.

Others

- Object recognition, complex NN training
- HD maps required for scene understanding
- Solve all complex driving scenarios
- $15k+ sensor stack
- Only viable for robo-taxis and luxury cars
- Reduce safety outliers to acceptable level

Ghost

- Physics-based NNs, no object recognition required
- No HD maps are required, scene discovered in real-time
- Start with commuting (highway & leader) miles, expand
- General-purpose cameras, radar, processors
- Every car can be autonomous
- AI Algorithms that are mathematically proven to be trained on a complete data set

By re-thinking the problem, Ghost is creating a complete AI-driven L4 autonomy stack that runs on industry-standard cameras and radar and low-power and scalable compute to deliver attention-free driving.

KineticFlow™ – a Universal Neural Network for Perception

Given the inherent "long tail" challenges of traditional image-based object recognition neural networks, Ghost is solving the obstacle and scene recognition problem by starting with the fundamental laws of physics. The core insight is that there are a set of physics properties in the world that dictate how all objects behave, how they move, and how light bounces off them. Instead of trying to train a neural network on the infinite ways that the nearly infinite objects in the world *look*, Ghost trains networks on physics-based properties of objects — properties that are observable, finitely representable, and universal. A second insight was to focus on video instead of static images. Ultimately, driving is about understanding the relative motion of all actors in a scene, thus Ghost uses video to detect and model motion rather than just attempting to recognize objects.

The result is KineticFlow, Ghost's universal, physics-based neural network for vision-based perception that is designed to deliver many unique benefits:

- **Detection without recognition.** Ghost detects objects and road information that range in size from tires and car parts to signs, signals, and lane markers to bikes, cars, motorcycles, and large semi-trucks at long distances. Traditional approaches to object avoidance start with recognition and identification, some use lidar while others reason about the probable size of the object by identifying it and comparing to a database with object sizes in order to estimate distance. Ghost simply detects the presence of objects and uses stereo vision, object expansion and radar techniques to estimate distance, effectively eliminating edge cases and increasing accuracy.

- **Training without human data labeling.** In the endless race to train the long tail of image recognition networks, many vendors are moving to millions or billions of frames of labeled training data, and often employ thousands of humans to verify auto-labeling processes. Ghost's networks benefit from training data sets that are orders-of-magnitude smaller than those required by traditional image-based object recognition. And Ghost's training data is fully auto-labeled with the physics properties of every pixel in the entire scene auto-verified for accuracy. This approach enables Ghost to quickly and cost-effectively re-train the network, as well as train for new camera makes, resolutions, angles, or additional sensors. Our novel approach to ingesting data into our data center allows us to build neural network datasets that are mathematically validated to offer homogenous data distribution allows us to compute what data is missing, if any, for any training set.

- **Stereo and mono-vision techniques for near and far perception.** KineticFlow uses a combination of stereo vision and mono vision to determine object distance, size, velocity, time to impact, and motion paths. Stereo vision uses disparity maps between two camera images and can precisely locate objects even when neither one is in motion (i.e., stop-and-go traffic). A variety of mono-vision techniques analyze the expansion rates of the various objects in a scene and across time to determine motion and observe their physics properties to disambiguate moving obstacles from stationary obstacles and/or road features. These techniques are performed together in KineticFlow and are self-reinforcing, enabling higher accuracy than using stereo or mono algorithms alone.

- **Low CPU requirements for in-car execution.** The universal nature of KineticFlow not only means that Ghost neural networks are safer and easier to train, but also that they consume far less compute. Because high-compute autonomous features can drain EV battery range or affect gas consumption in an ice vehicle, a low-compute autonomous system encourages the market for EVs by marrying the fuel and cost-saving benefits of EVs with the time-saving benefits of attention-free driving.

- **Supplemented with additional overlay networks for Radar, Surround Vision, and Scene Understanding.** As KineticFlow universally detects safety-critical objects and distinguishes between other objects in the scene, parallel Ghost neural networks detect and independently measure the distance to objects with radar, as well as leverage a dedicated visual scene canonicalization neural network to discover all relevant lane and semantic information necessary for driving in the scene.

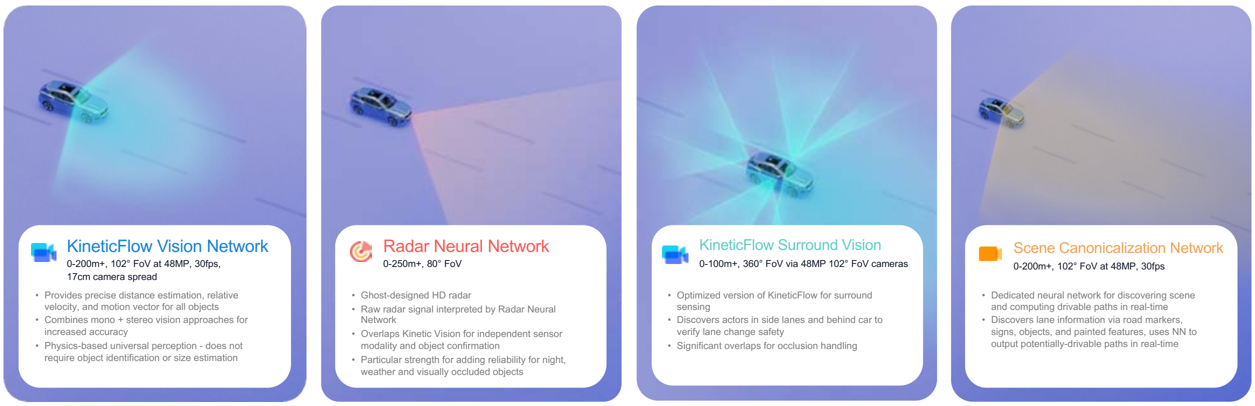

Ghost Perception Pipeline

Ghost implements a real-time perception pipeline that takes inputs from vision, radar, and other car sensors and processes this data through neural networks to create information for driving. Traditional driving neural networks function as depicted below/left, where obstacles are both detected and identified, shown by identification bounding boxes. Once identified, computer vision techniques are used to reason about size, location, and movement. In contrast, KineticFlow detects all the various planes in a given scene (i.e., relatively flat surfaces of objects) based upon their physics properties. Ghost detects the road and sky and then discovers and disambiguates lane markers from the road plane as the video network discovers and analyzes the motion of planes and groups those moving together into objects. The result is complete detection of every object, obstacle, and piece of road information in the scene at near and far distances, with accurate information of distance and motion characteristics, all discovered in just a few frames of video (~1/3 of a second for full scene and obstacle detection).

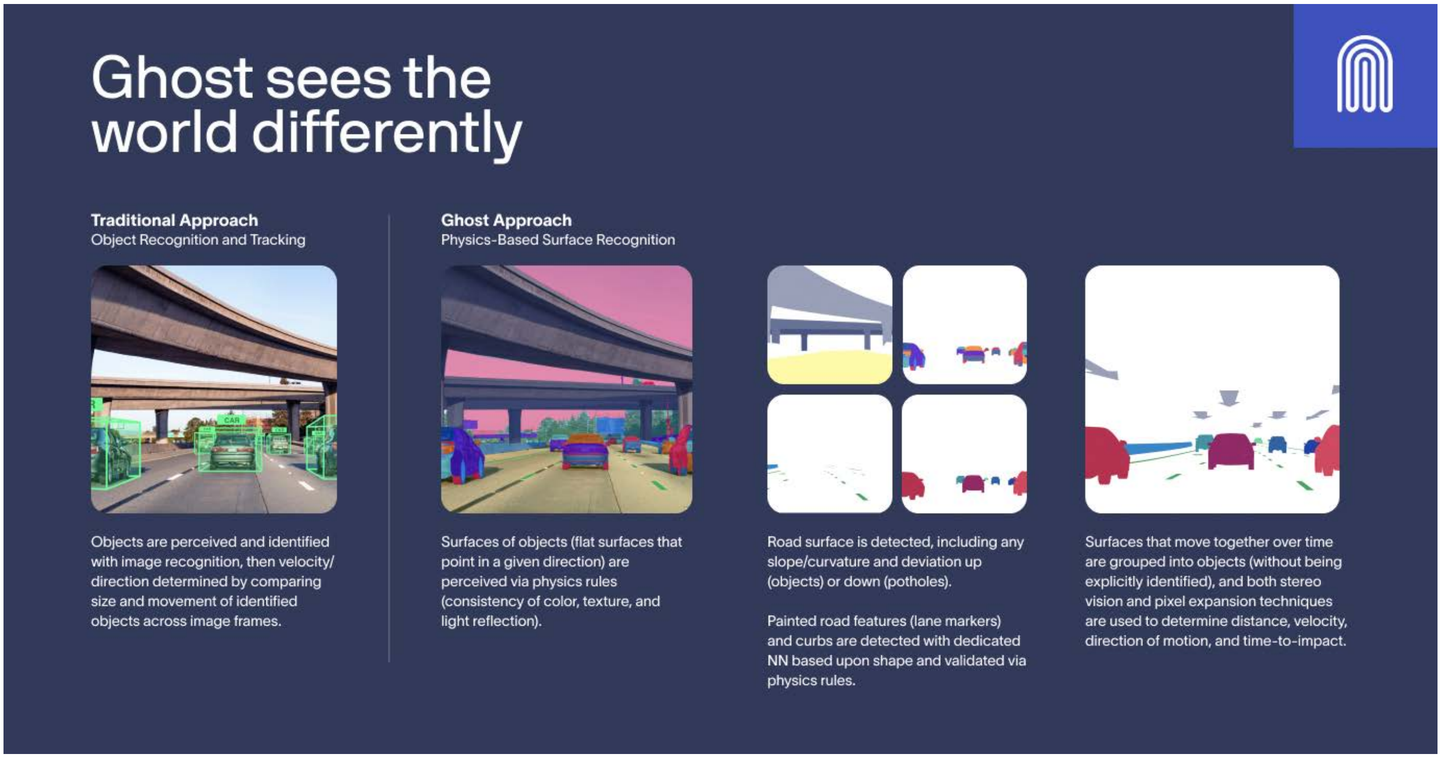

Sensors Built for Scale, AI, and Redundancy

It has been well-documented that the three major forms of autonomy sensors (camera, radar, and LiDAR) each has different capabilities, limitations, power profiles, and cost structures. And while it's tempting to think "the more number and types of sensors the better," the reality is that hardware and software complexity is sometimes the enemy of safety and reliability, and that system cost has a huge impact on the applicability of any technology – scaling down to every car has an important impact on the societal improvement to safety that can be delivered.

Ghost designed the Autonomy Engine's Reference Hardware sensor stack leveraging a combination of surround HD cameras and HD radar based upon the following design considerations:

- **Complete perception.** The combination of surround high resolution image sensors with high resolution radar is ideally suited to tackle the challenges of perception, from percieving at distances of hundreds of meters required for highway driving speeds, to detecting and determining the position of small objects, to operating in inclement weather. Each sensor modality reinforces and complements the other. They provide redundant detections and measurements and excel at different use cases, such as higher resolution and color perception (camera) and low light, challenging weather, and occlusions (radar).

- **Optimized for AI.** Ghost is introducing novel advances in AI processing of camera and radar data streams, both independently and together as a complete system. Modern high resolution image sensors and high resolution radar systems can provide rich, raw data feeds that have not been pre-processed, but instead provide the maximum information density for AI to interpret. These techniques are advancing at rapid rates, and in fact are enabling both camera images and radar to be used in parallel as “first class” sensors, rather than just using one to verify the primary findings of the other.

- **Redundant perception.** Many L2 systems today only leverage single sensors for various uses (a single forward LiDAR or long-distance camera, for example), and as such they simply require driver takeover in the case of sensor failure or occlusion. As a L4 system, Ghost must be designed end-to-end for redundancy across failures, which requires both sensor hardware redundancy (i.e., two front cameras with the same FoV) and sensor modality redundancy (i.e., both the front radar and camera can measure distance completely independently).

- **Scale and ongoing innovation.** Once the fundamental bar for safety has been met, scalability becomes a critical design criterion because cost directly impacts the breadth of adoption and thus the magnitude of societal benefit. Camera + radar is the right choice technically for safety to deliver complete and redundant perception, and it’s also the right choice for scale. Volume economics are driving massive annual improvements in camera and radar sensors and processors, propelled by investments in the cell phone industry and the global pace of AI research.

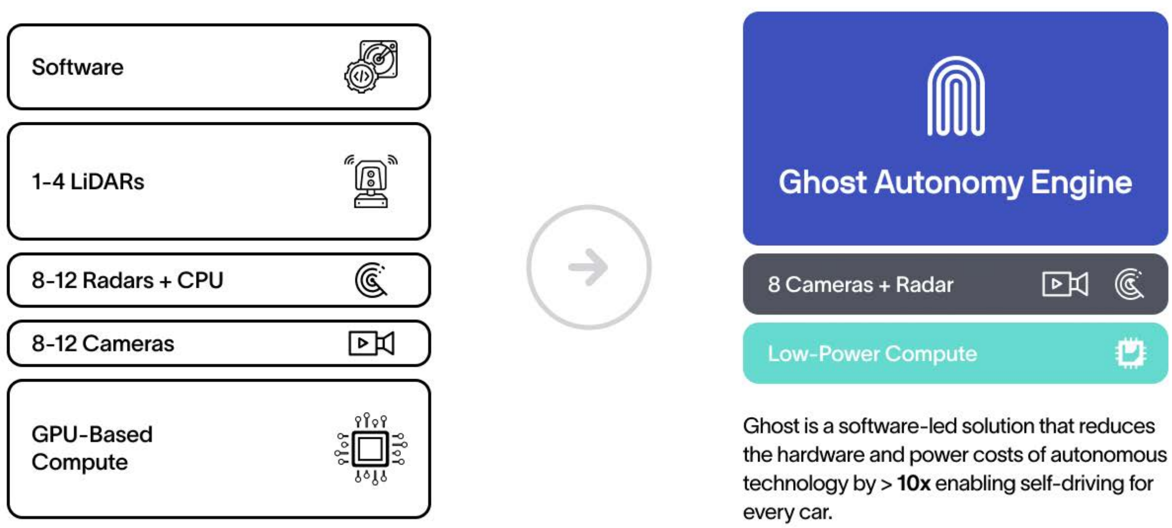

Ghost is a software-led solution that reduces the hardware and power costs of autonomous technology by > **10x** enabling self-driving for every car.

Finally, some may ask, “why not add LiDAR?” The Ghost Autonomy Engine is a customizable software platform that enables flexibility of sensor placement and configuration, and we specify a baseline reference sensor suite required for safe operation of the vehicle. We collaborate with our OEM partners to customize vehicle-specific packaging and sensor design from this hardware reference architecture. When evaluating today, LiDAR has several challenges which precluded us from specifying it as part of the reference suite, including high cost, high power consumption, mechanical fragility, poor inclement weather performance, difficulties in integrating attractively and aerodynamically in a car, and challenges to make redundant given its cost/power profile. Ghost successfully uses cameras, radar, and AI to achieve the perception performance required to drive safely without LiDAR, but we constantly review the capabilities and roadmaps of all sensor types for consideration in future generations.

Attention-Free Driving Starts on the Highway

Self-driving systems must be constrained during development and initial deployment to prove successful first uses as a basis for safe expansion. For some, this means urban driving at low speeds or in defined geo-fenced domains. Others focus on long-haul trucking. Yet others are growing through L2 assistance features with driver monitoring. Each of these approaches presents its own technology, safety, and market challenges.

Ghost is focused on building a self-driving system for mass-market vehicles, democratizing autonomy for those who need it most – consumers who drive long distances and would most benefit by recouping their time. Almost all (95%) of commuting time is spent on the highway or on suburban/rural feeder roads. To capture this benefit, Ghost has reinvented self-driving to be more fundamentally safe, to be more cost-effective, and to be more energy-efficient. We believe that not only does this represent the most attractive miles for consumers to automate, but it also serves to strengthen the market adoption of energy-efficient vehicles. Self-driving features are most likely to be expected as part of electric vehicles (EVs), and by coupling the EV+AV value proposition, OEMs will be able to accelerate consumer adoption of both technologies leading to both safety and environmental societal benefits.

The Ghost Autonomy Engine:
A Full Software Stack, Reference Hardware, Cloud, and AI Training Infrastructure

Ghost is building a complete stack to deliver L4 autonomy that OEMs can consume in part or full. The core Ghost IP is software, shipped as an integrated end-to-end reference platform:

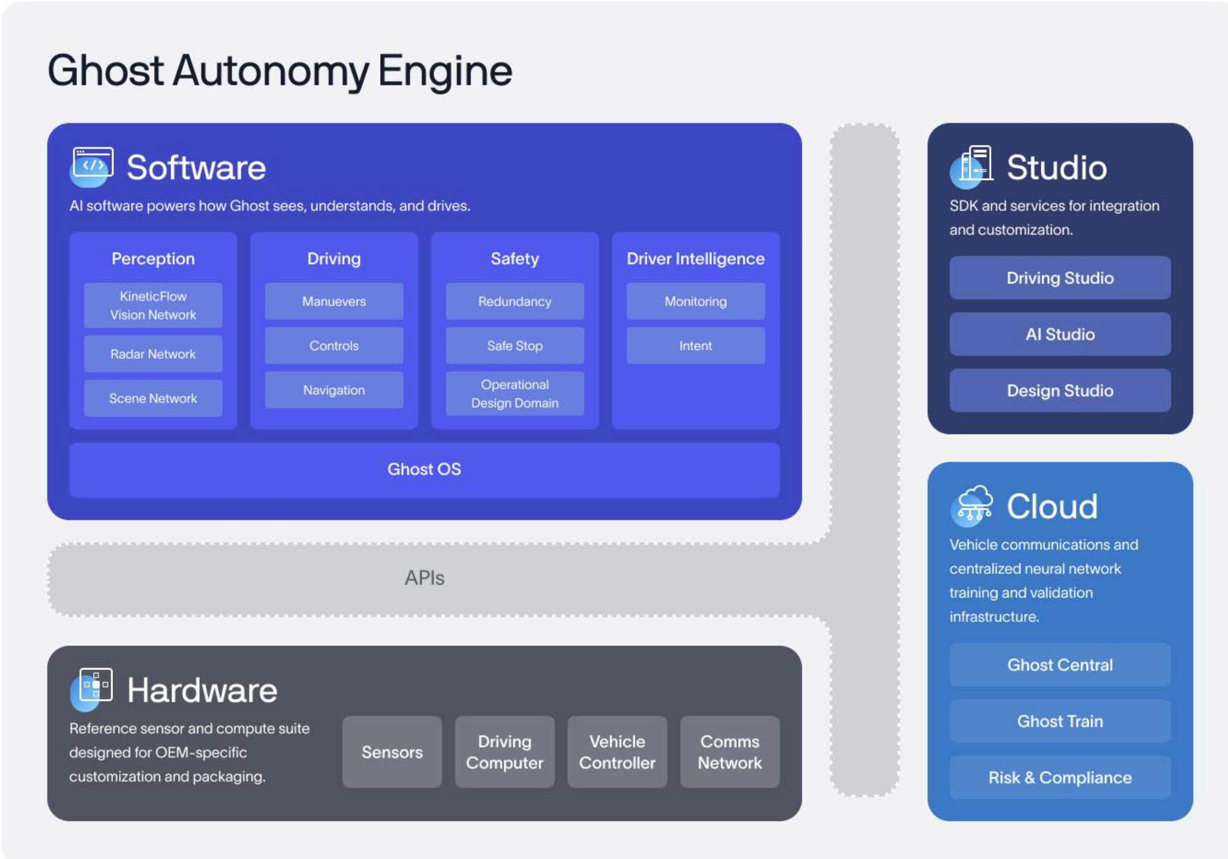

- **Ghost Software** is the autonomy brains of Ghost, designed to be embedded into and customized for OEM vehicles with the following major capabilities:

    - **Perception** is Ghost's combination of vision, radar, and scene understanding neural networks that take the inputs from 360° surround sensors and turn them into information for driving. Ghost Perception can detect stationary and moving objects and obstacles, determining their distance, location, direction of travel, and velocity, as well as road feature, lane, sign, and vehicle indicator cues.

    - **Driving** is Ghost's driving and routing program that takes discovered actors and drivable paths from Ghost Perception, and computes and executes driving maneuvers to optimize the route to a destination for safety and comfort.

    - **Safety** implements full-stack redundancy, ensuring that Ghost can continue driving even in the face of multiple failures. Safe Stop is capable of bringing the car to a complete and safe stop executing Ghost's Minimum Risk Condition, independent of failures of sensors, compute, or software.

    - **Driver Intelligence** is a suite of driver monitoring neural networks that use camera and control inputs to determine driver intent and to verify safe driver control during driving handovers.

- **Reference Hardware.** Ghost has designed a full auto-grade hardware stack of redundant surround-view cameras, radar, low-power and cost compute, vehicle integration, and communications for OTA updates and video clip collection for training data. This hardware platform can be purchased from Ghost, licensed as a reference design, and/or Ghost can work with Tier 1 suppliers to co-design appropriate hardware and integrate Ghost with your existing packaging.

- **Ghost Cloud** is Ghost's centralized cloud service that allows for bi-directional communication with Ghost-equipped vehicles via Ghost Connect. This enables OTA updates and uploads of video clips and logs from cars, and stores data in mechanisms compliant with global privacy regulations. Ghost Train is Ghost's proprietary environment for neural network training and training data auto-labeling. With Ghost Train, Ghost is able to quickly evolve Ghost's core neural networks with additional features, as well as re-train Ghost NNs to be compatible with specific sensors across multiple car models and generations. Partnering with Ghost doesn't deliver a one-time solution baked into silicon, but instead a constantly improving solution delivered as a subscription to consumers as software OTA.

- **Ghost Studio** is a combination of APIs, SDKs, and professional services to enable OEMs to integrate, customize, and extend the Ghost Autonomy Engine for every vehicle in their fleets. Driving Studio enables the development of custom vehicle-specific driving maneuvers, and interfaces with vehicle-specific controls. AI studio enables development of custom AI models to deliver AI perception-driven features. And Ghost's Design Studio works with OEMs to deeply embed the Ghost user experience within the OEM's HMI and car controls. With Ghost Studio, each OEM implementation of Ghost can be unique to your brand.

Your Ideal Autonomous Software Partner

Ghost's business is to help automakers bring autonomy to market and capitalize on the opportunities of the software-defined car. The Ghost Autonomy Engine platform, partnership model, and economic model are all designed to support OEMs:

- **Built for deep integration and customization.** The best Ghost experience is one that is fully integrated into a vehicle, with a single HMI, a seamlessly packaged suite of sensors, and a single mapping experience. Via Ghost Studio, Ghost is designed to be completely customized to your brand and enables customization across your models to suit the personality of each car. Ghost Studio has a full set of low-code/no-code APIs, making it simple for OEMs to build extensions and new features that drive further differentiation.

- **Maintaining the driver relationship.** While other self-driving companies aim to replace the driver, Ghost is *for* the driver, and appreciates that sometimes drivers want to drive. Ghost also recognizes that natural, fluid, and predictable interactions with the driver are essential for establishing trust. Ghost is designed to understand the intentions of the driver and take-over at will, a magical new experience that deepens the bond between driver and machine.

- **Targeting a full ADAS replacement.** Ghost's first version will deliver L4 highway driving, including the requisite safety ADAS features for the highway domain. Ghost's roadmap is to complete the full set of ADAS features for every driving domain, enabling a single stack for ADAS and L4 driving.

- **Verified software.** It is safety-critical to have a formal process for documenting and verifying capabilities, safety, and reliability at all layers. Ghost has built formal verification into both its software runtime and neural network training processes, enabling safety to be formally verified, trackable, and auditable over time.

- **Enabling new business models.** Ghost's cost-effective reference hardware platform enables it to be integrated into a wide range of cars. Software capabilities can then be made available via up-sell and/or subscription. The continuous improvement of self-driving capabilities and features gives consumers a clear reason to subscribe. With high customer demand for a self-driving subscription feature, self-driving can be a base "hook" to encourage customers into the subscription business model which can be broadened to other features and revenue streams over time.

Conclusion: Bring L4 Autonomy Across Your Fleet

Ghost has built an amazing team committed to helping you evolve the driving experience with autonomy. Our software engineers, neural network developers, mathematicians, user experience designers, hardware engineers, and program and process managers are driven by a passion to work in partnership with leading automakers to bend the safety curve of driving and take the integration of human and machine to the next level. Ghost's full stack of software, reference hardware, neural network training infrastructure, and customization and integration capabilities make us the ideal partner to differentiate your brands and compete for tomorrow's tech-savvy customer. We have a unique, AI-centric approach that is economically viable for widespread deployment, and our physics-based video neural networks are a new and innovative technological approach that delivers verifiable safety for L4 autonomy. We look forward to partnering together to re-imagine the driving experience.

Please contact us at partners@driveghost.com to explore working together.