SOFTWARE & ARCHITECTURE

Software-defined motion.

Inside DDX Dynamo - the software stack between simulation telemetry and the actuators, and why owning every stage of it is the real advantage.

> A motion platform is only as honest as the commands it receives. Speed, force, and travel all matter - but no actuator can recover a cue that arrived late, was filtered away, exaggerated, or built from incomplete data. The difference between motion that merely moves and motion that lets a driver read the car is decided in software.

VERSION 3.0 / SEPTEMBER 16, 2026

DRIVEN DYNAMIX

drivendynamix.com

A motion platform is only as honest as the commands it receives. Speed, force, and travel all matter - but no actuator can recover a cue that arrived late, was filtered away, exaggerated, or built from incomplete data. The difference between motion that merely moves and motion that lets a driver read the car is decided in software.

Driven DynamiX owns every performance-critical stage between the simulation and the actuators: telemetry reconstruction and predictive vehicle-state estimation, motion cueing, platform kinematics, and the real-time firmware inside our custom motion controller. No third-party motion engine sits in the critical path - and there is no layer we cannot measure, attribute, and improve.

Summary of results

- DDX Precision and DDX Prime operate within an approximately 2.5-3.5 ms telemetry-to-motion latency envelope, measured from raw simulation telemetry arriving at Dynamo to accelerometer-detected physical motion onset.
- Measured at the median in iRacing: input, physics, and telemetry take about 21 ms, and the DDX software, hardware, electrical, and mechanical stack adds roughly 3.5 ms on top of that, for a total of 24.5 ms from physical input to motion onset. Conventional interpolation smoothing would add approximately 17 ms more - 42 ms in total.
- DDX PreActiv delivers up to 2.7x the motion-cue fidelity of an optimized linear extrapolator, with a fraction of the harshness - fidelity being how closely the delivered cue matches the car's real motion in shape and timing, in an internal comparison on held-out captures.

From simulation to motion: the DDX software stack

1. Telemetry ingest and reconstruction

Every simulation reports the car differently: update rates from 60 Hz to 400 Hz; jitter, dropped frames, inconsistent channels, and the occasional invalid value. DDX Dynamo sanitizes and reconstructs the stream - cadence correction, discontinuity rejection, and a tuned blend of interpolation, extrapolation, and DDX PreActiv prediction - into one clean, consistent vehicle state delivered to the rest of the stack 1,000 times per second.

2. Motion cue construction

Surge, sway, heave, pitch, roll, and yaw are a poor cueing architecture. Most motion software constructs them as independent channels. Real vehicle dynamics are coupled: track orientation, body rotation, acceleration, and weight transfer happen together; changing the order of rotations and translations changes the result. Weight transfer through a banked turn is not "roll plus sway"; it is one 3D event.

Dynamo cues natively in full 3D. Slip angle, weight transfer, road detail, orientation, and grip are one coherent solution; complex motion is faithful and natural. Good cueing is the right information, at the right time, in the right relationship to every other cue - shaped to preserve useful detail within the platform's physical envelope.

3. Motion-controller command construction

Motion cues describe what the cockpit should do. Dynamo then solves the actual platform geometry to determine what every actuator must do to produce it - together, at 1 kHz - applying inverse kinematics, travel and joint limits, coordinated velocity and acceleration limits, and system-level safety constraints.

No fixed mechanical pivot defines how a DDX platform must feel. Its virtual pivot points and motion behavior live in the mathematics, where they can change by vehicle, by display configuration, by application - or by software update.

4. Real-time controller firmware

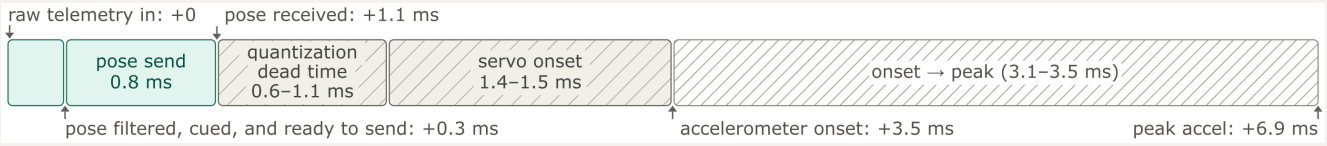
Motion commands travel over a tight 1 kHz link to our custom controller, where in-house firmware coordinates actuator motion and haptics at an approximately 56 kHz integration tick. Servo state and measured load are continuously monitored - managing motion limits, command continuity, faults, and emergency-stop behavior in hard real time.

DDX PreActiv: attacking the latency left inside the simulation

Making our own stack fast solves only the latency we control; the rest lives upstream, in the simulation itself.

Measured at the median in iRacing: input, physics, and telemetry take about 21 ms, and the DDX software / hardware / electrical / mechanical stack adds roughly 3.5 ms on top of that, for a total of 24.5 ms from physical input to motion onset. Run that same telemetry through conventional interpolation smoothing and you add another 17 ms - 42 ms in total.

> **The 3.5 ms is ours. The other 38.5 ms is the part every motion system inherits.**



Where the DDX-controlled share goes, stage by stage on DDX Precision: from raw telemetry arriving at Dynamo, through cue construction, transmission to the controller, and servo response, to accelerometer-detected motion onset at approximately 3.5 ms.

DDX PreActiv changes the problem from reaction to prediction: motion and force feedback respond to where the car is going, not where it was. The cues most critical to controlling a car at the limit reach net-zero effective latency - the whole 42 ms, gone.

DDX PreActiv combines recent telemetry, current vehicle state, driver inputs—including low-latency direct-input channels where supported—and learned vehicle physics to predict the vehicle state at multiple future time horizons. Dynamo uses those predictions to align motion cues in time with the driver's inputs, directly attacking simulator latency.

PreActiv predicts the whole car, not one channel: velocity and acceleration, orientation, per-wheel tire slip and suspension travel, steering-shaft torque, wheel-lock and wheelspin probabilities, and more. Each prediction carries calibrated uncertainty, so the system can shorten lead the moment confidence falls.

DDX PreActiv delivers up to 2.7x the motion-cue fidelity of an optimized linear extrapolator, with a fraction of the harshness. Motion-cue fidelity here is how closely the delivered cue matches the car's real motion in shape and timing, measured on held-out captures the predictor never trained on, and normalized so that an optimized linear extrapolator - measured by the same harness - scores 1.0. The multiplier grows with how far ahead the cue is placed; PreActiv measured above the extrapolator at every horizon we tested. And what it predicts, it renders faithfully: the sharpness you feel is the car's - neither exaggerated nor softened. DDX PreActiv is patent-pending, and it ships on every DDX machine.

One predicted vehicle state, more than one output

Two side-channel modules keep every output aligned with the same vehicle-state model:

- VR motion compensation removes cockpit movement from headset tracking using the platform's commanded motion - the one signal we know perfectly, because we authored it - so the virtual world holds still while the simulator moves beneath the driver.
- Simucube force-feedback augmentation. PreActiv outputs can be delivered as cues to supported steering and pedal hardware. Through SC-API, supported Simucube 3 wheelbases and Simucube ActivePedal hardware can provide net-zero effective-latency steering and pedal cues. The same path can add predicted steering torque, tire-slip, wheel-lockup, and wheelspin cues to augment force feedback in simulations with limited native steering feel, including many off-road and rally simulations. Motion and force feedback tell the same story, at the same moment.

We build hardware to embody the software

Our hardware is excellent because the software demands it. A fast, detailed signal exposes flex, backlash, weak servo control, irregular command cadence, and thin safety monitoring - flaws that slower, more heavily filtered systems can conceal. We build our own motion hardware and electronics wherever existing products are not good enough to express what Dynamo can do.

But the machine's character is not frozen into its mechanics. Hardware provides force, stiffness, bandwidth, and control authority; software decides how they are used. With no fixed pivots and no motion behavior baked into the linkage, the same platform gains new cueing strategies, new simulations, new applications, and better prediction - with only a software update.

> **Hardware is fixed the day it ships. Software compounds.**

The DDX advantage

One team owns the complete path from telemetry packet to actuator movement - so every delay can be attributed, every behavior tested, every layer evolved together. Our measured stack sits inside the approximately 2.5-3.5 ms telemetry-to-motion envelope after telemetry arrives. DDX PreActiv then attacks the simulator delay before it arrives - the part other motion systems simply inherit. And because the machine is software-defined, every improvement reaches simulators already in the field.

Today's consumer simulation titles model vehicle physics with extraordinary accuracy. What they lack is a motion system fast and faithful enough to turn that physics into training. DDX preserves rather than delays, filters, or exaggerates that information: low internal latency, 1 kHz cadence, fine-detail cueing, and PreActiv's time-alignment of the cues that matter most. Weight transfer, grip, slip, and road information arrive with the timing and detail required for real driver development.

The result is professional-level training capability - once reserved for Formula 1 teams, OEM programs, and massive simulator budgets - at a fraction of the cost. DDX does not replace every OEM engineering, HIL, or ADAS workflow; it makes professional driver-training cues accessible without lab-scale infrastructure.

Three companion reports are to follow. "Simulator Latency" measures each stage of the chain described here, from button contact to accelerometer-detected platform motion. "Motion Bandwidth" measures how much of a commanded cue survives the trip to the driver, and at which frequencies. "PreActiv User Study" measures what cue alignment and latency reduction are worth to drivers. To feel the difference now, book a demo at drivendynamix.com.

Study note: Latency figures are medians on one DDX Precision configuration, measured from physical button contact, or from raw telemetry arriving at Dynamo, to the first accelerometer-detected platform acceleration; the stage-by-stage protocol is published in the forthcoming "Simulator Latency" report. The 42 ms figure describes a conventional interpolated pipeline; DDX's measured median on that title is 24.5 ms, of which the DDX-controlled stack is roughly 3.5 ms. PreActiv fidelity and harshness are internal measurements from DDX's cue-metric harness on held-out captures, against an optimized linear extrapolator measured identically - no third-party benchmark exists for this metric class; harshness is delivered jerk relative to a real-vehicle reference, not a comfort rating. Net-zero effective latency means the predicted cue is time-aligned to its modeled vehicle event, not that transport, wheelbase, pedal, or human perception has zero latency; availability depends on supported hardware and configuration.

Source revision: ddx-kb@6a53dd86b8b52d2f1ecbd8292a3b3a7171d6c966
CLM-0026 / CLM-0028 / CLM-0029 / CLM-0031 / CLM-0032 / CLM-0034 / CLM-0062 / CLM-0063 / CLM-0065 / CLM-0126 / CLM-0127