

---

## From Simulation to Real-World Robotics on the Renesas RZ/V2H

---

*Gjergji Luarasi, MSc, Staff System Architect for Robotics Systems Solutions, UX Division, System Architecture, Renesas Electronics Corporation*

*Baoan Du, MSc, Principal Applications Engineer, UX Division, System Solution Technology, Renesas Electronics Corporation*

*Jiaru Han, MSc, Applications Engineer, UX Division, System Solution Technology, Renesas Electronics Corporation*



### Abstract

As quadruped robots and mobile manipulators transition from research to real industrial, inspection, and service applications, the need for robust locomotion control, real-time perception, and energy-efficient embedded compute has become increasingly critical. Achieving high-performance locomotion on embedded platforms remains a complex challenge due to dynamic environments, strict latency requirements, and the computational demands of reinforcement learning (RL), MPC-based whole-body control (WBC), and model-based estimation.

The RenRoboHIL (Renesas Robotics Hardware-in-the-Loop) pipeline addresses these challenges by providing a complete simulation-to-real (Sim2Real) evaluation and deployment environment for quadruped locomotion, built on the Renesas RZ/V2H RDK platform. The pipeline enables safe prototyping, high-fidelity simulation, real-time hardware verification, and final deployment on physical robots using the DRP AI3 heterogeneous accelerator for AI and control workloads.

This white paper introduces the architecture, workflow, evaluation setup, and preliminary outcomes of the RenRoboHIL pipeline, highlighting how the Renesas RZ/V2H MPU enables scalable, power-efficient, and real-time robotics applications at the edge.

### Purpose and Goals

The RenRoboHIL project is designed to provide a structured and quantitative evaluation framework for advanced robotic locomotion workloads on Renesas embedded compute platforms. Its primary objective is to bridge the gap between simulation-based algorithm development and validated real-world deployment on production-class MPUs and SoCs.

Specifically, the project aims to:

- Systematically evaluate representative locomotion control use cases, including classical model-based control, model predictive control (MPC), whole-body control (WBC), and learning-based approaches on current and next-generation Renesas MPU and SoC platforms.
- Assess the feasibility of the RZ/V2H platform as a real-time control backbone for legged robots, with particular focus on quadruped locomotion control by reinforcement-learning-based and neural control policies operating under strict latency and determinism constraints.
- Quantify compute performance and partitioning requirements across CPU and DRP-AI3 accelerators and compare heterogeneous execution against traditional CPU-only implementations for high-frequency robotics control loops.
- Evaluate the suitability of DRP-based AI acceleration for emerging robotic workloads like reinforcement learning inference.
- Establish a reproducible Sim2Real workflow using a unified messaging and software interface, enabling the same locomotion stack to be executed consistently across simulation environments and physical robotic systems.

Through these goals, RenRoboHIL serves as a reference framework for validating real-time performance, power efficiency, and scalability of Renesas embedded platforms for next-generation robotics applications.

### Evaluation Setup

The evaluation environment for the RenRoboHIL pipeline is designed to closely replicate real-world robotic locomotion workloads while enabling safe, repeatable, and instrumented testing on embedded hardware. The setup combines a high-fidelity physics-based simulation with a real-time embedded compute platform, connected through a unified communication interface that remains consistent across simulation and real deployment phases.

### Hardware Platform

At the core of the evaluation setup is the Renesas RZ/V2H Robotics Development Kit (RDK), shown in Figure 1, which serves as the primary real-time computation and control platform. The RZ/V2H MPU integrates quad-core Arm® Cortex®-A55 CPUs, dual-core Arm® Cortex®-R8 CPUs, and one Arm® Cortex®-M33, along with the heterogeneous DRP-AI3 accelerator, enabling concurrent execution of real-

## From Simulation to Real-World Robotics on the Renesas RZ/V2H


---

time control logic, estimation, and AI-based locomotion policies. The RDK provides Ethernet connectivity, low-level peripheral interfaces, and sufficient memory bandwidth to support high-frequency robotics workloads under realistic timing constraints.

The RZ/V2H executes the full locomotion control stack, including:

- Policy inference (reinforcement learning or neural control models)
- State estimation and filtering
- Whole-body or joint-space control
- Safety supervision, command arbitration, and saturation handling

By running the same software stack throughout simulation and deployment phases, the platform enables direct evaluation of real-time performance, computational load, memory utilization, and thermal behavior under representative operating conditions.



**Figure 1:** On the left is the RZ/V2H RDK used as the target compute platform for the RenRoboHIL pipeline. On the right is the Unitree Go2 quadruped robot used as the reference robotic platform.

To validate the Sim2Real pipeline on physical hardware, we use the Unitree Go2 quadruped (Figure 1) as the reference robot. Go2 is a commercially available 12-degree-of-freedom quadruped platform with torque-controlled actuators, onboard sensing, and a low-level interface designed for external real-time control. The RZ/V2H communicates with Go2 over Ethernet using DDS messaging via the Unitree SDK2 / `unitree_go` low-level interface, which provides joint-level commands and state feedback. The evaluation focuses on low-level control mode, in which the built-in high-level motion services on the robot are disabled.

Together, the combination of the RZ/V2H RDK and the Unitree Go2 quadruped forms a representative and production-relevant hardware setup for validating real-time embedded locomotion pipelines, while maintaining a clean separation between robot dynamics, control computation, and AI acceleration.


### Simulation Environment

A PC-hosted MuJoCo-based simulator is used to model a quadruped robot with realistic dynamics, contact interactions, and sensor feedback. MuJoCo provides accurate rigid-body dynamics and contact modeling, making it well-suited for validating legged locomotion algorithms before physical deployment.

In the baseline configuration:

- The simulator runs on a standard x86 PC.
- The virtual robot model includes articulated legs, joint actuators, contact sensors, and an IMU.
- Terrain conditions can be adjusted to evaluate locomotion robustness under varying scenarios.

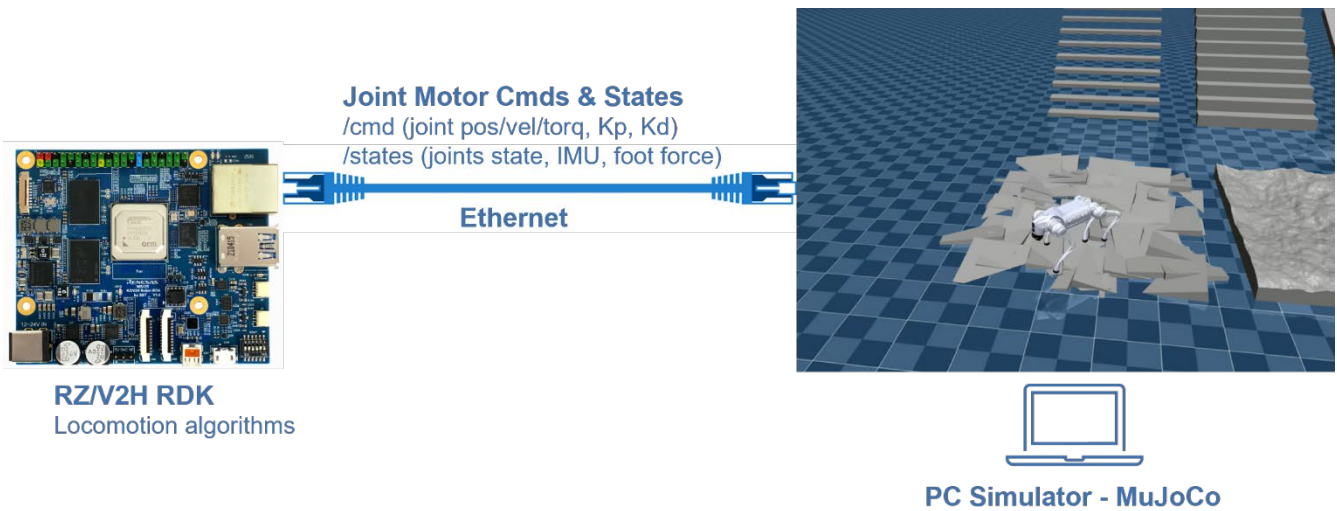
Figure 2 shows a screenshot of the MuJoCo-based simulation environment, where a Go2 robot is instantiated in proximity to various terrains that can be used for testing the locomotion algorithms.



**Figure 2:** Screenshot of the MuJoCo simulation environment highlighting the Go2 reference robotic platform and the different terrains used for evaluation.

Crucially, the simulator does not execute the locomotion controller internally. Instead, all control logic runs on the RZ/V2H RDK, ensuring that compute load, execution latency, and control loop timing reflect embedded deployment conditions even during simulation.

### Communication and Data Exchange



**Figure 3:** Schematic of the communication interface between the compute system (RZ/V2H RDK) and the simulation environment.

A Data Distribution Service (DDS) -based messaging interface over Ethernet is used to connect the RZ/V2H RDK with the PC simulator or the reference robot. This interface forms the backbone of the Sim2Real workflow and enforces a clean separation between control computation and simulated or real robotic platforms.

As shown in Figure 3, the messaging layer supports deterministic, low-latency streaming of joint commands and sensor feedback, including:

- Command signals (RZ/V2H → Simulator/real robot):
  - Target joint positions or velocities
  - Desired motor torques
  - Joint-level control gains (Kp, Kd)
- State feedback (Simulator/real robot → RZ/V2H):
  - Joint positions and velocities
  - Estimated motor torques
  - IMU data (linear acceleration and angular velocity)
  - Foot contact states and foot force estimates

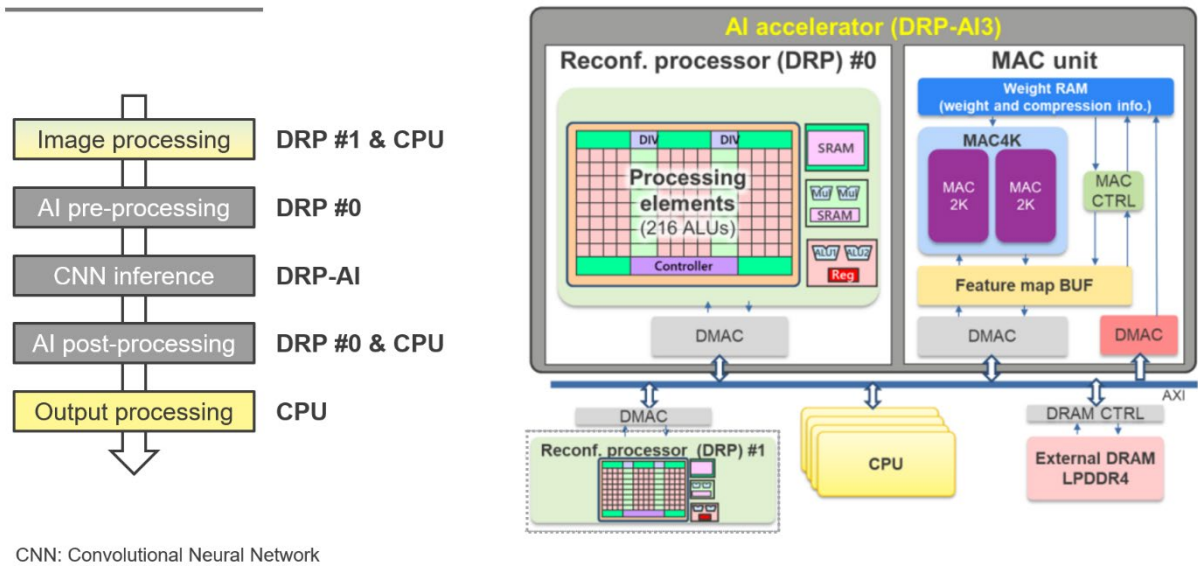
By standardizing the data interface, the same locomotion code can be executed unmodified when transitioning from simulation to real hardware. This significantly reduces integration effort and minimizes discrepancies between simulated and physical behavior.

### DRP AI3 Features for Robotics

The RZ/V2H integrates the next-generation DRP AI3 accelerator, which provides high-throughput INT8 inference and support for pruned lightweight AI models, Dynamically Reconfigurable Processor (DRP)

## From Simulation to Real-World Robotics on the Renesas RZ/V2H

blocks for low-latency non-AI control tasks, and a heterogeneous pipeline enabling parallel CPU/DRP/AI co-execution as shown in Figure 4. Robotics workloads such as locomotion often require policy inference (RL or neural MPC), state estimation from sensors, local mapping or terrain detection, planning and whole-body control loops at 200-1000 Hz. The DRP AI3 architecture accelerates these by offloading appropriate components to AI or DRP engines while keeping high flexibility logic on CPU cores.



**Figure 4:** Block diagram schematic of DRP-AI3 highlighting the processing steps for AI pipelines.

## RenRoboHIL Sim2Real Pipeline

The RenRoboHIL pipeline (Figure 5) is executed in three stages, ensuring safe and progressive deployment.




Figure 5: Schematic of RenRoboHIL Sim2Real Pipeline

### Step 1 - Simulation (PC)

In the first stage, the locomotion algorithm executes on the RZ/V2H RDK while the robot dynamics are emulated on a host PC using MuJoCo. The simulator acts as a stand-in for the physical platform by consuming command messages and returning sensor-equivalent feedback, allowing the embedded controller to be exercised under realistic timing conditions.

Over the Ethernet messaging interface, MuJoCo receives the control stream on the /cmd topic (joint targets together with per-joint gains Kp and Kd) and publishes the corresponding state stream on the /state topic (joint states, IMU measurements, and foot-force/contact signals). This configuration provides a fully safe environment with a fast iteration cycle, enabling tuning and debugging of locomotion behaviors without physical risk.

### Step 2 - Real Robot, Phase 1 (Bridged Signals)

In the second stage, the controller is exercised against a physical quadruped (for example, a Unitree-class (Go2) robot) while retaining a bridged communication path. Commands produced on the RZ/V2H are forwarded through a bridged DDS interface - RZ/V2H → RK3588 → robot - so that the robot receives the

same logical command topics used in simulation, but via an intermediary compute node that mirrors common integration architectures.

Real sensor data is streamed back into the RZ/V2H in closed loop for evaluation, making it possible to validate end-to-end timing, actuator limits (e.g., torque saturation), sensor noise sensitivity, and the fidelity of foot-contact/force signals under real interactions. This phase is used to confirm hardware compatibility and to ensure that the signal path and interface behavior remain realistic before moving to direct peripheral integration.

### **Step 3 - Real Robot, Phase 2 (Direct Signals)**

In the third stage, the RZ/V2H transitions to direct, low-level integration with the robot by communicating over peripheral interfaces such as RS485, SPI, and I<sup>2</sup>C. Joint commands and state feedback are exchanged without an intermediate bridge, and IMU inputs and motor-driver interactions are integrated directly into the embedded control stack. This enables full embedded locomotion deployment with the RZ/V2H operating as the primary real-time controller.

### **Heterogeneous Architecture Benefits for Locomotion**

#### **1. Low-jitter real-time control via DRP**

The DRP (Dynamically Reconfigurable Processor) is well-suited to deterministic, low-jitter execution of time-critical kernels that sit directly in the locomotion feedback loop. Offloading repeatable computations such as estimation filters, kinematics, and terrain-processing primitives to the DRP helps stabilize control-loop timing and reduces latency variance, which is essential for consistent foot placement, contact handling, and high-frequency stabilization.

#### **2. High-throughput AI via DRP-AI3**

For learning-based components, the DRP-AI3 accelerator provides high-throughput and energy-efficient inference, enabling modern locomotion stacks to run neural policies and perception-adjacent models without overloading the general-purpose CPUs. Typical candidates include reinforcement-learning (RL) policy inference, terrain classification, and learned motion-planning modules, where consistent inference latency and low power directly translate into higher control rates and better battery and thermal margins on mobile platforms.

#### **3. Flexible CPU control loops**

The CPU cores remain the most effective place for orchestration and control logic that benefits from software flexibility, branching, and rapid iteration. In practice, this includes contact-state switching and mode management, higher-level control logic and arbitration, and safety layers such as limit enforcement, fault handling, and supervisory checks that must integrate across multiple subsystems.

# Evaluation Results

### Sim2Real Stability

Initially, basic balance controllers for force control were evaluated for their performance on the RZ/V2H's CPU. They use PD control on body pose errors to compute and map forces to the feet. The robot actions used for testing include simple gaits, walking, trotting, standing up, lying down, etc. The CPU workload showed acceptable values around 40% at the control loop rate up to 200Hz. However, these models are not enough for most modern robots.

Secondly, an OCS2-based NMPC-WBC controller was evaluated. It is a traditional locomotion control method, widely used by research platforms. Dynamic gaits (trotting, bounding, bipedal gait) showed good stability. The downside of such a method is that it is computationally heavy, and its mathematical structure limits GPU acceleration. As a result, it led to full CPU utilization at control loop rate 25Hz, although further optimization is possible.

Lastly, neural locomotion models trained in MuJoCo demonstrated stable Sim2Real transfer when deployed on the RZ/V2H computing stack. Across the evaluated scenarios, end-to-end latency remained below 10 ms, satisfying the timing requirements for high-frequency locomotion control at 15% CPU utilization after optimization. Higher control frequencies (up to 8×) can be achieved if the neural network policy is deployed on the DRP-AI3 accelerator.

## Summary

The RenRoboHIL Sim2Real pipeline provides a unified, safe, and high-performance workflow for evaluating quadruped locomotion on Renesas embedded platforms. Powered by the RZ/V2H RDK and its DRP-AI3 heterogeneous architecture, the pipeline:

- Enables seamless transition from simulation to real hardware
- Supports modern RL and AI-driven locomotion methods
- Ensures deterministic real-time control at low power
- Provides a scalable approach for robotics developers targeting edge compute platforms

Renesas will continue to expand capabilities, tooling, and performance in upcoming RZ/V series MPUs to support increasingly sophisticated robotics applications.

## References

- [RZ/V2H](#): Quad-core Vision AI MPU Delivering High-Efficiency AI Inference and High-Responsive Real-time Control
- [DRP-AI](#): Renesas' proprietary AI accelerator that combines high AI inference performance with low power consumption

RENESAS ELECTRONICS CORPORATION AND ITS SUBSIDIARIES (“RENESAS”) PROVIDES TECHNICAL SPECIFICATIONS AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for developers skilled in the art designing with Renesas products. You are solely responsible for (1) selecting the appropriate products for your application, (2) designing, validating, and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. Renesas grants you permission to use these resources only for development of an application that uses Renesas products. Other reproduction or use of these resources is strictly prohibited. No license is granted to any other Renesas intellectual property or to any third party intellectual property. Renesas disclaims responsibility for, and you will fully indemnify Renesas and its representatives against, any claims, damages, costs, losses, or liabilities arising out of your use of these resources. Renesas' products are provided only subject to Renesas' Terms and Conditions of Sale or other applicable terms agreed to in writing. No use of any Renesas resources expands or otherwise alters any applicable warranties or warranty disclaimers for these products.

(Rev. 1.0 Mar 2020)

## Corporate Headquarters

TOYOSU FORESIA, 3-2-24 Toyosu, Koto-ku, Tokyo 135-0061,  
Japan  
<https://www.renesas.com>

## Trademarks

Renesas and the Renesas logo are trademarks of Renesas Electronics Corporation. All trademarks and registered trademarks are the property of their respective owners.

## Contact Information

For further information on a product, technology, the most up-to-date version of a document, or your nearest sales office, please visit:  
<https://www.renesas.com/contact-us>

© 2026 Renesas Electronics Corporation. All rights reserved.  
Doc Number: R01WP0033EU0100