ASX ANNOUNCEMENT

4 November 2026

Independent expert takes Chipforge from written requirements to working FPGA hardware

Highlights

- **An independent expert used Chipforge platform to take written requirements through to working FPGA hardware, without assistance from the Chipforge team and without manual edits to the generated code.**
- **Internal platform testing on NVIDIA's public CVDP benchmark scored 14 to 28 percentage points higher than published results for GPT-5, Claude 4 Sonnet and the MAGE research agent on the same tasks.**
- **Chipforge is in pre-production release and delivering supported projects with its engineers. The broader product release, Chipforge Studio, is targeted for Q1 2027.**
- **AMD Vivado and Altera Quartus workflows are complete, giving customers a choice of FPGA tool chains.**

Pathkey.AI Ltd (ASX: PKY) (Pathkey or the Company) confirms Chipforge validation by an independent expert. The Chipforge platform was used to take written design requirements through to working FPGA hardware, without assistance from the Chipforge team or manual changes to the generated code. Designing custom hardware normally passes through several specialist steps, each with its own tools and people. Chipforge brings those steps into one platform. Chipforge is in pre-production release and is delivering supported projects with its engineers while the broader product release is targeted for Q1 2027.

Chipforge connects the work involved in designing a semiconductor chip, testing it, fixing faults and preparing it to run on a chip. An engineer sets the requirements and reviews the results. The platform generates the chip code, tests its behaviour and revises the design using test results and the engineer's instructions. It then works with established chip-design tools to prepare the design for a programmable chip.

The demonstrated process uses an FPGA, short for field-programmable gate array. FPGA is a chip full of programmable logic, already connected through a network of wiring. What it doesn't have yet is a job. Load a design and it sets what that logic does and how it's connected, so it becomes exactly the chip you need. Wipe it and load another. Engineers can load a design, test it on the chip and change it without having to manufacture a new chip. Programmable chips can be used in finished products as well as to test designs before committing to custom silicon.

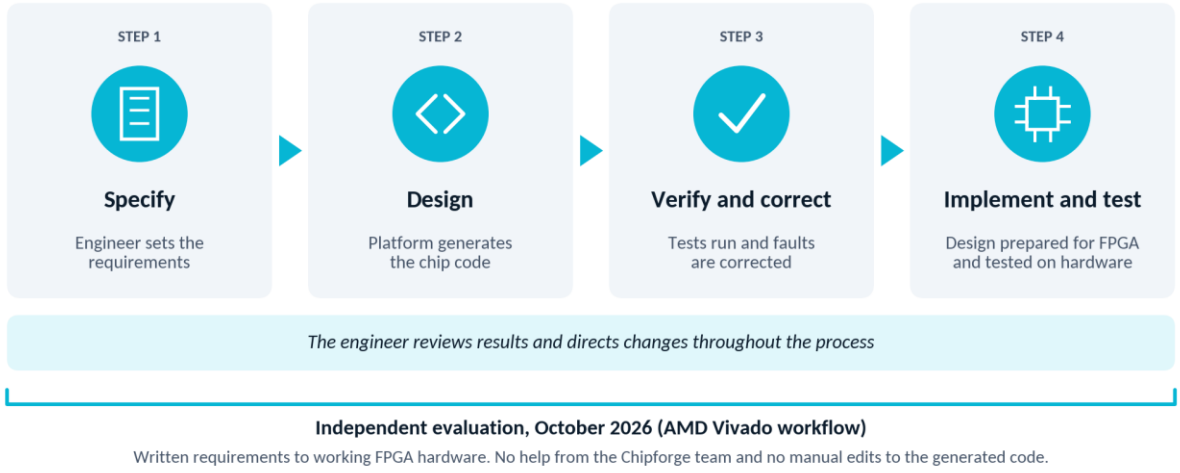


Figure 1. Chipforge FPGA design workflow and the scope of the independent evaluation. Source: Chipforge.

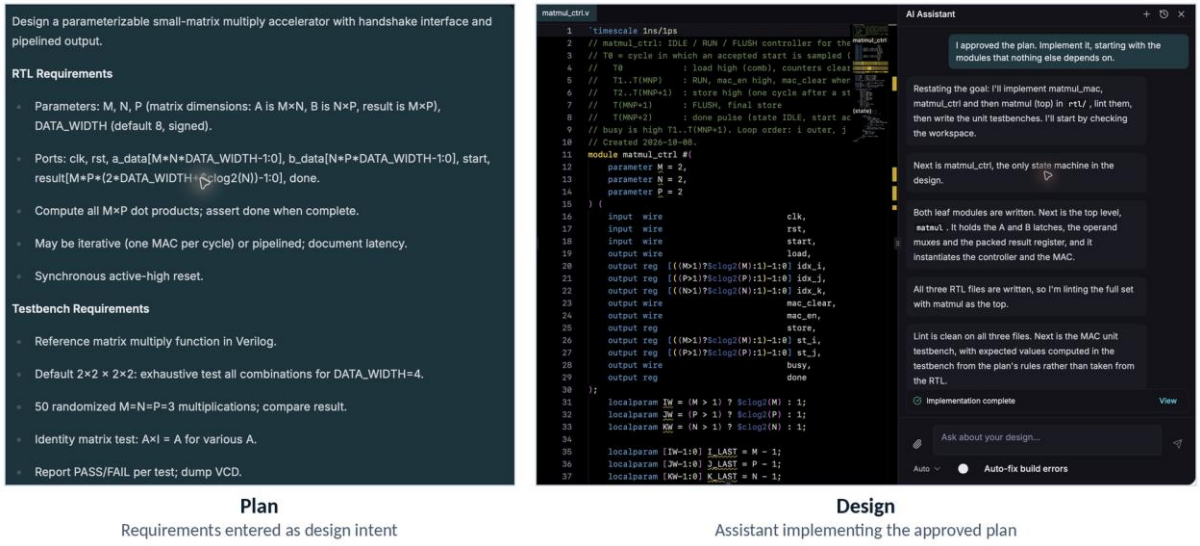


Figure 2. Chipforge Studio plan and design views, from an illustrative matrix multiplier example.

Independent hardware evaluation

The Company engaged Dr Smitha Kavallur Pisharath Gopi, a Senior Lecturer in the College of Computing and Data Science (CCDS) at Nanyang Technological University (NTU) in Singapore, to independently verify the capability of the Chipforge platform to deliver working FPGA designs from written specifications. Dr Gopi used the platform to take agreed hardware designs from written specification to working operation on an FPGA. Dr Gopi entered the design problems as design intent. Chipforge planned the design, generated the chip code, built the tests, ran the simulations, revised the design where tests failed and produced the file that programs the chip. The expert deployed it to physical hardware, and it ran as specified.

The Chipforge team provided no assistance. No manual edits were made to the generated code. Where Dr Gopi wanted changes, she described them to the platform and the platform implemented the changes.

The evaluation covered the AMD Vivado workflow and confirmed delivery of a functioning platform utilising the Chipforge intellectual property and that the product delivers core features necessary for end-users to experience the product's primary value proposition, as demonstrated through real-world use, testing and validation in a working environment. It assessed the agreed designs, rather than certifying the platform for every commercial application. Altera Quartus integration is also complete but was outside the scope of this evaluation.

Dr Gopi holds a PhD from NTU. Her independent evaluation report, dated 23 October 2026, records successful implementation of a Digital Frequency Monitor and AES-128 encryption over UART on a Digilent Nexys A7-100T FPGA board.

Recent platform development

Since the acquisition in June 2026, Chipforge has connected the FPGA design workflow and built the systems needed to support customer projects.

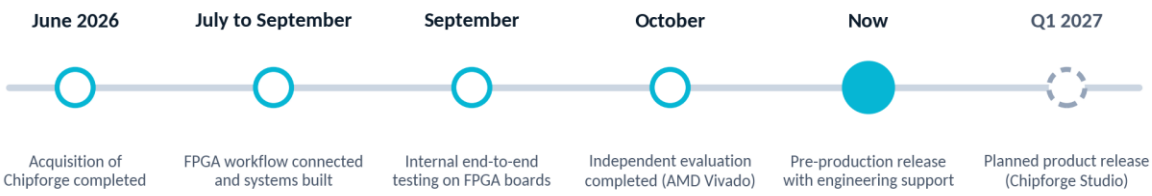


Figure 3. Chipforge progress since the acquisition..

The AMD Vivado workflow now runs from design intent through hardware code generation, testing and verification to the file that programs the FPGA. Altera Quartus integration is also complete.

Chipforge has programmed and tested designs on FPGA boards, covering individual chip components and more complex combined designs. This internal work is separate from the independent evaluation.

The team has also built software for planning designs and managing the platform. The production environment has software release controls, performance monitoring and alerts to help the team support customer work.

Chip design CVDP benchmark results

Chipforge tested the platform internally on 262 problems across four categories in CVDP (Comprehensive Verilog Design Problems), NVIDIA's public chip-design benchmark. The tasks involve creating a chip design from written requirements, completing unfinished code, modifying an existing design and fixing bugs.

The platform generates or edits the Verilog chip code and submits it to an open-source simulator. The simulator runs the benchmark's hidden tests to check whether the design behaves as required. A problem passes only if every test passes. There is no partial credit, human grading or AI self-assessment. Chipforge's results are first-pass. The published results count a problem as solved if any of five attempts passes.

Figure 4 sets the Chipforge results beside published results for other systems on the same tasks.

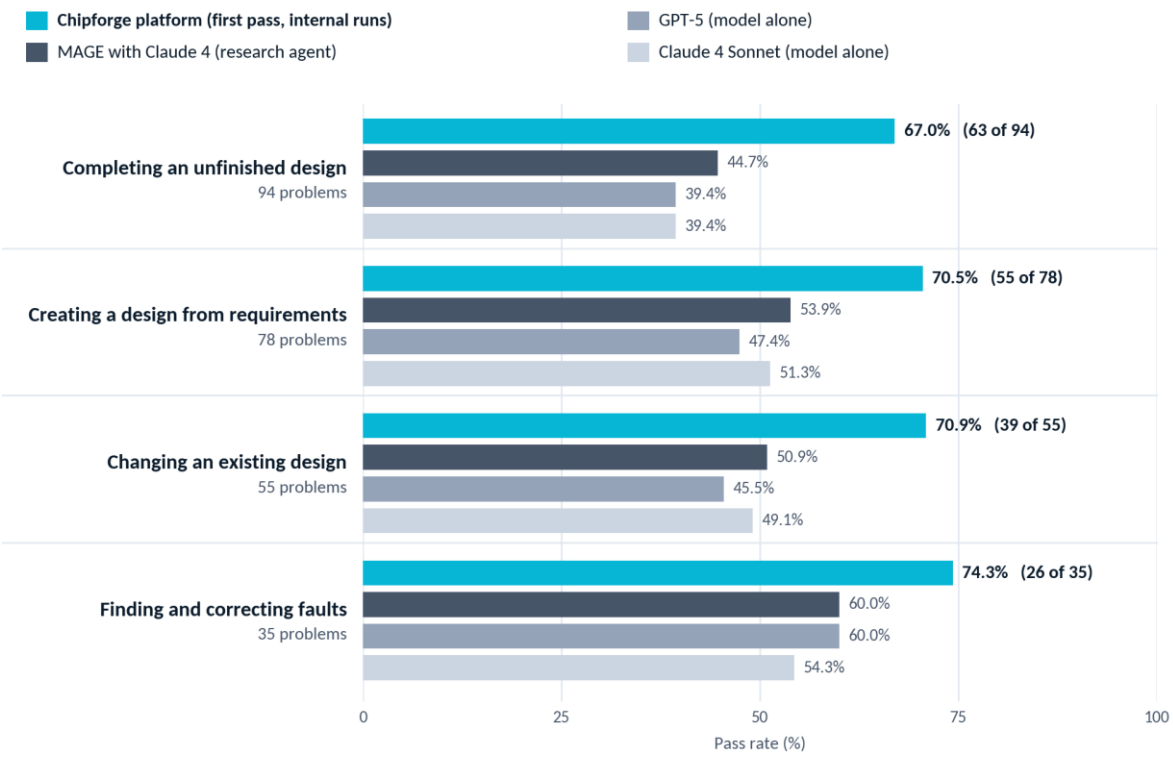


Figure 4. CVDP pass rate by task. Published results for other systems are from NVIDIA ACE-RTL research (arXiv:2602.10218, Table 1) and count a problem as solved if any of five runs passes.
Source: Chipforge internal platform runs, with task populations corresponding to CVDP v1.0.2. Benchmark background: Pinckney et al., Comprehensive Verilog Design Problems (arXiv:2506.14074). Percentages are reproduced as reported by Chipforge.

A general-purpose AI model writes code, but a mistake in chip code shows up only when the code is simulated or run on hardware. Chipforge adds the engineering steps around the model. It tests the code in a simulator, reads the failures, revises the design and tests again, then prepares the result for the chip.

The debugging result represents 26 of 35 problems solved. The tests show how the platform handles common design and correction tasks. Customer projects will establish what this means for engineering time, cost and delivery.

These are internal simulation results and have not been independently verified. They do not measure complete customer projects or operation on physical hardware. NVIDIA developed the benchmark but has not validated or endorsed Chipforge.

Customer projects and revenue

Chipforge is targeting work that requires specialist engineers to move between design, testing and debugging tools. Keeping those steps within the platform is intended to reduce handovers and rework. Customers can start with a specific design or fault, using Chipforge alongside their existing tools.

A traffic-light controller is a simple example. The customer specifies the sequence and timing, and Chipforge develops and tests the chip design before an engineer reviews it for use on a programmable chip. This is an illustration, not a customer project.

Initial engagements are fixed-price projects, with the scope, success criteria and engineering support agreed upfront. Successful delivery can lead to further projects, usage fees and software licences, with recurring revenue dependant on customers continuing to use the platform.



Figure 5. How an initial fixed-price project can lead to further work.

Chipforge's first commercial contract is a ten-week proof-of-value program, this first phase has been delivered. The next stage depends on the client being ready to integrate the work into its mechanical solution. This customer project is separate from the independent evaluation.

Pathkey Group CEO Andrew Farnsworth said:

"I'm really excited by the results of this independent review. Dr Gopi took a set of design problems and entered them as design intent. Through direct prompting, the platform guided the expert through planning and specification, then code development and the verification cycle, including creating the test benches and running the simulations. From there it moved through synthesis, generated the bitstream, deployed it to the physical hardware and ran the tests successfully. All of that happened with no intervention or support from the Chipforge team, which is very encouraging to see.

We are now working with customers, and we are no longer just talking about what the platform will do. We can show them what it does, which is a huge step for us.

We will be turning that capability into paid projects across Asia Pacific, including Singapore, Malaysia, Vietnam, Japan and Australia, wherever the opportunities lie."

Product development priorities



Figure 6. Current capability and development priorities.

The team is completing the initial release of Chipforge Studio, the software engineers will use to access the platform. Work includes testing, customer access, documentation and making it easier to get started. Supported FPGA projects can proceed while this work continues.

Current FPGA tool support includes AMD Vivado and Altera Quartus, with Lattice and Microchip support next. Wider tool support would allow Chipforge to address more customer hardware requirements. Vendor references do not imply partnership or certification.

Extending the platform

The longer-term roadmap extends to custom silicon, or application-specific integrated circuits (ASICs). These chips are manufactured for a specific purpose and require further engineering

before production. Chipforge can already support parts of ASIC design and digital verification. Physical implementation and the handoff to manufacturing are still in development.

Chipforge is also fine tuning and training an AI model for chip design and verification, and developing the ability to run within the customer's own computing environment. This would let customers keep proprietary designs and engineering data within their own infrastructure, without their design needing to touch an outside AI model. These capabilities remain in development and are separate from the current cloud FPGA offering.

Commercial progress will be measured through further paid engagements, work accepted by customers and repeat platform use. Product development and customer adoption remain subject to technical results, commercial terms and available resources.

Performance Rights milestone

The independent verification of the Chipforge platform's capability satisfies the vesting condition for a total of 74,999,997 Class A Vendor Performance Rights and 13,500,000 Class C Incentive Performance Rights. These Performance Rights were issued pursuant to the Company's acquisition of Chipforge Pte Ltd first announced on 29 April 2026 and shareholder approval was provided at the company General Meeting held 17 June 2026.

This announcement has been authorised for release to ASX by the Board of Pathkey.AI Ltd.

ENDS

investors@pathkey.ai

About Pathkey AI

Pathkey.AI Ltd (ASX: PKY) is an Australian technology company focused on developing proprietary AI powered large language model solutions for data intensive industries. TrialKey is focused on clinical trial design and optimisation. Chipforge is developing an AI-powered chip design platform for semiconductor engineering. Both TrialKey and Chipforge platforms are built around an agent-based AI architecture combining large-language-model functionality with neuro-symbolic reasoning to solve complex design optimisation problems. In each platform, AI agents generate candidate solutions, evaluate them against defined objectives and iteratively refine them through a structure reflection look using reinforcement driven optimisation.

Forward looking statements

This announcement contains forward-looking statements about product development and releases, recruitment, commercial launch, additional tool chains, secure deployment, model development, custom silicon, customer delivery and commercialisation. These statements reflect current intentions and assumptions and are subject to technical, commercial, staffing and funding risks. Actual outcomes and timing may differ materially. Development priorities do not guarantee product releases, customer acceptance and agreements recurring revenue or cash receipts. The independent evaluation assessed the product against agreed designs and criteria and is not certification of the whole platform. Pathkey does not undertake to update forward-looking statements except as required by law.