2026 DESIGN AND VERIFICATION™ DVCON CONFERENCE AND EXHIBITION UNITED STATES

SANTA CLARA, CA, USA

March 2 - 5, 2026

Broadening the Adoption of Hardware-Assisted Verification with Next Generation Emulation Appliance

Rohan Ganpati, Lance Tamura, Michael Young

cadence®

Dynamic Duo Introduction

Lance Tamura

Cadence Third Generation Dynamic Duo

Enterprise emulation and prototyping platforms

Next-Generation Hardware	*Cadence custom processor + AMD VP1902* *High Capacity: scales to **48 Billion Gates*** ***1.5X** higher performance*
Optimized Validation Throughputs	***Fastest** transition between platforms* *Pre-silicon hardware and software debug*
Unified Compiler and Interfaces	***Unified** front-end and modular compiler* ***Shared** physical and virtual interfaces*

cadence

Dynamic Duo III

Enterprise emulation and prototyping platforms

Unified compile and virtual/physical interfaces

Palladium® Z3

3-5X faster
Quick bring-up

Protium™ X3

Cadence® Z3
Emulation Processor

AMD VP1902
Adaptive SoC

build → run → debug → build → run → debug → build → run → debug

build → run → debug → run

#1 for total TAT

- 1.5x faster than Z2
- 2x more capacity per rack than Z2
- Scalable to 48BG
- Mixed-Signal and 4-state support
- Natively integrated power analysis

#1 for speed

- 1.5x faster than X2
- 3.3x more capacity per rack than X2
- Scalable to 48BG
- Unified compile & testbench link with Z3
- One shot compile technology

Hardware Debug Focus ← → Software Debug Focus

Dynamic Duo System Studio

Introducing **System Studio** for the small designs or IP

New configurations to address the market dynamics of advanced and mature node segments

Palladium® Z3

Protium™ X3

Large enterprise demands rack-based server that can scale to tens of billions of gates in datacenter environment

Palladium® Z3
System Studio

Protium™ X3
System Studio

New

Small design / IP team needs form factor that can be deployed as easy as a desktop server with no special power or cooling requirements

Why Dynamic Duo III System Studio

Design starts projection by technology node

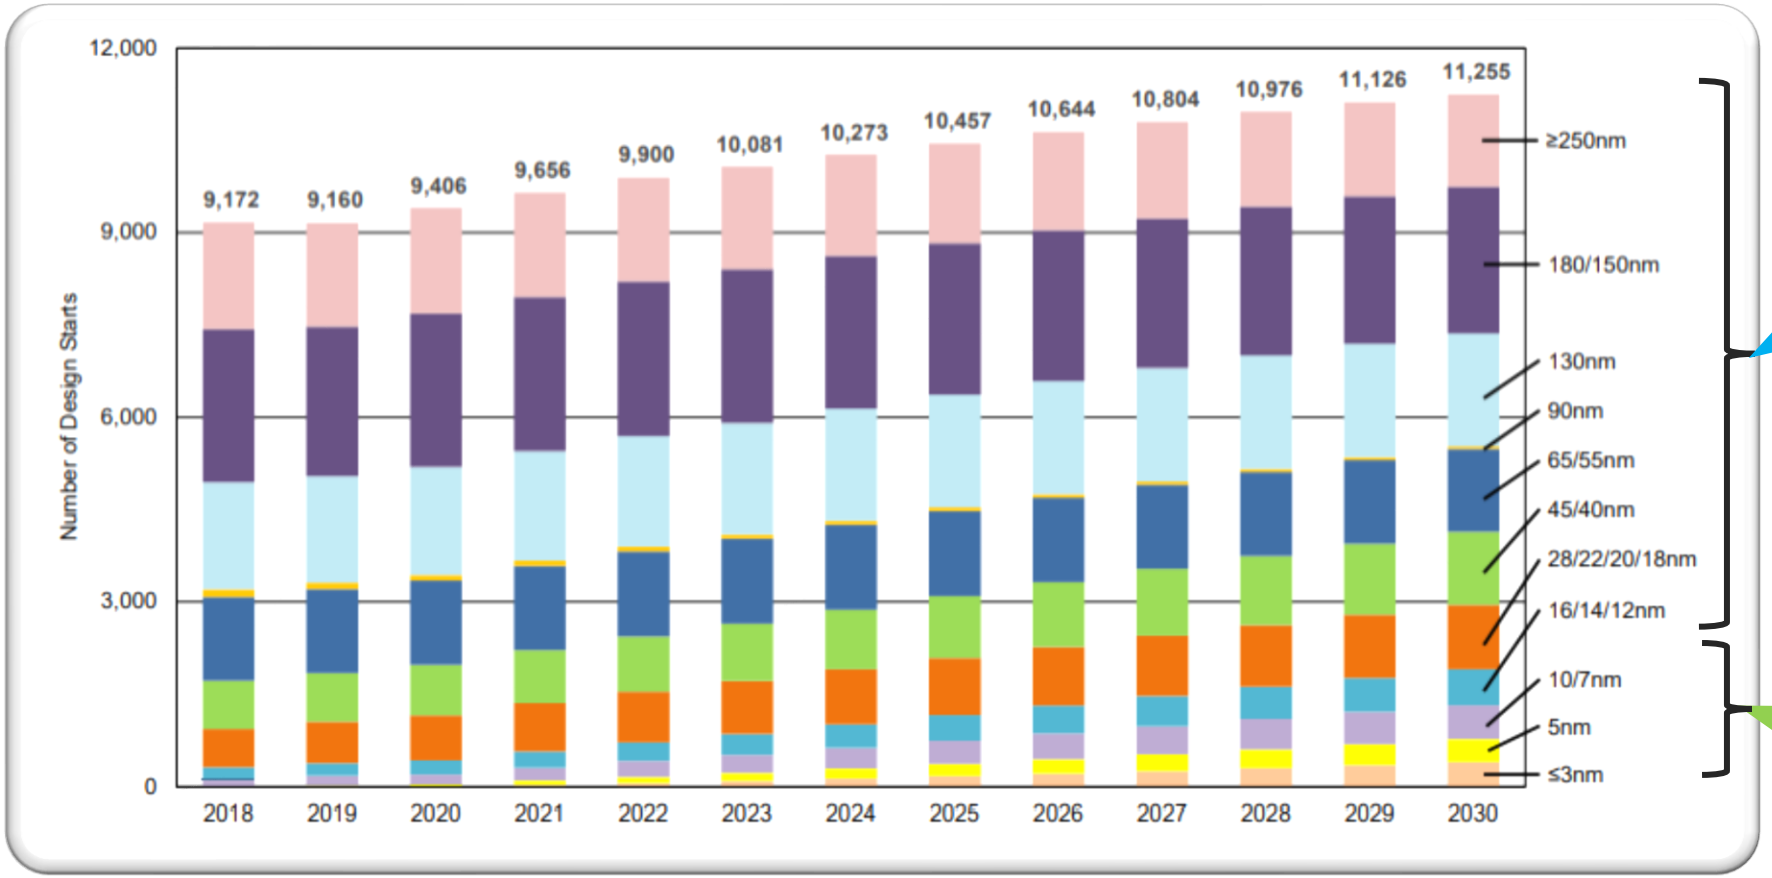

Mature Node Segments
- IoT and LP devices
- Sensors and ECU
- Industrial silicon
- Big A and Little D

Advanced Node Segments
- Hyperscale / AI
- EV / Graphics
- 5G / Networking

- By 2030, ~3/4 of design starts are projected to be using design nodes of 28nm or higher
- The need for small design verification acceleration is just as important for those design nodes

Source: IBS Design Activities

Dynamic Duo System Studio

Rohan Ganpati

Workload throughput matters even for small designs

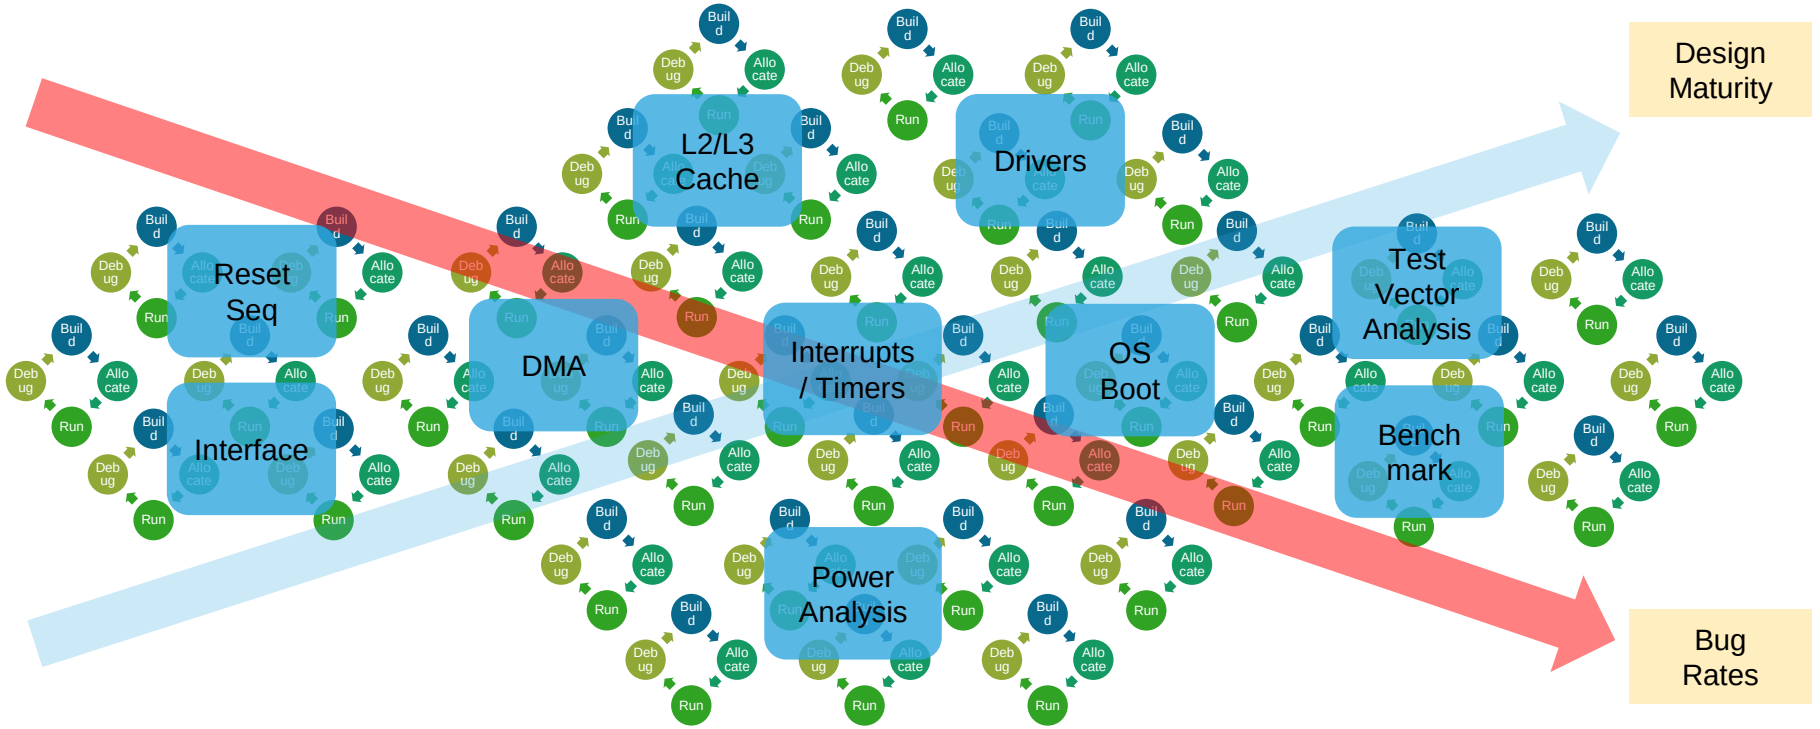

- How often will this loop be repeated over the course of a project?
- What type/size of workloads and users will be using emulation?
- How many workloads will be required to verify and validate your design?

Workloads

Palladium Z3 System Studio Overview

Off-the-shelf emulation appliance for design teams

- High-level Features
  - Stand-alone emulator with capacity of 128 million gates
    - Targets IP, sub-system or smaller ASIC designs
    - 8 domains of 16 million gates each
    - Supports up to 8 concurrent users per unit
    - Supports up to 4 target ICE ports

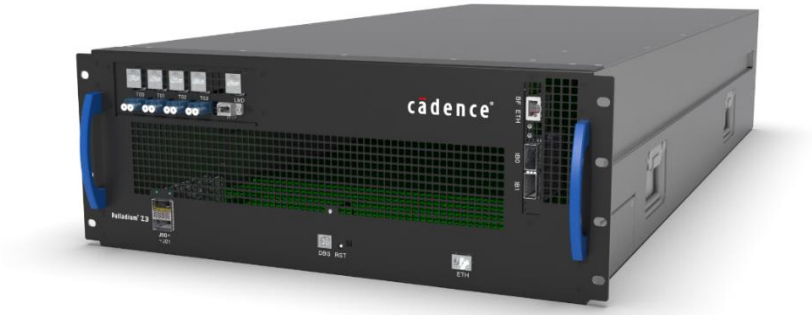

Palladium Z3
System Studio
Standalone 4U form factor

- Environment – Easy to deploy
  - Easy to relocate and maintain – no internal assembly required
    - Set up in a rack or on a lab table
    - Powered from typical lab outlets: 110/220V 15/20A
    - Air-cooled system with front-to-back airflow

- Software and Ecosystem
  - Compatible with the same software verification tools and flows as the Z3 GXL series

Protium X3 System Studio Overview

Enterprise prototyping for small SoCs and IPO

- High Level Features
  - Stand-alone system with 250M ASIC gates
    - 6 concurrent users, each 42M ASIC gates
  - 3-5x faster than Palladium Z3
  - Easy compile and bring-up through software automation
- Environment – Easy to deploy
  - Uses 110V input from standard wall outlets
  - Set up on lab bench or standard 23” rack, 4 RU height
  - Ideal for distributed teams requiring hands-on access
- Software and Ecosystem
  - Shared front-end Palladium – guaranteed congruent behavior
  - Powerful software-oriented debug capabilities
    - Hardware and virtual CPU debugger interfaces
    - Memory backdoor access
    - Record & Replay – for issue reproduction and isolation
  - Compatible with Cadence HW apps and solutions

Protium™ X3

System Studio

AMD VP1902 FPGA

Palladium and Protium Usage Example

Palladium is most efficient in compile, run, and debug loops

Enabling many iterations per day

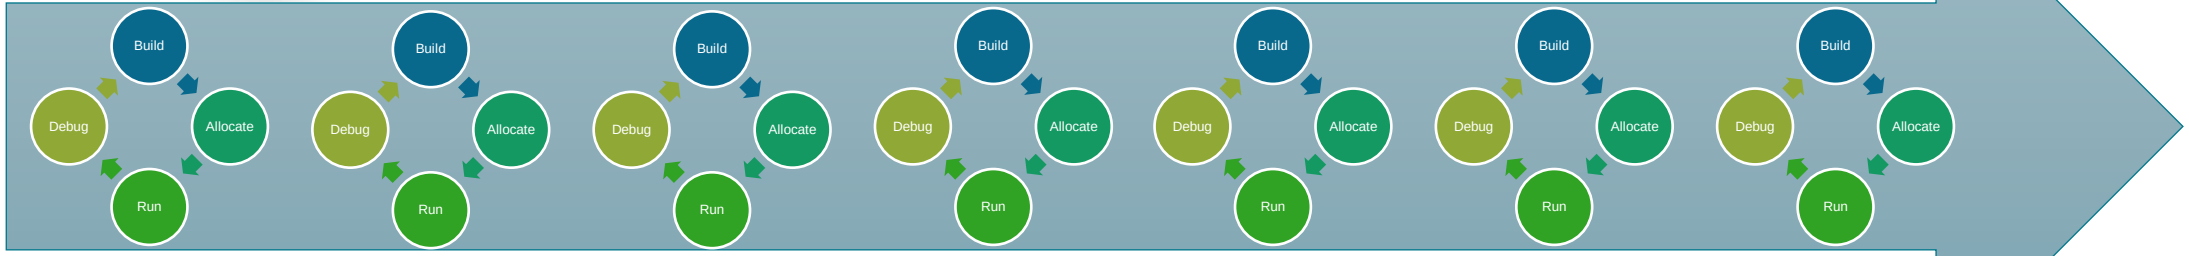

Protium is most efficient for long-running parallel runs on replicants

Enabling fast regressions, software development, and benchmarking

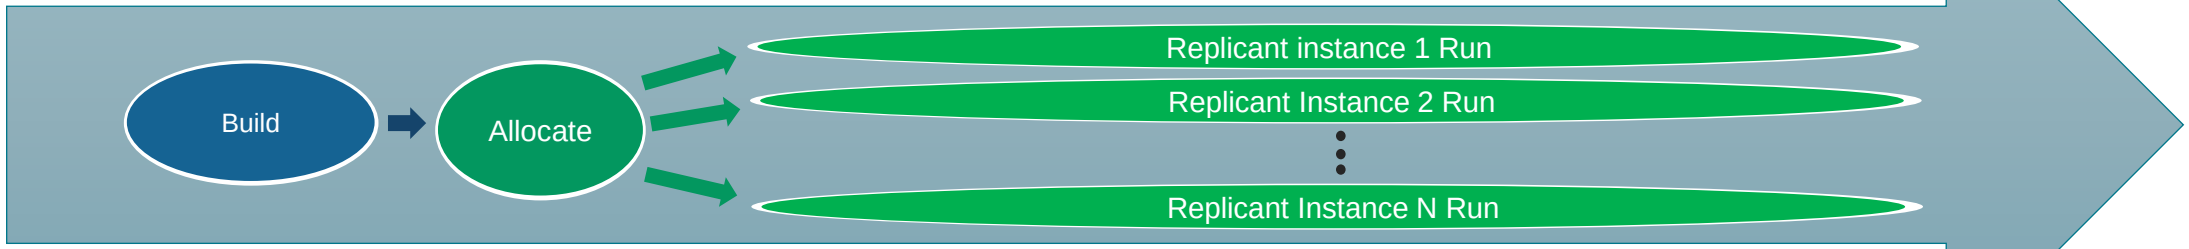

Dynamic Duo Debug Productivity

Palladium® Hardware Focused Debug

Protium™ Software Focused Debug

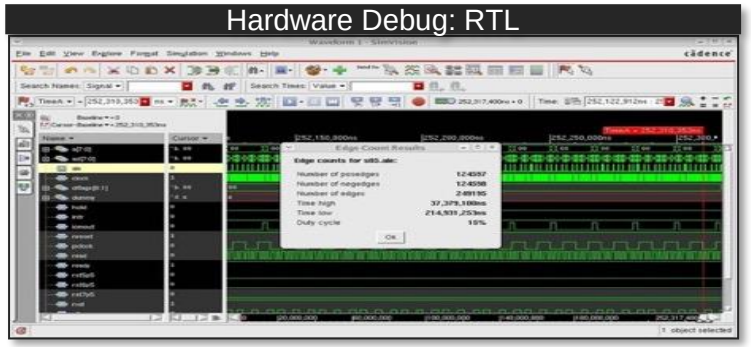

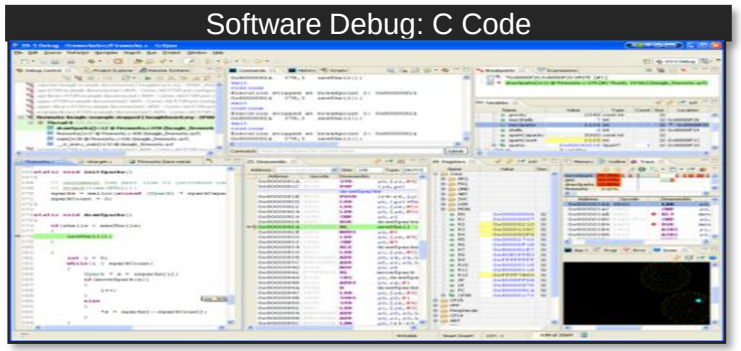

FullVision Waveforms

At-speed Streaming Probes

Code and Functional Coverage

Dynamic Probes

Dynamic RTL and SDL

Save and Restore

Force and Release

Assertions

Record and Replay

Memory backdoor read/write

FullVision Waveforms w/o predefining probes

At-speed System Probes

High-speed IO probes

Lightweight Assertions

Unique Palladium Debug

Unparalleled levels of productivity, without re-compiling, at speed

State Description Language

FullVision

Dynamic Probes

Record/Replay

Dynamic RTL

Offline Debug

Waveform streaming

Force & release

Save/Restore

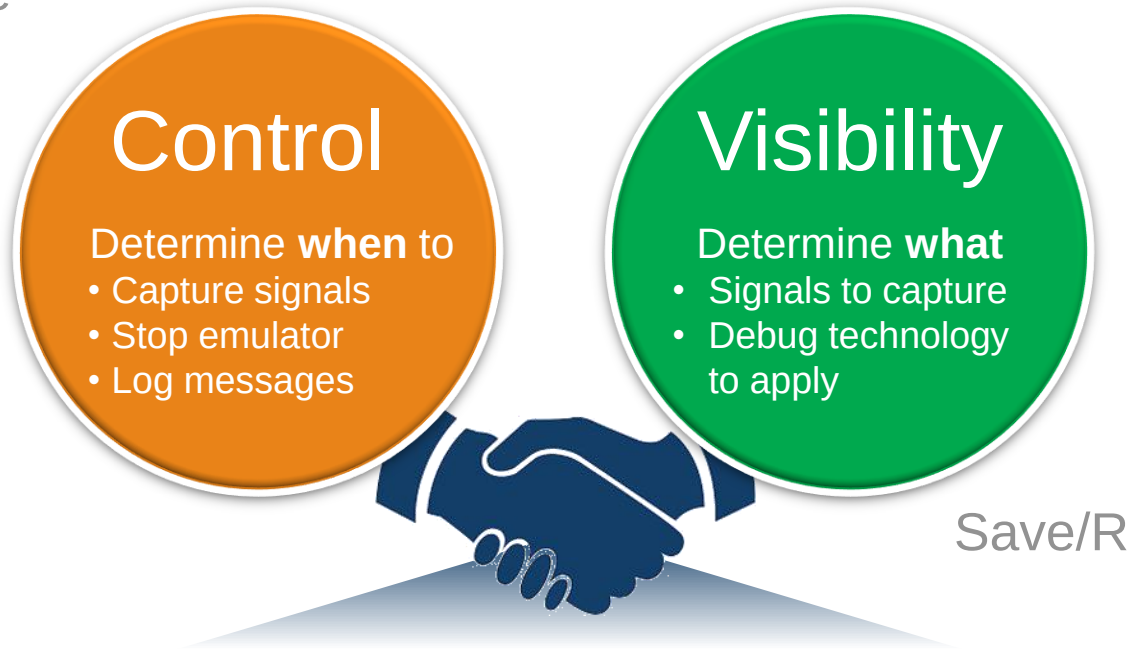

No re-compile

Re-compilation not required to gain visibility

@ Speed

No degradation to run-time speed

cadence

Unique Palladium Debug

Unparalleled levels of productivity, without re-compiling, at speed

FullVision View any design signal at full speed **without re-compilation**

Dynamic Probe Signals selected at run-time with large trace depth (up to 80M samples)

State Description Language (SDL) Specify multi-level complex trigger conditions. Use design events to detect scenarios, **without re-compile**

Dynamic RTL Create trigger condition using standard Verilog/VHDL during run-time **without re-compile**

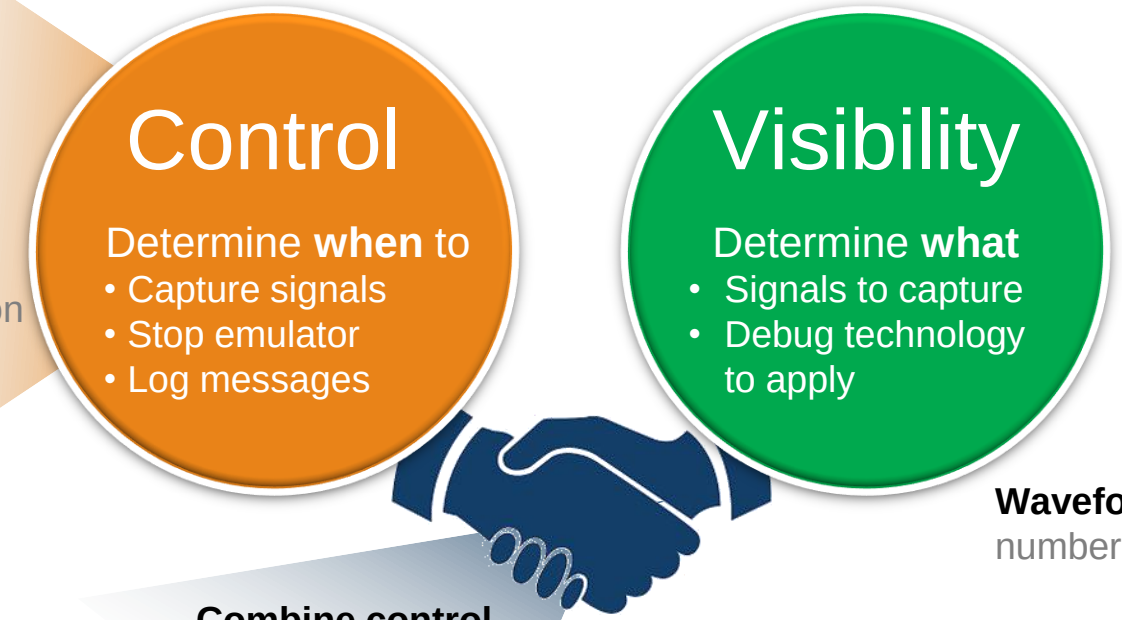

Record/Replay Complete trace set using replay, **without re-compiling**. Move forward & backward to debug time window of interest

Offline Debug Offline debugging on workstation to free up Palladium resource for other jobs

Waveform streaming Continuously view small number of signals at full rate **without re-compiling**

Combine control & visibility

- Triggering waveform capture
- FullVision & Waveform streaming
- Compiled Monitors with dynamic technologies (DRTL & SDL)

Save/Restore Re-start emulation run from previous state **without re-compiling**. Save time, avoid repetitive initializations/sequences

Force/Release Change design function during runtime. Set system conditions to analyze design behavior under un-modeled corner cases

Palladium Debug Example: FullVision

Controllability and visibility combinations

All Signals for All Time becomes unmanageable at Emulation Speeds

At Run Time (without recompile) add control to find the series of conditions of interest

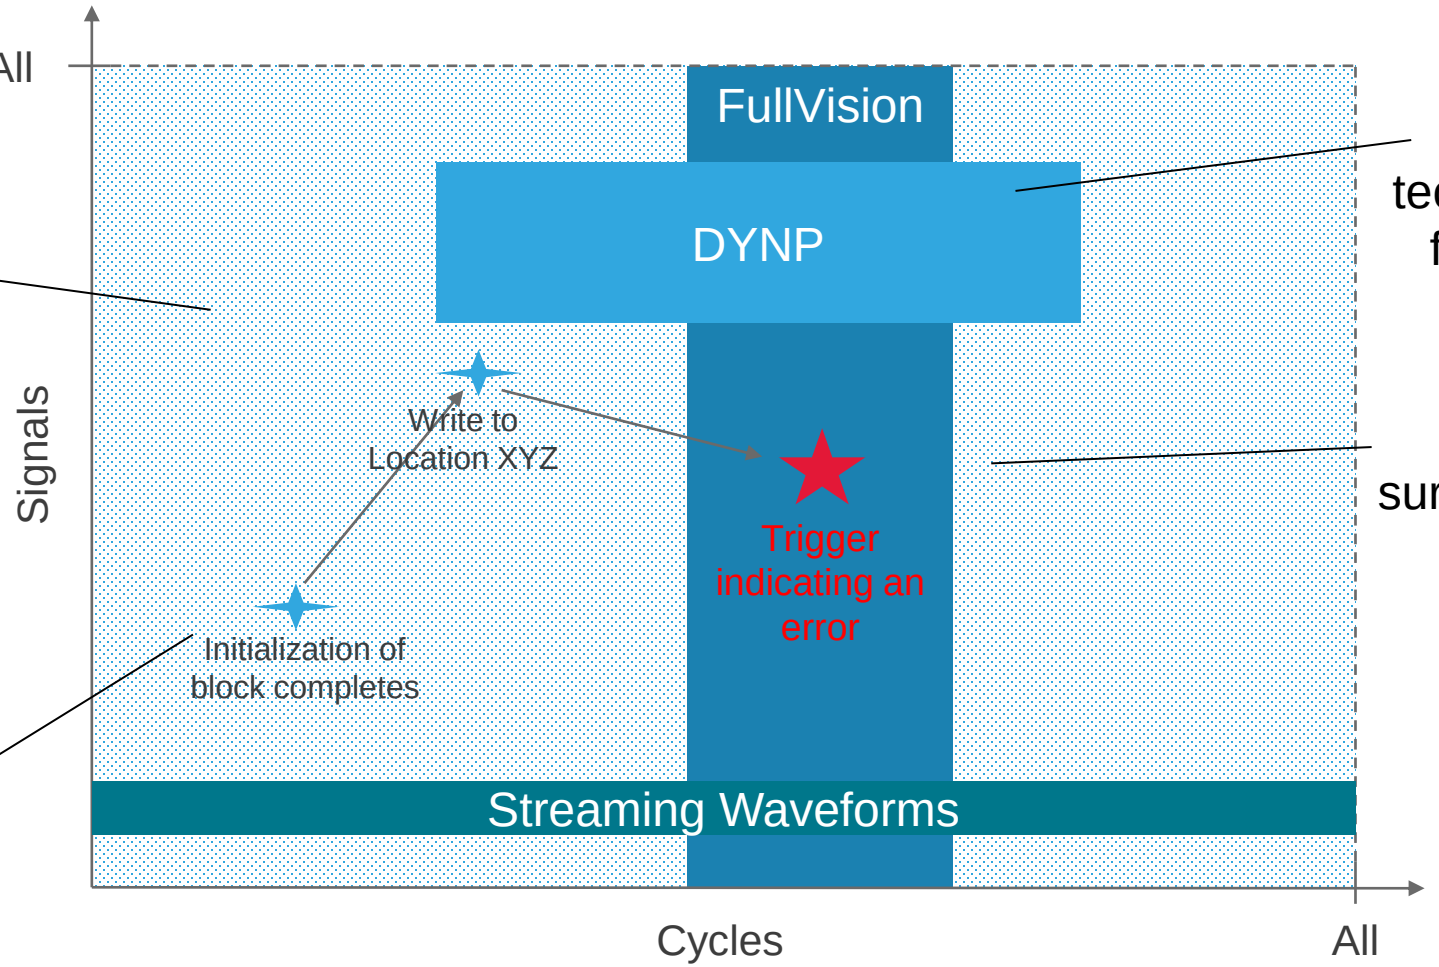

Use waveform technologies to optimize for depth and time to waveforms

Capture signals surrounding the area of interest

Palladium Debug Example: Record and Replay

Run once and replay any window of interest

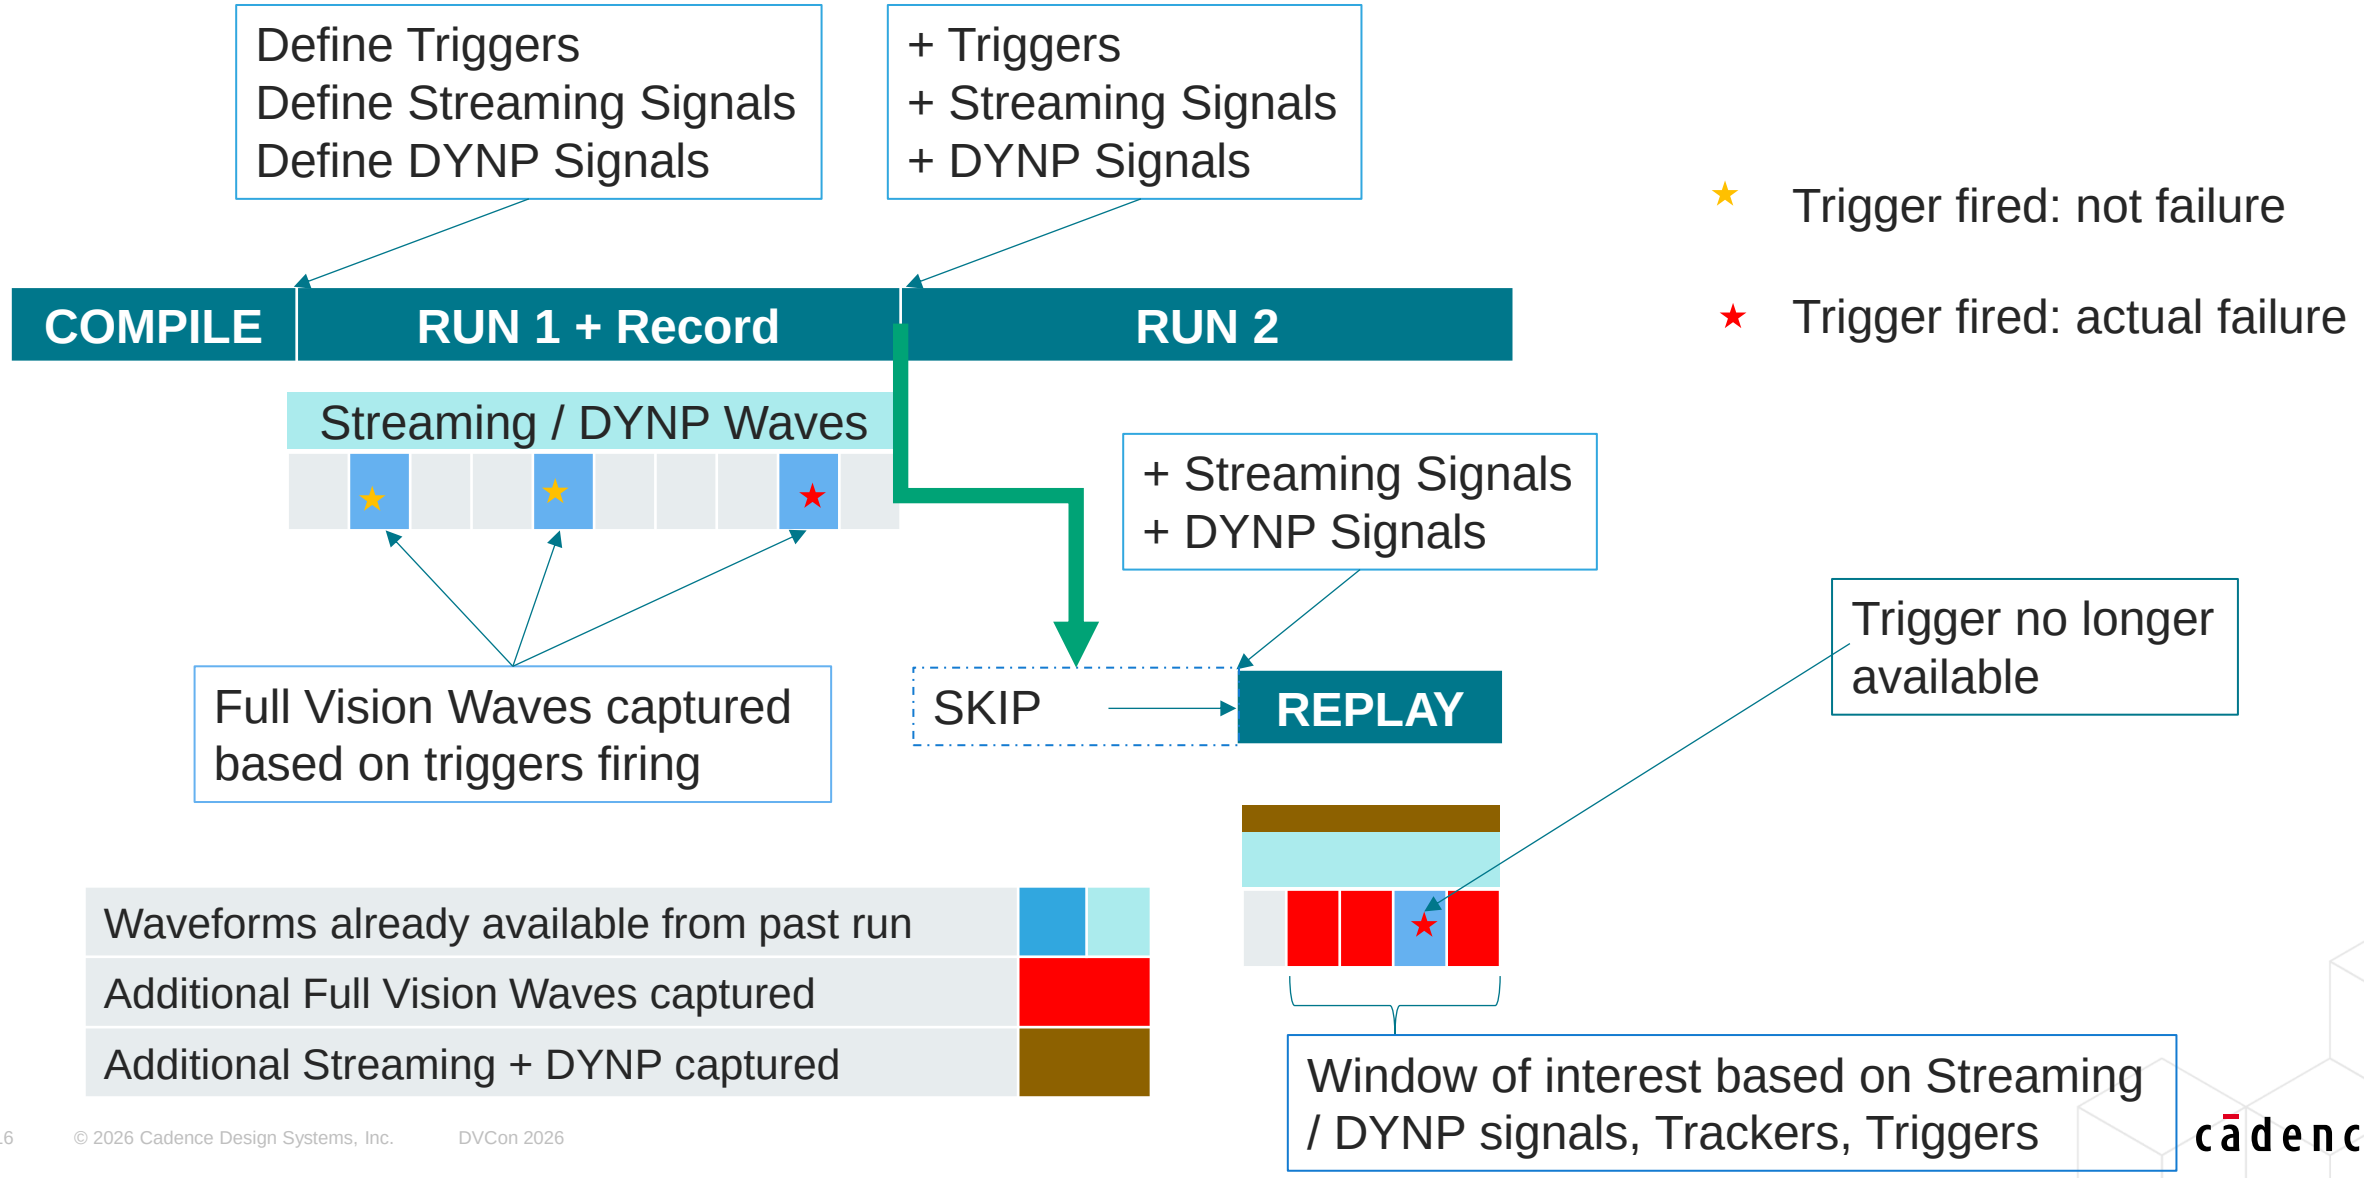

Palladium Use Models

Simulation acceleration, virtual and In-Circuit Emulation

Simulation Acceleration
TB-driven HW verification

Test Environment
Simulator
Test-Bench
UVM-A, C

Virtual Emulation
SW-driven HW verification
Interface-driven HW verification
HW/SW co-verification

Test Environment
Virtual Device Models
Application / Tester

Virtual Interfaces

Palladium Z3
System Studio

Embedded Device Models

Fast
Compile

FullVision
Debug

Coverage
Assertions

Multi-run
Modes

Low-power
DPA

Physical Interfaces

In-Circuit Emulation (ICE)
SW-driven HW verification
Interface-driven HW verification
HW/SW co-verification

Test Environment
Physical Devices
Application / Tester

Examples shown

Palladium APPs and Solutions

Fully compatible with HW APPs and Solutions

HW Apps and Solutions

Memory Model Portfolio

Ready-made memories for emulation and prototyping

SpeedBridge® and EDK

Off-the-shelf real-world HW interfaces (e.g., PCIe, USB)

Accelerated VIP and Virtual Bridge™

Virtual interface to IP, subsystem, and chip

Physical and Virtual Debug

Connect to third-party software debuggers

Next-Gen Power Analysis APP

95% Power sign-off accuracy*

4-State (0, 1, X, Z) Emulation APP

100X Accelerate 4-state faster*

Digital Mixed Signal Emulation APP

500X Accelerate DMS faster*

Palladium Safety Emulation APP

ISO 26262 Accelerate safety standards

* KPIs are based on actual customer examples; HW APPs are licensed and sold separately

cadence

Palladium Emulation Example

Performance and power analysis for IP to full-chip emulation

User Defined Parameters

IP Configurations to Full-chip Verification

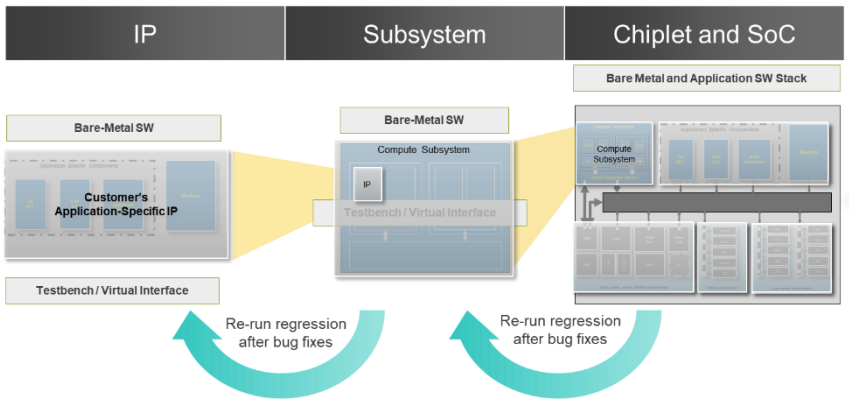

Power Profile Analysis with DPA 3.0: Average and Peak Power

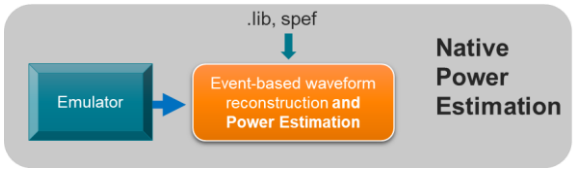

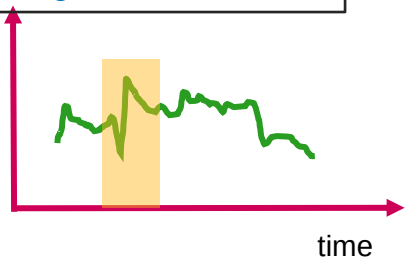

Performance analysis with System VIP and AVIP

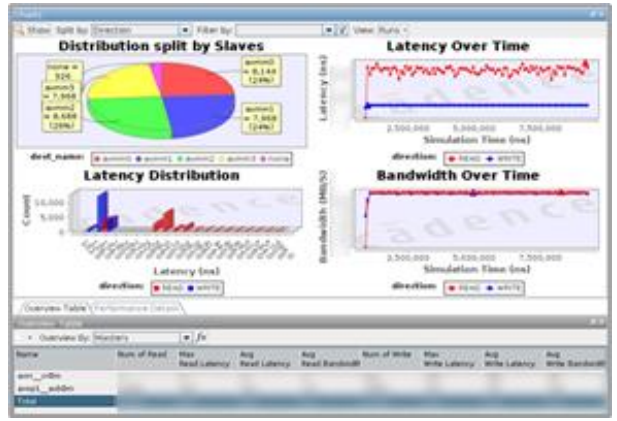

Bug hunting: eliminate hard to find bugs with system-level test

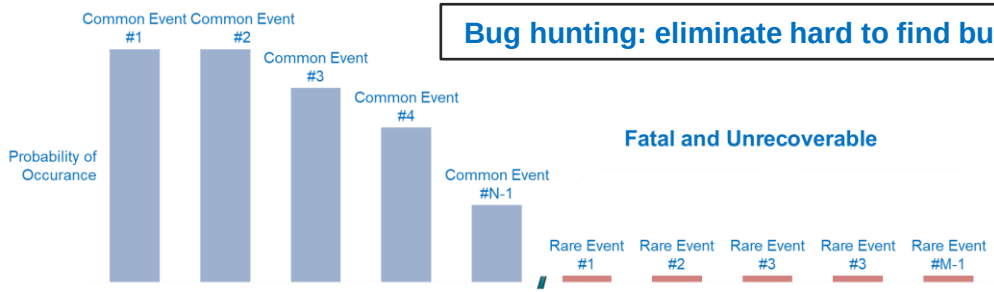

Palladium accelerates performance and power analysis at the same time from IP to Full-chip

Skymizer LPU Verification

William Wei

Speeding Skymizer LPU Verification with Palladium

William Wei | CMO, Skymizer

2026/03

We are Skymizer

- Founded in 2013, Skymizer specializes in AI chip IP and compiler technologies for edge inference.
- Headquartered in Taipei and Hsinchu, serving global semiconductor and device partners.
- Core IP includes the Language Processing Unit (LPU) series, optimized for low-latency, efficient on-device AI.

Agenda

- What is an LPU
- Verification Challenges of LPU
- Proposed Verification Flow
- Result of Works
- Conclusion

skymizer

What is an LPU (Language Processing Units)

- **Specialized for language models**

  Originating from Groq, the LPU is a dedicated AI processor built for transformer inference.

- **Distinct from NPUs/GPUs**

  LPUs prioritize deterministic, sequential dataflow and on-chip memory to optimize transformer workloads.

- **Evolve with Skymizer's IPs: EdgeThought ™ to HyperThought ™**

  Skymizer's LPU product line inherits this domain focus, enabling multi-modal, concurrent, and agentic AI on edge devices through compiler-hardware synthesis.

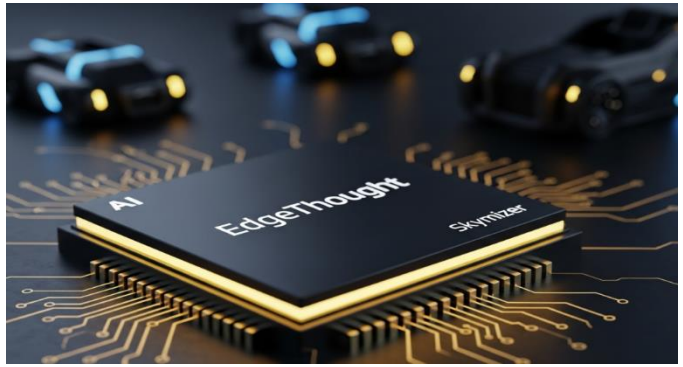

skymizer

Verification Challenge of LPU

- **Design Quality and Debugging**
  - Verification coverage
  - Large pattern set (long runtime)
  - Bug Visibility
- **SW Driver Development**
  - SOC bring-up
  - Skymizer LPU driver
  - IP driver (PCIE, LPDDR, …)
- **Performance analysis**

skymizer

Bringing Accelerated Computing to Digital Verification

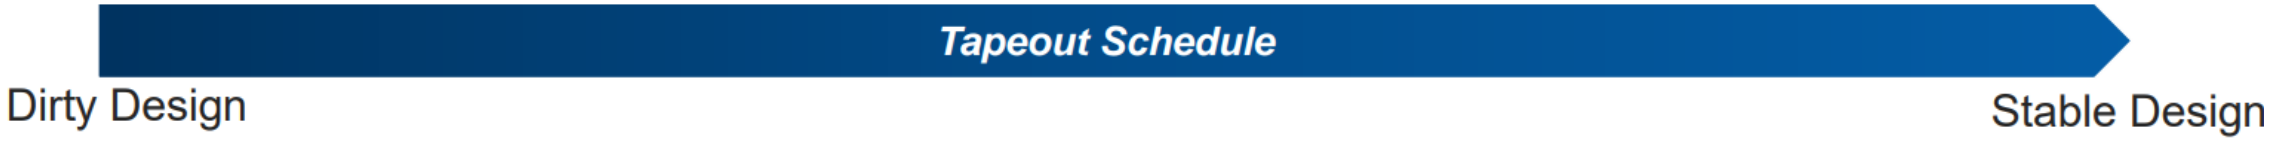

Logic Simulation
General-Purpose CPU

1-10kHz

Compile in minutes

Trace any signal at
any time at full speed

FPGA Prototyping
Development Board

MHz+

Compile in 3 to 6 months
and costs many $M

Debug via JTAG / scan

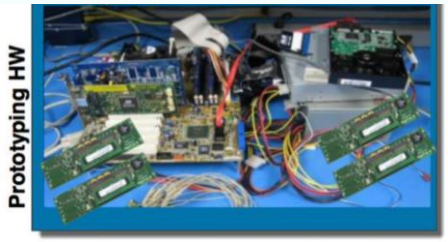

skymizer

Palladium Z3 System Studio Overview

A Processor-Based Emulator

SAME DEBUG & COMPILE

SMALLER FORM FACTOR

COMPATIBLE APP & SB/EDK

LOWER LAB REQUIREMENT

LOWER COST

- Typical lab outlets: 110/220V, 15/20A
- Air-cooled system
- Weight: 50KG
- Height: Standard 4U (17cm)

Best for IP, small designs, or secured area with no sharing allowed

skymizer

Cadence SpeedBridge

Connects emulation model to real world environment

- ✓ SpeedBridges are used to **bridge the speed** between Palladium and real-world devices
- ✓ Allowing users to emulate the design with **real-world applications** such as booting the operating system, transferring files, and displaying graphics /video

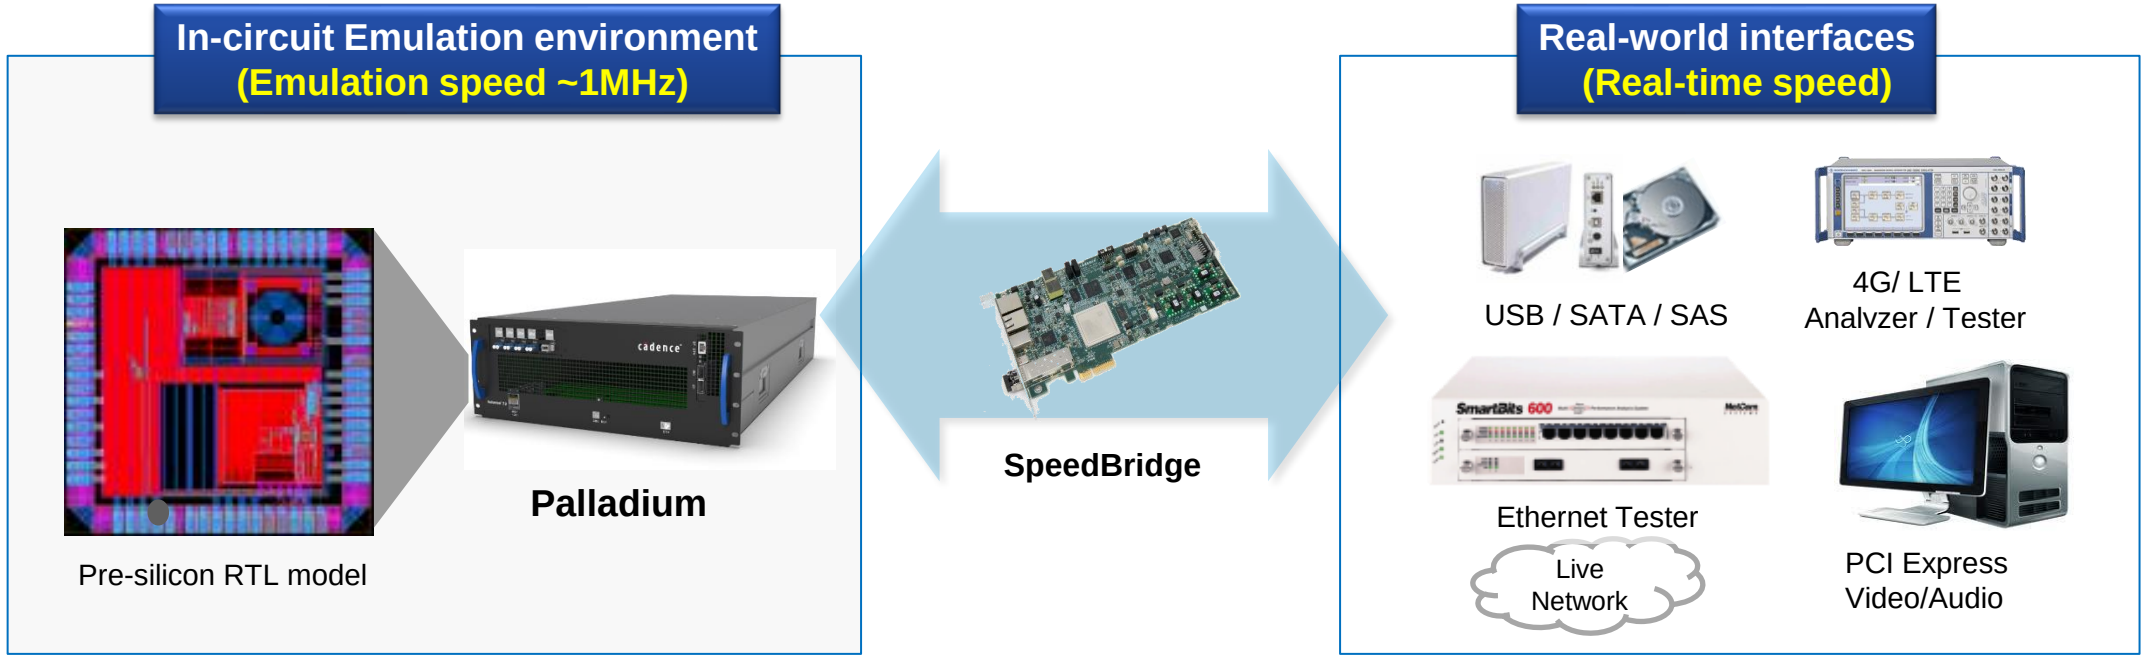

skymizer

Skymizer HTX301

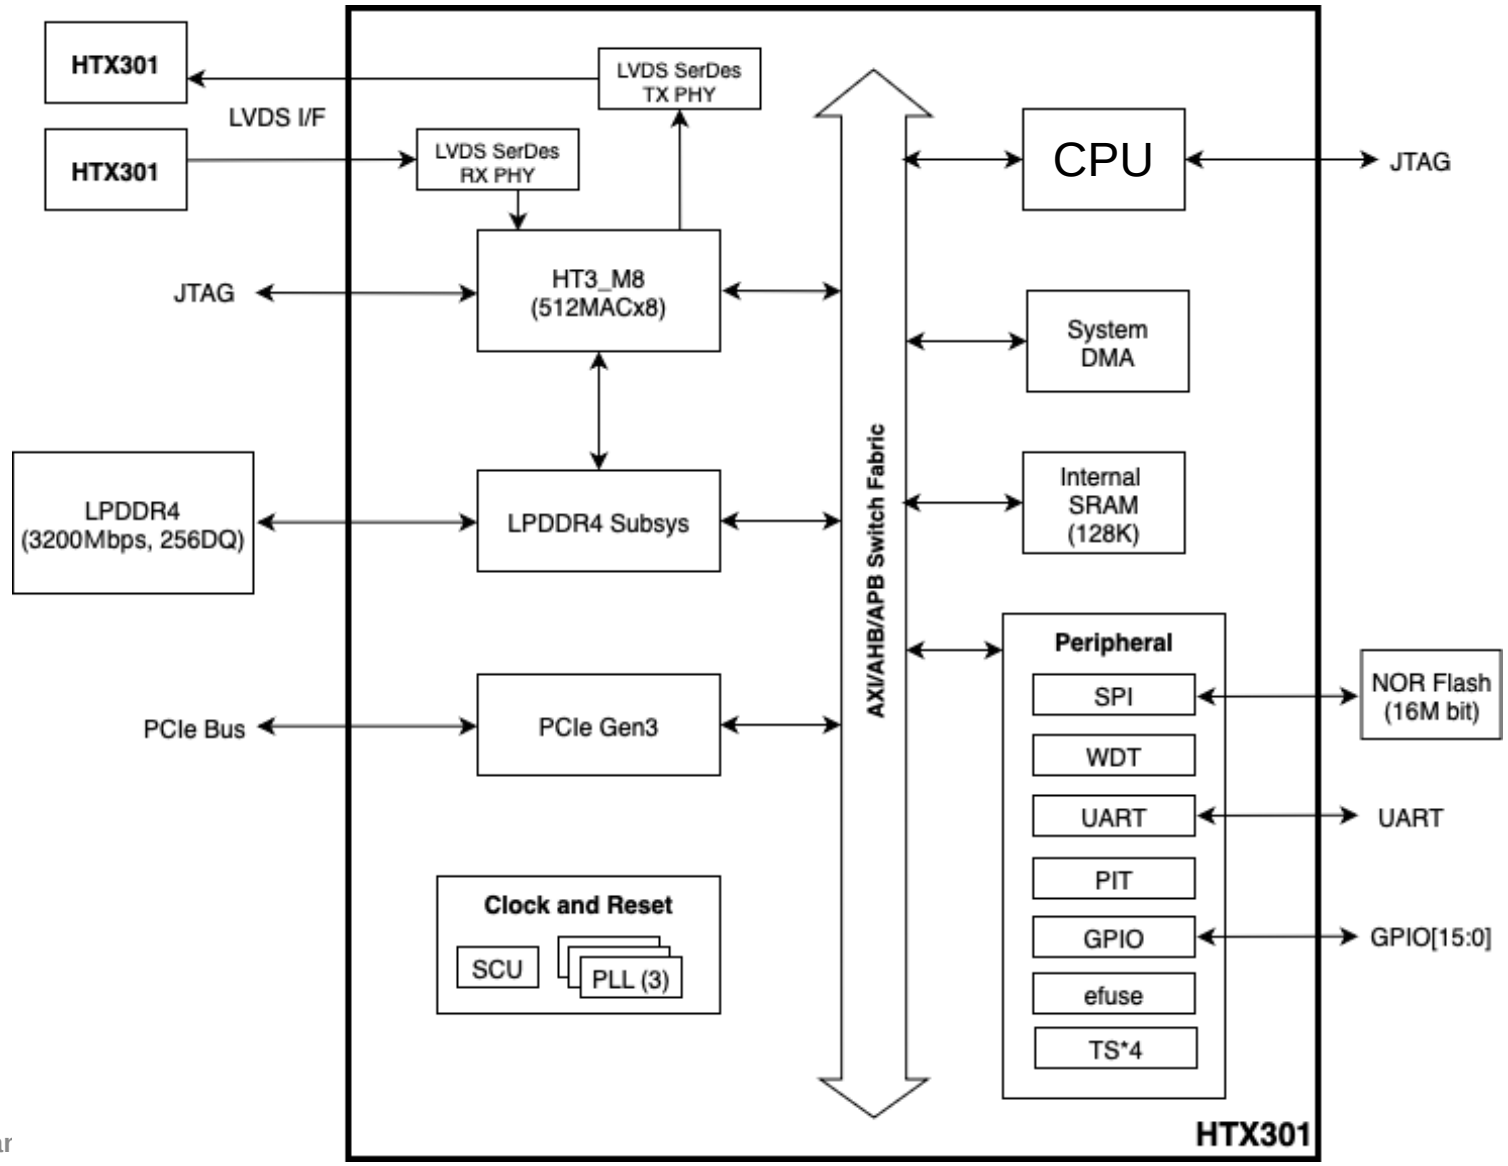

Current Verification Flow

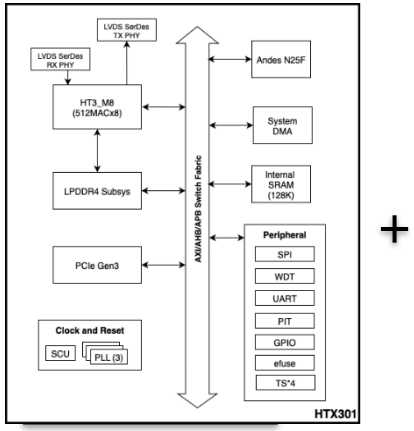

+ pcie_wrapper.v

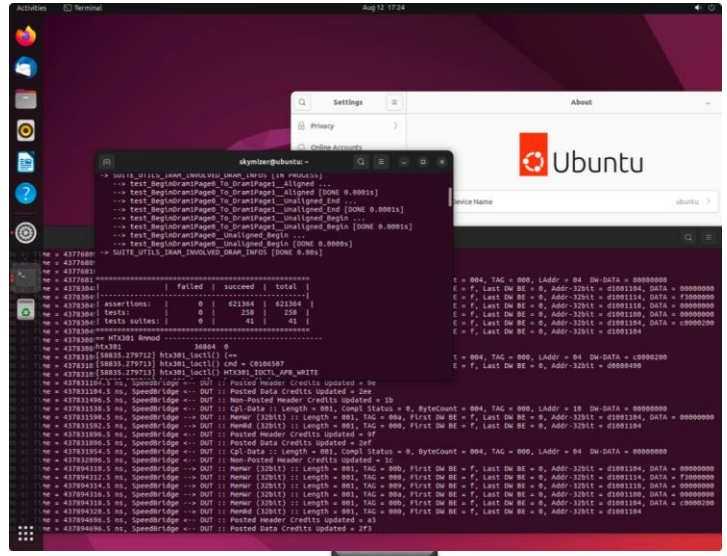

HD Optical

PCIE Gen5 SpeedBridge

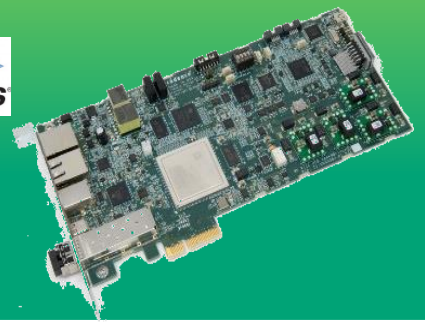

For EP DUT

skymizer

Current Verification Flow

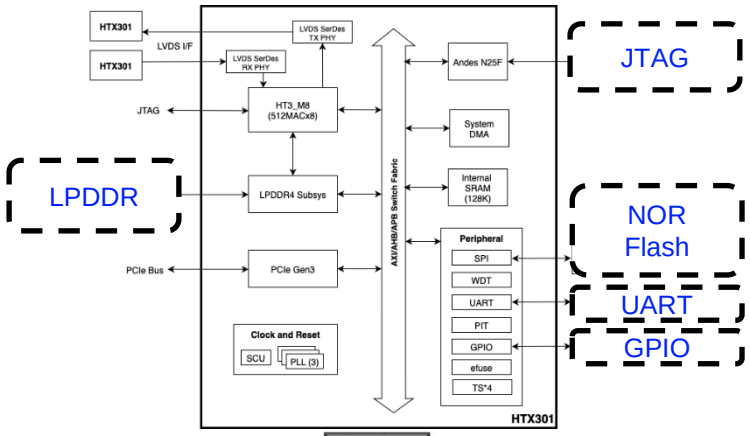

HD Optical

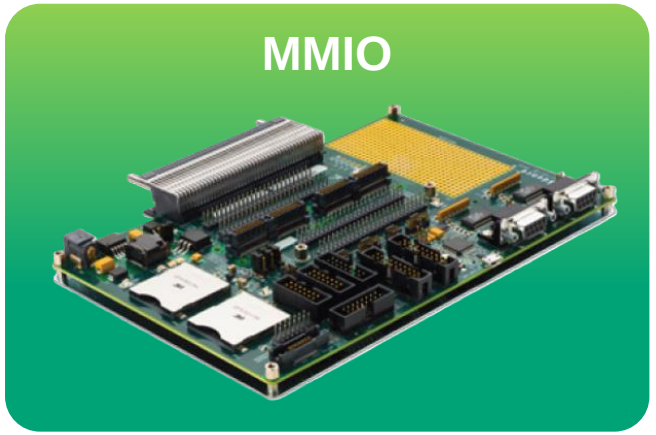

JTAG/UART

skymizer

Result of Works

- Token Throughput Rate (TinyStories260k , ignoring CI setup/teardown)
  - Simulator: 0.00247 TPS
  - Emulator: 0.11 TPS or **44.5x faster**
- Token Throughput Rate (Llama 3.2 1B, excluding ~10 mins CI setup/teardown)
  - Simulator: 4.45×10^{-6} TPS
  - Emulator: 0.00167 TPS or **374.4x faster**
- Compilation time
  - Palladium compile time ~= 10 mins or **72x faster**
  - FPGA RTL compile time ~= 12 hours (run synthesis and place-and-route to produce the FPGA bit file)

Scenario	Simulation Time	Emulation Time	Speed-up	Remark
TinyStories260k	13.5 Min	3 Mins	**4.5 times**	2 Tokens
Llama 3.2 1B	5.2 days	20 Mins	**374.4 times**	2 Tokens
Llama 3.1 8B	Expected to be 77 Days	36 Mins	**3080 times**	2 Tokens

skymizer

What We Achieved

Fast Bring-up

- ✓ Bring up on PZ1 cloud and seamlessly recompiled for Z3 on-prem
- ✓ Shift left PCIe driver development

Debug

- ✓ HW/SW co-debug with real traffic through PCIe SpeedBridge
- ✓ FullVision waveform dumping and SDL without re-compile

Acceleration

- ✓ Compiles in 10 minutes → multiple iterations per day
- ✓ Resulted in > 3000x performance improvement (simulator vs. Palladium on Llama-3.2-1B)

skymizer

Summary

- **Challenge** – Traditional RTL simulation of LPU/PCIe required weeks; software bring-up was blocked.
- **Approach** – Adopted Palladium Z3 + Z3SS flow: 10 min compile, cycle-accurate debug, FPGA co-verification for real-world I/O.
- **Outcome** – End-to-end verification time **reduced by > 90 %** and software bring-up **pulled in > 3 months**, with silicon-matched results.

"Palladium + FPGA parallel flow shrank LPU/PCIe verification from weeks to days and enabled pre-silicon software bring-up."

skymizer

Acknowledgements

Jun Wei Lin, Sr. Digital Design Engineer
Swimdi Chou, Software Engineer
Dave Wang, Software Engineer
Shawn King, Account Manager

Charles Hsieh, Product Engineer
Cheng-Yuan Hsueh, Principal Application Engineer
Arthur Lee, AE Director
Ted Chou, Account Technical Executive

skymizer

Boost deep learning accelerator
with compiler technology

Skymizer Taiwan Inc.

CONTACT US

E-mail sales@skymizer.com **Tel** +886 2 8797 8337

HQ 12F-2, No.408, Ruiguang Rd., Neihu Dist., Taipei City 11492, Taiwan

BR Center of Innovative Incubator, National Tsing Hua University, Hsinchu Taiwan

https://skymizer.ai

Summary

Dynamic Duo III: Ready to Scale Verification from IP to Enterprise

Consumer | Hyperscale | Mobile | Communications | Automotive | Aero/Defense | Industrial | Life Sciences

Vertical Solutions

Solutions and Apps

Memory Model Portfolio	SpeedBridge®/ EDK	Job Scheduler	Fault Emulation	4-state Emulation
VirtualBridge™ / AVIP	Virtual Test and Virtual Debug	Helium™ Virtual and Hybrid Platform	Power Analysis 3.0	DMS-RNM Emulation

HW Platforms

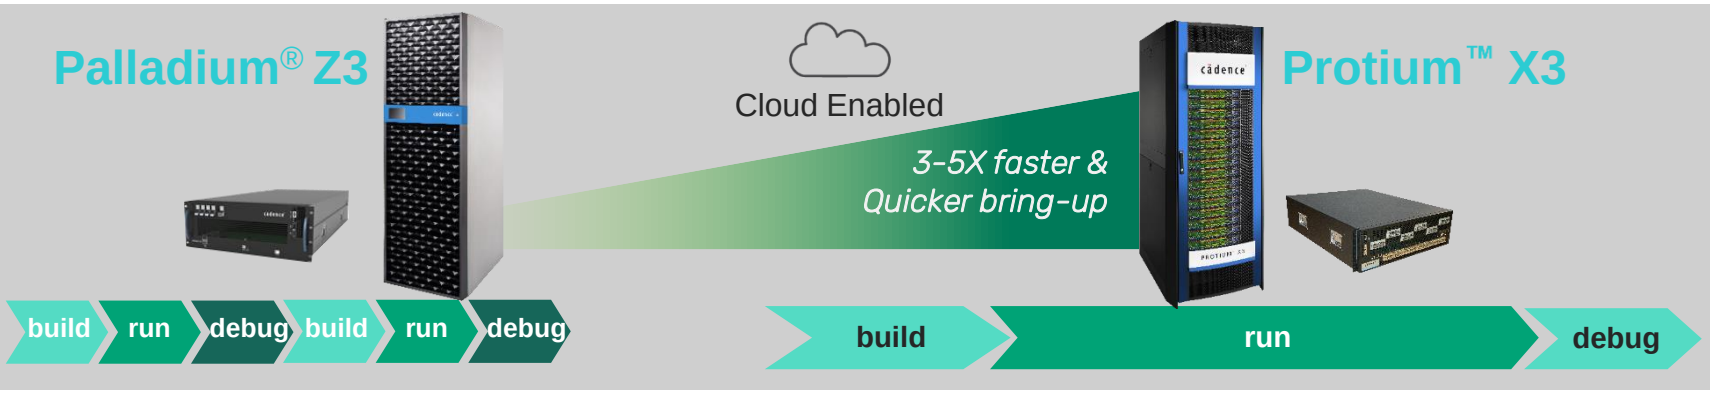

Acceleration Engines