



RAMTERA

PRODUCT DOCUMENTATION · PHYSICAL VERIFICATION (LVR)

RAMTERA-LVR-BM

Benchmark Methodology

How correlation and performance are measured — for our numbers and yours

August 2026 · Ramtera PRE engineering



LVR Benchmark Methodology

Verification-tool benchmarks are only as trustworthy as their method. This document fixes the method LVR uses — in internal qualification, in customer pilots, and in published papers — for the two questions that matter: does it find the same violations (correlation), and how fast does it run (performance). The rules below apply to Ramtera's own numbers first.

2 axes: correlation and performance	2 levels: rectangle and rule	declared environment for every figure	matched deck versions or no comparison
--	---------------------------------	--	---

Correlation methodology

Correlation runs LVR and the incumbent's known-good (golden) results on the identical design and the identical deck version, then compares at two levels. Rectangle level: each golden violation geometry must be matched by an LVR geometry (recall), and each LVR geometry must correspond to a golden one (precision). Rule level: a rule counts as clean only if its full violation set matches. Both levels are reported; a high rectangle score with rule-level misses tells a different story than the reverse, and hiding either would flatter the result.

Comparison hygiene

Rule	Why it exists
Deck versions must match exactly	different deck revisions legitimately produce different violations; comparing across revisions is excluded from correlation, not counted as tool error
Result caps are honored and disclosed	decks cap per-rule violation counts; a capped golden set can only be compared against a capped run (white paper WP-04)
Waivers excluded from both sides	correlation measures the engines, not the waiver states
Every mismatch gets a root cause	mismatches are triaged to a named cause — input difference, deck semantics, or tool defect — and tool defects go on the fix queue with the number attached

Performance methodology

Runtime is measured on a declared environment (CPU model, core count, memory, storage class, OS, compiler), on a fixed design and deck, as a thread sweep: the same run at increasing thread counts, reporting wall time and peak memory at each point. Single-number speedups without their sweep and environment are not published. Shared-filesystem effects are measured and attributed separately from engine scaling — the reference cloud runs showed storage, not the engine, as the first bound (WP-12).

WHAT IS AND ISN'T PUBLISHED — Ramtera does not publish correlation metrics derived from foundry reference data — those numbers live in certification dossiers under NDA. Public correlation claims use open PDKs or customer-permitted data only. In your pilot, you run this methodology on your designs and keep the numbers.

About Ramtera

Ramtera Inc. (Delaware, USA), with R&D at Ramtera Design Automation India Pvt Ltd (Bangalore), builds electronic design automation software for physical verification and physical implementation. LVR is Ramtera's physical verification engine (DRC, LVS, ERC, XOR, parasitic extraction); PRE is Ramtera's physical implementation environment (floorplanning, placement, clock tree synthesis, routing, timing analysis and optimization), built on industry-standard command conventions so existing flow scripts and operator skills transfer directly. The company holds granted patents and registered copyrights covering its engines.

Contact

Web	www.ramtera.com
Inquiries	sameer@ramtera.com
Engineering	Bangalore, Karnataka, India

This document reports measurements from Ramtera's engineering environment (Ubuntu 24.04, GCC 13.3, Tcl 8.6) unless stated otherwise. Product behavior and figures are subject to change; no compatibility, certification, or endorsement by any third party is stated or implied.