

PART 13 • SECOND EDITION • SEPTEMBER 2026 (FIRST PUBLISHED JULY 2026)

Commissioning & Go-Live

The building is finished. Nothing is proven. Commissioning is the only window where you may break your own data center on purpose — every test you skip runs later anyway, at full revenue exposure.

LLAMA 3 • UP TO 16,384 H100S

1 every ~3 hr (419 in 54 days)

One unplanned interruption roughly every three hours during a 54-day snapshot of the run

Llama 3 records the operational cost of interruptions. It does not identify inadequate commissioning as their cause; test the failure modes your acceptance boundary can expose.

META • AS-OF 2024-07

THE COUNTER-PRESSURE

~\$14M/mo

60 MW, one month late: lost revenue, overruns and penalties (modeled, STL Partners)

This is the force that truncates Level 5 testing. Every corner it cuts does not disappear — it relocates onto a live cluster, at the worst possible time.

STL PARTNERS, PREVENTING MULTIMILLION DOLLAR DATA CENTRE LOSSES THROUGH REPORTING · AS-OF 2025-03

THE REFRAME

Commissioning turns design into evidence on two interlocked tracks—the facility and its cluster—before integration defects reach revenue load.

Building standards leave cluster-specific link errors, burn-in coverage and workload outcomes to the acceptance contract. Write numerical limits and observable state gates before testing.

ACT I

The program: evidence, not vibes

A commissioning script is a contract for observable results. Freeze numerical limits, state gates and witnesses before testing; a disputed pass bar becomes a change order during energization.

THE LADDER

L1–L5: each level tests a wider integration boundary

LEVEL	BOUNDARY TESTED	CLUSTER-TRACK ANALOGUE
L1 — Factory acceptance	One piece of equipment meets spec before it ships	Rack integration test and firmware baseline at the integrator
L2 — Site acceptance	Equipment arrived undamaged, installed correctly	Inventory, cabling and optics verification, link-up checks
L3 — Pre-functional	One system works on its own, energized	Single-node burn-in and DCGM diagnostics
L4 — Functional performance	One system performs across its full operating envelope	Point-to-point fabric bandwidth and latency, rail by rail
L5 — Integrated Systems Test	Required systems behave together at the approved load and fault boundary	Platform collectives, training/inference and recovery evidence

THE PROGRAM'S DEFINING DECISION

Commit to an evidence-complete Level 5 IST — with a risk-based test method

L5 integrated proof

- Accepted test boundary complete; load banks staged at scale
- Every subsystem already at L4 sign-off
- Scripted campaign: inspection, injection, simulation/HIL, load banks, and approved live tests
- Expensive, schedule-hungry — the first thing to slip

Stop at L4 functional

- Every system proven alone, across its envelope
- The integration seams are never exercised together
- Cheaper and faster — on paper
- Residual risk moves to a live, revenue-bearing cluster

SO WHAT

Skipping L5 does not remove the integration risk — it reschedules the same faults onto the first real training run, where the load is irreplaceable GPUs and the audience is your customer.

THE RESILIENCE TARGET

~25–40% facility

The capital premium for Tier IV fault tolerance over Tier III concurrent maintainability

A dated McKinsey ratio, not a quote. The topology decision sets the failure-mode test matrix; a checkpointed training cluster may gain more from faster detection and recovery.

SAVILLS RESEARCH, EUROPEAN DATA CENTRES, MAY 2024 · AS-OF 2024-05

THE CLUSTER TRACK'S CONTRACT

The gates the standards won't write for you

1e-12

A legacy ibdiagnet default for SwitchX and ConnectX-3/-4 only: write your own BER gate

NVIDIA, IBDIAGNET INFINIBAND FABRIC DIAGNOSTIC TOOL USER MANUAL V2.13.0, PDF EXPORTED NOVEMBER 20, 2023. · AS-OF 2023-11

90% vs 96% scenario

Illustrative 90%-versus-96% sensitivity — contract a measured goodput gate

GUIDE ANALYSIS · AS-OF 2026-09-08

ACT II

The building on trial — and the load that lies

Facility acceptance runs on load banks. But an AI cluster's defining signature is not the magnitude of its load — it is the change in it, and the heat path it takes. No load bank reproduces both.

JULY 10, 2024 · NORTHERN VIRGINIA

~1,500 MW

Data-center load dropped in one six-fault, 82-second sequence — now a NERC Level 3 alert

Fleet-scale demand leaving the grid as sites transfer to backup is what electrical acceptance must design against: clean transfer outside the ride-through envelope, staying connected inside it.

NERC, INCIDENT REVIEW: LARGE LOAD LOSS · AS-OF 2024-07

FIXTURE CAPABILITY DECIDES COVERAGE

What each load proves, electrically and thermally

LOAD APPLIED	ELECTRICAL REALISM	THERMAL REALISM	CANNOT DEMONSTRATE
Resistive bank	kW at unity PF; specified switched steps	Air or liquid sink as specified	Reactive load; unrepresented waveform and die/contact behavior
Reactive bank	kVA/PF and specified switched steps	Air or liquid sink as specified	Unrepresented harmonics, slew and die/contact behavior
Programmable bank	Measured step, slew and repetition envelope	Specified heat split and loop connections	Untested spatial and software timing effects
Real workload	Actual tested software, rank map and waveform	Actual die/contact path and operating point	Risks irreplaceable accelerators; other workloads and untested protective faults remain outside coverage.

THE CENTERPIECE

Prove utility-loss behavior at an approved boundary with a scripted, risk-assessed method and explicit abort criteria.

Use injection, simulation/HIL, and staged test load first. A live utility interruption is exceptional work requiring owner, utility, AHJ, OEM, insurer, and CxA approval.

THE NEW L4 GATE

Power smoothing is now an acceptance criterion

65 J/GPU

In-shelf energy storage per GB300 NVL72 GPU — half the PSU volume is now capacitance

NVIDIA DEVELOPER BLOG · AS-OF 2025-07

~400 J/GPU

Vera Rubin on-rack reservoir: NVIDIA's published design figure, not a measured value

NVIDIA TECHNICAL BLOG, NVIDIA VERA RUBIN POD: SEVEN CHIPS, FIVE RACK-SCALE SYSTEMS, ONE AI SUPERCOMPUTER · AS-OF 2026-03

up to 30%

NVIDIA GB200/Megatron demonstration of GB300 shelf features; qualify the L4 limit

NVIDIA, HOW NEW GB300 NVL72 FEATURES PROVIDE STEADY POWER FOR AI, JULY 28, 2025. · AS-OF 2025-07

THE ISLAND PROBLEM

~90 GW announced

**Behind-the-meter generation, largely gas,
announced by mid-2026 — against ~2 GW operating**

Behind the meter you are not a customer of a grid — you are a grid. Commissioning is the only point where you prove the island holds voltage and frequency when the utility lets go.

CLEANVIEW, MICHAEL THOMAS, BYPASSING THE GRID, PUBLIC REPORT · AS-OF MID-2026 (PUBLISHER'S OBSERVATION WINDOW)

WHY THE LOAD BANK LIES ABOUT THE LOOP

132 kW nominal TDP: 115 kW liquid + 17 kW air

HPE GB200 NVL72 heat split at 132 kW nominal rack TDP: liquid and residual-air paths need separate acceptance

Liquid banks load the loop and worst branch before GPUs arrive. Real nodes then confirm die/contact behavior at the installed operating point; both tracks share the release gate.

HPE, NVIDIA GB200 NVL72 BY HPE DATA SHEET, OCTOBER 2026 EDITION. · AS-OF 2026-10

COOLING ACCEPTANCE – THREE DISTINCT BOUNDARIES

Prove the pipe, the coincident duty and the actual protective response

EVIDENCE	TEST BOUNDARY	RELEASE CONSEQUENCE
Pressure integrity	Adopted code and isolated pipe scope; OEM-qualified components	A cold hold leaves hot cycling and surge for separate tests
Fluid quality	Named approved fluid, materials and branch samples	Under-flushing becomes a cold-plate replacement campaign
CDU and branch duty	Liquid-bank heat plus coincident residual-air duty	Flow alone does not prove temperature or die margin
Interlock and recovery	Injected input; actual heat reduction, isolation and reset	A BMC acknowledgment alone leaves the actuator unproven

KEEP PRODUCT BOUNDARIES SEPARATE

Three cooling receipts, three different scopes

≤60 μS/cm (or within 10–20 μS/cm of procedure water); <5 NTU OCP row-manifold flush end-point, not a coolant release limit OCP, GUIDELINES FOR PRE-COMMISSION PREPARATION OF TECHNOLOGY COOLING SYSTEM · AS-OF 2025-03	45 °C / 65 °C QCT GB200: separate maximum inlet / return limits, not an operating pair QCT, QOOLRACK STAND-ALONE ADVANCED LIQUID COOLING FOR NVIDIA GB200 NVL72 SYSTEMS · AS-OF 2025-04	~140–220 L/min (water); ~144–230 L/min PG25 HPE liquid load: guide water-property flow derivation, not an OEM flow requirement HPE GB200 NVL72 DATA SHEET · AS-OF 2026-10
---	---	---

ACT III

The machine on trial

A GPU fabric does not fail loudly at commissioning. It fails quietly — one marginal optic, one mis-cabled rail at a time — and every defect you don't screen out at layer 1 comes back as a straggler mid-training-run.

BOTTOM-UP, IN ORDER – THE LAYER ORDER IS NOT NEGOTIABLE


The fabric gates, and what skipping each one costs

GATE	PASS BAR	COST OF SKIPPING
Physical layer	Post-FEC BER within the named interface limit; pre-FEC margin stated; zero flaps through line-rate soak	A marginal optic becomes a production straggler and days of forensics
Topology	Discovered adjacency graph matches the cabling map exactly — before any collective runs	A cross-patched rail passes BER, collapses bandwidth, and mimics a tuning problem
Point-to-point	Named endpoint/operation/message: qualified payload and completion tails; freeze commands and fault budgets	A slow path hides in the collective average; localization is lost
Time sync (PTP)	Every node inside a sub-μs offset bound — under load and in GNSS-loss holdover	Uncorrelatable telemetry; one non-PTP hop defeats the best grandmaster

THE BATHTUB CURVE **Burn-in buys down infant mortality before the failure curve flattens**

Burn-in ends at the contracted stop rule, not when the clock runs out

Schematic curve; phases not to scale.



Stop rule (in the contract): discovery rate below an agreed threshold over an agreed exposure.

The soak runs thermal and power stress, fabric and HBM pattern tests, checkpoint/restore and straggler tests.

For scale: SemiAnalysis reported ~7 days typical MTBF for one 512-H100 cluster at a top-tier operator.

Schematic · Source: SemiAnalysis (Oct 2024) + AMD · as of October 2026

A new cluster failing often is not broken — it is burn-in working. The soak exists so those failures happen on your clock, against safe load, instead of mid-run on the customer's.

NODE ACCEPTANCE, BY THE NUMBERS

Buy the failures on your own clock

≥3–4 weeks

October 2024 factory hot burn-in recommendation; not observed fleet settling time

SEMIANALYSIS, AI NEOCLOUD PLAYBOOK AND ANATOMY · AS-OF 2024-10

**~1 fault per 1,000 devices; exposure
unstated**

Meta: about one SDC fault per thousand devices; exposure period unstated

META ENGINEERING, HOW META KEEPS ITS AI HARDWARE RELIABLE, 22 JULY 2025. · AS-OF 2025-07-22

SIGN THE MEASURED RESULT

90% vs 96% scenario

**The 90% vs 96% scenario is not an acceptance bar —
sign the named reference-run result**

Use the scenario to test economics only. Contract the reference workload, productive-time and event accounting, observation window, recovery policy, minimum result, margin, and retest rule.

GUIDE ANALYSIS · AS-OF 2026-09-08

TRAINING AND INFERENCE ACCEPT DIFFERENT SERVICES

Freeze the workload, then test the failed state

PROFILE	REQUIRED EVIDENCE	HOLD CONDITION
Training/recovery	Unique state, capture, drain, durable commit and validated restart	Torn generation, lost-work or restart budget missed
Inference	Offered requests, correct completions and TTFT/decode/completion tails	Errors hidden from denominator or any tail limit missed
NVIDIA / AMD	Supported diagnostic manifest and NCCL/RCCL workload tuple	Required tests skipped or inaccessible
TPU / Trainium service	Exposed health, workload and replacement/recovery traces	Provider repair assumed to prove customer recovery

ACT IV

The ramp and the handover

A fixed bank hides synchronized cluster swings; a switched bank proves its measured waveform. Go-live adds actual software and heat coupling, so ramp through gates as each new operating point is qualified.

CANARY → PROXY → FULL LOAD

The soft-launch ladder: promote only when the gate clears

STAGE	FIRST REAL EXERCISE OF	PROMOTION GATE
Canary — one rack	End-to-end plumbing: power, cooling, fabric, telemetry flow	Required node tests pass; critical telemetry and interlocks verified
Partial-fleet proxy run	Collective traffic, storage, scheduler — and the first real power swing	Surviving power, liquid and air capacities hold; measured step envelope passes
Full synchronous load	Full heat flux and the full-amplitude synchronized swing	Training or inference contract, recovery and ORR pass; no failed path hidden in an average

UPTIME INSTITUTE, 2025 · ORGANIZATION-LEVEL DENOMINATOR

~40%

Organizations reporting a major human-error outage in the prior three years

About 40% of surveyed organizations reported this history. It is not a share of all outages; it supports a hard Operational Readiness Review without manufacturing an incident denominator.

UPTIME INSTITUTE ANNUAL OUTAGE ANALYSIS 2025 · AS-OF 2025

THE THROUGH-LINE

Commissioning puts failures on your clock. Leave an integration fault untested at the approved boundary, and its cost returns with interest on the customer's run.

TAKE IT HOME

What to do Monday

- 01 Engage an independent CxA at design; freeze numerical/state gates and witnesses before soft gates become change orders.
- 02 Contract evidence-complete L5 IST: method authority, prerequisites, hold/abort points, restoration, and substitute evidence.
- 03 Buy fixture capability: kW, kVA/PF, switching waveform and air/liquid heat path; log what each leaves unproven.
- 04 Sequence mechanical Cx and GPU bring-up: emulators prove their envelope; staged workloads confirm the product and operating point.
- 05 Set accept / RMA / re-soak thresholds before the soak starts, and hold them when the schedule screams.
- 06 Accept the training/inference contract and separate gates; hand operations the configuration/performance fingerprints as t=0.
- 07 Gate go-live on an Operational Readiness Review with a named owner who can say not yet.

|| THE DEFINITIVE GUIDE TO
AI Data Centers

This deck is the map. The guide is the territory — every number here is dated, sourced, and kept current in the full chapters.

- Ch 13.1 The L1–L5 program, governance, and the two-track model
- Ch 13.6 Integrated Systems Testing and the load-bank realism gap
- Ch 13.8 Burn-in soaks and the accept / RMA / re-soak thresholds
- Ch 13.10 The staged ramp, the readiness gate, and handover

READ THE FULL PART, WITH EVERY SOURCE AND DERIVATION, AT [AIDATACENTERGUIDE.COM](https://aidatacenterguide.com)

SOURCES

Every figure in this deck traces to the guide's dated register (1/4)

Meta (Llama 3 Herd of Models) · as-of 2024-07 ·
<https://arxiv.org/abs/2407.21783>

STL Partners, Preventing multimillion dollar data centre losses through reporting (report summary and 2025-03-27 press release) · as-of 2025-03 · <https://stlpartners.com/research/preventing-multimillion-dollar-data-centre-losses-through-reporting/>

Construct & Commission; BMP MEP; CxPlanner · as-of 2025 ·
<https://constructandcommission.com/level-5-data-center-commissioning-step-by-step-guide/>

Savills Research, European Data Centres, May 2024; publisher report. Quoted secondary estimate; not measured project evidence. · as-of 2024-05 ·
<https://pdf.euro.savills.co.uk/european/european-commercial-markets/spotlight-eu-data-centre---may-2024---final.pdf>

NVIDIA, ibdiagnet InfiniBand Fabric Diagnostic Tool User Manual v2.13.0, PDF exported November 20, 2023. · as-of 2023-11 ·
<https://docs.nvidia.com/ibdiagnet-infiniband-fabric-diagnostic-tool-user-manual-v2-13-0.pdf>

Guide analysis – stipulated sensitivity scenario; no claim of an industry measurement. · as-of 2026-09-08

NERC, Incident Review: Large Load Loss (2025-01-08), pp.7-8; NERC comments to FERC, Docket EL25-49 (2025-04-23), Attachment A (April 17 presentation), PDF p.17 · as-of 2024-07 ·
https://www.nerc.com/globalassets/our-work/reports/event-reports/incident_review_large_load_loss.pdf

SOURCES

Every figure in this deck traces to the guide's dated register (2/4)

NVIDIA Developer Blog / LITEON · as-of 2025-07 ·
<https://developer.nvidia.com/blog/how-new-gb300-nvl72-features-provide-steady-power-for-ai/>

NVIDIA Technical Blog, NVIDIA Vera Rubin POD: Seven Chips, Five Rack-Scale Systems, One AI Supercomputer (2026-03-16) · as-of 2026-03 · <https://developer.nvidia.com/blog/>

NVIDIA, How New GB300 NVL72 Features Provide Steady Power for AI, July 28, 2025. · as-of 2025-07 ·
<https://developer.nvidia.com/blog/how-new-gb300-nvl72-features-provide-steady-power-for-ai/>

Cleanview, Michael Thomas, Bypassing the Grid, public report; explicitly describes mid-2026 operating capacity and gives no publication day. Verification on September 8, 2026 does not create a source date. · as-of mid-2026 (publisher's observation window) · <https://cleanview.co/reports/behind-the-meter-data-centers>

HPE, NVIDIA GB200 NVL72 by HPE data sheet, October 2026 edition. · as-of 2026-10 ·
<https://www.hpe.com/psnow/generateDDS/NVIDIA%20GB200%20NVL72%20by%20HPE%20data%20sheet-PSN1014890104MYEN.pdf>

OCP, Guidelines for Pre-Commission Preparation of Technology Cooling System (TCS) Row Manifolds in Liquid Cooled Data Centers, v1.0 (2025-03-17) · as-of 2025-03 ·
<https://www.opencompute.org/documents/ocp-document-submission-guidelines-for-pre-commission-preparation-of-technology-cooling-system-tcs-row-manifolds-revq-pdf-1>

SOURCES

Every figure in this deck traces to the guide's dated register (3/4)

QCT, QoolRack Stand-Alone Advanced Liquid Cooling for NVIDIA GB200 NVL72 Systems · as-of 2025-04 · https://blog.qct.io/wp-content/uploads/2025/04/QCT-Qoolrack-Stand-Alone_Advanced-Liquid-Cooling-for-NVIDIA-GB200-NVL72-Systems.pdf

HPE GB200 NVL72 data sheet (PSN1014890104MYEN, October 2026), p.2: 132 kW total, 115 kW liquid, 17 kW air; OCP OAI System Liquid Cooling Guidelines §1.3 (2023-03-03 update) for the 7.5-12 °C rise and PG25 flow rule; guide water-property calculation. · as-of 2026-10 · <https://www.hpe.com/psnow/generateDDS/NVIDIA%20GB200%20NVL72%20by%20HPE%20data%20sheet-PSN1014890104MYEN.pdf>

Keysight, Validating 800GE and 1.6TE Interconnects · as-of 2026-08-24 · <https://www.keysight.com/us/en/assets/3125-1163/application-notes/Validating-800GE-and-1.6TE-Interconnects.pdf>

Engineering at Meta · as-of 2024-02 · <https://engineering.fb.com/2024/02/07/production-engineering/simple-precision-time-protocol-sptp-meta/>

SemiAnalysis, AI Neocloud Playbook and Anatomy · as-of 2024-10 · <https://newsletter.semianalysis.com/p/ai-neocloud-playbook-and-anatomy>

Meta Engineering, How Meta keeps its AI hardware reliable, 22 July 2025. · as-of 2025-07-22 · <https://engineering.fb.com/2025/07/22/data-infrastructure/how-meta-keeps-its-ai-hardware-reliable/>

SOURCES

Every figure in this deck traces to the guide's dated register (4/4)

Uptime Institute Annual Outage Analysis 2025 · as-of 2025 ·
<https://uptimeinstitute.com/about-ui/press-releases/uptime-announces-annual-outage-analysis-report-2025>

SOURCES

About these figures

Figures are resolved by reference from the guide's claim register at build time (Second edition · September 2026 (first published July 2026)) — a number appears here only as it appears, dated and sourced, at aidatacenterguide.com. Newer data may exist: check the live register.