LVDC: Building an Open and Evolving Ecosystem

Jason Adrian - Systems Architect, Google
Jared Huntington - Power Architect, Nvidia



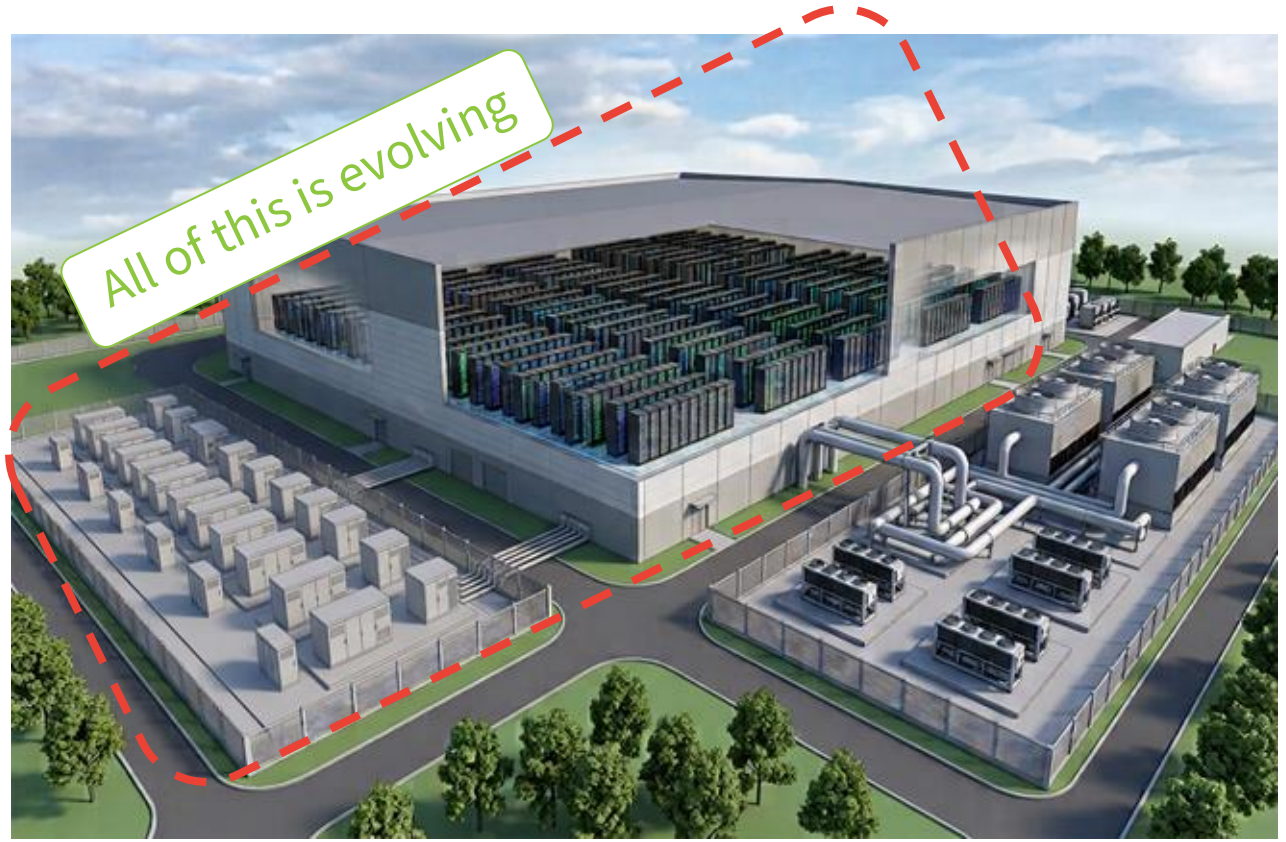
Jason



Jared

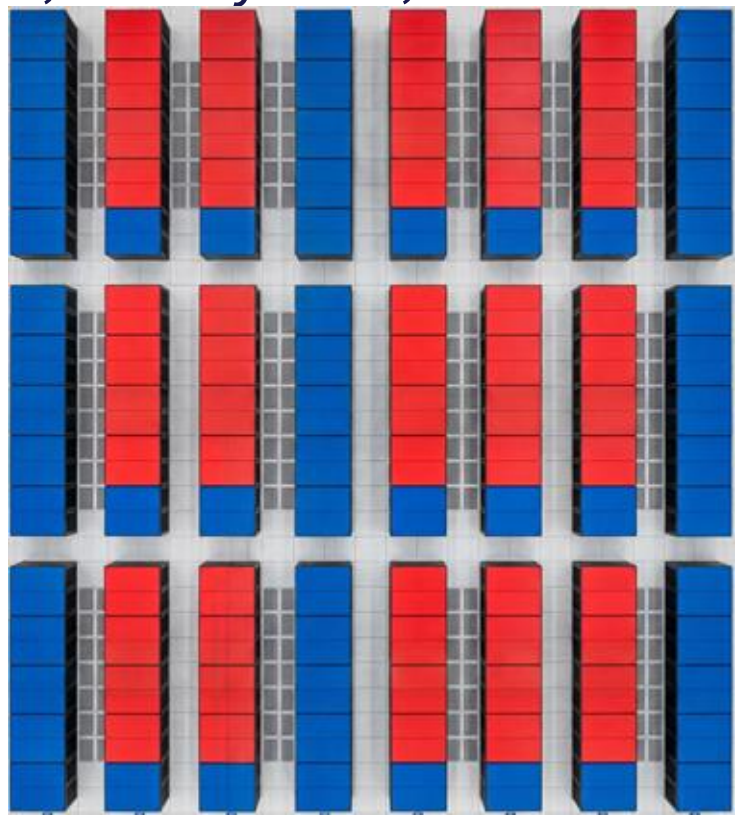
LVDC Revolution

- LVDC revolution impacts **everything** from the grid to the chip
- ML racks already transitioning to LVDC out of necessity
- LVDC enablement via sidecar tax is growing - 1MW racks drive imperative to native LVDC data centers!
- Entire ecosystem of interoperable products needed, and must scale to 10’s of GW rapidly
- How do we handle non-ML racks and other AC equipment?



The Big Picture

- 80%+ of the data center power is xPUs, but there are many other IT devices:
  - Compute, Storage, Scale-up/out & Core Networking, CDUs, Facility Fans, etc



- How to handle a mix of DC-native and AC loads and sources?
- Racks often need to be distributed across rows for power utilization, network cable length requirements, etc
- Splitting the data center into AC and DC powered rows would create complex placement challenges and fixed configurations

Breaking Points For In-Rack Power

What drives LVDC adoption?

As rack power reaches **≥100kW**, in-rack power conversion & backup consume up to **35%+** of available space:

- **AC:DC Power Conversion:** Takes 10–15% of rack capacity
- **N+N redundancy**: Increases power conversion 2x
- **In-Rack BBUs:** Consumes an additional 10–15% of space.

Moving to LVDC **offloads power conversion, freeing space for maximum xPU / accelerator density.**

Above 600kW: In-rack AC/DC conversion becomes completely unviable.

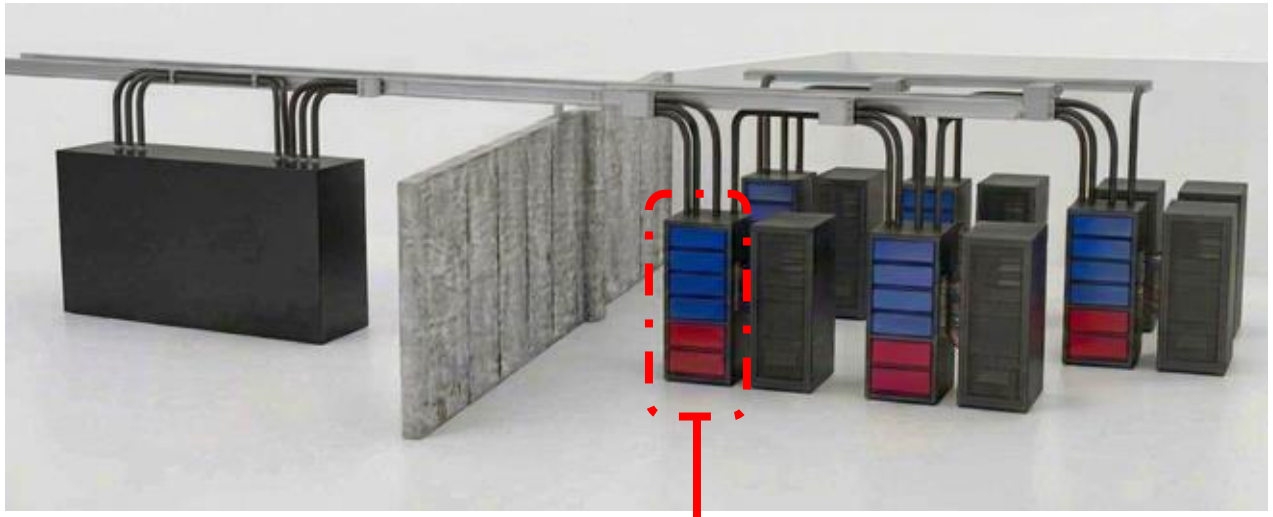
Should all racks move to LVDC?

xPU adoption of LVDC is obvious and under way, but what about everything else?

- High density networking and some dense compute evaluating LVDC - hitting the rack power wall
- Storage and low-density racks could remain AC (forces AC rows in future data centers, or AC inverters) or simply replace the AC:DC power shelves with 800VDC->54VDC power shelves
- To ensure the most **durable** and **fungible** data centers we need to support a rapidly evolving ecosystem of products, we should unify around LVDC distribution to the IT rows.



Moving Power Conversion Out of the Data Hall





Powering Compute, Storage & Networking

Standard Distribution

48-54 VDC Busbars

Use 48-54 VDC busbars for anything < 100 kW if possible to ensure optimal efficiency and compatibility.

Adapter Flexibility

Modular Power Shelves

Allows power shelves to flexibly provide required adaptation across generations:

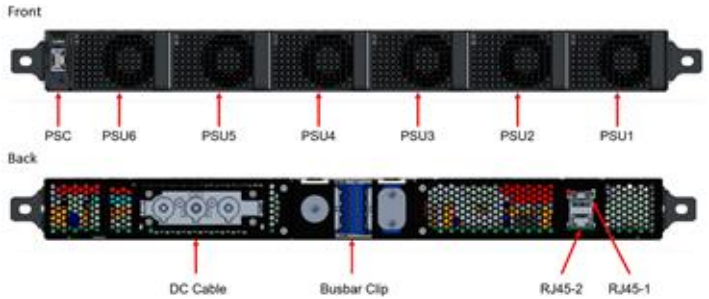
- **Existing Shelves:** Optimized for 415 / 480 VAC
- **New DC/DC Shelves:** Built for 800 VDC architectures

High-Power Loads & Density

2x DENSITY

New designs can consider adapters or 800V direct to trays for high power requirements.

DC:DC Power Shelves: Doubling power density
Increases capacity from **55 kW → 95 kW per RU**

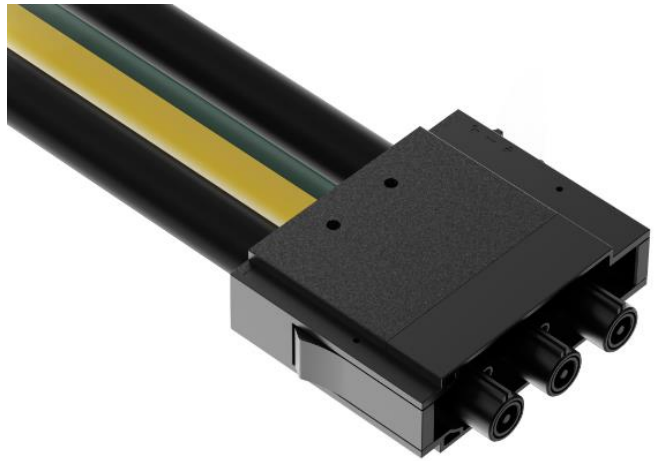


Rack Power Architecture Options

415V AC to 54V DC	800V DC to 54V DC
PAYLOAD 54V Busbar	**PAYLOAD** 54V Busbar

Making everything else LVDC

- Reuse the same standard power whips where possible
- Move CDUs and fans to DC motors - use same input connector and distribution as DC:DC power shelf
- Standard 800 VDC DC/DC bricks are available now
- PDB reference designs are available for cost optimization



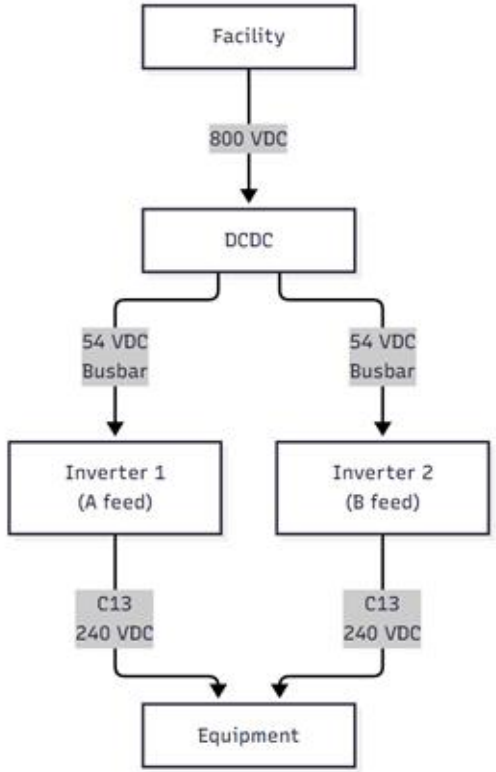


Enabling the last puzzle pieces

- Inverters as a last resort
- 800 VDC -> 54 VDC -> 240 VAC
  - DC/DC shelf + 1RU converter - not elegant but ok for small components
- 800 VDC -> 240 VAC - Might be some uses
  - Crash cart plugs, controls equipment, lighting, etc…

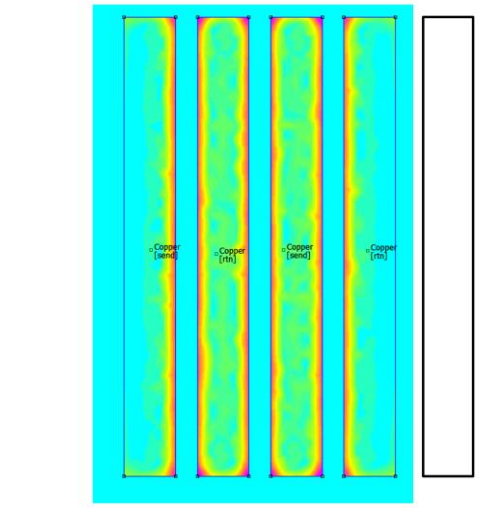


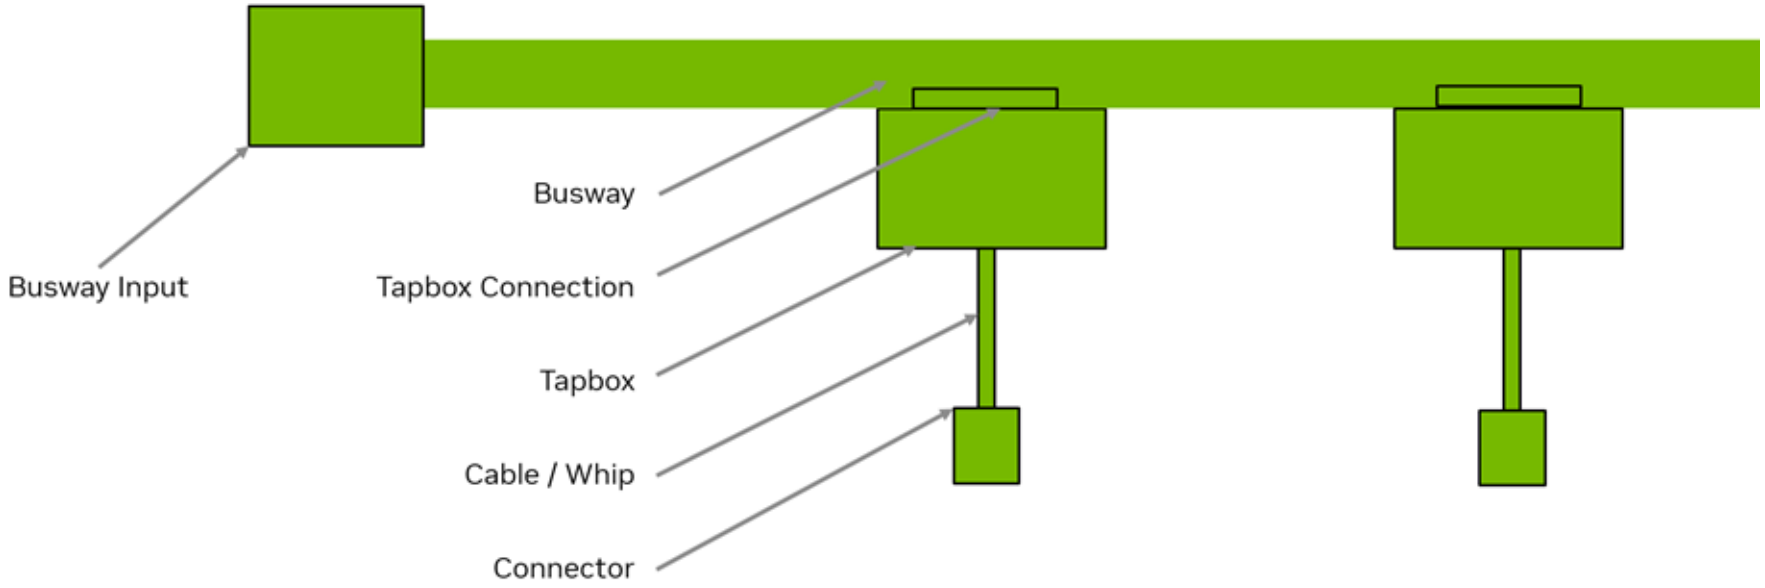




Universal Busway

- Not having interoperability between tapbox and busway causes issues
- If the busway and tapbox don’t land at the same time the datacenter can’t be powered on.
- If common solution can’t be converged on then RPP will be used
- Need to standardize to limited set of connectors to enable rack fungibility





Journey from R&D to Scale

While the LVDC ecosystem does not yet fully exist, Google and Nvidia are driving an open ecosystem from innovation to mass deployment.

01

Tech Incubation

LVDC architecture, SST testing, distribution, protection

Completed

02

Single Lineup Validation

Full single lineup testing (2026)

03

Datacenter Validation

Data center-scale validation (2027–2028)

Upcoming

04

Mass Deployment

Mass deployment of GW's of LVDC native (2029+)

Future Scale

Hyperscalers and Data Center Facility builders are designing and establishing the ecosystem.

Call to Action & Creating Open Standards

The LVDC transition impacts everything from chip to grid, and the pace required to make this change requires open collaboration, debate, and open standards.

- Google, Meta, Microsoft, NVIDIA are introducing new OCP workstreams to standardize:
  - Rack interface connectors
  - DC rack whips
  - DC tapbox
  - DC:DC power shelves
  - DC:DC power conversion modules
- Asking for all consumers and producers of the LVDC technology ecosystem to partner on open high level specifications for interoperability and velocity. Join us!

Thank You!

