

EXECUTIVE BRIEF

PROTECTING THE INVESTMENT

Five principles for repeatable,
resilient **GPU** infrastructure





Executive Summary

GPU infrastructure has changed the economics of the data center. The primary design question is no longer how efficiently a room can cool an average rack. It is how reliably each high-value rack can be protected, cooled, monitored, and deployed across many locations without redesigning the surrounding facility every time.

DDC Solutions moves the operating boundary from the room to the rack. The S-Series Rack Containment Unit (RCU) creates an individualized environment around the IT equipment, combining controlled airflow, NEMA 3R-certified protection against falling water and external liquid ingress, integrated monitoring, and optional rack-level fire suppression. Direct-to-chip liquid cooling (DLC) can be integrated when required, while air-side capacity remains available to manage residual heat and support a controlled response to liquid-cooling degradation.

THE STRATEGIC SHIFT

Standardize the rack environment, then adapt the building around a repeatable interface rather than engineering a unique cooling architecture for every site.

The Five Principles

	Principle	Business Outcome
1	Contain the rack, not the room	Control the IT environment and protect it from external falling-water exposure.
2	Control each rack, not the aisle average	Replace shared HAC behavior with rack-specific airflow and visibility.
3	Isolate fire response at the asset	Add a localized response layer while preserving building life-safety systems.
4	Use a hybrid thermal platform	Preserve CDU choice and air-side resilience as DLC requirements evolve.
5	Deploy deterministically across sites	Convert recurring facility custom engineering into a repeatable RCU deployment model.



① Contain the Rack, Not the Room

Rack containment versus open-air data centers

Open-air data centers expose rack performance to the behavior of the entire room. Supply-air distribution, bypass, recirculation, neighboring loads, room pressure, and facility-level controls all influence the environment seen by the servers. The result can be a large mechanical system that still manages conditions as a room average rather than at the point of heat generation.

An RCU changes the control boundary. The rack becomes a defined thermal and physical zone with its own managed airflow path. Cooling is applied where the heat is produced, and the surrounding room can be treated as neutral space rather than as the primary air-delivery mechanism.

Protection from external water exposure

The S-Series enclosure adds a protection function that an open rack does not provide. Its NEMA 3R certification establishes protection against falling water and external liquid ingress in the certified configuration. In a data center, this creates a physical barrier between the GPU equipment and exposure associated with overhead leaks, condensate, or sprinkler discharge. NEMA 3R should be positioned as documented ingress protection, not as a claim that the enclosure is waterproof or suitable for immersion.

Why the boundary matters

- Predictability: rack inlet conditions are managed locally rather than inferred from room averages.
- Water protection: a NEMA 3R-certified enclosure helps prevent falling water and external liquid from reaching the IT equipment.
- Density independence: high-density racks do not have to be diluted across the room to protect lower-density neighbors.
- Brownfield utility: capacity can be added rack by rack without making room-level airflow the sole constraint.

EXECUTIVE IMPLICATION

The RCU is not an accessory inside the data center. It is the thermal, physical, and water-protection envelope for the compute asset.



② Control Each Rack, Not the Aisle Average

Individualized air containment versus hot aisle containment

Hot aisle containment (HAC) improves airflow separation, but it remains a shared architecture. Multiple racks discharge into a common contained zone, and the performance of that zone depends on the combined rack mix, facility air handlers, pressure balance, ceiling or return path, and commissioning of the room as a system.

Individualized rack containment shortens and isolates the air path. Each RCU can respond to the load and airflow needs of the equipment it contains instead of relying on a uniform aisle strategy for nonuniform workloads.

Criterion	HAC	Individual RCU
Control boundary	Shared aisle and room systems	Single rack environment
Response to mixed loads	Balanced across the aisle	Adjusted at each rack
Monitoring context	Often room or zone level	Rack-specific thermal visibility
Water exposure	Open rack remains exposed	NEMA 3R-certified external ingress protection
Expansion model	Rebalance the shared system	Add a repeatable rack module

EXECUTIVE IMPLICATION

HAC is principally an airflow-efficiency strategy for a shared space. Rack containment is a control and protection strategy for an individual asset. The distinction becomes more important as rack densities diverge, hardware values rise, and deployment schedules leave less time for iterative room balancing.



③ Isolate Fire Response at the Asset

Rack-level suppression versus building-level suppression

Building fire detection and suppression remain essential life-safety and code-compliance systems. They protect the occupied structure and should not be characterized as optional or replaced by the RCU. The design opportunity is to add a more localized layer of response around the rack.

An available integrated clean-agent system within the S-Series can detect and respond inside the enclosure. This creates a distinct first-response zone at the asset while the building system remains the overarching protection layer.

Layered protection model

- Rack layer: localized detection and suppression within the enclosure.
- Building layer: code-required detection, alarm, evacuation, and facility-wide suppression.
- Operational layer: rack-level monitoring and alarms that support intervention and incident review.

THE RISK PRINCIPLE

Localize the initial response where practical, while preserving the independent building life-safety system required by code and the authority having jurisdiction.

Why localization has financial value

A rack-level event should not automatically become a room-level asset event. Local suppression can help constrain the incident domain, reduce exposure to adjacent equipment, and create a clearer boundary for recovery. Actual system performance, agent selection, interfaces, and jurisdictional acceptance must be confirmed for each project.



④ Use a Hybrid Thermal Platform

DLC flexibility, CDU choice, and air-side resilience


The RCU should not be positioned as air cooling instead of liquid cooling. It is a rack platform that can support air-cooled, liquid-cooled, and hybrid IT. This allows the cooling architecture to follow the server requirement rather than forcing every deployment into a single CDU or coolant strategy.

Preserve choice as liquid architectures evolve

- Integrate DLC when the IT platform requires it, without discarding enclosure, monitoring, airflow management, or protection functions.
- Evaluate CDU technologies against project requirements including pressure architecture, controls, coolant compatibility, service model, facility-water interface, and OEM support.
- Avoid making one CDU technology a universal prerequisite unless the commercial or engineering scope explicitly requires it.

Use air-side capacity as a resilience layer

Even liquid-cooled GPU systems retain air-cooled components and reject a residual share of heat to the rack environment. If the technology cooling system (TCS), supply fluid network (SFN), or CDU degrades, increased rack airflow can help remove residual and radiated heat, limit the rate of temperature rise, and support an orderly workload reduction or shutdown.



CRITICAL BOUNDARY

Air-side capacity may extend the response window during a liquid-cooling fault. It does not replace required direct-to-chip cooling and should not be represented as sustaining full liquid-cooled compute load after loss of coolant flow.

Executive implication

The value is optionality plus fault tolerance: select the liquid technology that fits the project, retain active management of the air-cooled heat fraction, and provide operators with a more controlled transition when the liquid system is impaired.

⑤ Deploy Deterministically Across Sites

A repeatable RCU platform versus a unique data center at every location

Traditional high-density deployments often begin by redesigning the room: air-handling capacity, plenums, containment, controls, suppression interfaces, structural clearances, commissioning procedures, and operating sequences are tailored to each building. Repeating that work site by site increases engineering effort, schedule risk, and performance variability.

A deterministic RCU deployment model standardizes the thermal, protection, and mechanical environment immediately surrounding the rack. The facility still must provide qualified power, water, structural support, controls integration, service clearance, and code compliance. The difference is that these become defined interfaces to a repeatable product rather than inputs to a new room-cooling concept at every location.

Financial implications

- Lower site-specific engineering exposure through a consistent rack-level design basis.
- Shorter and more repeatable commissioning through standardized sequences, sensors, alarms, and acceptance criteria.
- Phased capital deployment because capacity can be installed with the racks instead of building the full room around a forecasted end state.
- Improved portfolio consistency across enterprise, colocation, edge, and brownfield sites.
- Clearer comparison of cost per deployed rack and cost per supported kilowatt across locations.

Closing Statement

The investment case for rack containment is larger than cooling. It is the case for placing a repeatable protection, thermal-management, and monitoring boundary around each high-value compute asset. Compared with an open room, the RCU isolates the rack and adds NEMA 3R-certified protection against falling water and external liquid ingress. Compared with HAC, it individualizes control. Compared with building-only suppression, it adds a localized response layer. As DLC evolves, it preserves technology choice while maintaining air-side management and a more graceful fault-response path. Across a portfolio, it turns site-by-site infrastructure reinvention into a deterministic deployment model.

THE DECISION

Do not design every building as a one-off solution for the next GPU generation. Standardize the rack environment so power, water, liquid cooling, and facilities can evolve around a consistent deployment platform.





Protect the asset. Stabilize the environment.
Deploy with confidence.