REFERENCE DESIGN #505

BUILT TO POWER DATA CENTERS AT SCALE

Eaton 1GW Power Architecture Meets Trane Thermal Management

A collaborative reference design that supports strong performance and faster deployment.

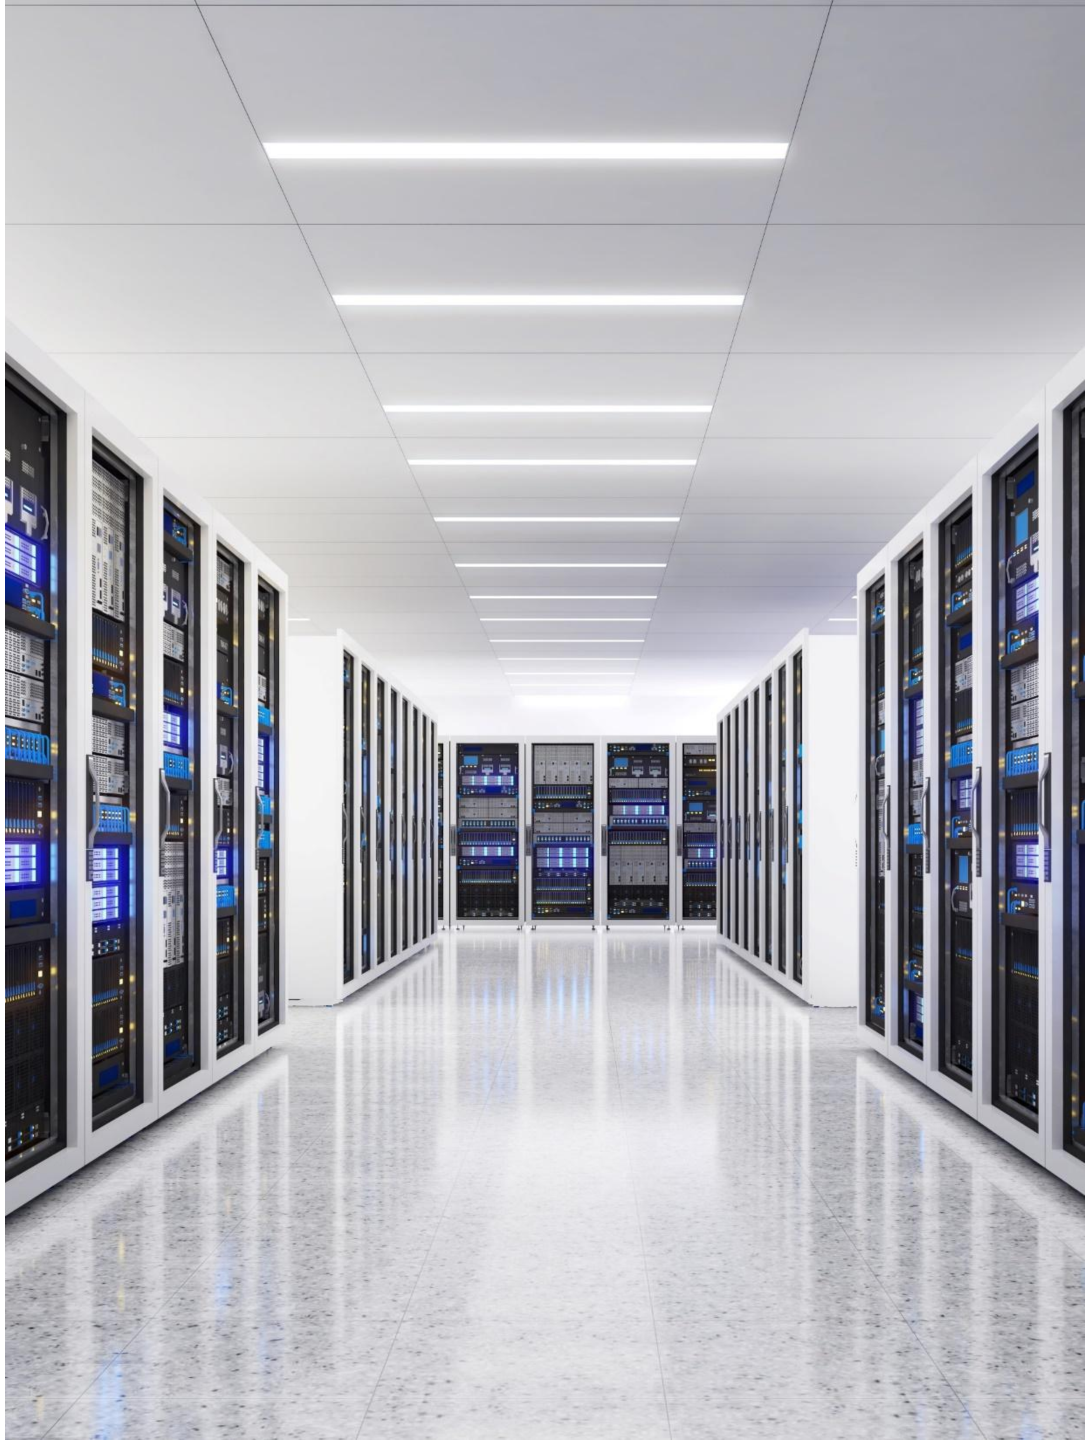

EXECUTIVE SUMMARY

AI data centers are rapidly scaling, with higher rack densities, faster build timelines, and tighter dependence on constrained electrical grids. At gigawatt scale, coordinated reference designs and power-and-cooling co-engineering become essential for predictable performance, high availability, and strong ROI. This white paper outlines a joint Eaton–Trane approach to aligned architectures for AI factories.

Eaton's electrical design maps to Trane's Continuum Rubin DSX Reference Design #501, built around eight repeatable 102.4 MW compute modules. Together, they create a scalable framework for gigawatt deployments that reduces risk, speeds time-to-revenue, and supports phased growth across greenfield or brownfield sites.

Eaton structures AI factory power into three domains:

- IT power, where traditional 480 V AC persists while 800 V DC gains traction for higher efficiency.
- Whitespace cooling power, where expanding liquid cooling distribution units require dedicated ride-through strategies.
- Centralized facility cooling, where medium-voltage distribution efficiently serves higher capacity chiller and compressor motors with controlled starts and predictable fault behavior.

Key trends drive this shift: 800 V DC can boost efficiency by 3–5%; CDU capacity has jumped from ~100 kW to 10 MW in five years; and medium-voltage distribution dramatically cuts motor current (e.g., 1,420 A at 480 V to 164 A at 4.16 kV for a 1,500 HP pump). Together, these illustrate the need for integrated, pre-engineered power-and-cooling architectures for gigawatt-scale AI infrastructure.

EXECUTIVE SUMMARY

In the combined reference design, 4.16 kV distribution feeds outdoor chiller plants, using Eaton Ampgard® motor control for large motors. Distributed cooling loads run at 480 V, and indoor CDU pods now receive Eaton 9395 UPS backup as critical loads. Emerging DC options—such as medium-voltage solid-state transformation and 480 V AC–to–800 V DC systems—outline a path as DC adoption grows. This layered approach boosts efficiency, limits fault exposure, and supports repeatable, campus-scale deployments.

Connectivity and controls play a central role. Trane's thermal controls coordinate chillers, loops, and CDUs, while Eaton Brightlayer™ solutions extend analytics and control across the electrical system. Together, they enable coordinated, data-driven operations that help improve maintenance, extend equipment life, and sustain availability.

Ultimately, future AI data centers will rely on co-designed power and thermal systems, tightly connected assets, and active lifecycle management. At gigawatt scale, success depends on power, cooling, and intelligence operating as one. The Eaton–Trane aligned designs provide a blueprint for more efficient, more available AI factories with higher compute output per unit of energy.

NVIDIA DSX is a data-center-scale AI infrastructure platform that provides a turnkey foundation for building and running large GPU-accelerated AI factories. It integrates compute, networking, and software into a validated, scalable system optimized for training and inference of large models. By combining DGX systems, high-performance networking, and NVIDIA software stacks, DSX simplifies deployment, boosts utilization, and ensures predictable performance as AI workloads grow.

INTRODUCTION

Data center power demand has moved far beyond the traditional tens-of-megawatts scale. Hyperscale AI campuses are progressing toward gigawatt-class facilities, driven by racks consuming 120–600 kW today and trending toward megawatt levels. Forecasts indicate that by 2030, one in five campuses will exceed 1 GW, rising to nearly one in three by 2035[1].

At these loads, replicating legacy power designs is no longer viable. Medium-voltage distribution and 800 V DC architectures are becoming essential enablers of future growth.

Cooling requirements are rising even faster. Liquid-cooled AI systems now depend on CDUs delivering 2–5 MW each and large centralized cooling plants, where a 100 MW compute block can impose 20 MW of motor load and pumps up to 1,500 HP.

Despite these industrial-scale demands—more like semiconductor fabs or midsized cities—many data centers still rely on 480-V rack-level architectures built for 5–20 kW racks. These systems have been pushed to their limits but were never intended for megawatt-class compute, multi-megawatt cooling, or gigawatt-scale campuses.

1. Bloom Energy, *2026 Data Center Power Report: When Power Defines Growth—How Power Availability Is Reshaping the Data Center Industry* (San Jose, CA: Bloom Energy, January 2026), 6.

INTRODUCTION

This design outlines Eaton’s reference power architecture for next-generation AI data centers, built around three interdependent domains:

- IT power in the whitespace, supported by traditional 480 V AC and emerging 800 V DC architectures for higher-density compute.
- Whitespace cooling power, where high-capacity CDUs introduce new critical electrical loads within the data hall.
- Centralized facility cooling, where medium-voltage distribution is required to reliably serve large chiller, pump, and motor systems.

The paper then describes the architectural principles needed for stability and availability at scale—covering UPS roles, medium-voltage motor control, and fault containment—and applies them in a case study aligned with the Trane Continuum Rubin DSX 1 GW AI Factory Reference Designs.

It concludes with how integrated power-thermal communications, controls, monitoring, and lifecycle services strengthen long-term operational resilience. Together, these elements show how coordinated power and thermal architectures improve deployment speed, efficiency, and predictability for gigawatt-scale AI data centers.

THE SCALE OF AI POWER INFRASTRUCTURE

IT power within the whitespace – racks to SuperPODs

Consider the trajectory. NVIDIA's DGX GB200 NVL72 rack, shipping today, consumes approximately 120 kW. The forthcoming GB300 NVL72, in production in 2025, pushes to approximately 140 kW. The Vera Rubin NVL72, arriving in the second half of 2026, will consume 190–230 kW per rack depending on power profile. And the Kyber-based Rubin Ultra NVL576, expected in 2027, targets approximately 600 kW per rack with Nvidia acknowledging that megawatt-class racks are on the visible horizon beyond that.

A single DGX GB200 SuperPOD (eight NVL72 racks) requires approximately 1 MW. A Kyber-era SuperPOD will approach 5 MW. Individual campuses are being planned at 100 MW to 1 GW. An illustration of a POD architecture is shown at right.

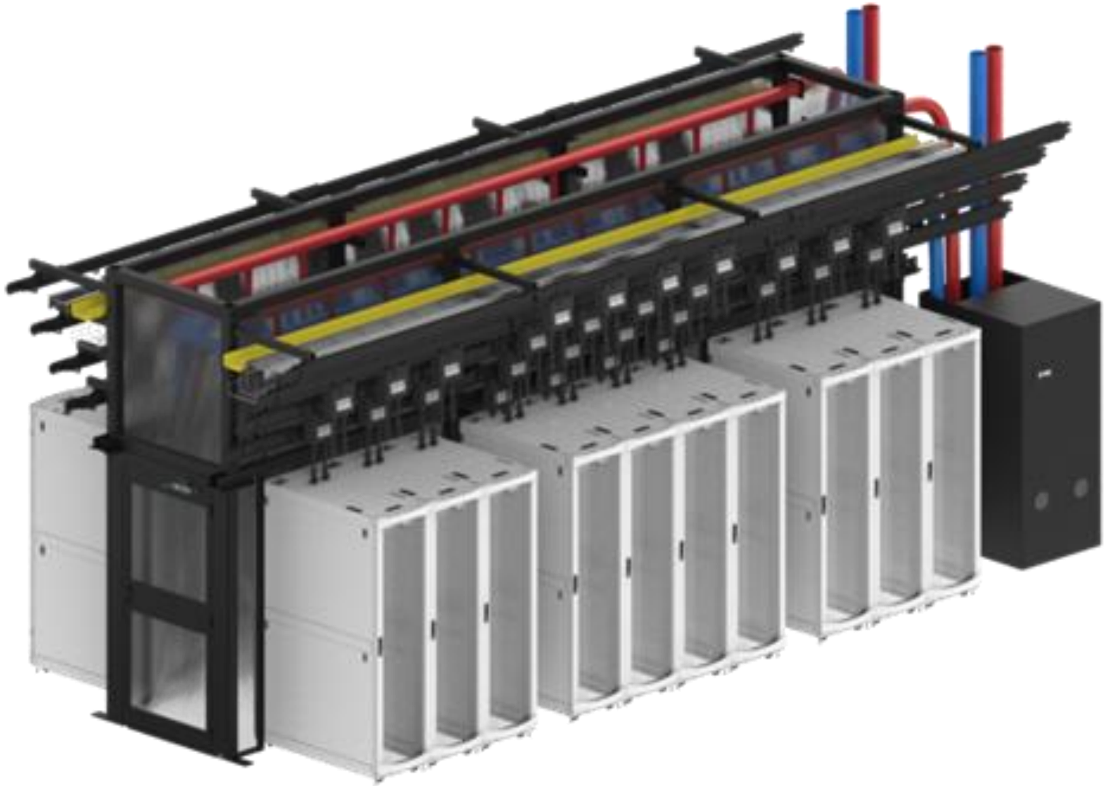

THE SCALE OF AI POWER INFRASTRUCTURE

IT power within the whitespace – racks to SuperPODs

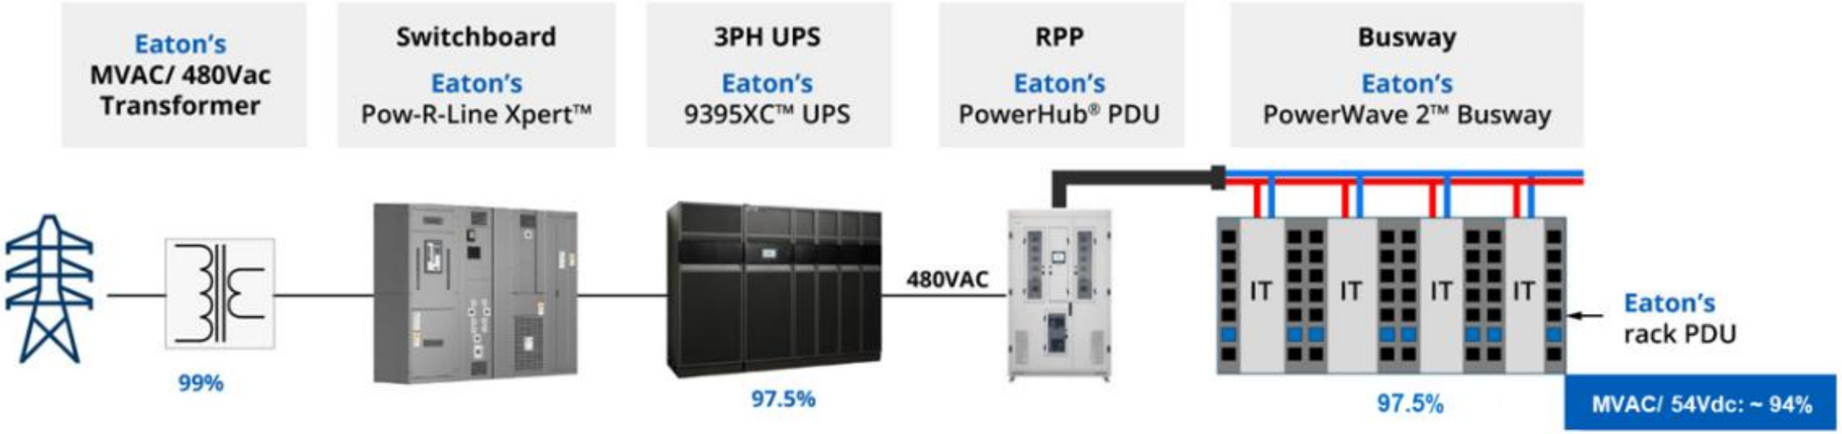

Conventional AC architectures remain the dominant method for powering data center whitespace. In this setup, medium-voltage service (typically 11 kV) is stepped down to 480 V AC, distributed through low-voltage switchgear to centralized AC UPS systems, and finally converted to 48–54 V DC at the rack. Designed for legacy server loads, this approach benefits from a mature ecosystem and proven protection practices. Eaton continues to support it with low-voltage switchgear, breakers, and UPS platforms. However, as rack power rises into the hundreds of kilowatts, this architecture faces growing current, copper, and conversion-loss challenges.

THE SCALE OF AI POWER INFRASTRUCTURE

IT power within the whitespace – racks to SuperPODs

Compute load current analysis

The current draw for compute PODs can be calculated from the fundamental relationship governing conductor sizing and switchgear rating is straightforward: $I = P / (V \times \sqrt{3} \times PF)$ for three-phase AC, and $I = P / V$ for DC. For the calculations below, a power factor of 0.95 and three-phase distribution are assumed for AC loads. Table 2 shows the resulting line currents for representative compute load levels at 480 VAC and 800 VDC.

Load	Representative System	I at 480 VAC	I at 800 VDC	Conversion stages
120 kW	1 × GB200 NVL72	152 A	150 A	3–4 (AC) vs. 1 (DC)
200kW	1 × VR NVL72	253 A	250 A	3–4 (AC) vs. 1 (DC)
600 kW	1 × Rubin Ultra NVL576 (Kyber)	760 A	750 A	3–4 (AC) vs. 1 (DC)
1 MW	1 × SuperPOD (8 × NVL72)	1,267 A	1,250 A	3–4 (AC) vs. 1 (DC)
5 MW	~5 SuperPODs / training cluster or 1 Kyber-SuperPOD	6,334 A	6,250 A	3–4 (AC) vs. 1 (DC)
10 MW	~10 SuperPODs / data hall or 2 Kyber-SuperPOD	12,668 A	12,500 A	3–4 (AC) vs. 1 (DC)

At these power levels, every element of the electrical chain from the utility interconnection to the busbar feeding a pump motor must be evaluated through the lens of industrial power engineering to optimize efficiency, power quality and availability. An additional benefit of 800V DC architecture is that it requires only two wires instead and the reduction in copper cross section and weight savings of up to 34%.

THE SCALE OF AI POWER INFRASTRUCTURE

Whitespace cooling power architecture – CDU growth

Coolant distribution units have undergone a rapid capacity increase driven by the transition from air-cooled CPU servers to liquid-cooled GPU racks.

Period	Max CDU Capacity	Representative Product	Vendor	Context
2018–2021	30–100 kW	In-rack 4U–10U CDUs	Boyd	HPC / early GPU cooling
2022–2023	100–200 kW	Boyd in-rack	Boyd	First-gen AI rack cooling
2024	500 kW–1 MW	In-row L2L CDUs	Multiple	GB200 NVL72 era
2025 H1	1 MW	Trane CDU, LiquidStack CDU-1MW	Trane, LiquidStack	GB200/GB300 NVL72
2025 H2	2 MW	Eaton Boyd ROL4000; Google Project Deschutes CDU (mfg. Boyd)	Eaton Boyd Thermal, Liquid~~s~~Stack	VR NVL72 3°C ATD; 80 PSI; OCP open arch.
2025–2026	5–10 MW	LiquidStack GigaModular	Trane	Full 100 MW block

CDU capacity growth, 2018–2026. Maximum single-unit or modular-stack capacity. Sources: vendor product announcements, OCP Summit 2025, Trane 1 GW Reference Design.

CDU capacity has grown nearly 100× in under five years, from about 100 kW to 10 MW, driven by rising rack power density. A 100 MW hall using GB200 NVL72 racks needs roughly 830 racks at ~120 kW each, all requiring liquid-cooling heat removal. As rack power heads toward 600 kW and higher, CDU capacity will keep scaling.

Google's Project Deschutes—a fifth-generation, open-architecture 2 MW CDU released through OCP—adds redundant power feeds, 25-micron filtration, and IEEE 519-compliant low-harmonic VFDs, highlighting how CDUs are rapidly becoming high-tech, mission-critical infrastructure.

THE SCALE OF AI POWER INFRASTRUCTURE

Centralized facility cooling – the MV opportunity

The chiller plant is the largest concentrated electrical load in a data center cooling system. In the Trane 1 GW AI Factory Reference Design, each 100 MW cooling block uses 10 CenTraVac CVHH chillers of about 3,000 tons (10–11 MW thermal) with power feeds from 400 V to 12 kV, plus roughly 20 large chilled- and condenser-water pumps.

Parameter	480 VAC	4.16 kV	Reduction vs. 480V
Motor rating	1,500 HP (1,119 kW)	1,500 HP	—
Line current / motor	1,420 A	164 A	4.8× / 8.7×
10 motors (one loop)	14,200 A	1,640 A	4.8× / 8.7×
20 motors (both loops)	28,400 A	3,280 A	4.8× / 8.7×
Conductor sizing	Multiple paralleled runs/motor	Single run/motor	Dramatic reduction
Switchgear	LV MCC	Compact MV MCC	30–50% footprint reduction

Current analysis for these pump motors at two voltage levels. The implications for conductor sizing, switchgear footprint, and installation cost are significant.

At 480 V AC, twenty 1,500 HP motors draw 28,400 A total; at 4.16 kV, they draw just 3,280 A—an 8.7× reduction. Beyond smaller conductors, 4.16 kV distribution removes the need for multiple paralleled cable runs, cuts cable-tray volume, improves efficiency, and allows a single Eaton Ampgard MV motor control center to serve the full pump array with capacity to grow.

THE SCALE OF AI POWER INFRASTRUCTURE

Centralized facility cooling – the MV opportunity

Eaton Ampgard® for cooling plant applications

Eaton's Ampgard® medium-voltage motor control centers cover 2,300–15,000 V and motor loads from roughly 50 to 8,000 HP. They support FVNR, RVR, RVAT, and RVSS starting methods and integrate with MV drives such as Eaton's SC9000. The arc-resistant Ampgard® AR design, certified to IEEE C37.20.7-2017 (Type 2B), withstands arc-faults up to 50 kA for 0.5 s, enabling safer operation and energized maintenance in mission-critical chiller plants.

Eaton's medium-voltage Ampgard® starters provide reliable control and protection for large hyperscale loads such as chillers, pumps, and other plant equipment. With front access and coordinated protection using vacuum contactors, current-limiting fuses, and isolation devices, they support safe, predictable maintenance and reduce repair time. Paired with Eaton MV switchgear and breakers, they enable clean transitions, selective coordination, and consistent operation across large campuses.

Eaton Ampgard medium-voltage motor control (2.5, 5.0, 7.2 kV)

ARCHITECTURES FOR STABILITY & IMPROVING AI FACTORY AVAILABILITY

Availability as the primary design objective

At scale, the primary objective of power and cooling infrastructure is to sustain continuous IT availability while operating as a predictable and compliant grid participant. AI factories function as large, tightly coupled electrical systems in which disturbances can propagate rapidly across compute and cooling domains. Power architectures must therefore be designed from the IT load outward, ensuring that compute remains operational through grid events while complying with interconnection and grid code requirements in regions such as Europe and ERCOT served markets in Texas. Stability, power quality, and controlled fault behavior are foundational requirements for both operational uptime and regulatory acceptance.

ARCHITECTURES FOR STABILITY & IMPROVING AI FACTORY AVAILABILITY

Whitespace power architecture and critical load segmentation

Within the whitespace, UPS systems isolate sensitive IT loads from grid disturbances. Eaton's double-conversion UPS platforms, including the 9395 and 93PM, provide continuous voltage regulation, harmonic mitigation, and fast response to rapid compute load changes, helping maintain grid-compliant voltage, frequency, and ramp rates.

As liquid cooling becomes mandatory for high-density compute, whitespace cooling is now mission-critical. CDUs must stay powered to prevent rapid thermal spikes, so they are increasingly placed on dedicated UPS systems separate from IT UPSs. Eaton's UPS platforms provide coordinated ride-through and controlled recovery during brief grid events without creating sudden electrical step changes.

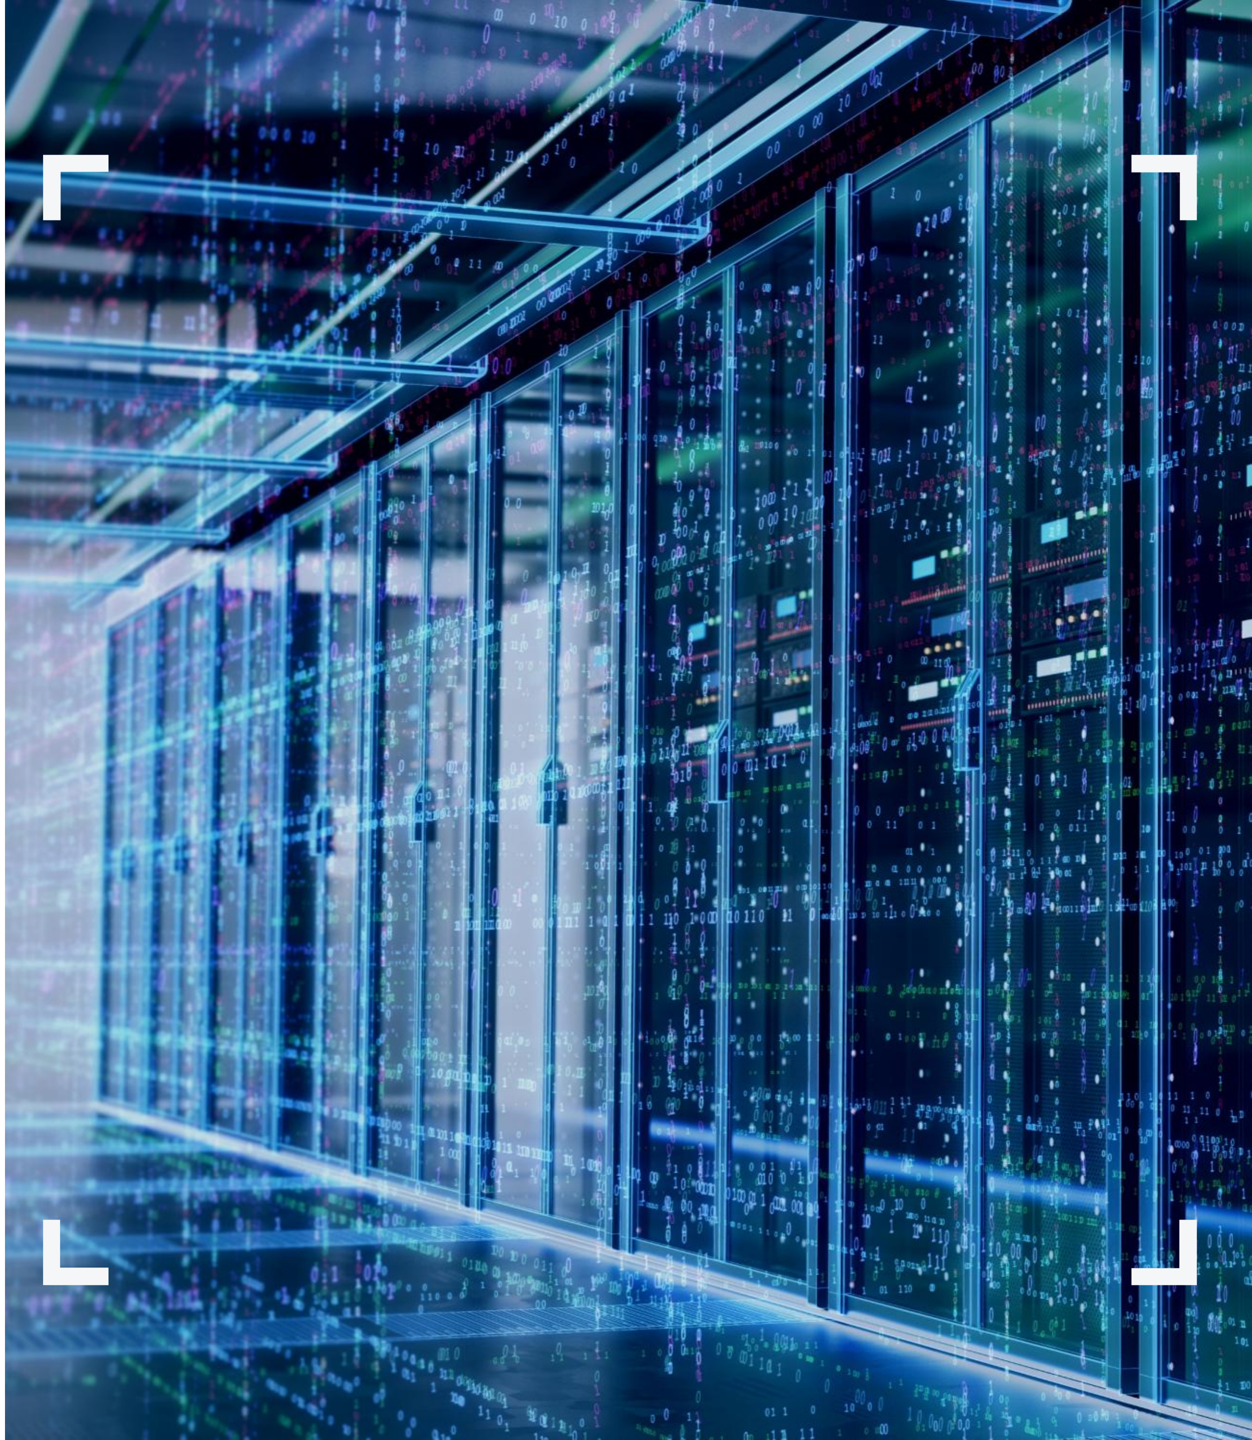

ARCHITECTURES FOR STABILITY & IMPROVING AI FACTORY AVAILABILITY

Facility-level stability and medium voltage power architecture

Facility-level stability depends on managing large motor loads in a grid-compliant, electrically robust way. Centralized cooling plants add high-inertia equipment, large starting currents, and complex fault dynamics that must be coordinated to prevent voltage sags or utility-side disturbances. Medium-voltage architectures are essential for lowering current and containing faults.

Eaton’s VacClad® medium-voltage switchgear provides strong fault isolation, arc-resistant protection, and selective coordination, keeping disturbances confined to a cooling block rather than impacting the wider campus or grid. Downstream, Ampgard® medium-voltage motor control centers enable controlled starting of large pumps and chiller compressors, and integrate with SC9000 drives for smooth acceleration, reduced electrical stress, and improved power quality. Together, these systems support compliance with European grid codes and ERCOT requirements for voltage stability and disturbance management.

ARCHITECTURES FOR STABILITY & IMPROVING AI FACTORY AVAILABILITY

Power quality, ride through, and grid interaction

Cooling systems are mission-critical, and uninterrupted power is essential for IT uptime. Eaton UPS systems provide clean, continuous AC power to both IT equipment and key cooling loads such as multi-megawatt CDUs. In hyperscale facilities, UPS ride-through is vital—it absorbs brief utility disturbances without starting generators, helping reduce wear and improving overall stability while protecting sensitive infrastructure.

Eaton's EnergyAware™ functionality expands the UPS role by enabling grid-interactive operation without compromising reliability. It allows operators to use UPS batteries for peak shaving, demand response, and frequency regulation while preserving energy for ride-through and outages. In sites with large, staged medium-voltage cooling plants, this helps maintain stable voltage and frequency, preventing disturbances from affecting cooling systems. With critical-load protection always prioritized, EnergyAware™ supports higher efficiency, sustainability, and lower total cost of ownership while reinforcing the electrical stability needed for uninterrupted cooling.

ARCHITECTURES FOR STABILITY & IMPROVING AI FACTORY AVAILABILITY

Coordinated system operation

These layers must operate as a coordinated system, not as isolated components. Eaton's portfolio enables power conditioning, fault containment, motor control, and monitoring to work together across IT, whitespace cooling, and facility infrastructure. By pairing UPS decoupling at the IT and CDU levels with medium-voltage distribution and motor control at the facility level, the architecture keeps large-scale AI infrastructure running through grid events while presenting a stable, compliant load to the utility. In this model, high AI-factory availability and strong grid stability are achieved through aligned power, cooling, and control systems.

Further information is available in the Eaton paper Good Grid Citizen available here: [Good grid citizen].

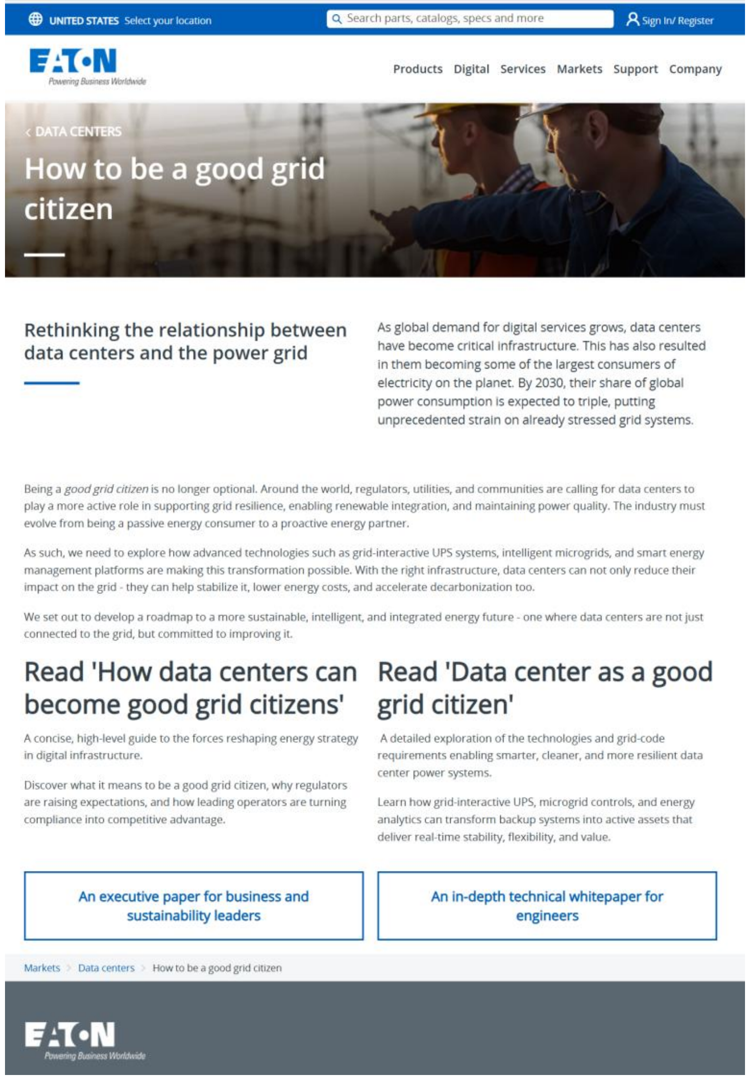

TRANE REFERENCE DESIGN #501 & #502

Introduction

This section applies the earlier architectural principles through a case study showing how Eaton's reference power architecture aligns with the Trane Continuum Rubin DSX 1 GW AI Factory Reference Design. The Trane model offers a large-scale example of how modular, high-power cooling systems are supported electrically through medium-voltage distribution, coordinated motor control, and power-quality measures that protect IT availability and grid stability.

Trane's publication High-Efficiency Cooling for the 1 GW AI Factory (Reference Design #501) defines a scalable, cooling-focused framework for gigawatt-class AI sites. It was later followed by Reference Design #502, which optimized the original cooling system design to increase efficiency, reduce peak cooling power requirements, and enable a smaller dry cooler array. Eaton's reference power architecture provides a complementary electrical framework, illustrated here by mapping it onto the Trane cooling topology. Together, they show how high-power electrical and large-scale thermal systems can be integrated into a coordinated, next-generation AI data center architecture.

TRANE REFERENCE DESIGN #501 & #502

Reference Design #501 and #502 is based on water-cooled chiller systems that offer significant performance and sustainability advantages over air-cooled alternatives, delivering higher efficiency, reduced system footprint, and improved thermal performance for large-scale AI data centers.

A dual-temperature chiller plant with dedicated medium- and high-temperature loops enables maximization of free cooling hours.

A dual-temperature swing chiller configuration allows redundant units to be flexibly assigned to either loop, reducing the total number of chillers required when full dual-loop redundancy is not necessary.

DESIGN OVERVIEW

Reference Diagram

- ✓ 1 GW mechanical cooling blueprint for AI factories
- ✓ Integrates air- and liquid-cooling systems
- ✓ Modular structure for scalability and serviceability

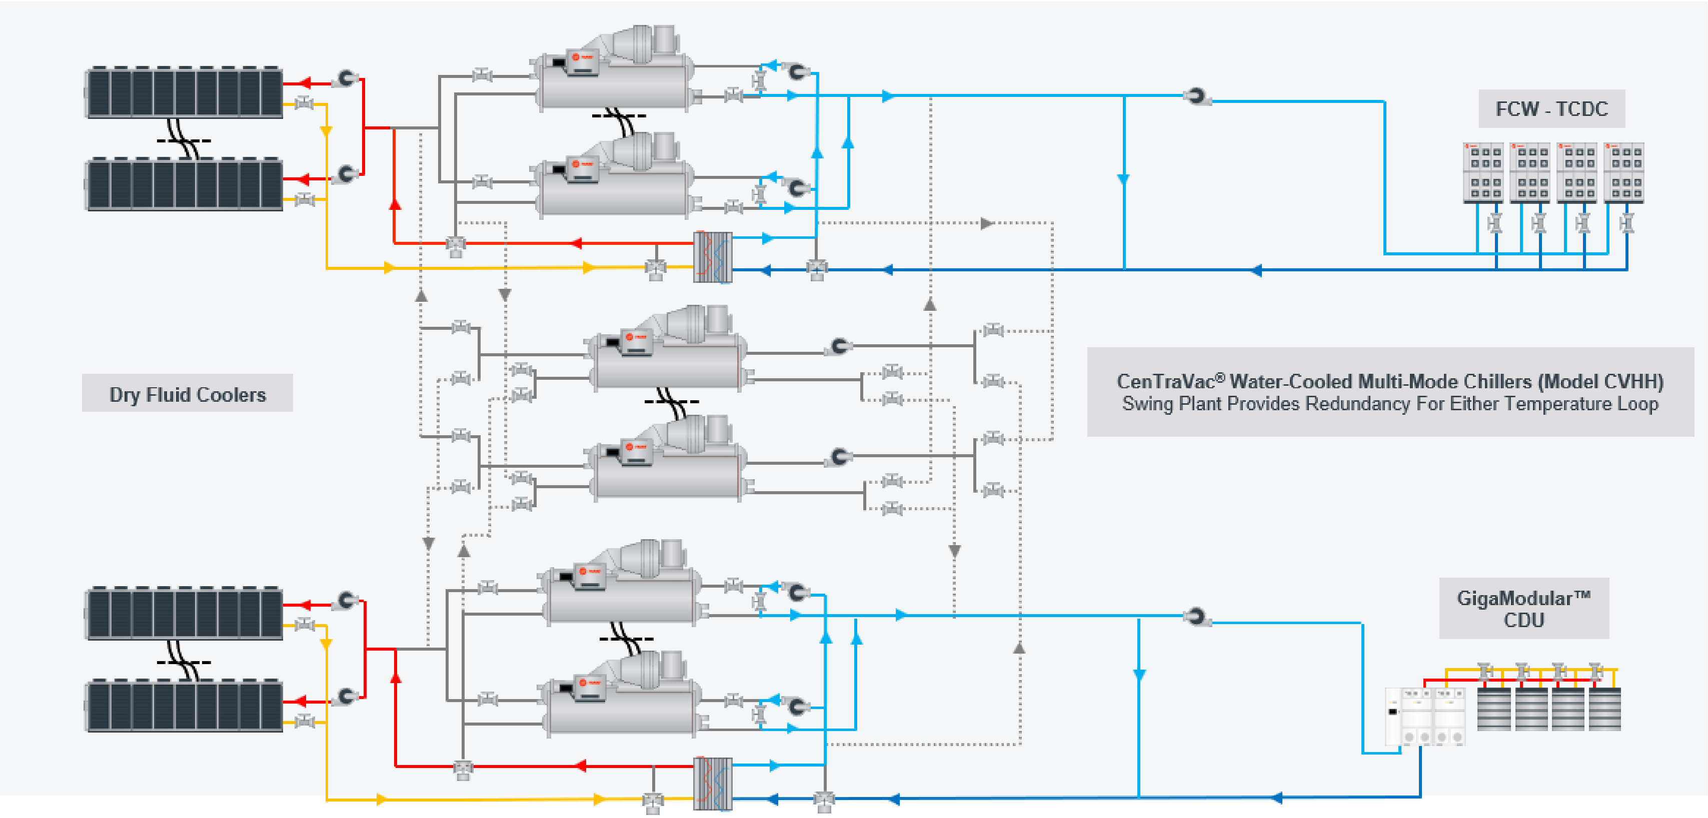

Trane Reference Design #502

TRANE REFERENCE DESIGN #501 & #502

Reference Design #501 divides the 1 GW AI factory into 102.4 MW cooling blocks and recommends high-capacity centrifugal chillers for their efficiency and reduced unit count. Fewer chillers shrink the footprint and simplify commissioning. While Trane's CenTraVac® line offers low-voltage models, higher-capacity units require medium voltage.

Most chiller components—heat exchangers, compressor housings, impellers—are unaffected by using 4,160 V power. Because current drops as voltage increases, medium-voltage chillers draw far fewer amps, allowing smaller wiring, lower copper use, and improved sustainability.

BEST PRACTICES

Block Sizing

Designing a single chiller plant for a 1 GW data center isn't practical due to the massive pipe, valve, and pump sizes required. A better solution is to divide the load into 102.4 MW partitions, typically using 48-inch piping based on system design.

- Each partition is served by 10 - 9.6 MW CDU pods and 1 - 6.4 MW CDU pod
- Each CDU has N+1 Redundancy
- Each totals 102.4 MW and includes 512 (200-kW) racks.
- 9.6 MW Deployment : 4 rows, 12 *200-kW racks per row
- 6.4 MW Deployment: 4 rows, 8 *200-kW racks per row
- Remaining site power reserved for airside cooling loop, additional IT loads and other site needs

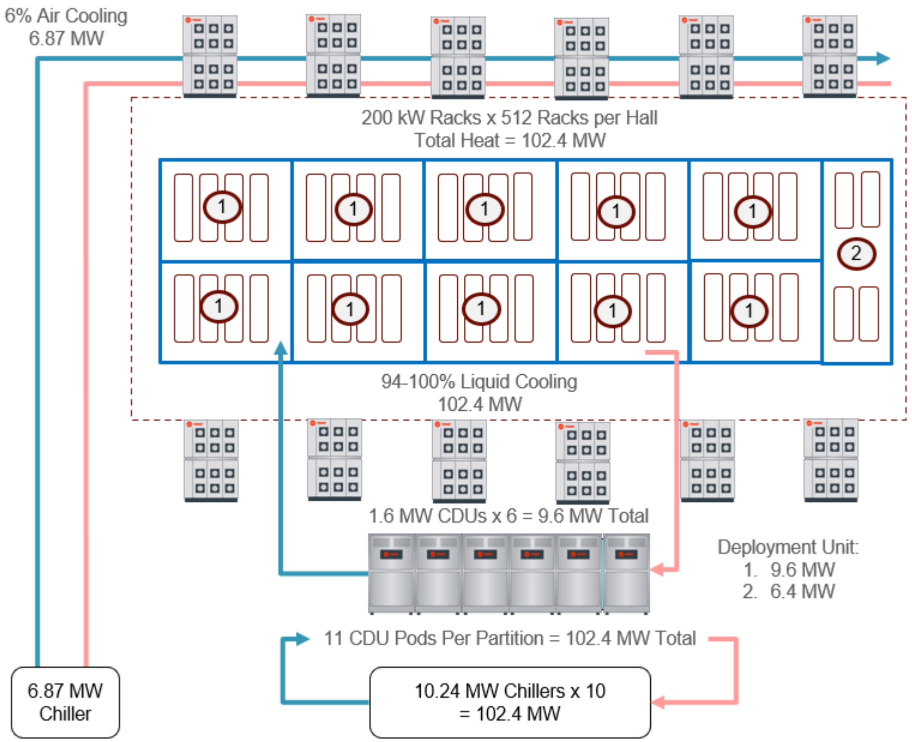

Trane 1 GW AI Factory Reference Design 17

Each 102.4 MW block still needs about ten large centrifugal chillers. At this scale, unloading is managed through staging and inlet guide vane control, enabling reduced-voltage starters instead of medium-voltage VFDs—cutting cost and complexity while maintaining high efficiency.

TRANE REFERENCE DESIGN #501 & #502

Higher voltage lowers amperage, so medium-voltage chillers use smaller wiring, reducing copper and improving sustainability. About ten centrifugal chillers are needed per 102.4 MW block. Staging and inlet guide vanes provide efficient unloading, allowing reduced-voltage starters instead of medium-voltage VFDs—helping to cut cost and complexity while maintaining efficiency.

Efficiency gains enabled by combining Trane cooling systems with Eaton power architecture include:

- ✓ Medium-voltage operation is ideal for Trane data center chillers above 1,000 HP. Supplying these motors at medium voltage can cut electrical equipment and installation costs by 15–30% and reduce copper use by more than 80% compared to low-voltage designs.

- ✓ Infrastructure is also simplified: eliminating MV-to-LV transformation stages and using smaller conductors reduces material needs and streamlines layouts.

- ✓ With fewer conversion steps and lower distribution losses, medium-voltage systems can improve overall energy efficiency by 1–5%, depending on site design and operating conditions.

1GW AI FACTORY COMPUTE QUADRANT LAYOUT

Designing a single centralized chiller plant to serve a full 1 GW data center is neither practical nor desirable.

Reference Design #501 & #502 adopt a distributed cooling architecture, partitioning the total load into 102.4 MW data hall blocks, each served by a dedicated cooling plant and associated infrastructure.

This modular strategy limits hydraulic and electrical fault domains, improves redundancy, and enables predictable expansion as additional compute quadrants are brought online.

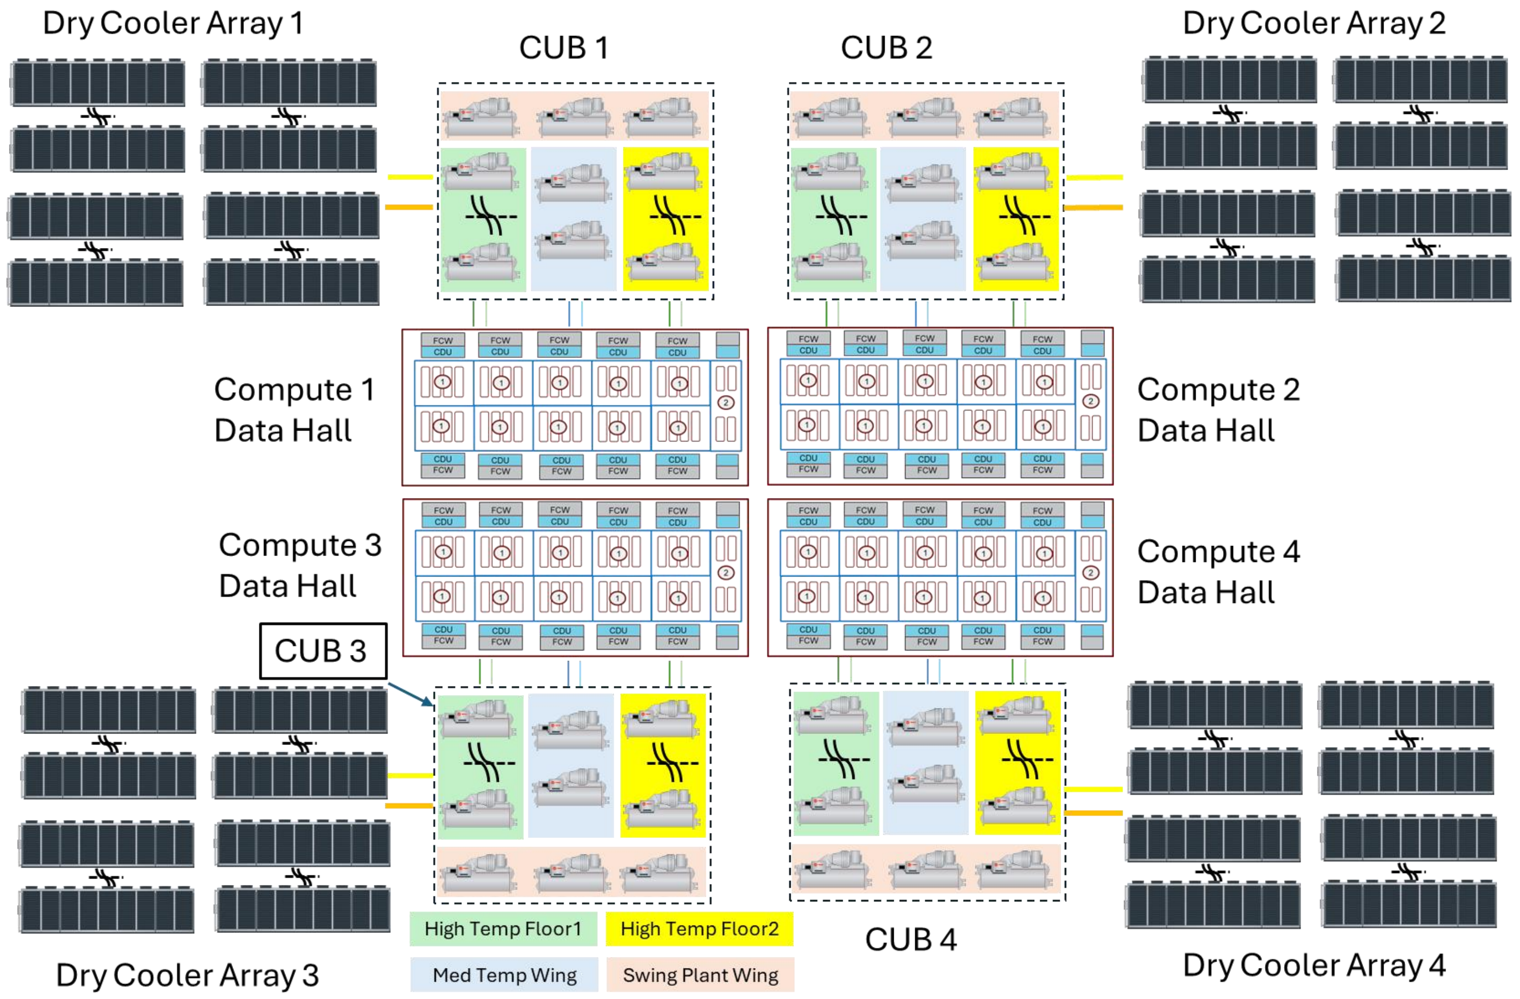

THERMAL MANAGEMENT LAYOUT FOR COMPUTE QUADRANTS

The figure shows a compute-cooling quadrant layout with four compute halls arranged around central utility buildings, each with its own chiller plant and dry coolers. This configuration naturally supports medium-voltage distribution to each cooling block, along with localized motor control, protection, and power-quality management. The result is a scalable, fault-tolerant topology suited for predictable gigawatt-scale deployment.

Each cooling block maps directly to a compute quadrant made of two stacked data-hall levels producing a combined 102.4 MW IT load. This one-to-one pairing creates a repeatable modular unit that can be replicated across the site, aligning compute density, cooling capacity, and electrical infrastructure while simplifying design and phased expansion.

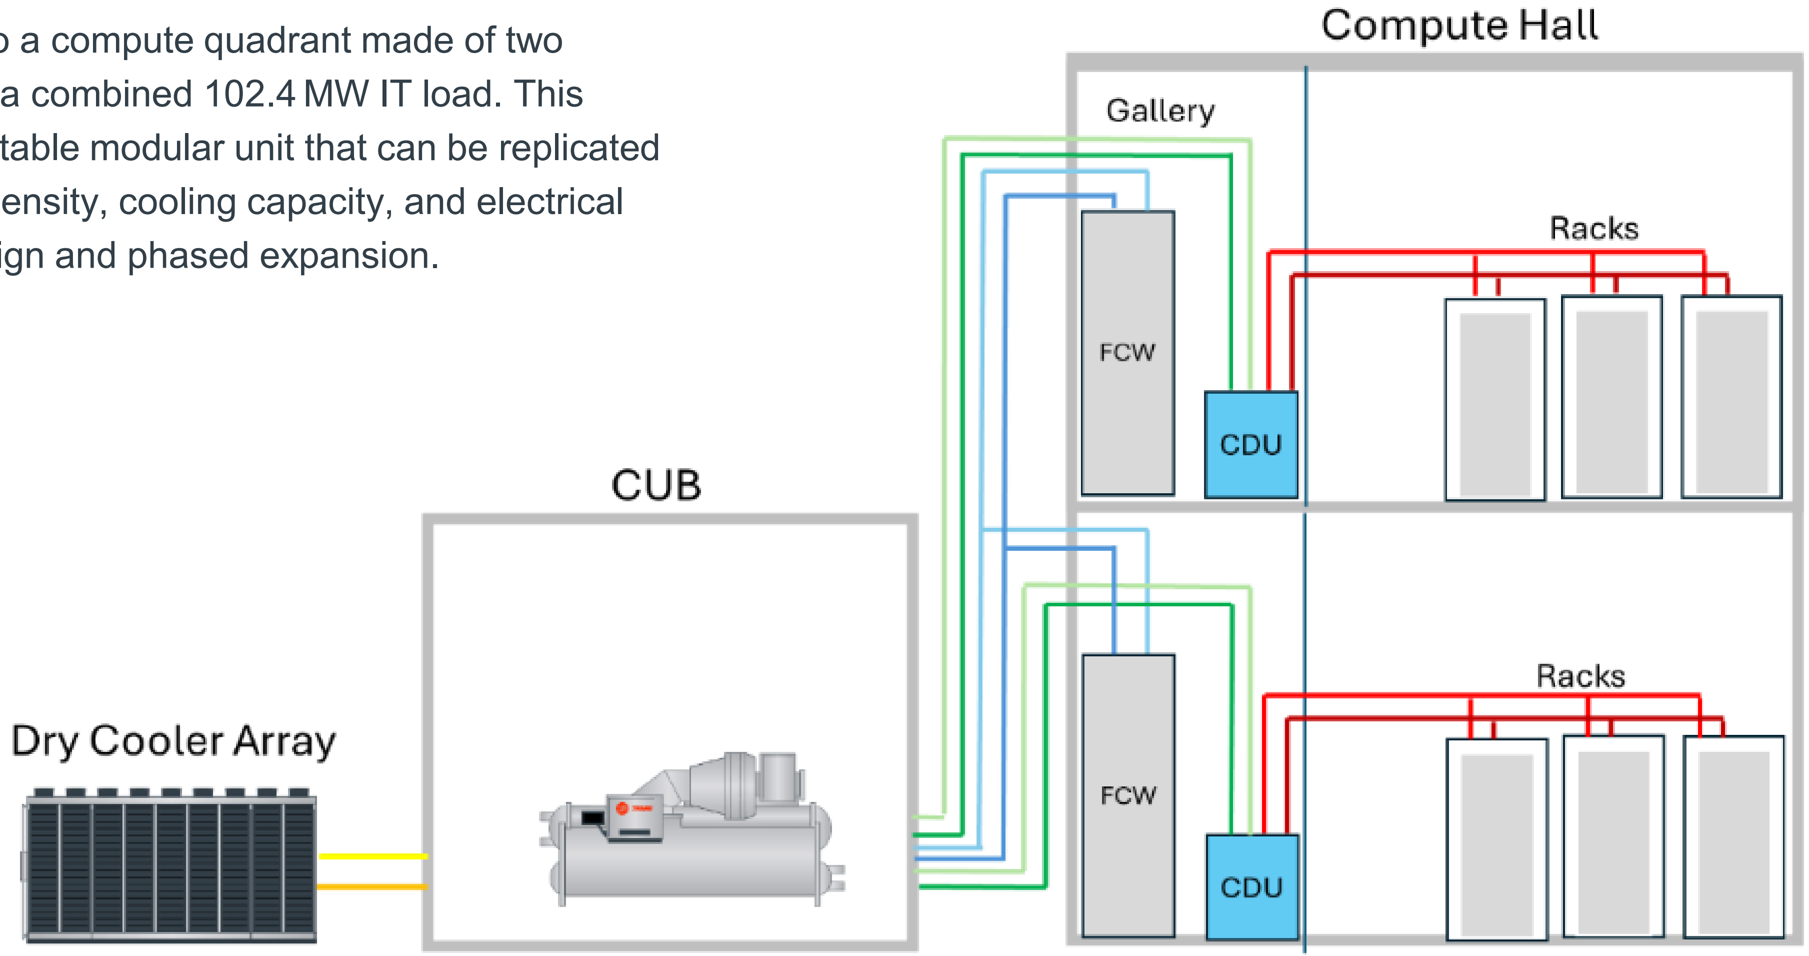

EATON COOLING PLANT POWER ARCHITECTURE

The electrical distribution system supports each 102.4 MW cooling block as its own repeatable domain, aligned with the modular layout of the 1 GW AI factory. Medium-voltage 4.16 kV distribution powers the outdoor chiller plants, efficiently serving large motor loads while reducing current, conductor size, and fault energy. Eaton Ampgard® medium-voltage motor control starts, protects, and coordinates the chiller motors, providing controlled operation and localized fault containment for high-inertia, mission-critical equipment.

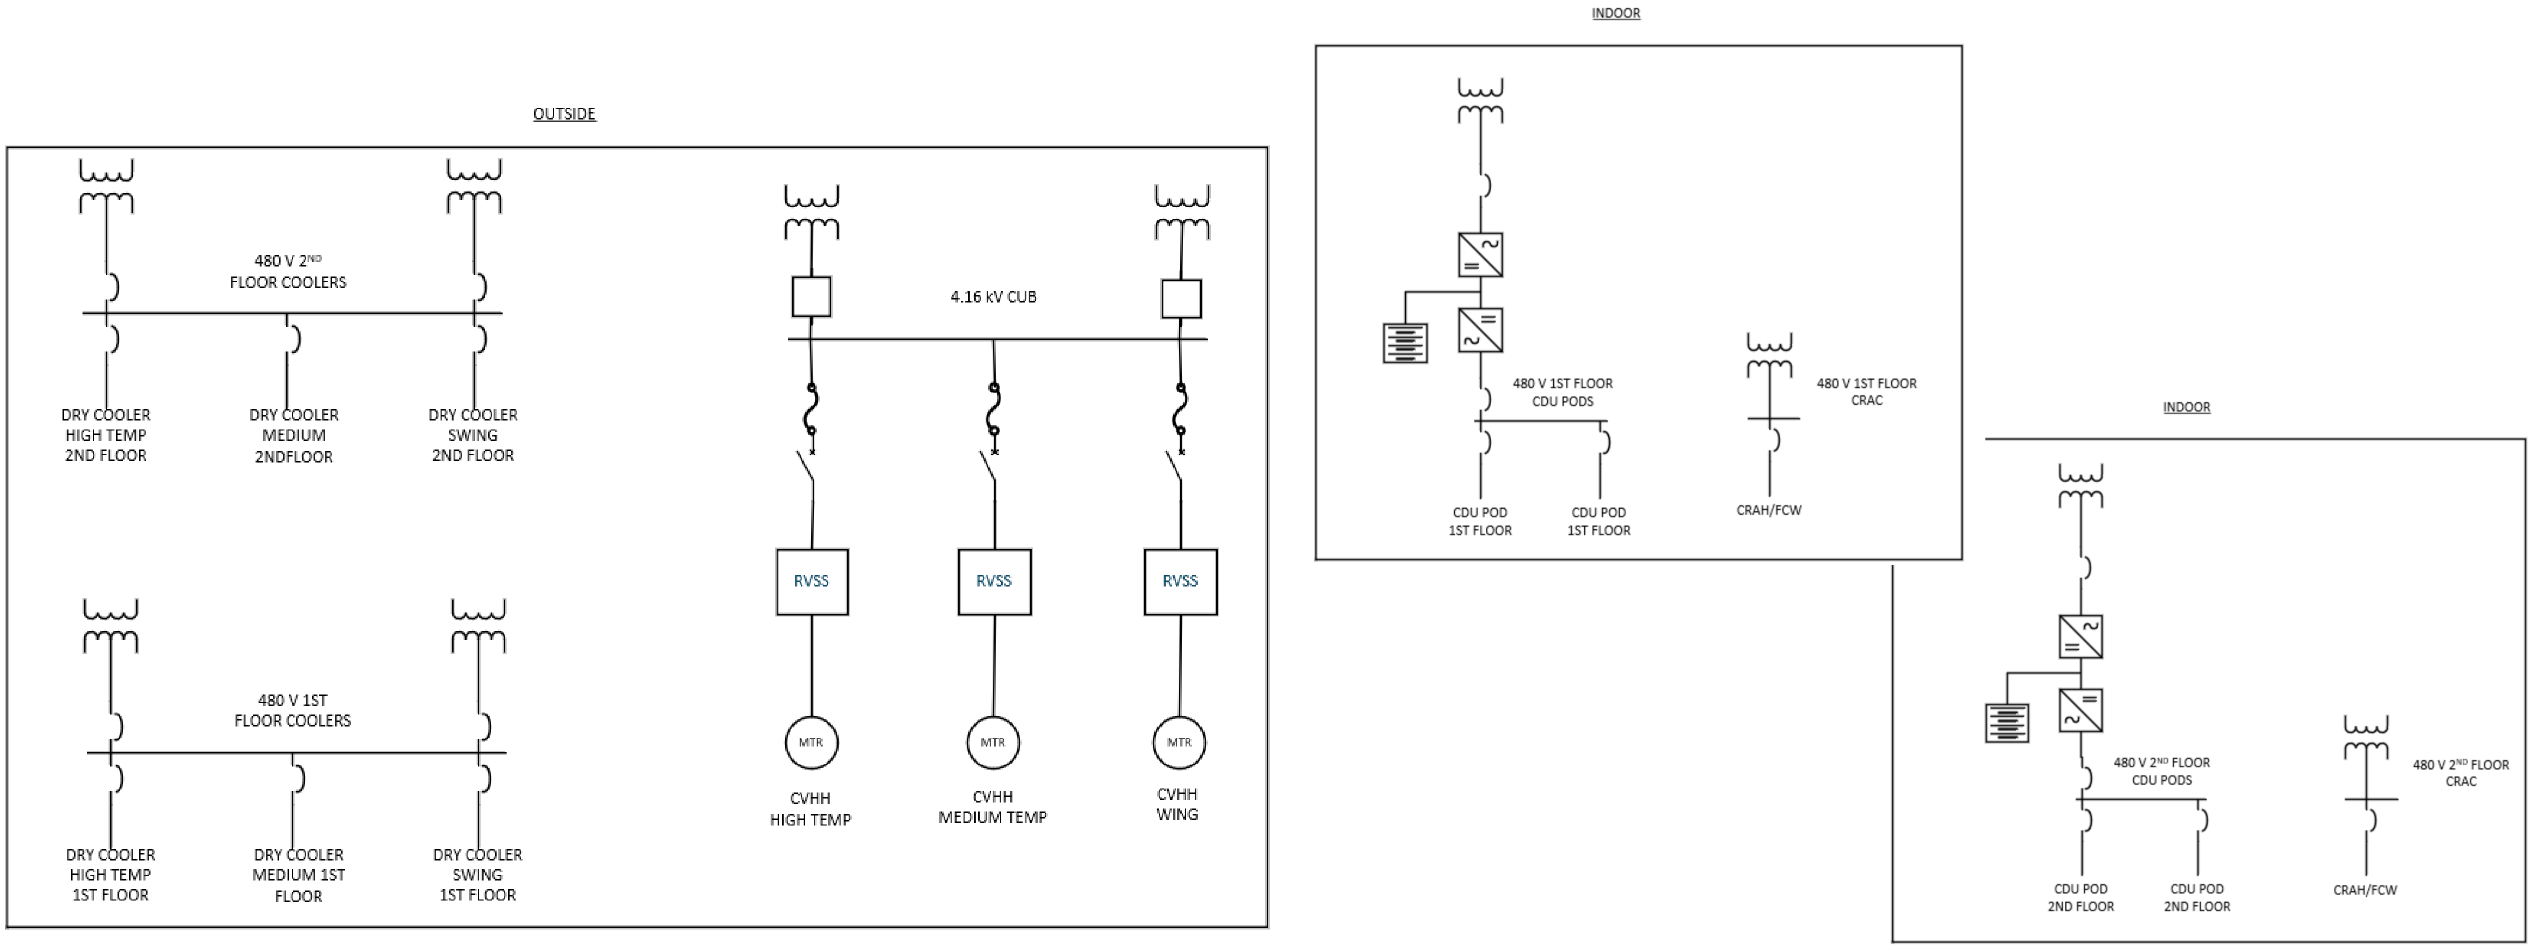

EATON COOLING PLANT POWER ARCHITECTURE

Downstream of the chiller plants, additional cooling equipment is served at low voltage. Outdoor dry coolers run from 480 V switchboards to match their distributed fan loads, and CRAH/FCW units inside the facility are likewise fed from 480 V distribution for standardized installation. Indoor CDU pods also operate at 480 V but are backed by a dedicated Eaton 9395 UPS, ensuring liquid-cooling continuity during brief power disturbances and protecting rack-level thermal stability.

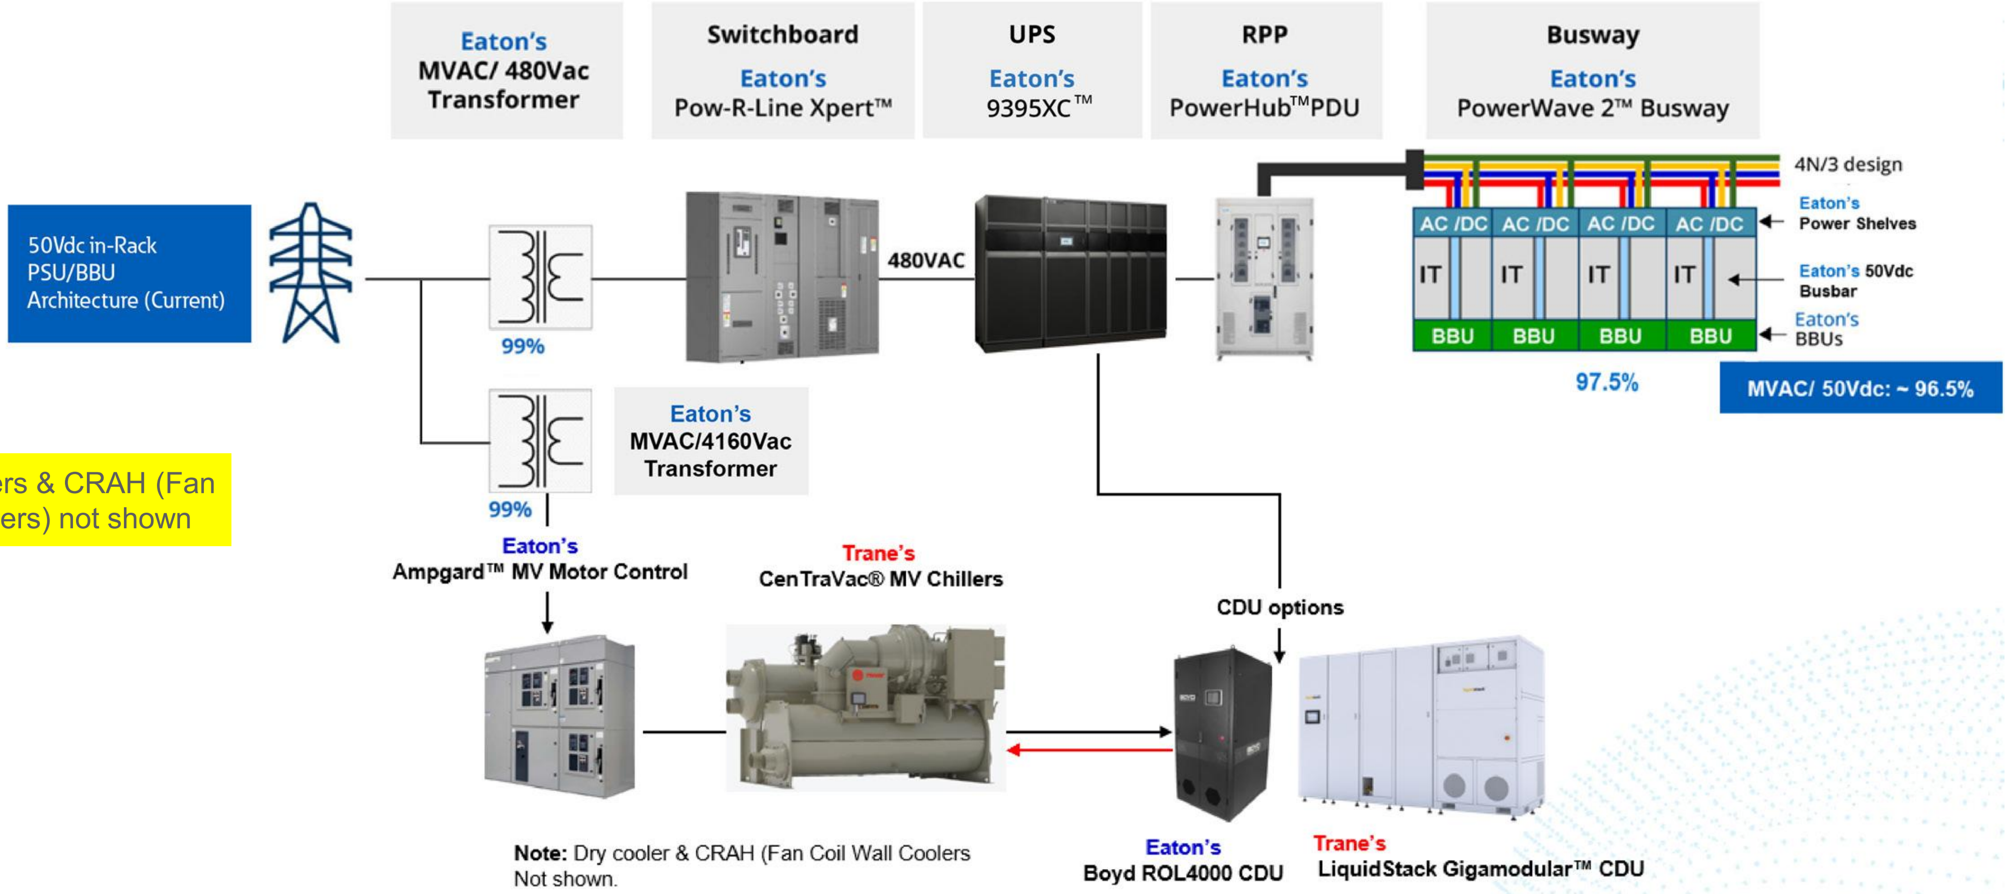

Note: Dry coolers & CRAH (Fan Coil Wall coolers) not shown

POWER-COOLING ARCHITECTURE LAYOUT FOR EATON-TRANE REFERENCE DESIGN

Product Types		Product Relationship and Specs		
Medium Voltage	Pad mount transformers	Oil filled transformer	Oil filled transformer	Oil filled transformer
	Generator (Siemens)	Natural gas generator	Natural gas generator	Natural gas generator
	SW Gear	MVS (Metal enclosed) switchgear	MVS (Metal enclosed) switchgear	MVS (Metal enclosed) switchgear
		GIS or RMU	GIS or RMU	GIS or RMU
	MCC	MCC (Mini voltage starter)	MCC (Mini voltage starter)	MCC (Mini voltage starter)
Chillers	Application and qty	CVHH (medium temp) qty 1	CVHH (high temp) qty 10	CVHH (swing) qty 3
	Cooling capacity	1938 Tons (6.87MW)	2912 (10.24MW) each	2912 (10.24MW) each
	MVoltage/Frequency	4160V, 60 Hz	4160V, 60 Hz	4160V, 60 Hz
	Circuit motor order RLA	172.9 A	243.4 A	243.4 A
	Circuit line-side MCA / MOP	216 A/ 230 A	304 A / 450 A	304 A / 450 A
Switchboards	Product line	Pow-R-Line Xpert (PRLX)	Pow-R-Line Xpert (PRLX)	
UPS, 3Ph	Product line	Eaton 9395XS	Eaton 9395XS	
	Voltage	480V, 3Ph	480V, 3Ph	
	Capacity	1125-2250kW	1125-2250kW	
CDU and CRAH/FCW	Product	CRAH/FCW	CDU	
	Capacity, qty	256kW, qty 30	9.6 MW, qty 10	
	V input	460V, 3Ph	460V, 3Ph	
	FLA / MCA / MOP	32.2 A / 33.5 A / 35 A Disconnect with manual motor protector per fan (6 total)	363 A / 385.5 A / 450 A Molded case breaker	
	Capacity, qty		6.4 MW, qty 1	
	FLA / MCA / MOP (w/ Molded CCB)		273 A / 295.5 A / 350 A	
Dry coolers	Application and qty	Medium temp (72.6F) plant: Qty 8 (per wing)	High temp (78.8f) plant: Qty 52 (per floor)	Swing plant: Qty 4 (per wing)
	V input	460V 3Ph	460V 3Ph	460V 3Ph
	FLA / MCA / MOP	155.2 A / 157.6 A / 175 A	155.2 A / 157.6 A / 175 A	155.2 A / 157.6 A / 175 A

CONNECTIVITY: POWER, COOLING, AND SERVICES

Integrated connectivity as a foundation for resilience

As AI data centers grow in power density and complexity, long-term resilience depends on tightly connected power and thermal systems. These layers can no longer operate independently—they must function as an integrated system that enables continuous availability, proactive health management, and predictable response to disturbances. Connected infrastructure lets operators shift from reactive operations to coordinated control, condition-based maintenance, and lifecycle optimization while preserving the deterministic behavior required for mission-critical sites.

Connectivity now goes beyond monitoring. It includes real-time communication between electrical and thermal assets, system-level controls that coordinate subsystems, and service platforms that turn operational data into actionable insight. When built into the reference architecture, these capabilities help strengthen uptime, reliability, and risk management as AI factories scale to gigawatt-class campuses.

BRIGHTLAYER™ DATA CENTER SOLUTIONS

Eaton's Brightlayer™ Data Center Solutions provide the digital foundation for connecting electrical infrastructure with monitoring, analytics, and control. Brightlayer enhances connected power assets with visibility, insights, and predictive capabilities that support the most critical operator outcomes: safety, risk reduction, efficiency, operational resilience, and sustainability.

By integrating data from UPS systems, switchgear, motor control, and power distribution equipment, Brightlayer enables continuous assessment of system health and performance. Operators gain real-time visibility into power quality, load behavior, and asset condition, enabling faster diagnosis of abnormal conditions and informed operational decisions. Predictive analytics support early identification of degradation trends, allowing maintenance to be planned proactively rather than in response to failures. Within high-value AI facilities, this capability is essential for protecting capital-intensive electrical assets and maintaining confidence in long-term availability.

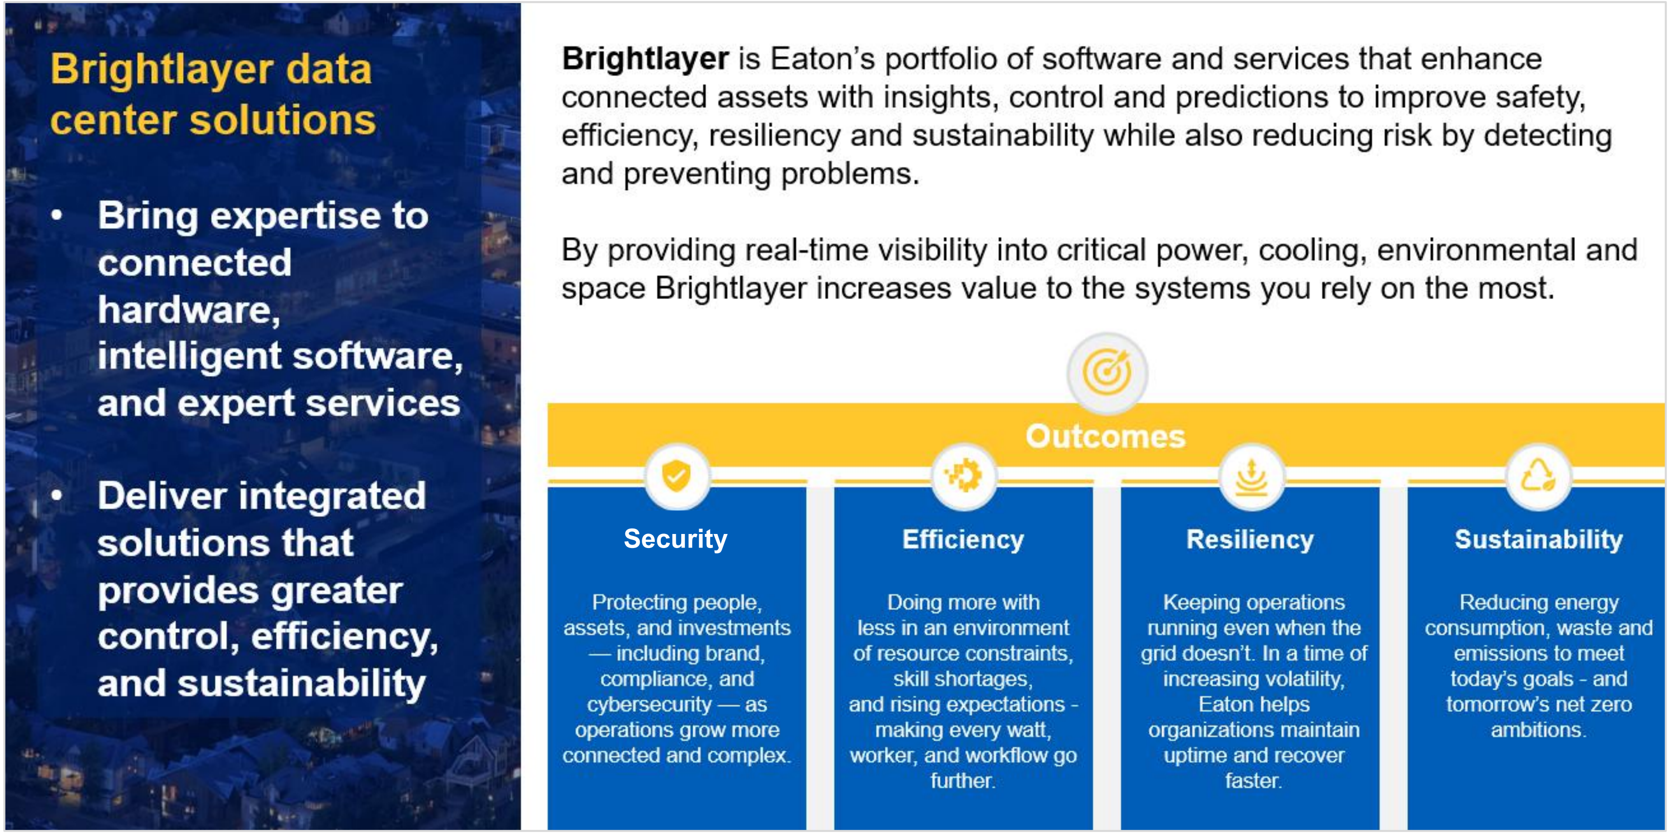

TRANE SYSTEM ORCHESTRATION FOR SMARTER THERMAL MANAGEMENT

Trane's full-stack controls and software portfolio deliver a complete, integrated, system-level approach. Built on the foundation of thermal equipment and scaled to the unique needs of the data center, these offerings help enhance reliability through proactive, real-time control, digital insights, and automated response to critical events, delivering energy improvements and smarter operation. From chiller plants to CDUs, today's thermal systems need to coordinate like a team. Trane enables seamless communication between pieces of equipment and across subsystems, with control and monitoring that turns complex infrastructure into synchronized performance.

System Control

Tracer® is a unified thermal management system that simplifies operations by handling chiller rotation, staging, and sequencing with built-in tools that boost reliability and enable data-driven performance.

AI Enhancements

Take optimization further with our AI-driven efficiency engine. **Trane® AI Control** blends decades of expertise with real-time data to continuously improve your thermal management system and capture every possible watt of savings.

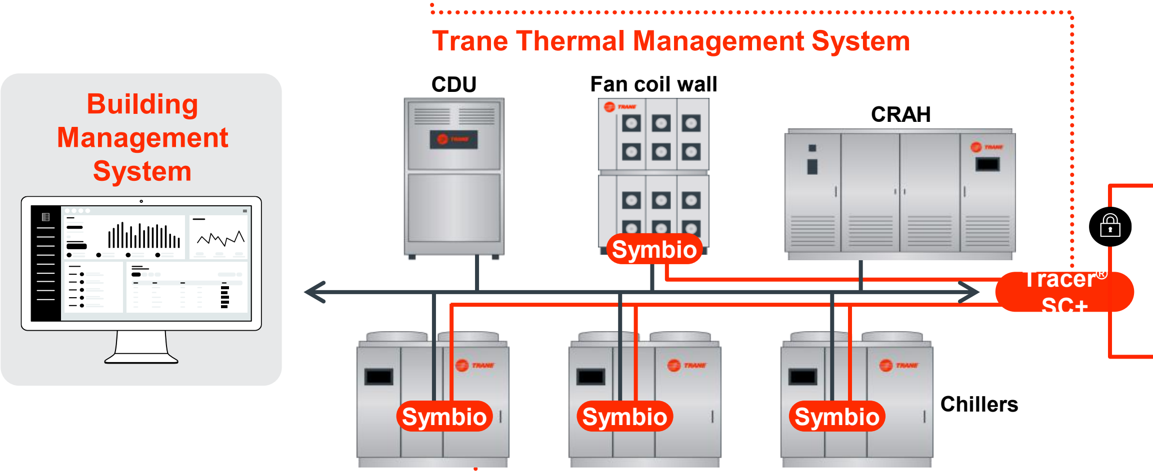

Digital Hub

Trane® Cloud

- Secure connectivity
- Remote service
- Analytics and diagnostics

Tracer® Synchrony®

- Chiller plant control
- Optimization strategies
- Equipment monitoring

Insights + Optimization

When power is limited, your cooling system must run at peak efficiency—every kilowatt saved means more racks online. **Trane Intelligent Services** uses whole-system data to deliver actionable energy-saving strategies.

Equipment Connectivity

Symbio® controls collect and translate key equipment data into the **Tracer SC+,** enabling data-driven service and optimization.

Predictive Maintenance

With your permission, we can securely connect to your equipment or Tracer SC+ to enable **Connected Mechanical Service**. By monitoring trends and performance data, we help identify issues early so your team or our technicians can act before failures occur. And when on-site support is needed, **Trane Service Agreements** and **local technicians** are ready to respond.

SYSTEM CONTROLS

Trane's approach is built around the key system approaches that drive long-term data center resilience:

- Connect
  - Thermal management system control synchronizes mechanical systems across the chilled water loop to improve performance while offering built-in redundancy and fault tolerance.
  - Integrated equipment unit controls deliver asset-level monitoring, uptime protection, and intelligent performance. Data across multiple connected units can be aggregated in tailored dashboards, providing insights into both daily operation and historical performance.
- Enhance
  - Chiller Plant Manager orchestrates chiller plant operation through chiller sequencing, rotation, and staging to improve energy efficiency, reduce strain, and extend asset life.
  - AI-enhanced control continuously fine-tunes system performance to improve efficiency and reduce energy use.
  - Optional distributed team control with dynamic leader reassignment reduces dependency on centralized BMS systems and sustains system coordination in the event of member failure.
- Maintain
  - Real-time monitoring and analytics help detect deviations before they become disruptions — with options for remote or rapid on-site resolution.
  - Intelligent Services combine analytics and expert insight to drive performance improvements, prioritize action, and free up operational capacity.

Trane's modular system architecture and secure open standard protocols support integration with other building systems including Eaton's Brightlayer™ intelligent power management solutions, enabling proactive system response to changes in load.

SERVICES, LIFECYCLE VALUE, AND ASSET PROTECTION

Eaton and Trane share a unified vision for data center services: to help deliver seamless, lifecycle-long support for data centers—helping to maximize uptime, reduce operational costs, and automate critical decisions.

Our combined mission is to embed service design considerations early—beginning at the reference design and digital twin stages—to optimize end-to-end data center performance across customers' power, electrical, and cooling systems.

SERVICES, LIFECYCLE VALUE, AND ASSET PROTECTION

Eaton and Trane’s professional Services, powered by a digital networking foundation, delivers intelligent energy and cooling management across hyperscale and colocation data centers.

- ✓ Advanced unit and system controls provide real-time actionable insights and full-system orchestration, monitoring, and management to deliver the best outcome.
- ✓ Actionable insights and analytics detect deviations before they become disruptions with remote or rapid on-site resolution.
- ✓ Professional service tools are fundamental, but it’s our experienced service teams and service model that truly set us apart.
- ✓ Our global network includes highly trained technicians, available 24/7—unmatched in scale and reach—to support your sites wherever they are located.

As AI factories grow in scale and asset value, the importance of these capabilities increases. Power and cooling systems represent a growing share of total capital investment, and their reliability directly governs compute availability. Integrated connectivity, controls, and services transform these systems from static infrastructure into actively managed assets, protecting investment value while enabling safe, efficient, and resilient operation over decades of service.

CONCLUSION

AI data centers are moving from megawatt sites to gigawatt-class campuses, making industrial-scale power and cooling the key constraints. This shift requires new architectures, not incremental extensions of legacy designs. Eaton’s reference power architecture, aligned with Trane’s cooling approach, shows how to meet these demands while staying adaptable to changing compute and grid conditions.

The future architecture is organized into three coordinated domains:

- ✓IT power, where 480 V AC remains viable and 800 V DC boosts efficiency by reducing conversion stages.
- ✓Whitespace cooling, where high-capacity CDUs become mission-critical and need dedicated UPS-backed ride-through.
- ✓Centralized facility cooling, where medium-voltage distribution is essential for large chiller and pump motors with controlled starts and strong fault containment.

The goal across all domains is the same: maximize IT availability and maintain predictable grid behavior through coordinated, high-quality power delivery.

Key trends underline this need. Shifting to 800 V DC can improve efficiency by 3–5%. CDU capacity has grown from ~100 kW to 10 MW in five years. Medium-voltage distribution cuts current dramatically—for example, from 1,420 A at 480 V to 164 A at 4.16 kV for a 1,500 HP pump. Together, these trends show that gigawatt-scale AI campuses require pre-engineered, integrated power-and-cooling architectures incorporating medium-voltage and high-voltage DC to accelerate deployment, improve reliability, and speed time-to-revenue.

CONCLUSION

Eaton supports this architecture with a portfolio spanning critical power, low- and medium-voltage distribution, motor control, and emerging 800 V DC systems.

Key elements in the reference design include Eaton's 9395 and 93PM UPS platforms, medium-voltage distribution and protection, VacClad switchgear, Ampgard motor control, and SC9000 drives, all enhanced by Brightlayer Data Center monitoring and analytics.

Together, these components create a scalable electrical foundation that ensures IT and cooling continuity and enables predictable, repeatable deployment across blocks.

CONCLUSION

Trane Reference Design #501 provides the cooling topology used to show how Eaton's power architecture applies to a modular, gigawatt-scale system.

- Modular scaling: The campus is built from ten repeatable 102.4 MW blocks, and Eaton's power architecture aligns to the same modular structure for predictable deployment.
- Medium-voltage backbone: 4.16 kV distribution feeds outdoor chiller plants, with large motors started and protected by Eaton Ampgard medium-voltage motor control.
- Low-voltage where appropriate: Outdoor dry coolers use 480 V switchboards to match distributed fan loads and simplify block-level standardization.
- Whitespace cooling: CRAH and FCW units inside the data hall also run from 480 V distribution, leveraging mature low-voltage practices.
- CDU continuity: Indoor CDU pods operate at 480 V and receive backup from a dedicated Eaton 9395 UPS to maintain liquid-cooling stability during power events.
- Evolution path: The design anchors on 480 V AC today while acknowledging emerging 800 V DC options as DC protection and components mature.
- Controls and digital integration: Trane thermal controls coordinate plant-level cooling, while Eaton Brightlayer extends monitoring and analytics across electrical systems.
- Services and lifecycle: Digital connectivity supports maintenance, monitoring, and lifecycle management at gigawatt scale.

Ultimately, success at this scale depends on power, cooling, and intelligence operating as one system. The aligned Eaton and Trane reference designs provide a coordinated blueprint for efficient, high-availability AI factories.

HELPFUL LINKS

Trane BIM and Selection Tools →

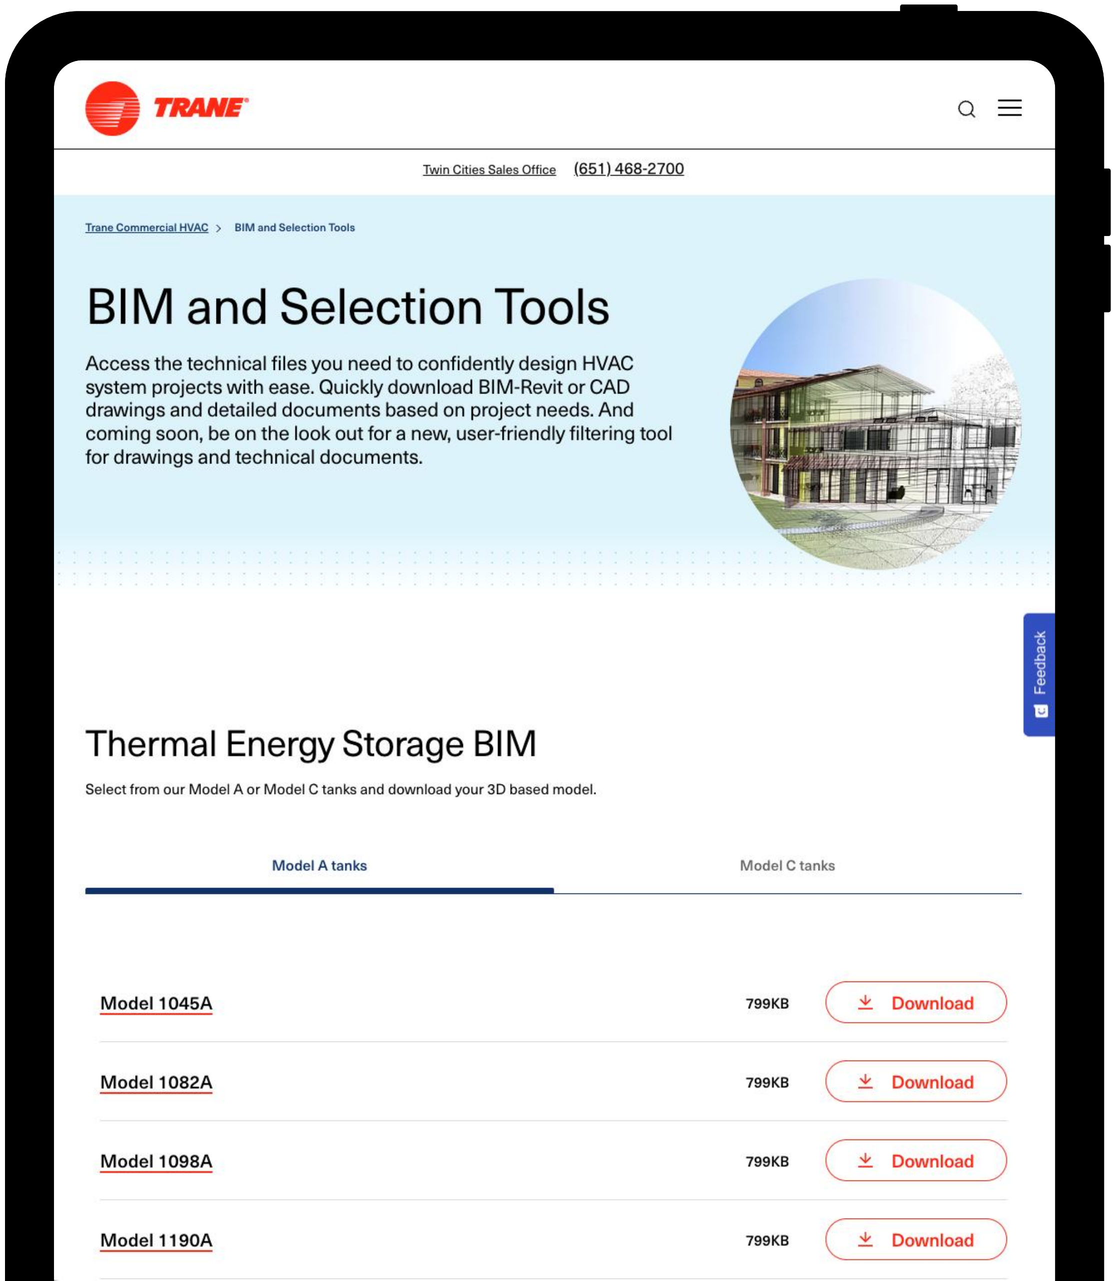

Trane.com/DataCenters

This is for informational purposes only and does not constitute professional advice. Trane believes the facts and suggestions presented here to be accurate; however, final design and application decisions are your responsibility. Trane disclaims any responsibility for actions taken on the material presented. This presentation is protected by U.S. and international copyright laws. Reproduction, distribution, display, and use of the presentation without written permission of Trane is prohibited.

DC-PRR005A-EN
07/22/2026

APPENDIX

THE SCALE OF AI POWER INFRASTRUCTURE

IT power within the whitespace – racks to SuperPODs

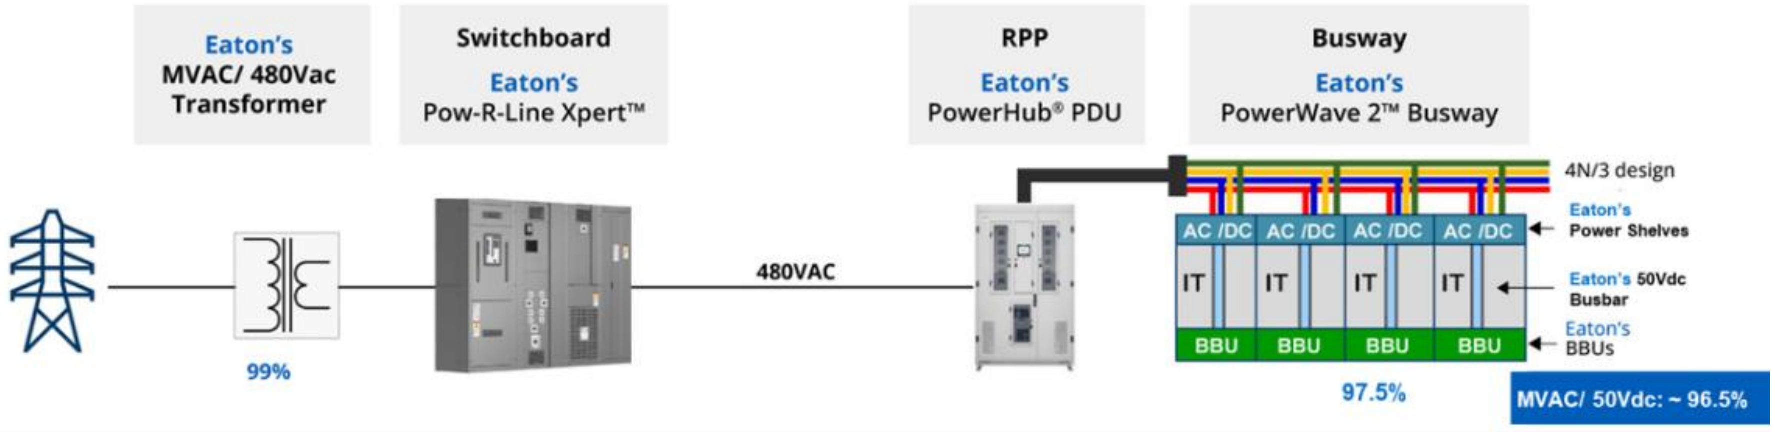

Hybrid AC architectures add localized DC conversion to support rising rack power while keeping 480 V AC distribution. Medium-voltage service is stepped down to 480 V AC, routed through low-voltage switchgear, and then converted to DC near the load using solutions like Eaton's sidecar designs. Racks receive DC power, with final conversion to 48 V or 800 V depending on the compute platform. Moving conversion out of the rack reduces AC-DC stages, boosting efficiency and easing thermal and space constraints. Eaton's UPS, sidecar, and low-voltage distribution products support this transitional approach, which bridges today's AC systems and future higher-voltage architectures.

THE SCALE OF AI POWER INFRASTRUCTURE

IT power within the whitespace – racks to SuperPODs

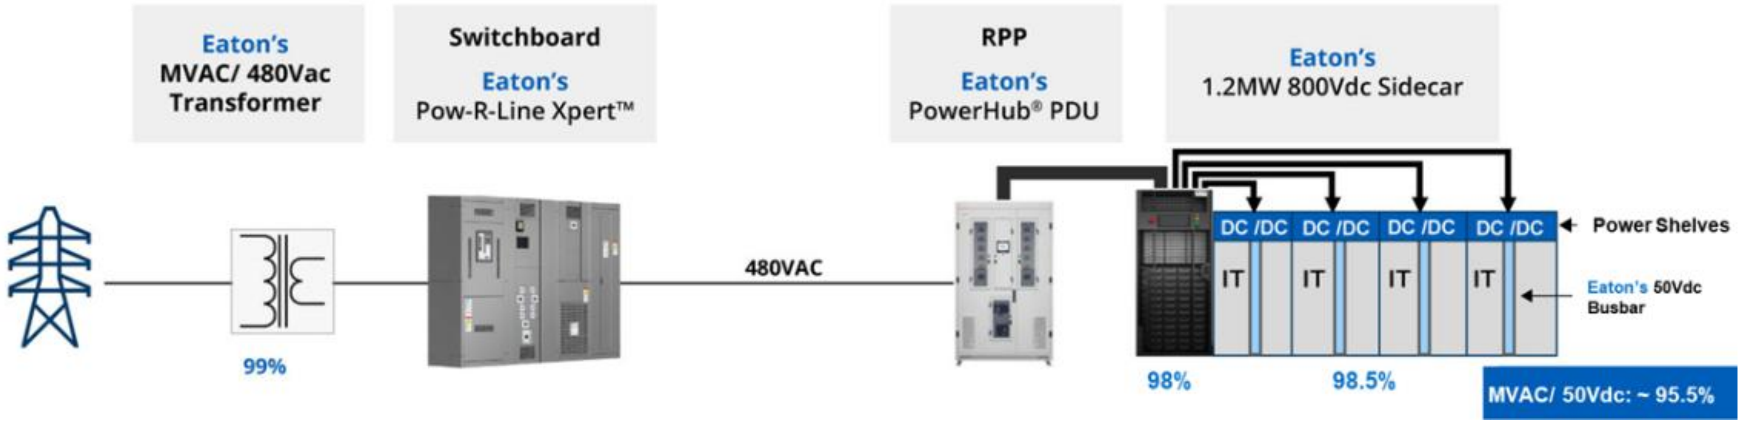

Direct DC architectures are the most scalable long-term solution for high-density AI whitespace. Medium-voltage service feeds a solid-state substation that converts directly to 800 V DC, which is distributed through the whitespace via DC busway. Final DC-to-DC conversion happens at the rack, reducing conversion stages and sharply lowering distribution current. This approach delivers maximum efficiency, power density, and scalability as racks approach megawatt levels. Eaton is advancing 800 V DC protection, solid-state breakers, transformers, and other DC-native components to support this shift.

THE SCALE OF AI POWER INFRASTRUCTURE

MVSST / 800VDC

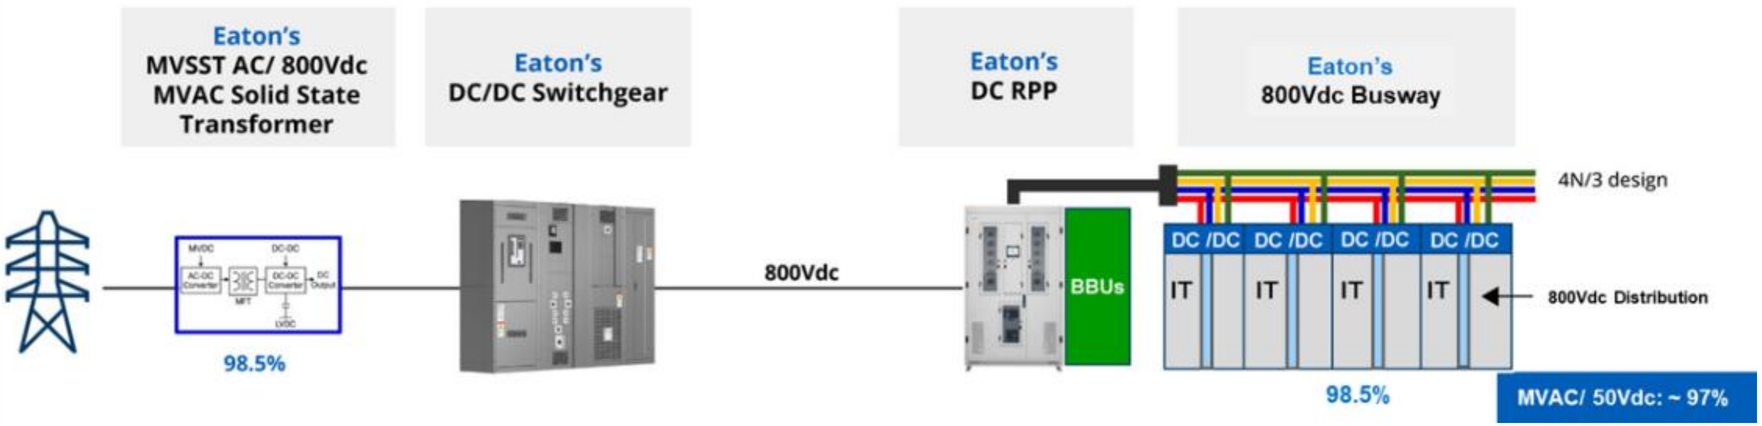

Direct DC architecture represents the most scalable long-term approach for powering high-density AI whitespace. In this model, medium-voltage service feeds a medium-voltage solid-state substation (MVSS) transformer that converts power directly to 800 V DC, which is then distributed throughout the data center halls to power Ai IT racks via DC busway. Final DC-to-DC conversion occurs at the rack, minimizing conversion stages and significantly reducing distribution current for a given power level. This architecture offers the highest potential efficiency, power density, and scalability as rack power approaches megawatt-class systems. Eaton is advancing the commercialization and development of 800 V DC protection technologies, solid-state circuit breakers (SSCBs), solid-state transformers (SSTs), Remote Power Panels, Busway and DC-native components to enable this transition.