DAIKIN

DATA CENTER COOLING

AI Factory Reference Architecture

NVIDIA DSX AI Factory Platform
Vera Rubin NVL72

60 MW

Reference D60-W: Water-Cooled Chillers with Dry Coolers



INTRODUCTION TO THE AI FACTORY INFRASTRUCTURE

A REFERENCE ARCHITECTURE FOR SCALABLE, HIGH-PERFORMANCE AI FACTORIES

Purpose of This Guide

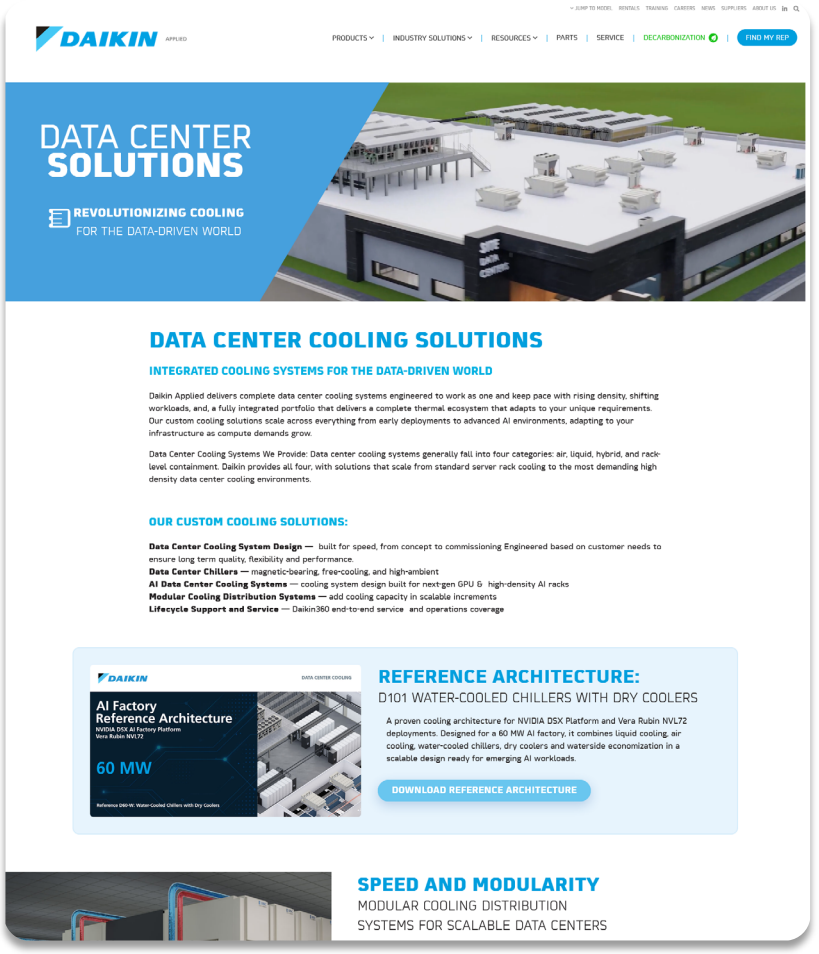
This guide presents Daikin's cooling reference architecture for project-specific integration with the NVIDIA DSX Platform and Vera Rubin NVL72, featuring a 60 MW AI factory across six deployable units that is scalable to meet the demands of growing and emerging technologies.

Integrated Cooling Architecture

This architecture combines liquid cooling for AI workloads with dedicated air cooling for residual heat loads, supported by water-cooled chillers, dry coolers and integrated waterside economizers.

What You'll Find Inside

The guide defines the design conditions, equipment configuration, resiliency strategy and operating modes to support project-specific engineering to bring a Total Thermal Value Stream to customers building AI Factories.



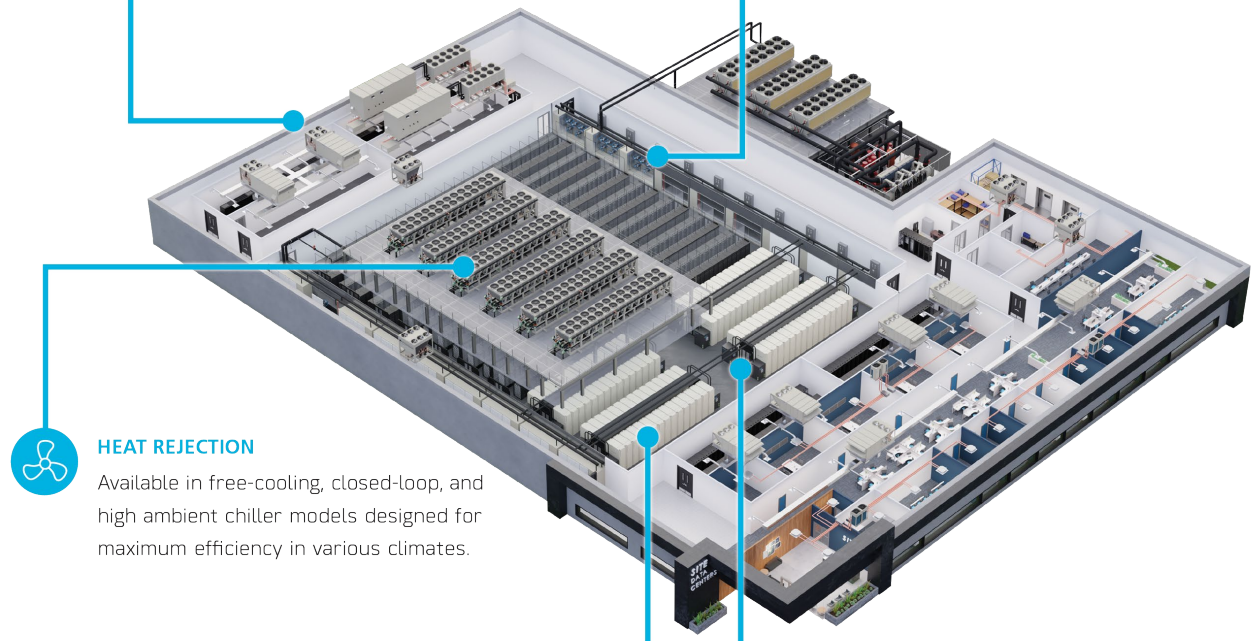
Learn More About Daikin Data Center Solutions

THE DAIKIN DATA CENTER INFRASTRUCTURE

Our data center cooling portfolio covers the full thermal path, from capturing heat at the chip to rejecting it outdoors. Cold plates, coolant distribution units, rack containment units, and heat rejection equipment work together as one system to meet your facility's cooling requirements.

- CDUs manage coolant flow and temperature, transferring heat from the IT cooling loop to the facility water loop while keeping the two circuits separate.
- Rack-level systems manage airflow and capture server heat, supporting air-cooled and hybrid deployments.
- Chillers and dry coolers provide facility cooling and heat rejection, selected to match water temperatures, cooling loads and site conditions.

Each solution is selected around rack density, water temperatures, local climate and capacity needs, connecting chip-level cooling to facility-wide heat rejection.



CRITICAL SPACE COOLING

Combine quiet operation, high efficiency, and application flexibility for critical operations from command and control spaces, to network rooms, office and ancillary spaces.

CRAH UNITS

Flexible configurations to fit any data center need deliver cool air to server racks while removing heat.

HEAT REJECTION

Available in free-cooling, closed-loop, and high ambient chiller models designed for maximum efficiency in various climates.

MODULAR DATA HALL SOLUTIONS

Deployable standalone or within data centers, these solutions deliver rapid setup, efficient cooling, and scalable, adaptable infrastructure.

IT HEAT REJECTION

Centralized, end-of-row, and in rack Coolant Distribution Units (CDUs) for traditional and modular data centers.

DESIGN BOUNDARY CONDITIONS

Ashburn, Virginia

60
MW IT Load

NVIDIA DSX AI Factory Platform
Vera Rubin NVL72
6 Deployable Units
768 reference positions

92% liquid cooling
8% RCU air cooling

Site and Weather

Parameter	Reference Basis
Location / Station	Ashburn, VA (WMO: 724030)
Summer Dry Bulb	39.8°C (103.6°F)
Summer Mean Coincident Wet Bulb	28.1°C (82.6°F)
Winter Dry Bulb	−12.9°C (8.8°F)

Temperatures

Circuit	Supply / EWT	Return / LWT
Facility Liquid	37°C (98.6°F)	47°C (116.6°F)
Rack Liquid	41°C (105.8°F)	51°C (123.8°F)
RCU Water	22°C (71.6°F)	33.11°C (91.6°F)
Condenser Water	44.8°C (112.6°F)	50.4°C (122.7°F)
IT Air Temperatures	30°C (86°F)	44°C (111.2°F)

DEPLOYABLE UNIT ARCHITECTURE

13.8 MW
DU reference capacity

10 MW
Plant operating allocation

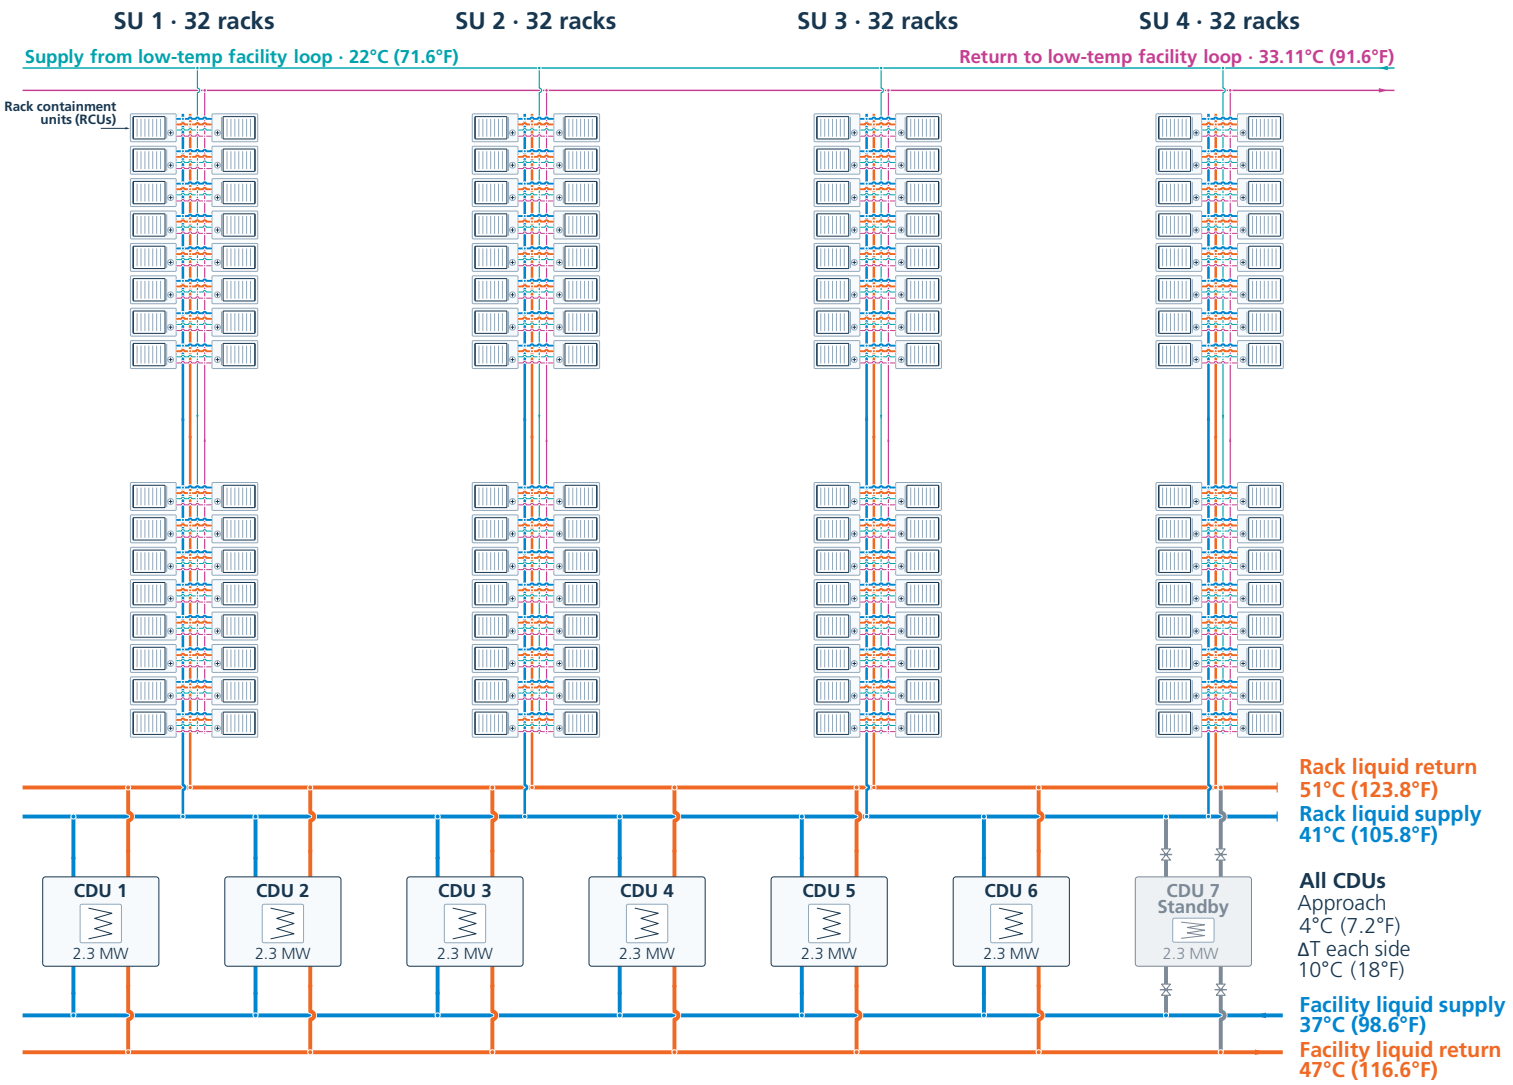
128 racks
4 Scalable Units

7 × 2.3 MW CDUs
6 duty + 1 standby

9.2 MW liquid cooling

0.8 MW RCU air cooling

Operating split: 92% liquid / 8% air



128 reference rack positions; compute/support allocation is project-specific.

AI FACTORY ARCHITECTURE

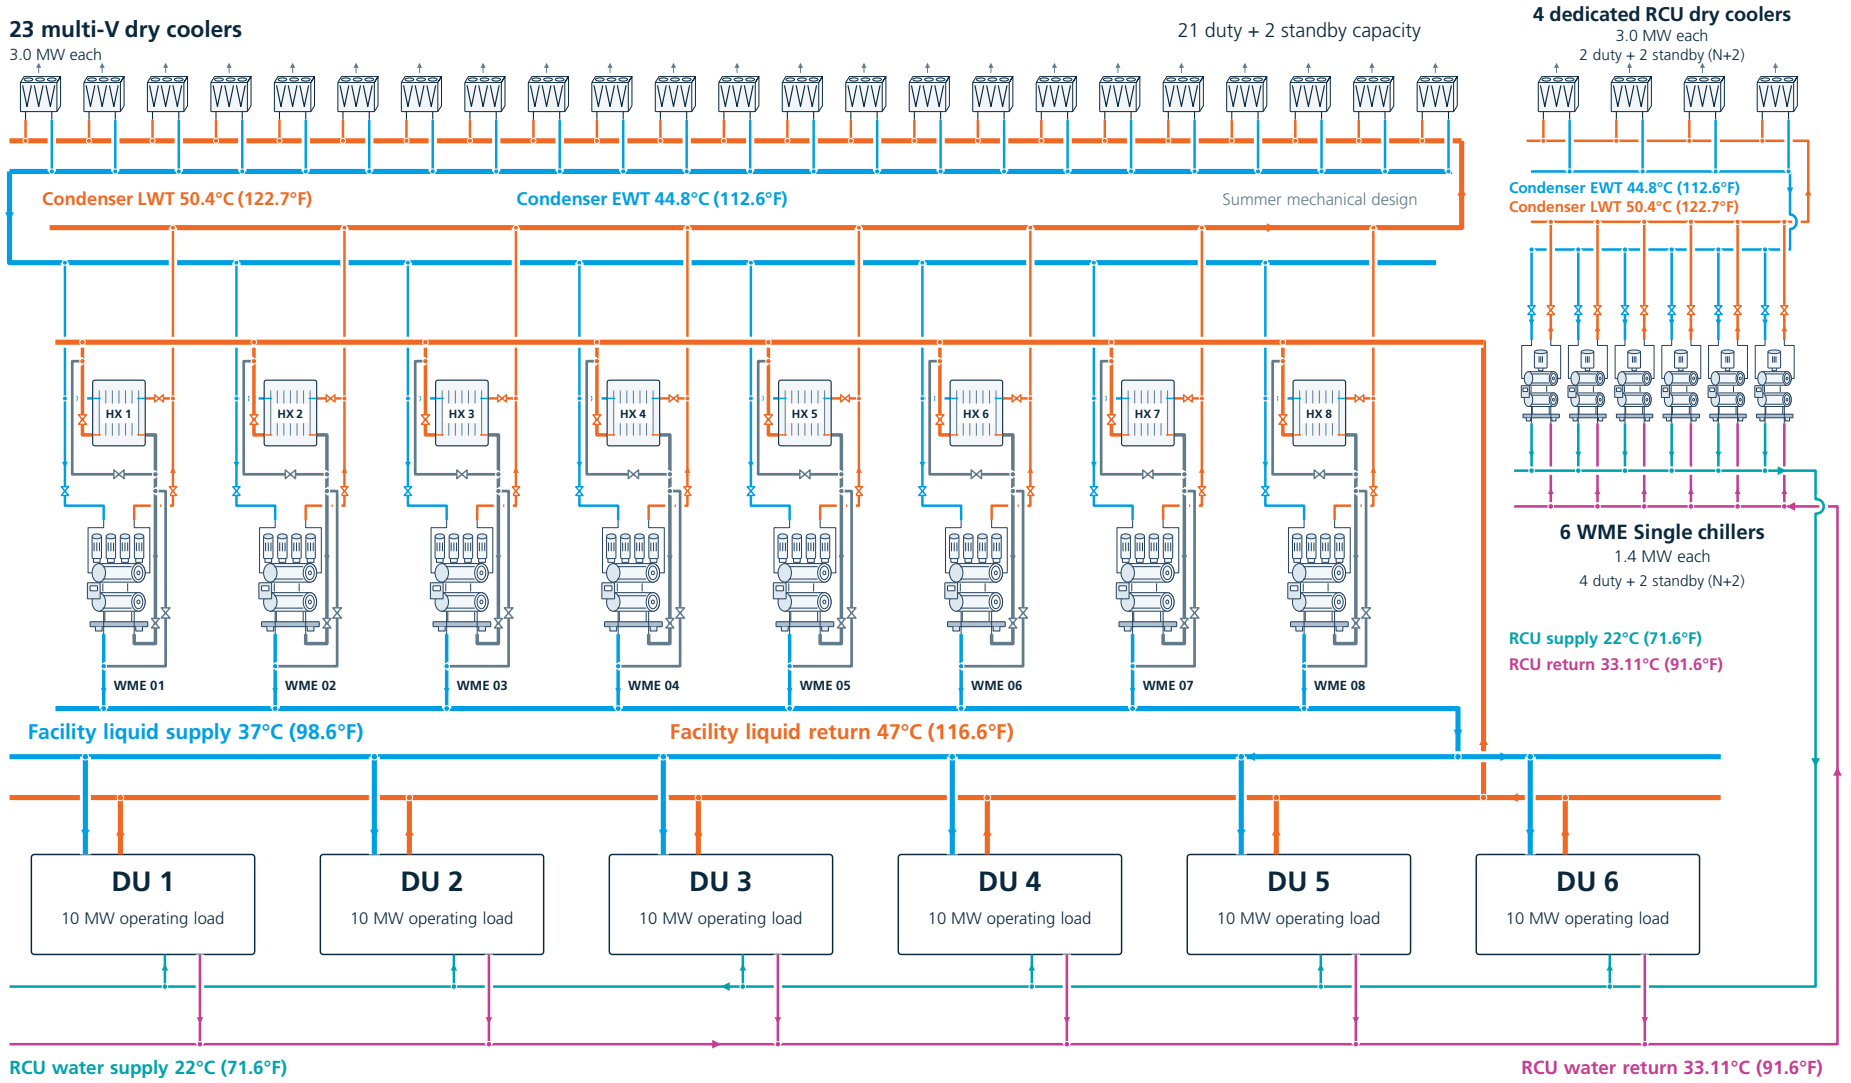
Six Deployable Units with Dual Temperature Loops & Multiple Operating Modes

Heat rejection plant supports 60 MW IT load; additional plant capacity is required to fully utilize the total DU thermal capacity.

Modes of Operation

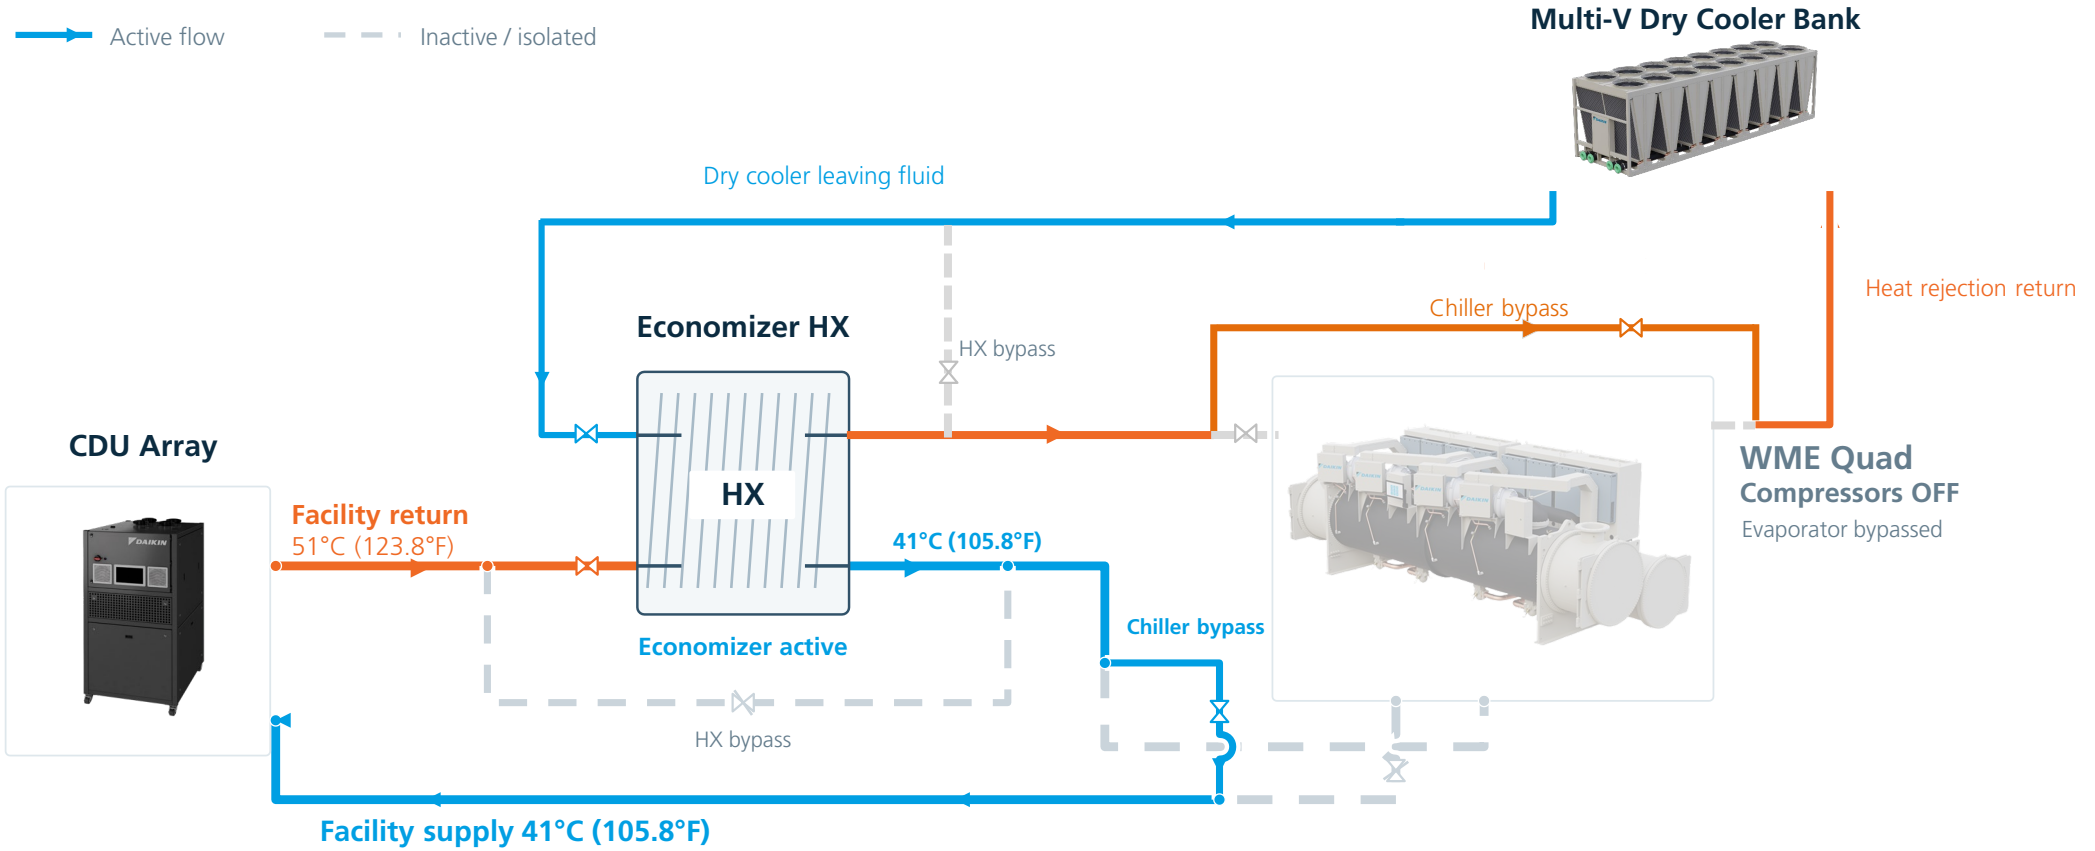
- Dry Cooling
- Hybrid Cooling
- Mechanical Cooling

8 WME Quad chillers
6 duty + 2 standby (N+2)
10 MW each



MODE 1: DRY COOLING

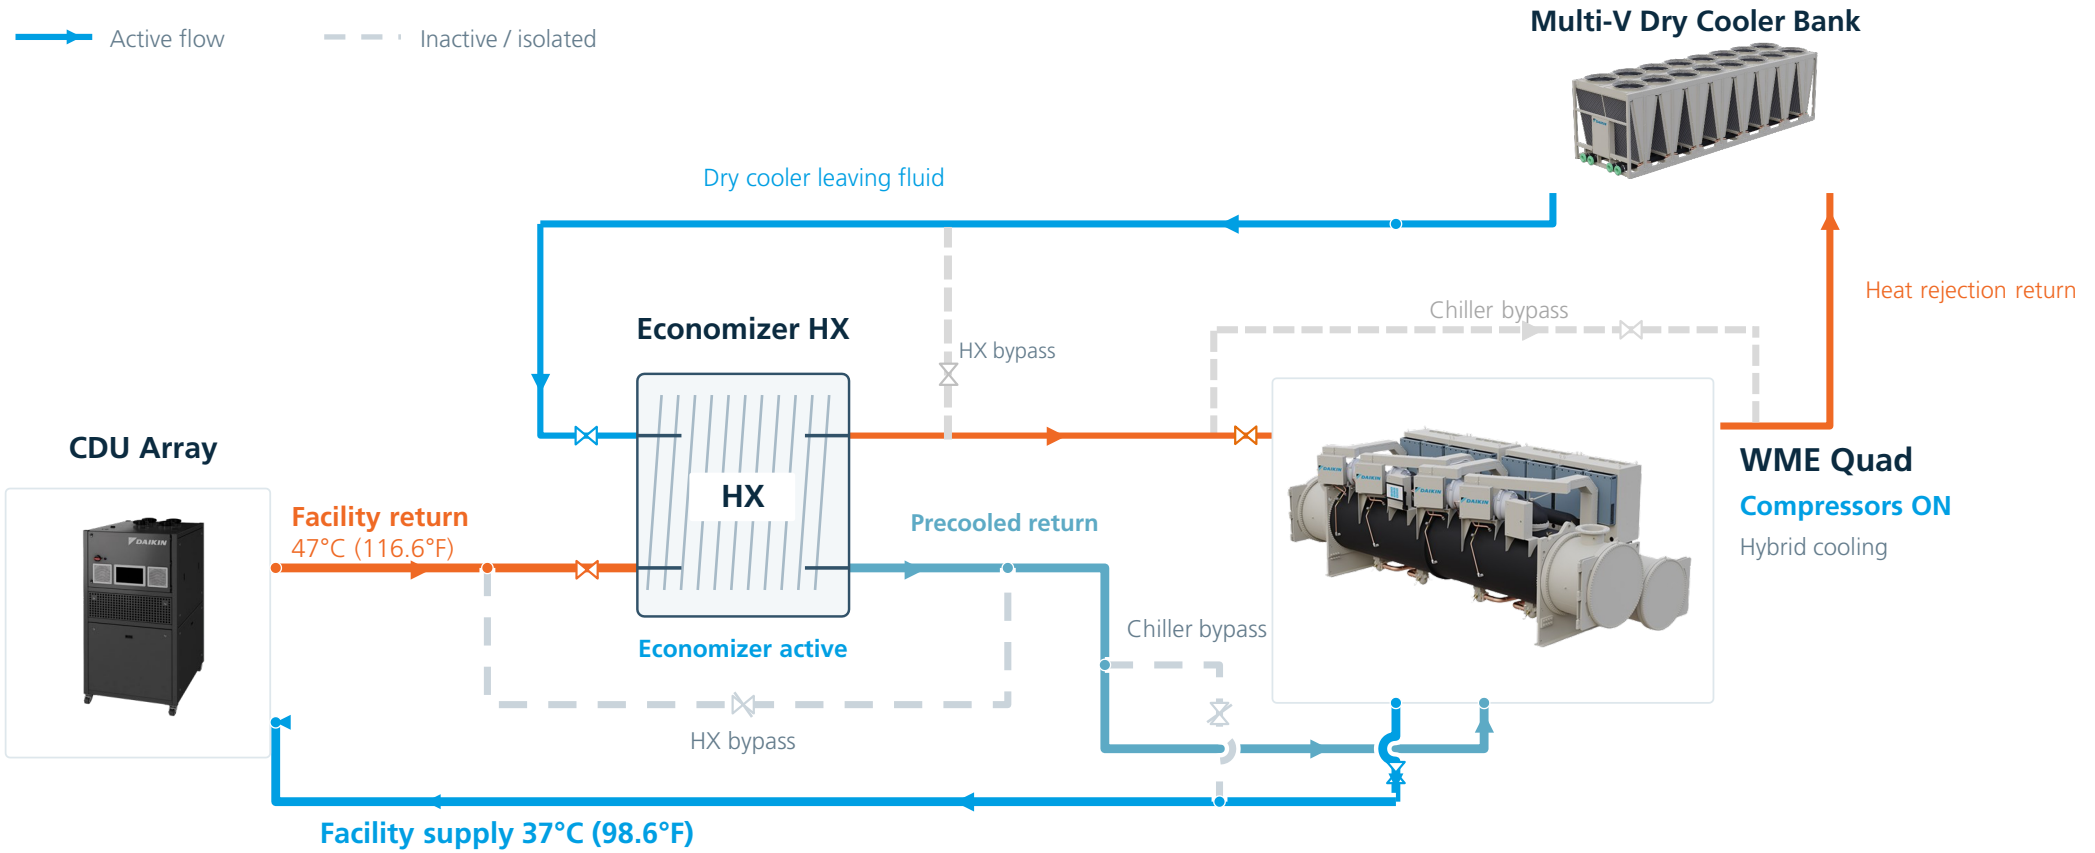
WME Liquid-cooling Plant
Dry cooler leaving fluid + HX approach ≤ 41°C (105.8°F)



Dry Cooling: Enabled when the dry cooler leaving-fluid temperature, plus the economizer HX approach, is at or below the facility supply-temperature setpoint. In this mode, the Chiller compressors are off. *During full economization only, facility supply resets up to 41°C, providing 45°C rack supply through the CDU. Hybrid and mechanical operation modes restore 37°C facility supply.*

MODE 2: HYBRID COOLING

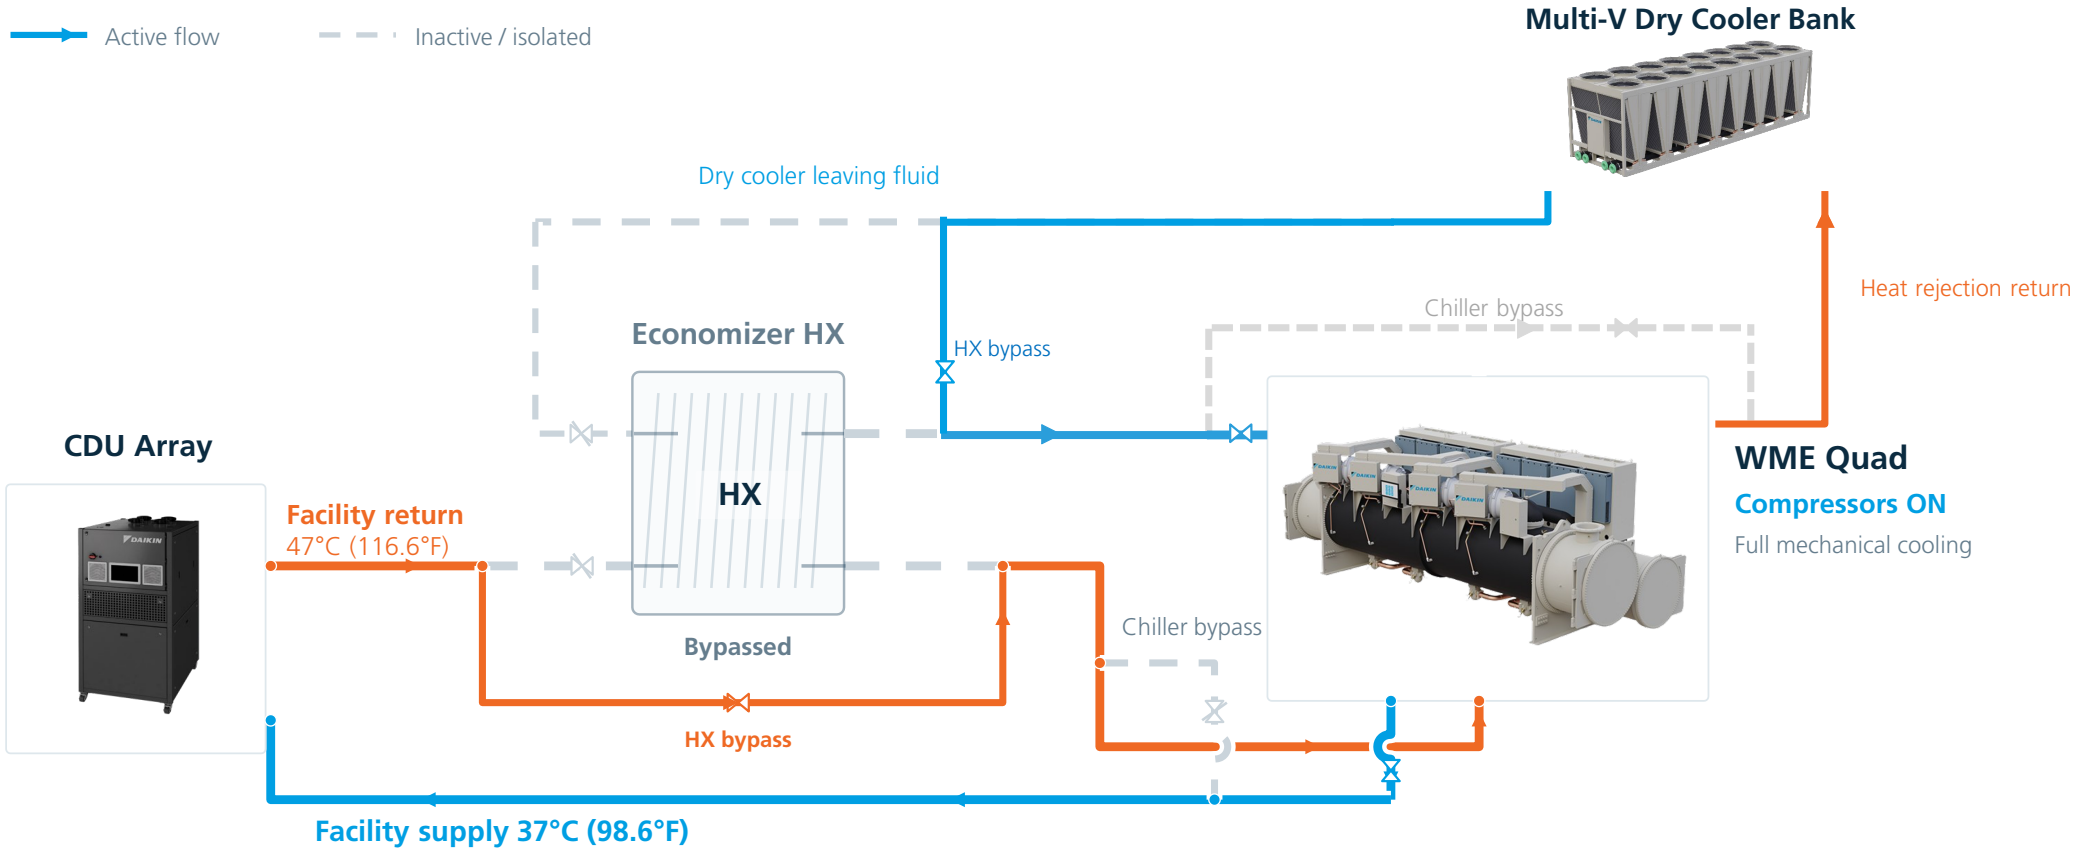
WME Liquid-cooling Plant
41°C (105.8°F) < dry cooler leaving fluid + HX approach < 47°C (116.6°F)



Hybrid Cooling: Enabled when the dry cooler leaving-fluid temperature, plus the economizer HX approach, is above the facility supply setpoint but below the facility return-water temperature. The economizer precools the return water, and chillers provide the remaining cooling to reach the supply setpoint, significantly reducing energy consumption.

MODE 3: MECHANICAL COOLING

WME Liquid-cooling Plant
Dry cooler leaving fluid + HX approach ≥ 47°C (116.6°F)



Mechanical Cooling: Enabled when the dry cooler leaving-fluid temperature, plus the economizer HX approach, is at or above the facility return-water temperature, leaving no useful economizer cooling. In this mode, the economizer is bypassed and the chillers maintain the facility supply setpoint.

60 MW Central Plant

Load (MW)	Per DU allocation	Central plant
Liquid	9.2	55.2
RCU air	0.8	4.8
Total	10	60

Plant operating basis: 92% liquid / 8% air
Existing 13.8 MW DU architecture retained

Deployable Unit Redundancy

- CDU 6+1
- RCU fans and coils: 2N design target

Central Plant Redundancy (N+2)

- WME Quad (High Temperature Loop) 6+2
- Dry Coolers (High Temperature Loop) 21+2
- WME Single (Low Temperature Loop) 4+2
- Dry Coolers (Low Temperature Loop) 2+2

DAIKIN DATA CENTER CAPABILITIES

How We Deliver For Data Centers

SOLUTIONS ENGINEERED FOR PERFORMANCE, RESILIENCE AND LONG-TERM VALUE

Unified Performance
Systems engineered to operate together deliver better cooling stability, efficiency, and scalability.

Resiliency by Design
Equipment capacity redundancy and coordinated controls support availability.

Future-Proof Flexibility
Air and liquid cooling paths designed to evolve with rising compute density.

Sustainability at Scale
Non-evaporative heat rejection, low-GWP refrigerants and heat reuse options reduce environmental impact.

Fostering Strong Partnerships
Taking ownership and driving measurable results through a proactive, hands-on approach focused on innovative, tailored solutions.

Faster Deployment
Reference architecture accelerates design, integration and commissioning.

Protective Maintenance Plans
Comprehensive plans designed to keep critical assets operating at their best.

24/7 Service & Support
From seasonal tune-ups, critical repairs, and system diagnosis to comprehensive project management and readily available, trained field service personnel for on-site support.

COOLANT DISTRIBUTION UNIT

Daikin Coolant Distribution Unit

Available 2027-H2



7 × 2.3 MW

Per Deployable Unit

6 duty + 1 standby

4°C approach

50 psi (3.45 bar)

Core Features

Controlled Rack Conditions
Regulates liquid supply temperature, flow and pressure.

Redundant Pumps & Filtration
Prioritizes continuous operational uptime by integrating robust failsafes to counter hardware failure.

Load Sharing & Operating-Hour Balancing
Daikin proprietary control logic to maintain system uptime and maximize efficiency.

Hot-swappable Components
Engineered as part of a maximum uptime design to allow maintenance and repairs without data center downtime.

Performance

- 25 µm, Beta 75
- Approach 4°C
- 1 MW to 3.5 MW
- Wetted-material compliant

RACK CONTAINMENT UNITS (RCU)

DDC S-Series Platform

CAPTURES HEAT AT EVERY RACK AND MOVES IT STRAIGHT TO THE WATER LOOP.



128

RCUs per Deployable Unit

Core Features

Protect the Investment
Fire suppression and water-loop protection built into every rack.

Density Without Hot Spots
Heat is captured at the rack, so the room stays neutral. No chasing heat around the hall.

Air + Liquid, Side by Side
Mix air, DLC and hybrid racks in one row. Residual heat is captured at the rack.

Grow Rack by Rack
Deploy in existing halls with power and chilled water. Add capacity as GPUs arrive.

Performance

- Liquid-to-Chip Ready
- Dynamic Density Control™
- Biometric Rack Access
- Telemetry + DCIM Integration

Magnitude® WME Magnetic Bearing Centrifugal Chiller

OIL-FREE READINESS FOR AI FACTORIES: LESS MAINTENANCE DURING FREE COOLING. NO OIL-RELATED STARTUP DELAYS. SUPPORTING RAPID COOLING RECOVERY WHEN DEMAND ARRIVES.



AHR INNOVATION AWARDS

8

WME Quad WC chillers

10 MW each
6 + 2 redundant

Core Features

Compact Footprint
Maximize usable space and simplify installation, helping to reduce costs while optimizing valuable floorspace.

Efficiency
Oil-free magnetic bearing technology helps operators reduce energy costs and improve PUE performance.

Dry-cooling Applications
Eliminate the need for cooling towers, reduce water usage, simplify infrastructure and support more sustainable cooling strategies.

Maximum Uptime
RideThrough® and RapidRestore® support cooling continuity and rapid recovery from power disturbances. Oil-free technology eliminates oil-related startup delays after extended standby.

Performance

- R-513A or R-515B
- 10 MW
- Oil-free Operation

Load Sharing Across All Available Chillers

Magnitude® WME Magnetic Bearing Centrifugal Chiller

YEAR-ROUND COOLING FOR THE DEDICATED RCU LOOP.
OIL-FREE TECHNOLOGY, LOAD SHARING AND RAPID RECOVERY SUPPORT RELIABLE 22°C (71.6°F) WATER SUPPLY.



6

WME Single WC chillers

1.4 MW each
4 + 2 redundant

Core Features

Compact Footprint
Maximize usable space and simplify installation, helping to reduce costs while optimizing valuable floorspace.

Efficiency
Oil-free magnetic bearing technology helps operators reduce energy costs and improve PUE performance.

Dry-cooling Applications
Eliminate the need for cooling towers, reduce water usage, simplify infrastructure and support more sustainable cooling strategies.

Maximum Uptime
RideThrough® and RapidRestore® support cooling continuity and rapid recovery from power disturbances. Oil-free technology eliminates oil-related startup delays after extended standby.

Performance

- R-513A or R-515B
- 1.4 MW
- Oil-free Operation

Load Sharing Across All Available Chillers

OIL-FREE TECHNOLOGY & AI FACTORY TOKENOMICS

01 Lower Standby Overhead

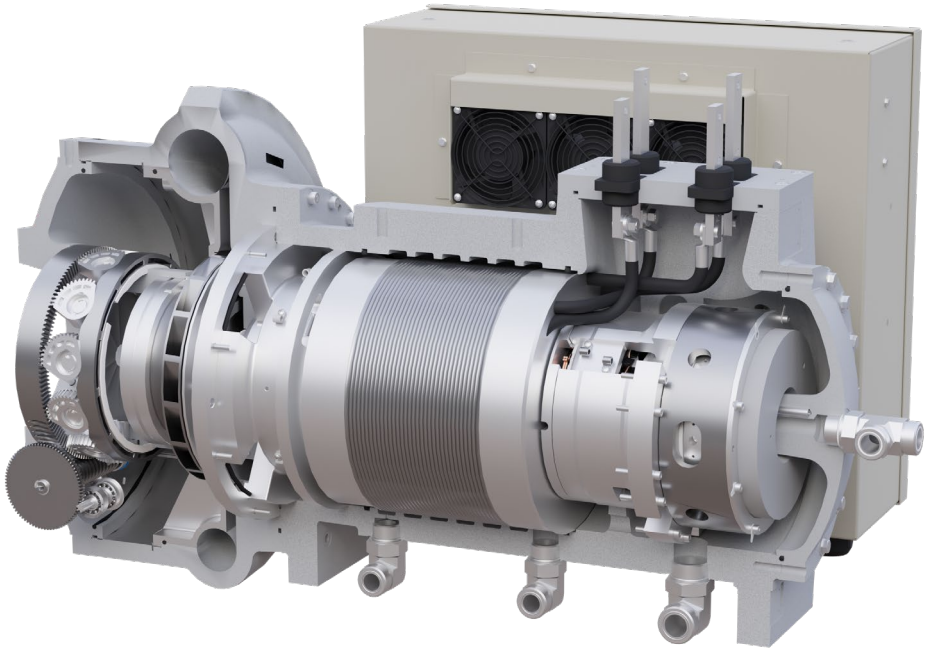
Dry cooling with less oil-system upkeep.

While dry coolers carry the load, oil-free chillers eliminate oil heaters, oil pumps and oil servicing—reducing standby energy use and maintenance costs to lower total cost of ownership (TCO) and support more AI tokens per dollar (Token/$).

02 Cooling Recovery That Protects Compute

Fewer startup barriers when cooling demand arrives.

After extended idle periods, oil-free technology removes oil-temperature and pre-lubrication startup requirements. This supports rapid cooling recovery and helps protect GPU availability and sustained token throughput.



WME-C Quad Motor Cooling Compressor

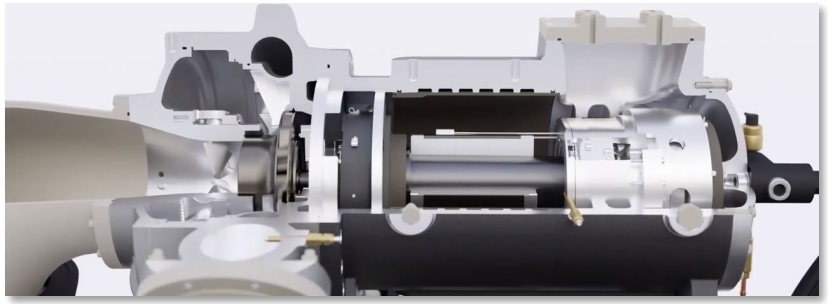
DAIKIN RAPIDRESTORE & RIDETHROUGH

03 RideThrough®

Maintains compressor rotation during a brief disturbance.

Within the supported disturbance envelope, the compressor continues rotating. When power returns, the drive resynchronizes with the spinning motor, avoiding a full restart. Available cooling during the event depends on the selected chiller, water flow and power conditions.

BRIEF DISTURBANCE
Rotation retained; drive resynchronizes

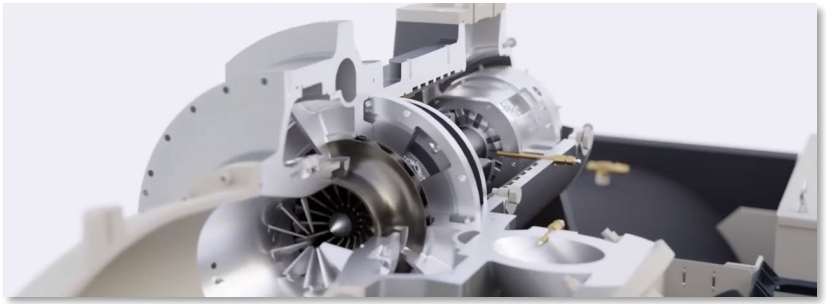


04 RapidRestore®

Compressor restart, followed by cooling-capacity recovery.

After an interruption that requires shutdown, RapidRestore expedites compressor restart once power, water flow and start permissives are restored. Cooling capacity then ramps back up. Restart time and capacity-recovery time are separate measures and depend on the selected chiller and operating conditions.

AFTER SHUTDOWN
Restart first; cooling capacity then recovers



MULTI-V DRY COOLER BANK

Data Center Dry Cooler

Available 2027-H2

HEAT REJECTION FOR THE LIQUID AND RCU COOLING PLANTS



27 Liquid: 21 + 2 standby

Air: 2 + 2 standby

3.0 MW nominal

Core Features

Water Conservation
Avoids evaporative heat rejection and cooling-tower make-up water.

Efficiency & Sustainability
Ideal for modern high-density AI and liquid-cooled server architectures.

Lower Maintenance
Avoids cooling-tower blowdown. Closed-loop water quality and treatment remain necessary.

Low Noise
Acoustic selection to meet the site requirement. Measurement distance and operating basis to be confirmed.

Performance

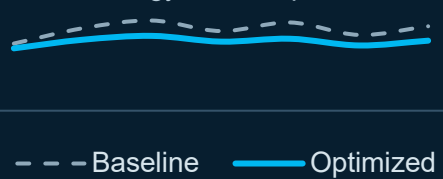
- Approach 5°C
- Noise 70 dBA
- Footprint
  8' 5" x 53' 11"
- Seismic certification

iDCM SUPERVISORY SOLUTION

LOAD SHARING, OPERATING-HOUR BALANCING AND THREE-MODE PLANT CONTROL

Site BMS / DCIM

Energy Consumption



iDCM Supervisory Layer

Master Controller 1



Master Controller 2



Optimization for energy efficiency

Advanced predictive maintenance

Primary IP network

Secondary IP network (redundant)





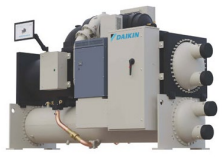




8 WME Quad



23 Dry Coolers



6 WME Single



4 Dry Coolers



42 CDUs



768 RCUs

AFTERSALES AND LIFECYCLE SERVICE

DAIKIN 360
BUILDING SOLUTIONS

A COMPLETE SOLUTION FOR YOUR BUILDING'S **ENTIRE LIFECYCLE**

Commissioning

Start-up support and recommissioning



Planned Care

Factory-trained maintenance, diagnostics and predictive care



Response & Repair

24/7 emergency availability, parts and repair support



Lifecycle Support

Modernization, replacement and temporary cooling



Daikin360 delivers a complete, integrated building solution across three core pillars:

HVAC & MECHANICAL
Service and maintenance for maximum performance

BUILDING AUTOMATION
Connected, intelligent building management

ENERGY & OPTIMIZATION
Efficiency, sustainability and long-term value