10.4 and 6.2 MW, Tier III, ANSI, Chilled Water, Liquid-Cooled AI Clusters (AMD Instinct™ MI455X GPUs)

Design Overview

Data Center Total IT Capacity
10,440 kW (greenfield)
6,150 kW (retrofit)

Target Availability
Tier III

Racks and Density
Total Racks: 62 (greenfield)
193 (retrofit)
Rack Power Density:
Support racks up to 45 kW/rack
AI racks up to 246 kW/rack

Data Center Overall Space
43,251 ft^2 (greenfield)
56,304 ft^2 (retrofit)

Regional Voltage and Frequency
480V, 60Hz

About this Design

- IT space and power distribution designed to accommodate AI clusters with density up to 246 kW per rack
- Design optimized to support liquid-cooled racks, with liquid-to-liquid coolant distribution units (CDUs) and adiabatic dry coolers
- Chilled water systems optimized for high water temperatures using *Uniflair FWCV* fan walls and *Uniflair XCAF* air-cooled packaged chillers
- Redundant design for increased availability and concurrent maintainability

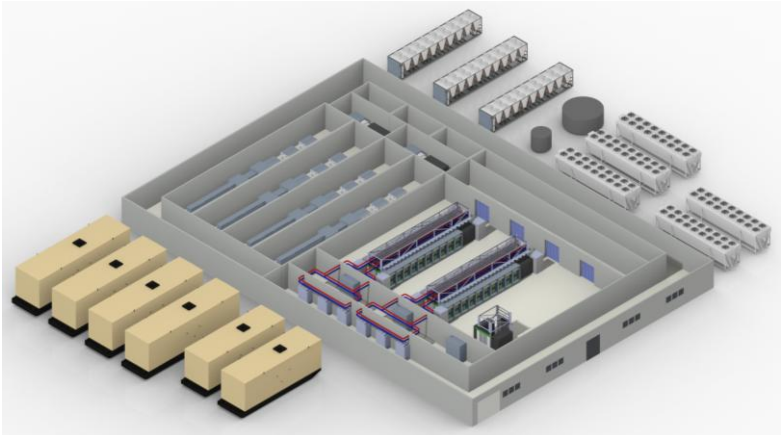
Introduction

AI applications and liquid cooling bring new challenges to data center design. Schneider Electric's data center reference designs help shorten the planning process by providing validated, proven, and documented data center physical infrastructure designs to address such challenges. This design focuses on the deployment of AI infrastructure, specifically with the *AMD Helios Rackscale Solution Powered by AMD Instinct™ MI455X GPUs*, in a single data hall. Two distinct scenarios are considered for the application: greenfield and retrofit.

For the greenfield situation, the data hall is purpose-built and optimized for two *AMD 2,880 GPU 20-Helios rackscale solution clusters* with a shared pod of support racks, using *Motivair MCDU-70* floor mounted liquid-to-liquid CDUs and adiabatic dry coolers. The air-cooled infrastructure considers both *Uniflair FWCV* perimeter fan-walls and hot-aisle containment or *Motivair M14* rear-door heat exchangers. Facility power and cooling design are optimized for capital cost, efficiency, and reliability.

For the retrofit situation, the data hall is existing and originally sized for four 2.5 MW power trains, with existing pods of 44 racks and average rack density of 20 kW/rack. Three pods are removed and replaced with 10 *AMD Helios rackscale solutions* hooked together via ethernet to create one (1) 720 GPU server. The retrofit scenario considers the same cooling architectures and equipment as the greenfield case.

The Reference Design includes information on four technical areas: facility power, facility cooling, IT space, and lifecycle software. These areas represent the integrated systems required to meet the design's specifications provided in this overview document.



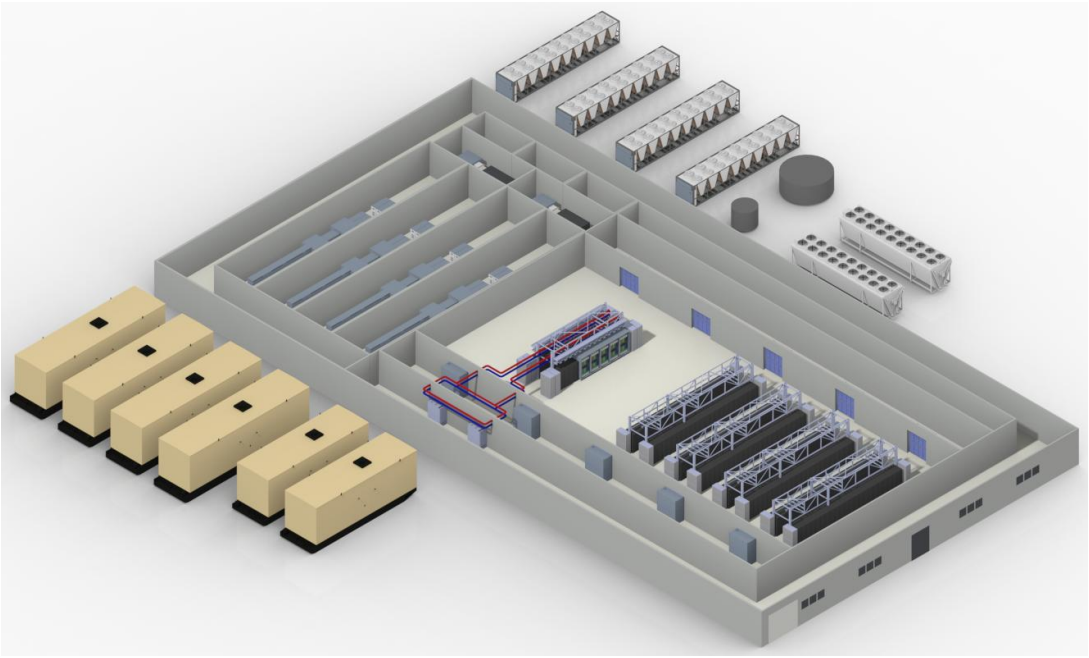
3D CAD Rendering of the data center for the greenfield situation and installed in Dallas, TX.

Annualized PUE at 100% Load (Greenfield)

City	Hot Aisle Containment (HAC)	Rear Door Heat Exchanger
San Francisco	1.124	1.121
Dallas	1.137	1.134
Dulles	1.125	1.122

Annualized PUE at 100% Load (Retrofit)

City	Hot Aisle Containment (HAC)	Rear Door Heat Exchanger
San Francisco	1.143	1.133
Dallas	1.185	1.175
Dulles	1.154	1.144



3D CAD Rendering of the data center for the retrofit situation and installed in Dallas, TX.

Facility Power – Greenfield

The facility power system supplies power to all components within the data center. In this concurrently maintainable electrical design, power to the IT space is supplied through four 3.75 MW powertrains. The four powertrains provide four makes three distributed redundant UPS power to the IT space, backed up by diesel generators. Each powertrain consists of a 5000-amp *QED-2* main switchboard feeding three 1250 kW *Galaxy VXL* UPS with 5 minutes of runtime in parallel and a 5000-amp *QED-2* distribution section. At this loading, the *Galaxy VXL* UPSs' overload capacity can manage *AMD Helios rackscale solutions* electrical design point (EDP) power peak. Downstream, these powertrains feed 1600A remote power panels that power the AI racks with four makes three redundancy and the support racks with 2N redundancy. Separately, in San Francisco and Dulles, two 2 MW powertrains feed the fan walls, chillers, and adiabatic dry coolers with 2N redundant power. Dallas requires a mechanical powertrain of slightly larger capacity at 2.5 MW. The mechanical powertrains also feed two 500 kW *Galaxy VL* UPS that provides critical power to the liquid-to-liquid coolant distribution units (CDUs) and facility water system pumps.

The facility power system is designed to support integrated peripheral devices like fire panels, access control systems, and environmental monitoring and control devices. Power meters in the electrical path monitor power quality and allow for predictive maintenance & diagnostics of the system. These meters also integrate with *EcoStruxure™ Power Monitoring Expert*.

Every component in this design is built and tested to the applicable IEEE or UL standards.

Further design details, such as dimensions, schematics, and equipment lists are available in the engineering package.

Design Options

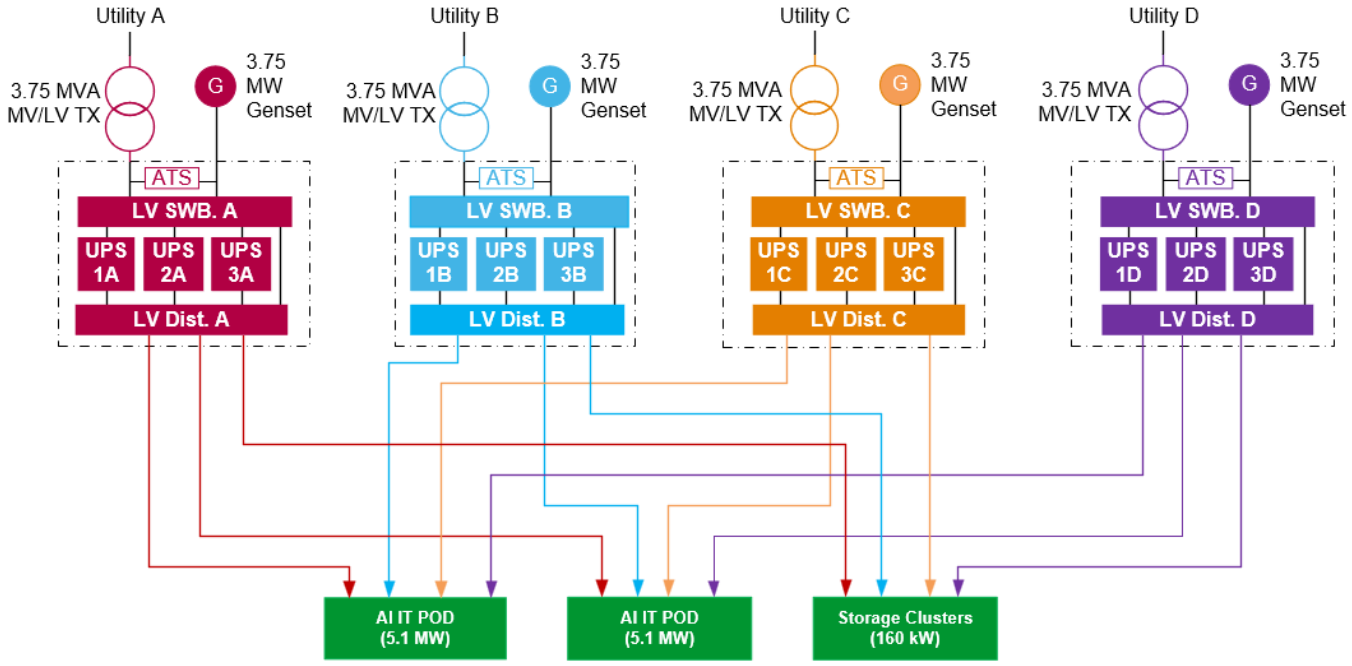
This reference design can be modified as follows without a significant effect on the design's performance attributes:

- Provision for load bank
- Change UPS battery type & runtime
- Add/remove/change standby generators:
  - Location & tank size

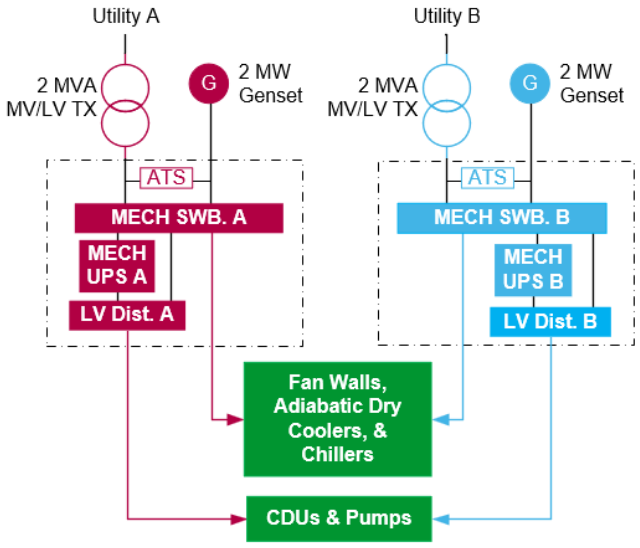
Facility Power Attributes for Greenfield

Name	Value	Unit
Total facility peak power (N, IT and cooling, HAC)	11,614 (San Francisco) 11,854 (Dallas) 11,750 (Dulles)	kW
Total amps (IT main bus, each)	5,000	A
IT space UPS capacity, per powertrain	3,750	kW
Total amps (Facility cooling bus, each)	2,500 (San Francisco, Dulles) 3,000 (Dallas)	A
Facility cooling UPS capacity	500	kW
Input voltage (IT main bus)	480	V
Switchboard kAIC (IT main bus)	65	kA
Generator redundancy (IT main bus)	Distributed redundant	-
IT power path	4 makes 3	-
IT space UPS redundancy	Distributed redundant	-
IT space UPS runtime @ rated load	5	minutes
IT space UPS output voltage	480	V
Input voltage (Facility cooling bus)	480	V
Switchboard kAIC (Facility cooling bus)	42	kA
Generator redundancy (Facility cooling bus)	2N	-
Facility cooling UPS redundancy	2N	-
Facility cooling UPS runtime @ rated load	5	minutes

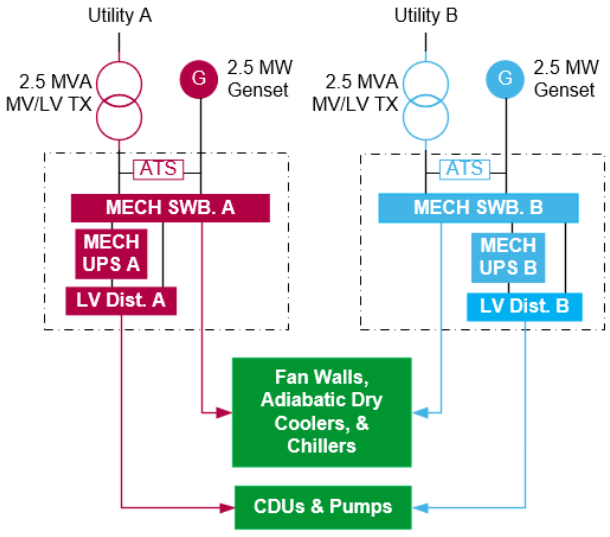
Facility Power Block Diagram – Greenfield



Mechanical Powertrain – San Francisco, CA & Dulles, VA



Mechanical Powertrain – Dallas, TX



Facility Power – Retrofit

The facility power system supplies power to all components within the data center. In this concurrently maintainable electrical design, power to the IT space is supplied through four 2.5 MW powertrains. The four powertrains provide four makes three distributed redundant UPS power to the IT space, backed up by diesel generators. Each powertrain consists of a 3000-amp *QED-2* main switchboard feeding two 1250 kW *Galaxy VXL* UPS with 5 minutes of runtime in parallel and a 3000-amp *QED-2* distribution section. At this loading, the *Galaxy VXL* UPSs' overload capacity can manage *AMD Helios rackscale solutions* electrical design point (EDP) power peak. Downstream, these powertrains feed 800A remote power panels that power the AI racks with four makes three redundancy and the support racks with 2N redundancy. To support existing IT racks, the powertrains feed 1000 kVA PDUs. Separately, two 2.5 MW powertrains feed the fan walls, chillers, and adiabatic dry coolers with 2N redundant power. The mechanical powertrains also feed two 200 kW *Galaxy VL* UPS that provides critical power to the liquid-to-liquid coolant distribution units (CDUs) and facility water system pumps.

The facility power system is designed to support integrated peripheral devices like fire panels, access control systems, and environmental monitoring and control devices. Power meters in the electrical path monitor power quality and allow for predictive maintenance & diagnostics of the system. These meters also integrate with *EcoStruxure™ Power Monitoring Expert*.

Every component in this design is built and tested to the applicable IEEE or UL standards.

Further design details, such as dimensions, schematics, and equipment lists are available in the engineering package.

Design Options

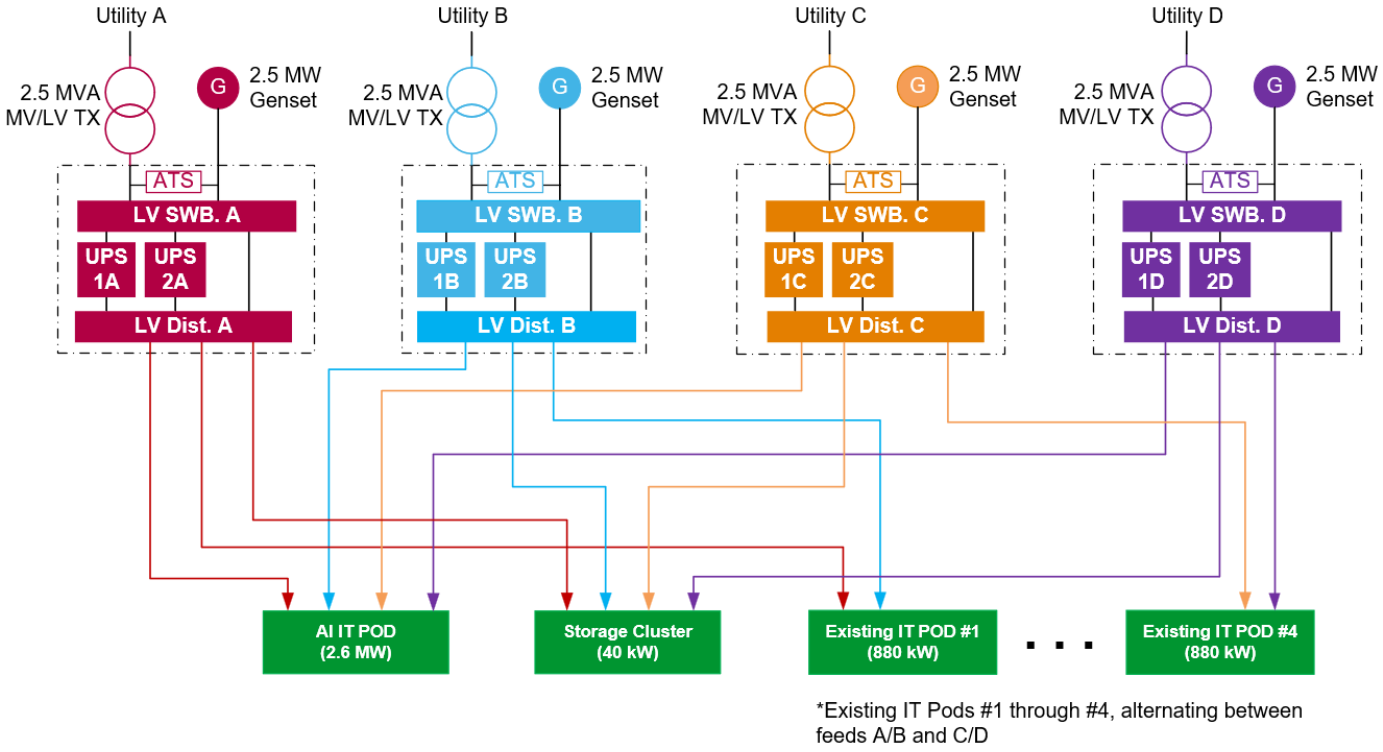
This reference design can be modified as follows without a significant effect on the design's performance attributes:

- Provision for load bank
- Change UPS battery type & runtime
- Add/remove/change standby generators:
  - Location & tank size

Facility Power Attributes for Retrofit

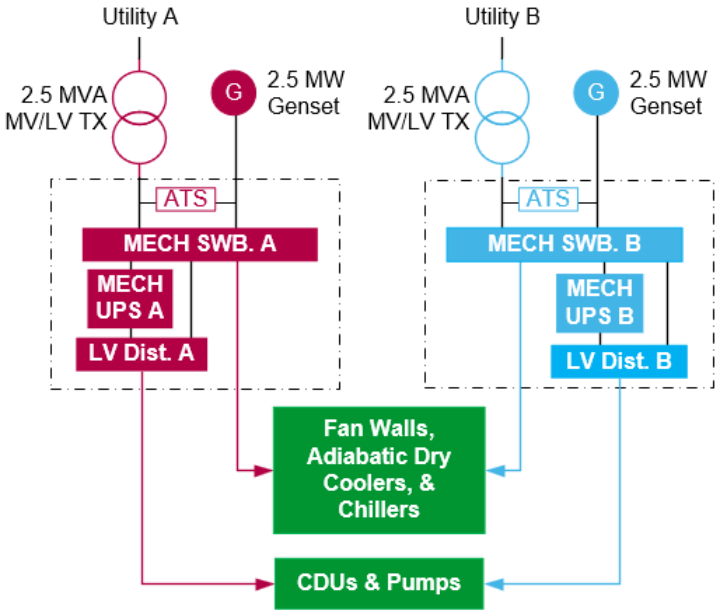
Name	Value	Unit
Total facility peak power (N, IT and cooling, HAC)	8,008 (San Francisco) 8,020 (Dallas) 7,903 (Dulles)	kW
Total amps (IT main bus, each)	3,000	A
IT space UPS capacity, per powertrain	2,500	kW
Total amps (Facility cooling bus, each)	3,000	A
Facility cooling UPS capacity	200	kW
Input voltage (IT main bus)	480	V
Switchboard kAIC (IT main bus)	65	kA
Generator redundancy (IT main bus)	Distributed redundant	-
IT power path	4 makes 3	-
IT space UPS redundancy	Distributed redundant	-
IT space UPS runtime @ rated load	5	minutes
IT space UPS output voltage	480	V
Input voltage (Facility cooling bus)	480	V
Switchboard kAIC (Facility cooling bus)	42	kA
Generator redundancy (Facility cooling bus)	2N	-
Facility cooling UPS redundancy	2N	-
Facility cooling UPS runtime @ rated load	5	minutes

Facility Power Block Diagram – Retrofit



*Existing IT Pods #1 through #4, alternating between feeds A/B and C/D

Mechanical Powertrain – All Locations



Facility Cooling

Hot Aisle Containment (HAC)

The facility cooling design features a dual path piping system optimized for data center efficiency. A chilled water loop integrates *Uniflair XCAF* chillers, with free cooling capabilities, to deliver 68°F (20°C) chilled water to fan walls in N+1 configuration. This lower temperature water loop handles the air-cooling needs of the data center.

Rear-door Heat Exchanger

A chilled water loop integrates *Uniflair XCAF* chillers, with free cooling capabilities, to deliver 68°F (20°C) chilled water to fan walls in N+1 configuration and rear-door heat exchangers in an N configuration. A reduced number of fan walls account for the heat dissipated in the white space by the equipment or the external environment. The rear-door heat exchangers account for most of the air-cooling needs of the data center. A centralized hot air duct system is not utilized because the rear-door heat exchangers exhaust cold air back to the room at a predetermined setpoint.

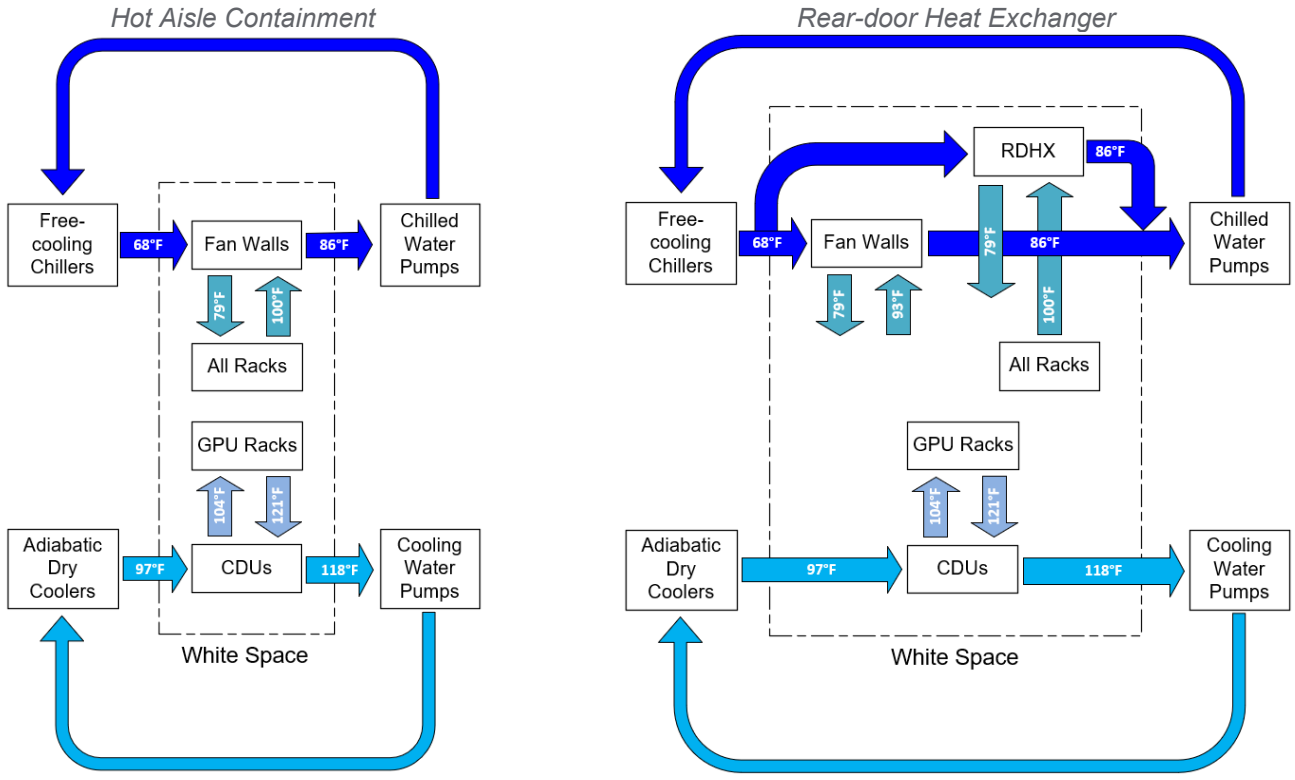
Common Architecture Elements

A separate high temperature water loop supplies 97°F (36°C) water to *Motivair Floor-Mounted Liquid-to-Liquid MCDU-70* used to cool the *AMD Helios rackscale solutions* that require liquid cooling. Outdoor heat rejection for this loop is supported by *Guntner GFD V-Shape* adiabatic dry coolers. With supply water temperatures approaching 104°F (40°C), dry coolers become useful heat rejection units in a wide range of climates. They can also help reduce energy and CapEx cost compared to chillers. The dry coolers use adiabatic assist to provide sufficient cooling for the data center when ambient temperatures exceed the dry bulb switchover temperature.

An integrated thermal storage system provides 5 minutes of continuous cooling at the maximum cooling capacity, in case of power outage, to allow the chillers and dry coolers to restart. The CDUs and facility pumps are on UPS power. More information on the fan wall, or rear-door heat exchanger, and CDU cooling architecture is detailed in the IT Space section of this document.

This design is instrumented to work with *EcoStruxure™ Building Operation* and *AVEVA Unified Operations Center*.

Facility Cooling Block Diagrams



Facility Cooling Attributes

Name		Greenfield	Retrofit	Unit
		Value		
Total max cooling capacity (N)		15,787 (San Francisco) 15,174 (Dallas) 15,502 (Dulles)	11,684 (San Francisco) 11,018 (Dallas) 11,617 (Dulles)	kW
Input voltage		480		V
Heat rejection medium		Water & propylene glycol mixture		-
Chiller redundancy		N+1		-
Dry cooler redundancy		N+1		-
CDU redundancy		N+1		-
Outdoor heat exchange		Dry cooler with adiabatic assist & packaged chiller		-
Chiller CW supply temperature		68 (20)		°F (°C)
Chiller CW return temperature		86 (30)		°F (°C)
Dry cooler CW supply temperature		96.8 (36)		°F (°C)
Dry cooler CW return temperature		118.4 (48)		°F (°C)
Storage tank size		20,000 (high temp. loop) 5,000 (low temp. loop)	10,000 (high temp. loop) 15,000 (low temp. loop)	gallon
Ride-through time		5		minutes
Outdoor ambient temperature range	San Francisco	33 to 101 (0.4 to 39)		°F (°C)
	Dallas	7.3 to 111 (-14 to 44)		
	Dulles	-3.1 to 104 (-20 to 40)		
Dry bulb switchover temperature	San Francisco	68.9 (20.5)		°F (°C)
	Dallas	81.8 (27.7)		
	Dulles	80.2 (26.7)		

Design Options

This reference design can be modified as follows without a significant effect on the design's performance attributes:

- Change storage tank size and subsequent ride-through time
- Reduce water consumption with different supply/return temperature for the dry coolers, different dry bulb switchover temperature, or the lack of adiabatic cooling
- Implement a single chilled water loop, configured with economizers, to serve both the air-cooled and liquid-cooled loads

IT Space – Greenfield

The IT space of this design consists of a single data hall. The data hall features up to forty 246 kW liquid cooled IT racks (modeled after *AMD Helios rackscale solution*) and 22 air cooled support racks with rack density of up to 45 kW. The liquid cooled racks remove 84% of the heat via liquid while 16% requires air. The racks are organized to two pods consisting of AI racks with some of the support racks, and one pod of the remaining support racks.

The AI pods are deployed with *Prefabricated Modular EcoStruxure™ Pod Data Center* to provide support for busway and/or cabling, and TCS piping. The AI pods may be deployed with either a 6 ft. wide ducted hot aisle containment for proper airflow to perimeter fan-walls via a common ceiling plenum, or without the hot aisle containment and rear-door heat exchangers for each rack. The support pod is deployed with *NetShelter Pod Containment* to provide support for busway and/or cabling. The same flexibility exists for the support pods, hot air may be captured with a ducted hot aisle and routed to perimeter fan-walls, or the hot air exhaust from each rack may be cooled by rear-door heat exchangers before return to the room.

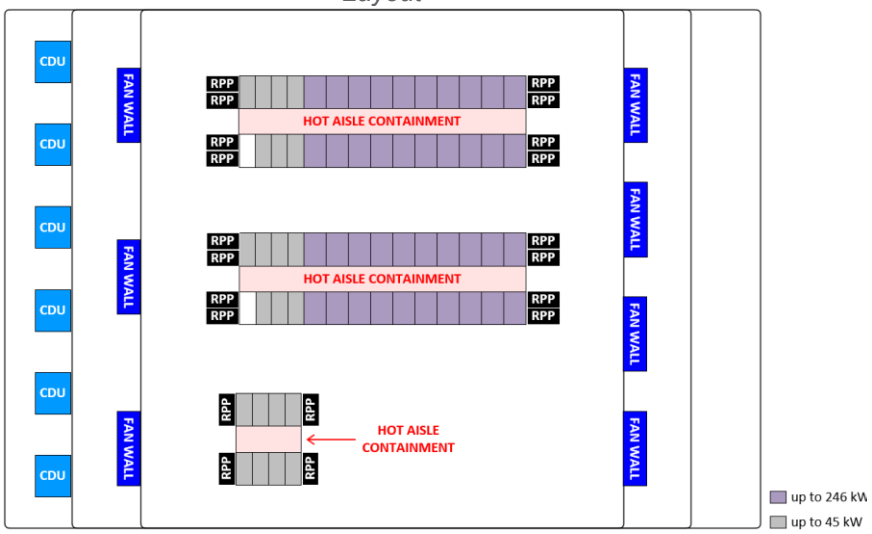
When implementing hot aisle containment, seven *Uniflair FWCV* chilled water fan walls supply conditioned air to the data hall in an N+1 configuration. Otherwise, with *Motivair M14* rear-door heat exchangers, two *Uniflair* FWCV chilled water fan walls supply conditioned air to the white space. Three *Motivair Floor-Mounted Liquid to Liquid MCDU-70s* provide precise liquid cooling to each AI pod with N+1 redundancy. This design utilizes a Technology Cooling System (TCS) loop optimized for efficiency and synchronized with *AMD Helios rackscale solution* infrastructure specifications. Prior to the adoption of any ASHRAE TCS classification and reference design, the loop must undergo exhaustive engineering analysis and high-fidelity thermal simulation. This validation is mandatory to ensure the thermal protection, operational reliability, and integrity of all liquid-cooled equipment within the data hall. Redundant piping systems across the data hall provides an alternate path for chilled water in case of cooling equipment failure or maintenance.

The support racks are *NetShelter SX* racks respectively configured with 2N 33 kW *NetShelter Power Shelfs*. The 246 kW AI racks are *Open Compute ORW* configured with 6+2 72 kW *HPRv2 Power Shelves*. Each row of the AI pods is powered by four *Hyper Solutions with Schneider Electric High Density 1600A RPPs*, providing A-, B-, C-, and D-side power to the row. The 246 kW AI racks have four 60 A feeds coming from each of the RPPs, for a total of 16 power feeds per rack from the RPPs – two for each power shelf. The *Hyper Solutions with Schneider Electric High Density 1600A RPP* includes *PowerLogic HDPM* for branch circuit power monitoring, *PowerLogic ION900T* for metering at input with transient capture, and *PowerPacT H-Frame* circuit breakers that can be configured with shunt trip units for those feeding liquid cooled racks. The pod of support racks is powered by the same RPPs associated with the AI pods.

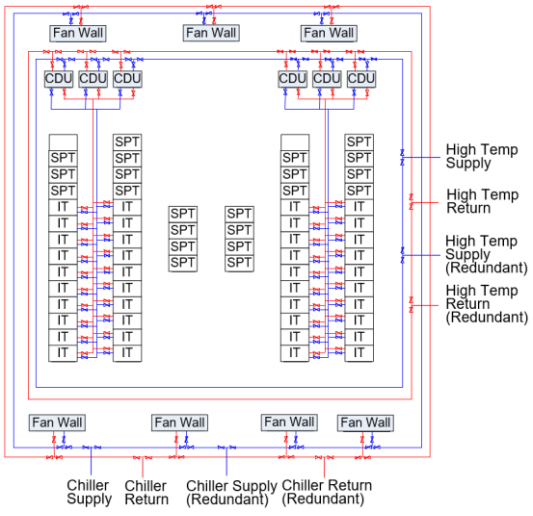
This design is instrumented to work with *EcoStruxure™ Building Operation* and *EcoStruxure™ IT Expert.*

IT Space Diagrams

Layout



Piping



Design Options

This reference design can be modified as follows without a significant effect on the design's performance attributes:

- Use *Uniflair FXCV* fan walls
- Use *NetShelter Aisle Containment*
- CRAHs can be selected instead of fan walls
- Use TCS loop surge tank to ensure TCS supply temperature stability
- AI rack power shelves may be configured as 5+5 (N+N)

IT Space Attributes

Name	Value	Unit
IT load	10,440	kW
Supply voltage to IT, AI racks	480	V
Supply voltage to IT, support racks	480	V
Rack power feed redundancy	Redundant	-
Number of 246 kW liquid-cooled racks	40	racks
Number of up to 45 kW support racks	22	racks
IT floor space	6,516	ft^2
CRAC/CRAH type	Fan wall or Rear-door heat exchanger & Fan wall	-
CRAC/CRAH redundancy	N+1	-
CRAC/CRAH supply air temperature	79 (26)	°F (°C)
CW supply temperature	68 (20)	°F (°C)
CW return temperature	86 (30)	°F (°C)
Containment type	Ducted hot aisle or None	-
CDU type	L2L	-
CDU redundancy	N+1	-
CDU CW supply temperature	96.8 (36)	°F (°C)
CDU CW return temperature	118.4 (48)	°F (°C)
TCS loop supply temperature	104 (40)	°F (°C)
TCS loop return temperature	120.7 (49.3)	°F (°C)

IT Space – Retrofit

The IT space of this design consists of a single data hall. The data hall features an existing setup of 176 racks with average density of 20 kW. These existing racks are organized into four pods of 44 racks each. The retrofit scenario considers that three of these 44-rack pods are removed and replaced by a single AI pod with up to ten 246 kW liquid cooled IT racks (modeled after *AMD Helios rackscale solution*) and 7 air cooled support racks with rack density of up to 45 kW. The liquid cooled racks remove 84% of the heat via liquid while 16% requires air.

The AI pod is deployed with *Prefabricated Modular EcoStruxure™ Pod Data Center* to provide support for busway and/or cabling, and TCS piping. The AI pod may be deployed with either a 6 ft. wide ducted hot aisle containment for proper airflow to perimeter fan-walls via a common ceiling plenum, or without the hot aisle containment and rear-door heat exchangers for each rack.

When implementing hot aisle containment, ten *Uniflair FWCV* chilled water fan walls supply conditioned air to the data hall in an N+1 configuration. Otherwise, with *Motivair M14* rear-door heat exchangers, two *Uniflair* FWCV chilled water fan walls supply conditioned air to the white space. Two *Motivair Floor-Mounted Liquid to Liquid MCDU-70s* provide precise liquid cooling to the AI pod with N+1 redundancy. This design utilizes a Technology Cooling System (TCS) loop optimized for efficiency and synchronized with *AMD Helios rackscale solution* infrastructure specifications. Prior to the adoption of any ASHRAE TCS classification and reference design, the loop must undergo exhaustive engineering analysis and high-fidelity thermal simulation. This validation is mandatory to ensure the thermal protection, operational reliability, and integrity of all liquid-cooled equipment within the data hall. Redundant piping systems across the data hall provides an alternate path for chilled water in case of cooling equipment failure or maintenance.

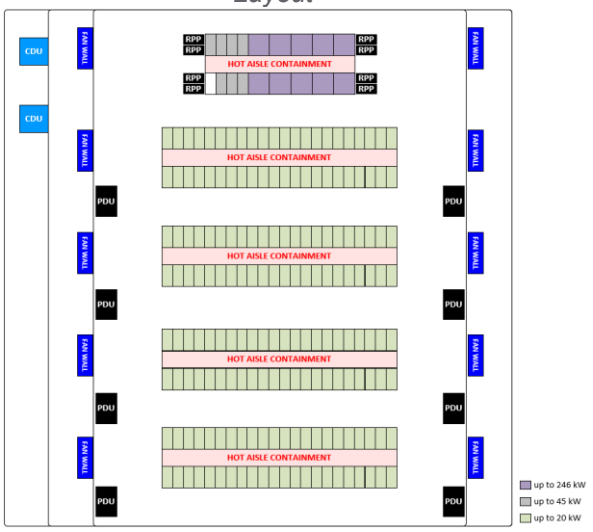
The support racks are *NetShelter SX* racks respectively configured with 2N 33 kW *NetShelter Power Shelfs*. The 246 kW AI racks are *Open Compute ORW* configured with 6+2 72 kW *HPRv2 Power Shelves*. Each row of the AI pods is powered by four *Hyper Solutions with Schneider Electric High Density 800A RPPs*, providing A-, B-, C-, and D-side power to the row. The 246 kW AI racks have four 60 A feeds coming from each of the RPPs, for a total of 16 power feeds per rack from the RPPs – two for each power shelf. The *Hyper Solutions with Schneider Electric High Density 800A RPP* includes *PowerLogic HDPM* for branch circuit power monitoring, *PowerLogic ION900T* for metering at input with transient capture, and *PowerPacT H-Frame* circuit breakers that can be configured with shunt trip units for those feeding liquid cooled racks.

The existing IT racks are configured with 1+1 60A *NetShelter* metered rack-mount power distribution units (rPDUs). Each row of the existing IT pods is powered by two *Hyper Solutions with Schneider Electric High Density 1200A* RPPs providing 2N redundant feeds. Two of the *1000 kVA Galaxy PDU* provide power to each pod, with each PDU responsible for the pair of RPPs at either end of each pod.

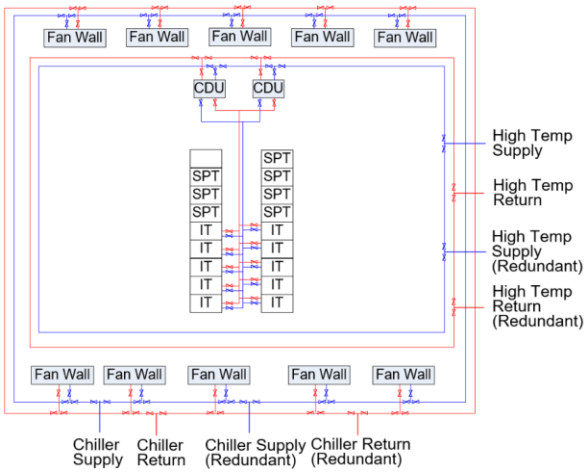
This design is instrumented to work with *EcoStruxure™ Building Operation* and *EcoStruxure™ IT Expert*.

IT Space Diagrams

Layout



Piping



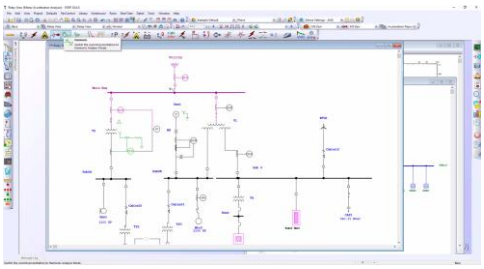
Design Options

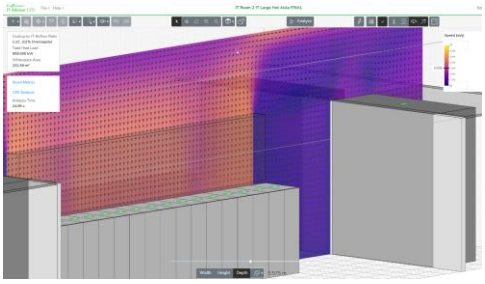
This reference design can be modified as follows without a significant effect on the design's performance attributes:

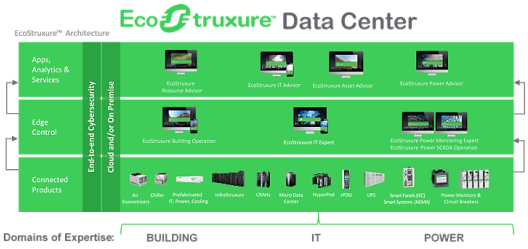
- Use *Uniflair FXCV* fan walls
- Use *NetShelter Aisle Containment*
- CRAHs can be selected instead of fan walls
- Use TCS loop surge tank to ensure TCS supply temperature stability
- AI rack power shelves may be configured as 5+5 (N+N)

IT Space Attributes

Name	Value	Unit
IT load	6,150	kW
Supply voltage to IT, AI racks	480	V
Supply voltage to IT, support racks	480	V
Supply voltage to IT, existing racks	415	V
Rack power feed redundancy	Redundant	-
Number of 246 kW liquid-cooled racks	10	racks
Number of up to 45 kW support racks	7	racks
Number of 20 kW IT racks	176	racks
IT floor space	10,336	ft^2
CRAC/CRAH type	Fan wall or Rear-door heat exchanger & Fan wall	-
CRAC/CRAH redundancy	N+1	-
CRAC/CRAH supply air temperature	79 (26)	°F (°C)
CW supply temperature	68 (20)	°F (°C)
CW return temperature	86 (30)	°F (°C)
Containment type	Ducted hot aisle or None	-
CDU type	L2L	-
CDU redundancy	N+1	-
CDU CW supply temperature	96.8 (36)	°F (°C)
CDU CW return temperature	118.4 (48)	°F (°C)
TCS loop supply temperature	104 (40)	°F (°C)
TCS loop return temperature	120.7 (49.3)	°F (°C)







Visit *EcoStruxure™ for Data Center* for more details.

Lifecycle Software

High-density AI clusters push the limits of data center facility infrastructure, so it's critical to leverage advanced planning and operation tools to ensure safe and reliable operations.

Planning & Design

Electrical Safety and Reliability: Due to the high amount of power supplied to an AI cluster, design specifications such as available fault current, arc flash hazards and breaker selectivity must be analyzed in the design phase. Applications like *Ecodial* and *ETAP* simulate the electrical design and reduce the chance of costly mistakes or even worse, injury.

Cooling: AI clusters are pushing the limits of what can be done with air-cooling. Modeling the IT space with computational fluid dynamics (CFD) helps spot issues including high pressure areas, rack recirculation, and hot spots. This is especially true when retrofitting an existing data center with an AI cluster. Schneider Electric's *EcoStruxure™ IT Design CFD* can quickly model airflow, allowing rapid iteration to find the best design and layout.

Operations

EcoStruxure™ is Schneider Electric's open, interoperable, integrated Internet of Things (IOT)-enabled system architecture and platform. It consists of three layers: connected products, edge control, and applications, analytics, and services.

EcoStruxure™ Data Center is a combination of three domains of *EcoStruxure™:* Power, Building, and IT. Each domain is focused on a subsystem of the data center: power, cooling, and IT. These three domains combined will reduce risks, increase efficiencies, and speed operations across the entire facility.

- *EcoStruxure™ Power* monitors power quality, generates alerts, while protecting and controlling the electrical distribution the electrical distribution system of the data center from the MV level to the LV level. It uses any device for monitoring and alerting, uses predictive analytics for increased safety, availability, and efficiency, while lowering maintenance costs.
- *EcoStruxure™ Building* controls cooling effectively while driving reliability, efficiency, and safety of building management, security, and fire systems. It performs data analytics on assets, energy use, and operational performance.
- *EcoStruxure™ IT* makes IT infrastructure more reliable and efficient while simplifying management by offering complete visibility, alerting and modelling tools. It receives data, generates alerts, predictive analytics, and system advice on any device to optimize availability and efficiency in the IT space.

There are several options for supervisory visibility and control. *AVEVA Unified Operations Center* can provide visibility at a site or across an entire enterprise.

Design Attributes

OVERVIEW		Value		Unit
		Greenfield	Retrofit	
Target availability		III		Tier
Annualized PUE		1.121 – 1.124	1.133 – 1.143	San Francisco, CA
		1.134 – 1.137	1.175 – 1.185	Dallas, TX
		1.122 – 1.125	1.144 – 1.154	Dulles, VA
Data center IT capacity		10,440	6,150	kW
Data center overall space		43,251	56,304	ft^2
Maximum rack power density		246		kW/rack
FACILITY POWER		**Value**		**Unit**
		Greenfield	**Retrofit**	
Total facility peak power (N, IT and cooling)	Hot Aisle Containment	11,614 (San Francisco)	8,008 (San Francisco)	kW
		11,854 (Dallas)	8,020 (Dallas)	
		11,750 (Dulles)	7,903 (Dulles)	
	Rear-door Heat Exchanger	11,584 (San Francisco)	7,944 (San Francisco)	
		11,823 (Dallas)	7,957 (Dallas)	
		11,720 (Dulles)	7,840 (Dulles)	
Total amps (IT main bus, each)		5,000	3,000	A
Input voltage (IT main bus)		480		V
Switchboard kAIC (IT main bus)		65		kA
Generator redundancy (IT main bus)		Distributed redundant		-
IT power path		4 makes 3		-
IT space UPS capacity, per powertrain		3,750	2,500	kW
IT space UPS redundancy		Distributed redundant		-
IT space UPS runtime @ rated load		5		minutes
IT space UPS output voltage		480		V
Total amps (facility cooling bus, each)		2,500 (San Francisco & Dulles) 3,000 (Dallas)	3,000	A
Input voltage (facility cooling bus)		480		V
Switchboard kAIC (facility cooling bus)		42		kA
Generator redundancy (facility cooling bus)		2N		-

Design Attributes (continued)

FACILITY COOLING		Value		Unit
		Greenfield	Retrofit	
Total max cooling capacity (N)	San Francisco	15,787	11,684	kW
	Dallas	15,174	11,018	
	Dulles	15,502	11,617	
Input voltage		480		V
Heat rejection medium		Water & propylene glycol mixture		-
Chiller redundancy		N+1		-
Dry cooler redundancy		N+1		-
CDU redundancy		N+1		-
Outdoor heat exchange equipment		Dry cooler with adiabatic assist & packaged chiller		-
Chiller CW supply temperature		68 (20)		°F (°C)
Chiller CW return temperature		86 (30)		°F (°C)
Fluid cooler CW supply temperature		96.8 (36)		°F (°C)
Fluid cooler CW return temperature		118.4 (48)		°F (°C)
Storage tank size	High temperature loop	20,000	10,000	gallon
	Low temperature loop	5,000	15,000	
Ride-through time		5		minutes
Outdoor ambient temperature range	San Francisco	33 to 101 (0.4 to 39)		°F (°C)
	Dallas	7.3 to 111 (-14 to 44)		
	Dulles	-3.1 to 104 (-20 to 40)		
Dry bulb switchover temperature	San Francisco	68.9 (20.5)		°F (°C)
	Dallas	81.8 (27.7)		
	Dulles	80.2 (26.7)		

Design Attributes (continued)

IT SPACE	Value		Unit
	Greenfield	Retrofit	
IT load	10,440	6,150	kW
Supply voltage to IT, AI racks	480		V
Supply voltage to IT, support racks	480		V
Supply voltage to IT, existing	Not applicable	415	V
Maximum rack power density	246		kW/rack
Number of racks	62	193	racks
IT floor space	6,516	10,336	ft^2
Rack power feed redundancy	Redundant		-
CRAC/CRAH redundancy	N+1		-
Containment type	Ducted hot aisle or None		-
CDU Type	L2L		-
CDU redundancy	N+1		-
CRAC/CRAH supply air temperature	79 (26)		°F (°C)
CW supply temperature	68 (20)		°F (°C)
CW return temperature	86 (30)		°F (°C)
CDU CW supply temperature	96.8 (36)		°F (°C)
CDU CW return temperature	118.4 (48)		°F (°C)
TCS loop supply temperature	104 (40)		°F (°C)
TCS loop return supply temperature	120.7 (49.3)		°F (°C)

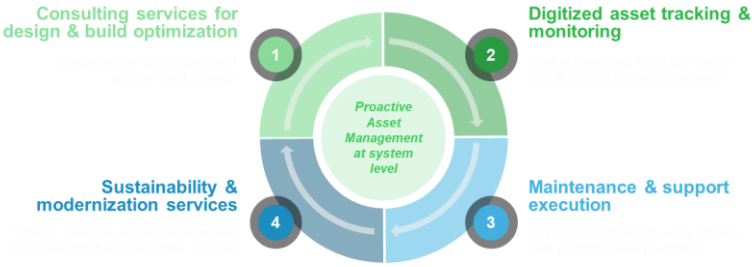
Schneider Electric Lifecycle Services

High-density AI clusters push the limits of data center facility infrastructure, so it's critical to leverage advanced planning and operations tools to ensure safe and reliable operations. To ensure cost-competitive and reliable data center operation over the data center's lifecycle, data center operators must adopt new ways of adapting to mounting operational complexity.

While AI Factories represent new opportunities for data centers, AI Factories will also bring about many challenges to data center operators. Data centers today are facing increasing OpEx burdens driven by maintaining critical power, cooling, and IT infrastructure. Meanwhile, staffing qualified engineers and technicians to oversee and service that infrastructure has gotten highly competitive, and all pieces of this infrastructure are set to be pushed to the cutting edge by AI workloads and variability. It is critical to build intelligent, reliable infrastructure and systems that can meet the challenges ahead and adapt to the next generation of data center computing.

Design and Build for the Entire Lifecycle of the AI Factory

Intelligent infrastructure starts at the design phase. While digitizing and monitoring critical assets is essential at every phase of a data center's lifecycle, ensuring digital-ready infrastructure at the outset of data center design avoids costly, disruptive shutdowns that may likely leave gaps in monitoring coverage. Designing with digital services in mind leads to a seamless system ready to deliver real-time insights from day one – meaning, a data center designed for the future.



EcoConsult

Whether a data center operator is using a brand-new facility, upgrading an existing facility, or repurposing a non-data center facility to be equipped to support an AI Factory, it is critical to leverage advanced planning and operations tools to ensure safe and reliable operations. EcoConsult is a comprehensive portfolio of best-in-class data center solutions, from hardware to software & services. Continuous reevaluation as the next generation of technology enters data centers is critical to ensuring success in the next phase of the data center evolution. Schneider Electric delivers an integrated engineering approach purpose-built for AI and high-density data centers, where power density, thermal constraints, and operational risk are amplified. **Design** and **system studies** validate electrical and thermal architectures upfront to support high-density GPU clusters, ensuring safety, power quality, coordination, and resilience before deployment. **Audits** establish a clear baseline of readiness for existing facilities being retrofitted for AI workloads, identifying risks tied to capacity, protection, and operational continuity. The **Electrical Digital Twin** connects design intent with real-world operations, providing a living model of the electrical network to manage rapid load growth, enable scenario analysis, and maintain reliability as AI demand scales.

Future Ready Data Center Life Extension with EcoFit

After evaluating existing infrastructure, addressing any outdated equipment should be a key priority. Outdated equipment and installations waste energy and increase operating costs, reducing efficiency. EcoFit helps data center operators solve these problems and modernize equipment and installations at the level needed to ensure a future ready data center. It enables high-performance digital technologies and supports your decarbonization journey by leveraging circularity services.

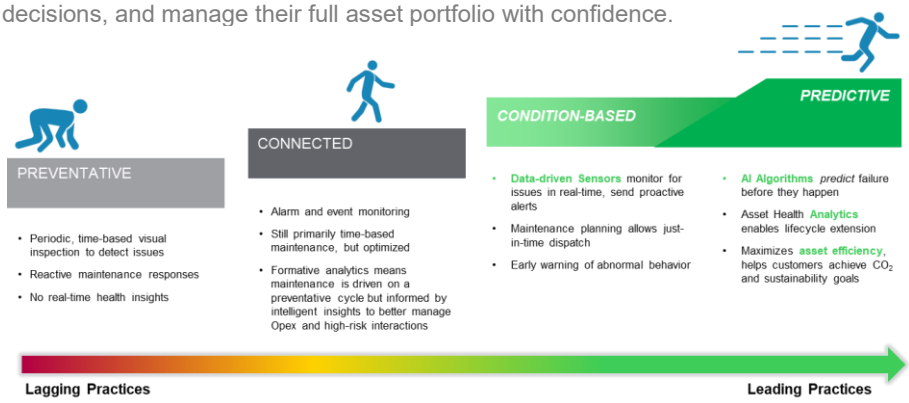
- *EcoFit Essential:* Upgrade your existing equipment with new features.
- *EcoFit Advanced:* Replace the core components with active and natively connectable new components.
- *EcoFit Replacement:* Get a comprehensive replacement with one of our latest equipment products/platforms/solutions.

Operate and Maintain with EcoCare

Get the most out of equipment installation through its entire lifecycle with EcoCare service plans for next-level predictive, preventive, and onsite maintenance. EcoCare from Schneider Electric is a subscription-based service plan offering 24/7 proactive remote monitoring, AI-powered predictive analytics through EcoStruxure, and expert on-site and remote support for electrical equipment and building management systems. EcoCare enables operators to revolutionize how they maintain their critical assets through condition-based maintenance.

Systemic AI-Driven Condition-Based Maintenance

Traditional time-based maintenance approaches are poorly suited for AI and high-density environments, where load profiles, thermal behavior, and failure modes change rapidly. Because time-based maintenance is disconnected from real-time asset health, it often results in over-maintenance, unnecessary interventions, and increased risk of human error. In contrast, AI-enabled condition-based maintenance is designed around continuous visibility into critical power and cooling infrastructure, using real-time data, advanced analytics, and strategic sensor placement to detect early failure indicators. This design-led approach enables proactive intervention before faults occur, reducing unplanned downtime, and allowing maintenance to be executed based on actual equipment condition rather than fixed intervals. A condition-based strategy leads to improving resilience, lowering risk, and optimizing operational efficiency in high-density AI data centers. The Health Index in EcoStruxure translates real-time asset data, AI analytics, and deep domain expertise into a single, intuitive view of equipment health across the entire data center. By combining simple, at-a-glance visibility with the ability to drill into asset-level detail, the Health Index enables engineers and technicians to prioritize risk, make informed maintenance decisions, and manage their full asset portfolio with confidence.



What are system-level analytics?

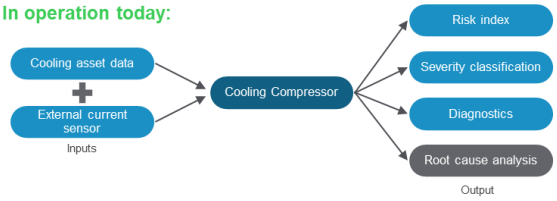
A data center is made up of systems within systems. Any solution monitoring critical assets within a data center needs to not only understand the performance of each individual asset but also the performance of its surrounding systems. Considering the entire system unlocks deeper understanding of critical data center infrastructure.

While this does mean evaluating enormous amounts of data, AI is enabling the change from an asset level approach to an end-to-end systemic approach. Schneider Electric's AI Hub is a global center of excellence that combines data science expertise with Schneider Electric's domain expertise to create valuable insights to data center systems. This leading hub looks to continue its intense growth to match pace with the need for technological & operational advancements within Data Centers.

Examples of system level analytics:



In operation today:



Data Aggregation and Collection

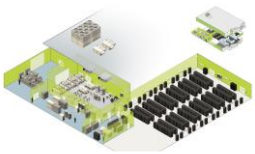
At the asset level, EcoStruxure aggregates data from a variety of intelligent sensors, including, thermal and environmental sensors for switchgear, partial discharge sensors, CDU controllers, oil transformer condition sensing, and leak detection for liquid cooling. Critical asset data can be collected through a variety of methods on-site. Schneider Electric solution architects and subject matter experts work with end users to ensure the necessary data is aggregated from the customer's BMS, EPMS, DCIM, and/or a corporate data lake. Collected data is then reviewed by AI-driven models and validated by Schneider Electric's subject matter experts, bringing decades of experience across power and cooling systems

Cybersecurity & Connectivity

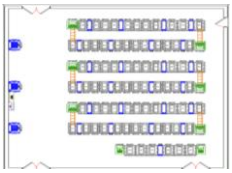
Cybersecurity is embedded by design across Schneider Electric's AI-ready CBM and analytics architecture, ensuring that the massive volumes of operational and performance data generated by high-density AI infrastructure can be securely collected, transported, and analyzed. From intelligent devices and sensors to edge and cloud platforms, security controls are integrated across every layer to protect system integrity, data availability, and performance.

All connected components and digital services are designed and maintained in accordance with ISO/IEC 27001 information security management practices and SOC 2 Type II compliance requirements for cloud-hosted CBM platforms. These frameworks govern our access control, data governance, & incident response methodologies to ensure protection of operational and performance data.

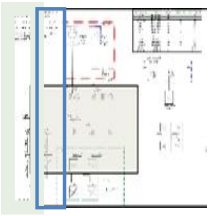
Get more information for this design:



3D spatial views



Floor layouts



One-line schematics



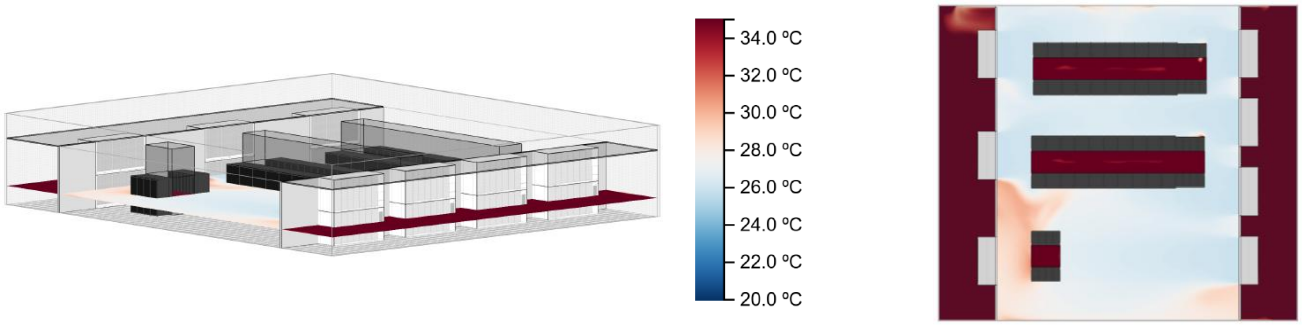
Bill of materials

Engineering Package

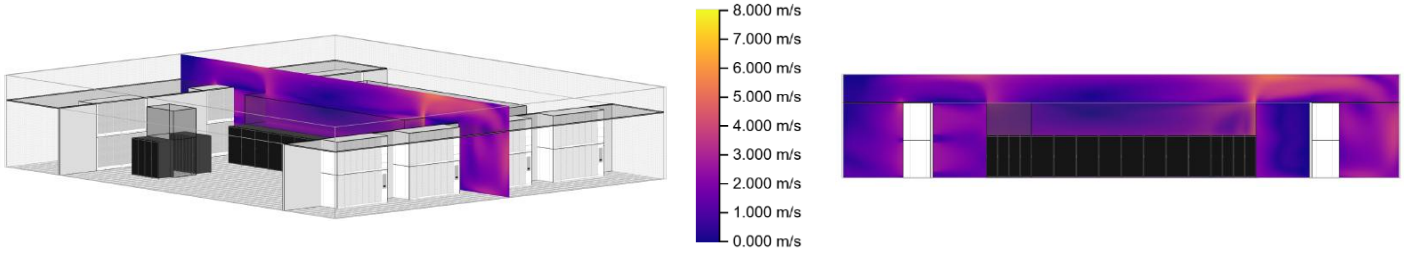
Every reference design is built with technical documentation for engineers and project managers. This includes engineering schematics (CAD, PDF), floor layouts, equipment lists containing all the components used in the design and 3D images showing real world illustrations of our reference designs.

Documentation is available in multiple formats to suit the needs of both engineers and managers working on data center projects.

Download the engineering package



Temperature distribution for the greenfield data hall with hot aisle containment as simulated by *EcoStruxure™ IT Design CFD* for Dallas, TX.



Air speed across the greenfield data hall with hot aisle containment as simulated by *EcoStruxure™ IT Design CFD* for Dallas, TX.

AMD
together we advance_AI

Email referencedesigns@se.com **for further assistance**