10-20 MW, Tier III, ANSI, Modular AI Factory Reference Design

Design Overview

Data Center IT Capacity
10,000 kW – 20,000 kW
4032 Vera Rubin NVL72 GPU

Target Availability
Tier III

Annualized PUE at Max-Q operating state
San Francisco, CA: 1.12
Dallas, TX: 1.13

Racks and Density
Total Racks:
- 56 AI racks
- 32 Networking racks

Rack Power Density:
- Networking racks up to 45 kW/rack
- AI racks up to 227 kW/rack

Data Center Overall Space
IT / White Space: 15,756 ft^2
Utilities Space: 17,762 ft^2
Total Deployment Unit Space: 33,518 ft^2

Regional Voltage and Frequency
480V, 60Hz

About this Design

- IT space and power distribution designed to accommodate AI clusters with density up to 227 kW per rack
- Design optimized to support liquid-cooled racks, with liquid-to-liquid coolant distribution units (CDUs) and fluid coolers with adiabatic assist
- 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 factories and liquid cooling bring new challenges to data center design, including unprecedented power density, thermal management requirements, and pressure to deploy capacity faster and more predictably. 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 these challenges. This design extends this approach through prefabrication, integrating key power, cooling, and compute white-space infrastructure into pre-engineered modular systems designed for factory assembly, testing, and repeatable deployment.

This design focuses on the deployment of high-density infrastructure for AI factories, specifically *NVIDIA's Vera Rubin NVL72* racks, within a single data hall. The data hall is purpose-built and optimized for four NVIDIA 1008 GPU Vera Rubin NVL72 clusters with a shared pod of networking racks, using *Motivair* floor-mounted liquid-to-liquid CDUs and fluid coolers with adiabatic assist. The data hall is also designed to enable various operating states of the AI racks to always maximize token generation on available infrastructure. Facility power and cooling design are optimized for capital cost, efficiency, reliability, and deployment flexibility. The modular architecture provides a repeatable building block that can help customers scale AI capacity while reducing the engineering and onsite integration required for subsequent deployments.

Reference Design 115 includes information on four technical areas: facility power, facility cooling, IT space with power and cooling, and lifecycle software and services. By coordinating these systems within a prefabricated, modular architecture we call Schneider Electric RapidScale, the reference design provides a practical blueprint for designing, manufacturing, deploying, and operating AI factory infrastructure at scale.

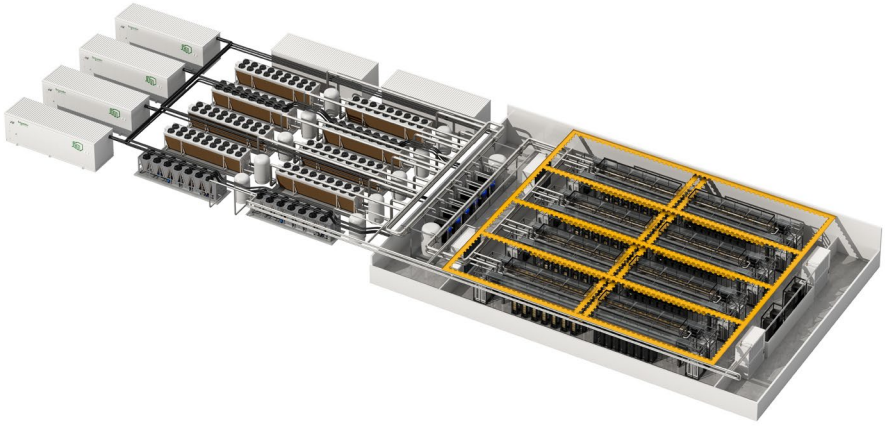

Data Hall Diagrams

Max-Q operating state, N+1 equipment active
11.4 MW IT capacity

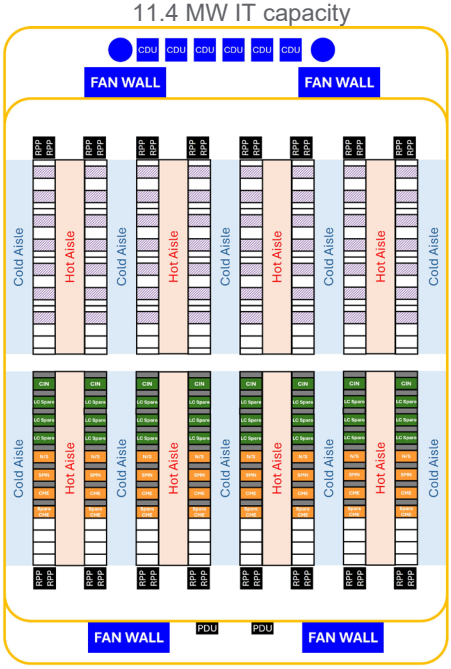

Max-P operating state, N equipment active
10.4 MW IT capacity

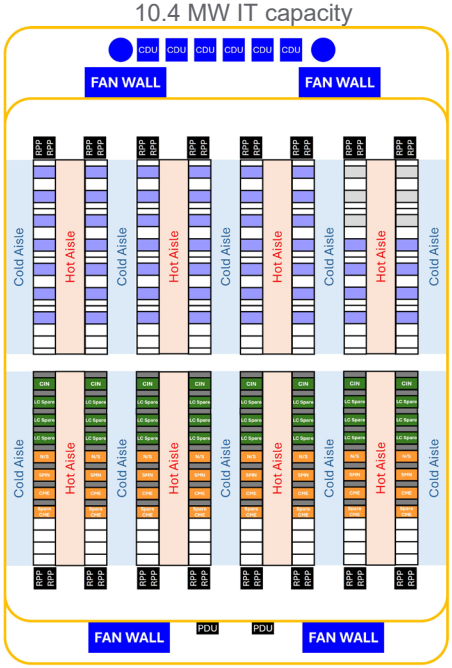

- Max-Q: ~188 kW
- Max-P: 227 kW
- Inactive
- Blank
- VCM
- N/S network: 15 kW
- SMN network: 15 kW
- CME nodes: 35 kW
- CIN network: 45 kW

Maximizing Data Hall Output

The AI racks have a range of operating points, this design is based around two of them: Max-P and Max-Q. Max-P is the maximum performance operating point of the AI rack generating the greatest number of tokens. For *NVIDIA's Vera Rubin NVL72* racks, it draws 227 kW at this operating point. Max-Q is the most efficient operating point of the AI rack on a token-per-Watt basis for the workload, drawing around 188 kW for *NVIDIA's Vera Rubin NVL72*. The data hall has been designed to dynamically operate the AI racks and leverage redundant capacity.

The normal operating state of this data hall can have all fifty-six AI racks operating at Max-Q by leveraging the redundant power and cooling equipment. This operating state also enables up to forty-six AI racks operating at Max-P, thereby fully utilizing the facility assets and maximizing token output. The AI racks have constant flow rate; it does not change with its operating point. Hence, the liquid-to-liquid cooling distribution unit (CDU) system is designed with redundancy for the maximum number of AI racks that can be active at a time.

The other operating state is where a smaller number of AI racks are operating, either up to forty-two AI racks at Max-P or up to fifty of the AI racks at Max-Q. This scenario is used if any redundant components (e.g., powertrain or cooling) are taken down for maintenance or in a failure scenario. Additionally, if the data center operator wants to always run with N+1 redundancy, then the GPU clusters can be downsized and always operated at Max-P. This is the basis of a standard data center design, as the electrical and air-cooling system is designed with redundancy for this operating state.

The networking racks are sized to the maximum number of AI racks possibly active. All networking racks are assumed to be active in both scenarios.

By dynamically operating the AI racks, this design enables the maximum IT capacity on the facility's power equipment. MaxLPS is another key operating point, representing the average rack power across multiple workloads, estimated at ~156 kW per rack. Like Max-Q, it can increase the number of active AI racks within the same power budget and support provisioning additional racks in the data hall or AI factory.

Operating States Attributes

Name	Max-P	Max-Q	MaxLPS
Maximum IT capacity active	10,414 kW	11,408 kW	9,616 kW
Operating point of AI racks	227 kW	188 kW	156 kW
Number of AI racks active	42	56	56
Total capacity of AI racks	9,534 kW	10,528 kW	8,736 kW
Total capacity of networking racks active	880 kW	880 kW	880 kW
Minimum number of active IT electrical powertrain	3 out of 4	4 out of 4	3 out of 4
Minimum number of active air-cooling equipment	3 out of 4 fan walls, 1 out of 2 chillers	3 out of 4 fan walls, 1 out of 2 chillers	3 out of 4 fan walls, 1 out of 2 chillers
Minimum number of active liquid-cooling equipment	6 out of 8 L2L CDUs	6 out of 8 L2L CDUs	6 out of 8 L2L CDUs

Facility Power

Facility Power Block Diagram

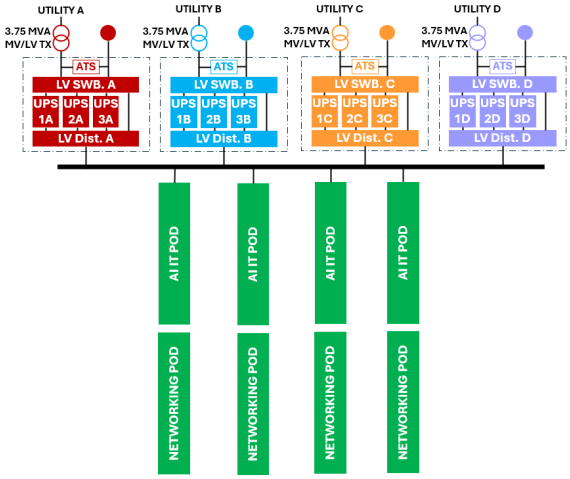

Mechanical Powertrain

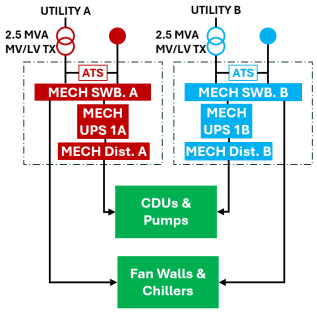

Design Options

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

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 RapidScale powertrains. The four powertrains provide 3+1 distributed redundant UPS power to the IT space. Each powertrain consists of a 5000-Ampere *QED-2* main switchboard feeding three 1250 kW *Galaxy VXL* UPS with 5 minutes of runtime in parallel and a 5000-Ampere *QED-2* distribution section. At this loading, the *Galaxy VXL* UPSs' overload capacity can manage *NVIDIA's Vera Rubin NVL72* racks' electrical design point (EDP) power peak. Downstream, these powertrains feed remote power panels (RPP) that power the AI racks or power distribution units (PDUs) that power the networking racks with 3+1 redundancy. Separately, two 2.5 MW RapidScale powertrains feed the fan walls, chillers, and fluid coolers with 2N redundant power. They also feed either a 750 kW *Galaxy VXL* UPS (Dallas) or 625 kW *Galaxy VXL* UPS (San Francisco) that provides critical power to the liquid-to-liquid coolant distribution units (CDUs) and facility water system pumps. The *Galaxy VXL UPS* can also be configured to provide Fault-Ride-Through capability.

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 required UL product standards. Further design details, such as dimensions, schematics, and equipment lists are available upon request.

Facility Power Attributes

Name	Value	Unit
Total facility peak power (N, IT and cooling)	13,250	kW
Total amps (IT main bus, each)	5,000	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	3+1	
IT space UPS capacity, per powertrain	3,750	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	A
Input voltage (Facility cooling bus)	480	V
Switchboard kAIC (Facility cooling bus)	42	kA
Generator redundancy (Facility cooling bus)	2N	
Facility cooling UPS capacity	750 (Dallas) 625 (San Fran.)	kW
Facility cooling UPS redundancy	2N	
Facility cooling UPS runtime @ rated load	5	minutes

Facility Cooling

Facility Cooling Block Diagrams

Facility Water System Flow

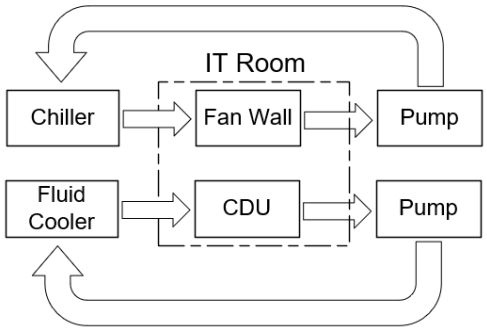

Facility Piping Diagram

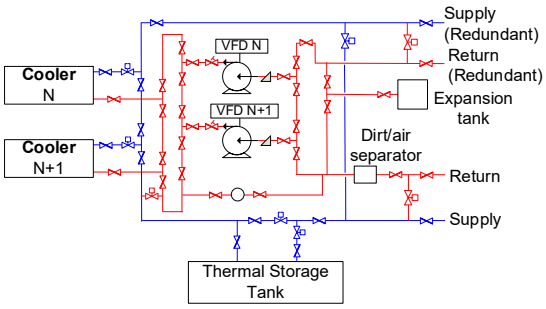

Similar piping architecture for chilled-water loop and fluid cooler loop

Design Options

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

- Change storage tank size
- Reduce water consumption with different operating points of fluid coolers

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 64°F chilled water to fan walls in N+1 configuration. This lower temperature water loop handles the air-cooling needs of the data center. A separate high-temperature water loop supplies 108°F water to RapidScale skidded *Motivair Liquid-to-Liquid MCDUs* used to cool *NVIDIA's Vera Rubin NVL72* liquid-cooled racks. Outdoor heat rejection for this loop is supported by *Guntner GFD V-Shape* dry coolers. With supply water temperatures approaching 108°F, 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 at temperatures higher than the dry-bulb switchover temperature.

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

This design is instrumented to work with any building management system, such as *EcoStruxure™ Building Operation.*

Further design details such as dimensions, schematics, and equipment lists are available upon request.

Facility Cooling Attributes

Name	Value	Unit
Total max cooling capacity (N)	16,247 (Dallas) 16,931 (San Francisco)	kW
Input voltage	480	V
Heat rejection medium	Water & propylene glycol mixture	
Chiller redundancy	N+1	
Outdoor heat exchange	Fluid cooler with adiabatic assist & packaged chiller	
Chiller CW supply temperature	64	°F
Chiller CW return temperature	78	°F
Fluid cooler CW supply temperature	100	°F
Fluid cooler CW return temperature	118	°F
Ride-through time	5	minutes
Outdoor ambient temperature range	7 to 110.8	°F
Dry-bulb switchover temperature	81.8 (Dallas) 69.0 (San Francisco)	°F

IT Space Diagrams

Max-Q operating state, N+1 equipment active
11.4 MW IT capacity

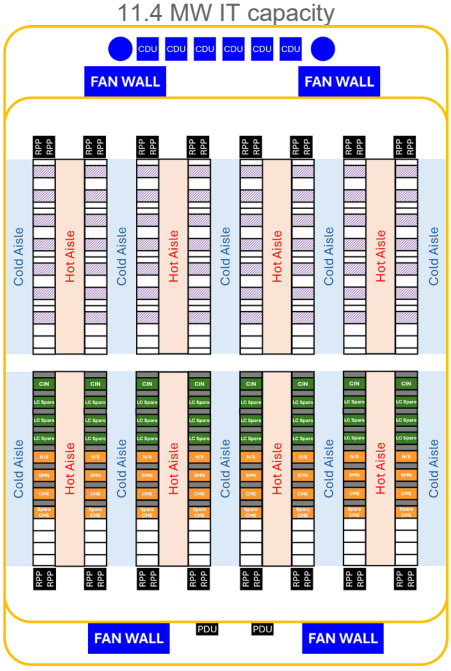

- Max-Q: ~188 kW
- Blank
- VCM
- N/S network: 15 kW
- SMN network: 15 kW
- CME nodes: 35 kW
- CIN network: 45 kW

Data Hall Piping

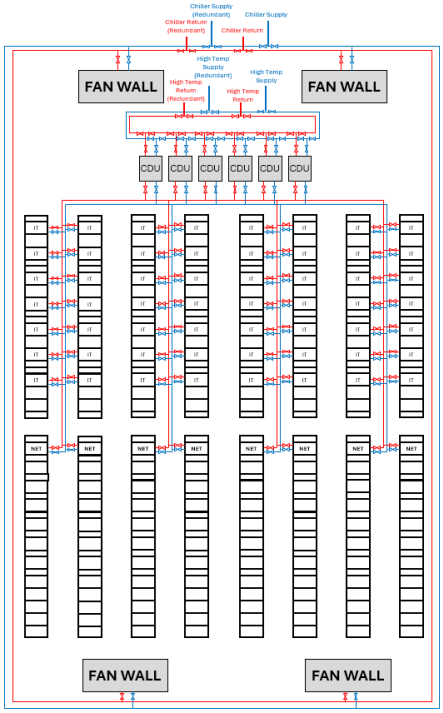

IT Space

The IT space of this design consists of a single data hall. The data hall features up to fifty-six 227 kW liquid-cooled IT racks (modeled after *NVIDIA*'s *Vera Rubin NVL72*), eight 15 kW air-cooled N/S and SMN network racks, eight 35 kW air-cooled CME nodes network racks, and eight 45 kW liquid-cooled CIN network racks. The liquid-cooled racks remove 96% of the heat via liquid while 4% requires air cooling. The racks are organized to one pod consisting of only AI racks and one pod of networking racks.

Each AI pod and networking pod is deployed with RapidScale IT Pods to provide a 6 ft. wide ducted hot aisle containment for proper airflow, RPPs, and cabling support, and TCS piping for the AI pods. Ducted hot aisles and a common ceiling plenum return hot air to the fan walls for cooling.

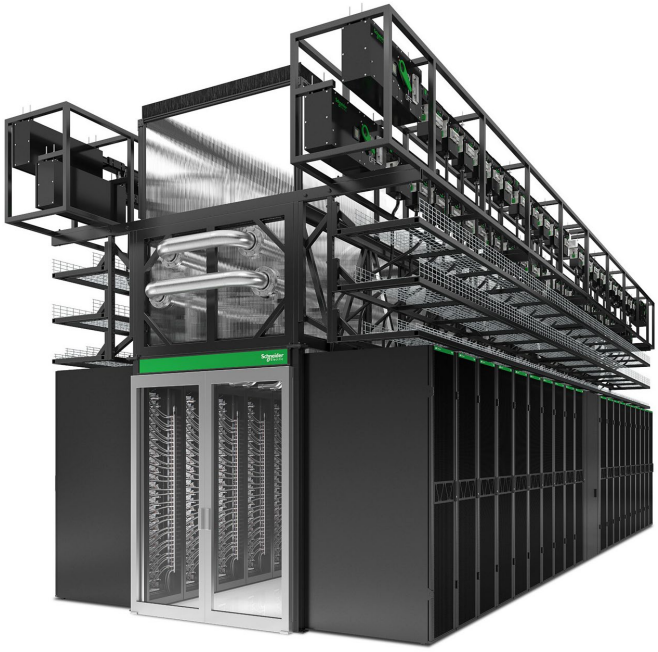

RapidScale IT Pod

Four *Uniflair FWCV* chilled-water fan walls supply conditioned air to the data hall in an N+1 configuration. Six RapidScale *Motivair Skidded Liquid-to-Liquid MCDUs* provide precise liquid cooling to all liquid-cooled racks in the data hall with N+1 redundancy. This design utilizes a Technology Cooling System (TCS) loop governed by ASHRAE S-45 class and synchronized with Vera Rubin 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.

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
- CRAHs can be selected instead of fan walls
- Variations in AI cluster configuration
- Use TCS loop surge tank to ensure TCS supply temperature stability

The 15 kW N/S and SMN network racks are *NetShelter Open Architecture Rack MGX 1.1* racks configured with 1+1 33 kW *NetShelter Power Shelves*. The 35 kW CME node racks and 45 kW CIN network racks are *NetShelter Open Architecture Rack MGX 1.1* racks configured with 2+2 33 kW *NetShelter Power Shelves*. The 227 kW AI racks are *NetShelter Open Architecture Rack MGX 1.2* configured with 3+1 110 kW *NetShelter 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 227 kW AI racks have a pair of 100 A feeds coming from each of the RPPs, for a total of eight 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 pods of networking racks are powered by two *Galaxy PDU 500 kVA* feeding eight *Galaxy RPPs*. The 15 kW SMN and N/S network racks have a pair of 60 A feeds from two of the four *Galaxy RPP*s for 2N redundant feeds. The 35 kW CME node and 45kW CIN network racks have two pairs of 60 A feeds from two of the four *Galaxy RPP*s for 2N redundant feeds.

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

IT Space Attributes

Name	Value	Unit
IT load	10,414 – 11,408	kW
Supply voltage to IT, AI racks	480	V
Supply voltage to IT, networking racks	415	V
Rack power feed redundancy	Redundant	
Number of 227 kW liquid-cooled racks	56	racks
Number of 15 kW N/S network racks	8	racks
Number of 15 kW SMN network racks	8	racks
Number of 35 kW CME nodes rack	8	racks
Number of 45 kW CIN network racks	8	racks
IT floor space	5,486	ft^2
CRAC/CRAH type	Fan wall	
CRAC/CRAH redundancy	N+1	
CRAC/CRAH supply air temperature	77	°F
CW supply temperature	64	°F
CW return temperature	78	°F
Containment type	Ducted hot aisle	
CDU type	L2L	
CDU redundancy	N+1	
CDU CW supply temperature	100	°F
CDU CW return temperature	118	°F
TCS loop supply temperature	108	°F
TCS loop return temperature	126	°F

Lifecycle Software

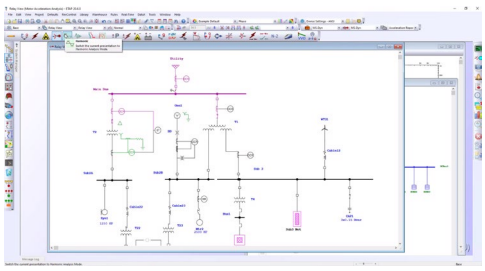

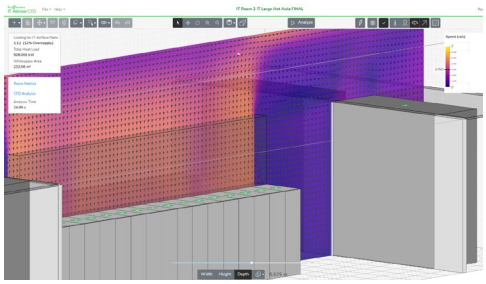

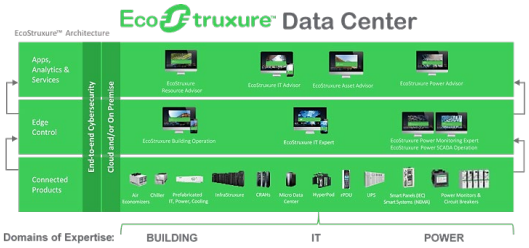

Visit *EcoStruxure™ for Data Center* for more details.

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, perform required power system analyses (e.g., short-circuit and arc flash), 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 system of the data center from the MV level to the LV level. It uses any device for monitoring and alerting, performs 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
Target availability	III	Tier
Annualized PUE at Max-Q operating state	1.11	San Francisco, CA
	1.12	Dallas, TX
Data center IT capacity	10,414 – 11,408	kW
Total Deployment Unit Space:	33,518	ft^2
Maximum rack power density	227	kW/rack
FACILITY POWER	**Value**	**Unit**
Total facility peak power (N, IT and cooling)	13,250	kW
Total amps (IT main bus, each)	5,000	A
Input voltage (IT main bus)	480	V
Switchboard kAIC	65	kA
Generator redundancy (IT main bus)	Distributed redundant	
IT Power path	3+1	
IT space UPS capacity, per powertrain	3,750	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	A
Input voltage (facility cooling bus)	480	V
Switchboard kAIC (facility cooling bus)	42	kA
Generator redundancy (facility cooling bus)	2N	
FACILITY COOLING	**Value**	**Unit**
Total max cooling capacity	16,247 (Dallas); 16,931 (San Francisco)	kW
Input voltage	480	V
Heat rejection medium	Water & propylene glycol mixture	
Chiller redundancy	N+1	
Outdoor heat exchange	Fluid cooler with adiabatic assist & packaged chiller	
Chiller CW supply temperature	64	°F
Chiller CW return temperature	78	°F
Fluid cooler CW supply temperature	100	°F
Fluid cooler CW return temperature	118	°F
Ride-through time	5	minutes
Outdoor ambient temperature range	7 to 110.8	°F
Dry-bulb switchover temperature	81.8 (Dallas); 69.0 (San Francisco)	°F

Design Attributes continued

IT SPACE	Value	Unit
IT load	10,414 – 11,408	kW
Supply voltage to IT, AI racks	480	V
Supply voltage to IT, networking racks	415	V
Maximum rack power density	227	kW/rack
Number of racks	56	AI racks
	32	Networking racks
IT floor space	15,756	ft^2
Rack power feed redundancy	Redundant	
CRAC/CRAH redundancy	N+1	
Containment type	Prefabricated, ducted hot aisle	
CDU Type	L2L	
CDU redundancy	N+1	
CRAC/CRAH supply air temperature	77	°F
CW supply temperature	64	°F
CW return temperature	78	°F
CDU CW supply temperature	100	°F
CDU CW return temperature	118	°F
TCS loop supply temperature	108	°F
TCS loop return supply temperature	126	°F

Schneider Electric Services for the AI Factory

Lifecycle services focused where AI Factories spend most of their life: in operation

Operation is the longest and most costly phase of the data center lifecycle, often extending 10 to 20 years, while design and construction are completed in months. In an AI Factory, high-density infrastructure makes electrical distribution, critical power, liquid cooling, controls, and IT operations closely interdependent throughout that period. EcoCare supports this phase through condition-based maintenance, drawing on the sensing and connectivity specified during design.

Assess and Design with EcoConsult

Early design decisions set an AI Factory's capacity, resilience, safety, efficiency, and maintainability for the life of the facility. EcoConsult brings Schneider Electric power, cooling, and data center expertise into design, turning AI workload requirements into an infrastructure strategy and surfacing constraints and tradeoffs early. For this reference architecture, its studies can assess:

- Electrical and cooling capacity, redundancy, and resilience
- Electrical safety, protection coordination, and power quality
- IT-space thermal performance and energy efficiency
- Maintainability, future growth, and monitoring requirements

Design sensing in at Day 0. Sensing, metering, and secure connectivity specified in the design are commissioned with the equipment, so they record assets' healthy baseline from startup, enabling predictive analytics and condition-based maintenance.

Simulate and Build with Digital Twins and Commissioning

Digital twins let teams validate design assumptions, uncover constraints, and optimize system performance before anything is built. The ETAP Electrical Digital Twin models AI Factory power from grid to chip, including dynamic load behavior at the chip level, and can feed a broader AI Factory digital twin that adds thermal and airflow (CFD) modeling. It supports:

- Load flow, capacity utilization, and future power requirements
- Short-circuit, protection coordination, and arc-flash studies
- Power quality and harmonics analysis
- Dynamic what-if scenarios that expose system constraints early

With the design validated, Schneider Electric services carry engineering intent into the installed facility through configuration verification, controls and integration validation, commissioning, and an accurate as-built digital baseline, creating continuity between design and operation and giving EcoCare its starting point.

Operate and Maintain with EcoCare

Once the facility is live, availability depends on seeing the condition of interdependent power and cooling systems before a problem reaches the AI workload. EcoCare, Schneider Electric's service plan for the operating phase, works from the sensing designed into the facility to move operators from reactive and calendar-based maintenance to a condition- and risk-based strategy. It also backs site teams with Schneider Electric experts as skilled staff grow harder to find.

EcoCare brings together data from electrical distribution, UPS systems, and power-quality metering, plus chillers, fan walls, CDUs, pumps, liquid cooling, and environmental sensors through integrations with BMS, DCIM, and/or SCADA systems, and turns it into action through:

- Monitor: 24/7 remote monitoring and expert interpretation through the Connected Services Hub
- Analyze: AI-powered analytics that spot early signs such as overheating or partial discharge and rank risk
- Recommend: maintenance and mitigation actions timed to asset condition
- Respond: SLA-backed technical support, emergency dispatch, and on-site service
- Improve: a customer success team, success plan, and training for site teams

Modernize and Upgrade with EcoFit

Over its operating life, an AI Factory will host several generations of AI hardware, pushing rack density, power demand, and cooling requirements higher. EcoFit extends the life of installed equipment and reduces its environmental footprint through Life Extension Essential and Advanced offers, with retrofit or replacement where that is the better investment. EcoCare condition data and maintenance history show what to modernize, when, and where investment returns the most.

Designing AI Factories for Predictive Maintenance

Digital Services and EcoCare turn connected infrastructure data into expert-guided service action.

Predicting failures takes more than monitoring: it takes a secure data path, analytics, expert interpretation, and field service. Schneider Electric Digital Services and EcoCare provide all four, but they can only work with what the facility measures, so for this reference architecture the design should include:

- Continuous thermal monitoring of MV and LV switchgear, busway, and transformer connections
- Partial discharge and environmental (temperature and humidity) monitoring for MV equipment
- Breaker diagnostics and power-quality metering on incoming and critical feeders
- Networked monitoring of UPS systems, batteries, and PDUs
- Cooling telemetry from chillers, CDUs, pumps, and fan walls through the BMS and DCIM
- A secure, segmented OT network with gateways and IEC 62443-certified devices
- Asset register and as-built configuration data captured at commissioning

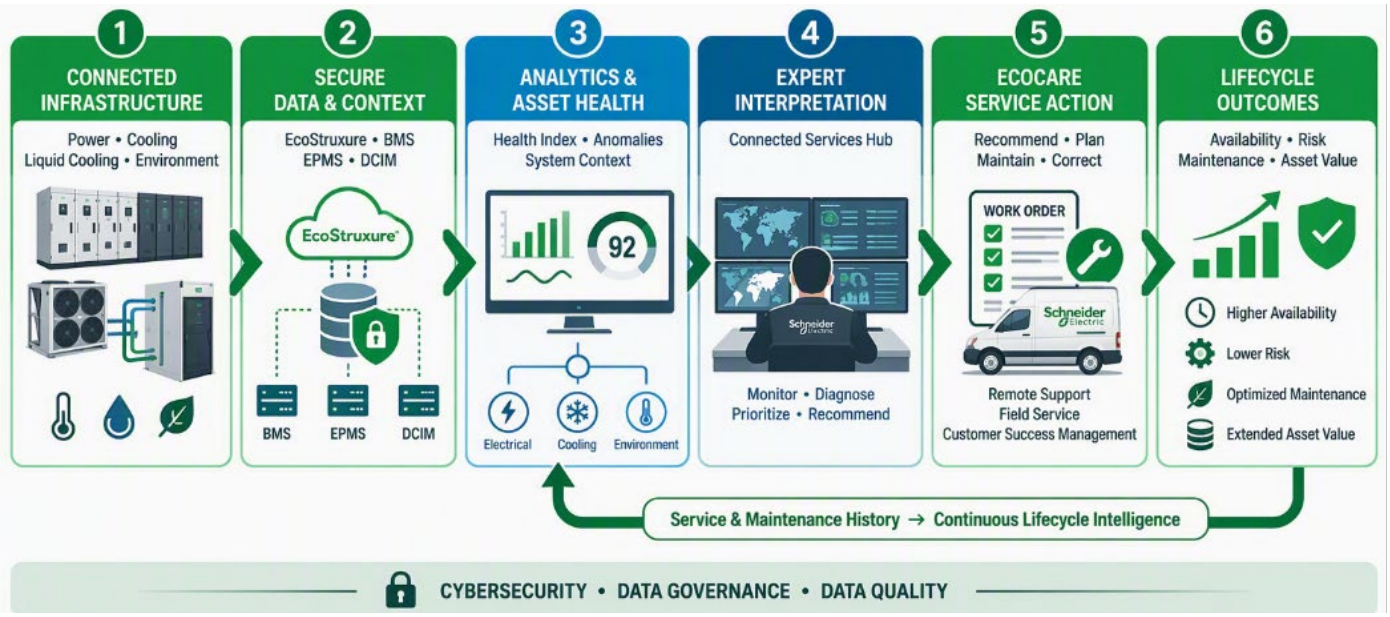

CONDITION-BASED SERVICE PROCESS	
1. Connected Infrastructure The best condition-based maintenance and AI-driven analytics start at design. Schneider Electric specifies the sensing, connectivity, and cybersecurity architecture across critical power and cooling systems, so connected assets deliver trusted data from commissioning onward.	**2. Secure Data & Context** Data from connected assets and applicable BMS, DCIM, and SCADA sources flows into Schneider Electric's secure cloud, where context turns isolated alarms into a view of equipment condition and performance.
3. Analytics & Asset Health Analytics flag abnormal conditions, emerging trends, and shifts in asset health. Outputs can include anomaly alerts, the Asset Health Index, the Maintenance Index, and asset- or system-level analysis showing what needs attention and why.	**4. Expert Interpretation** Connected Services Hub experts pair AI-driven analytics with Schneider Electric domain knowledge to validate alerts, diagnose root causes, assess potential impact, and recommend prioritized mitigation or maintenance actions.
5. EcoCare Service Action Service teams act on those recommendations, resolving what they can remotely and dispatching targeted on-site work before issues become failures. This results in fewer intrusive interventions, shorter MTTR, and more first-time fixes.	**6. Lifecycle Outcomes** This closed-loop process helps protect availability and resilience, reduce the risk of unplanned failures, focus technicians where they're needed, extend asset life, and base modernization on condition data.

Get more information for this design:

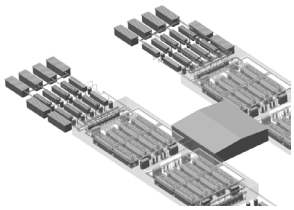

3D spatial views

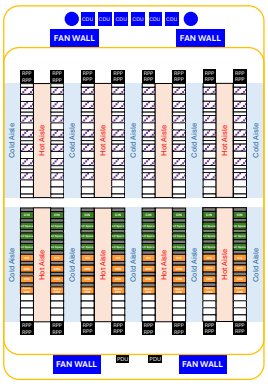

Floor layouts

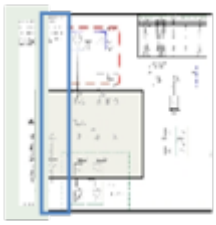

One-line schematics

Bill of materials

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.

Contact Schneider Electric for more information

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