10.4 – 11.4 MW, Tier III, ANSI, Chilled Water, Liquid-Cooled AI Clusters (NVIDIA Vera Rubin)

Design Overview

Data Center IT Capacity
10,414 kW – 11,408 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
37,391 ft^2

Regional Voltage and Frequency
480V, 60Hz

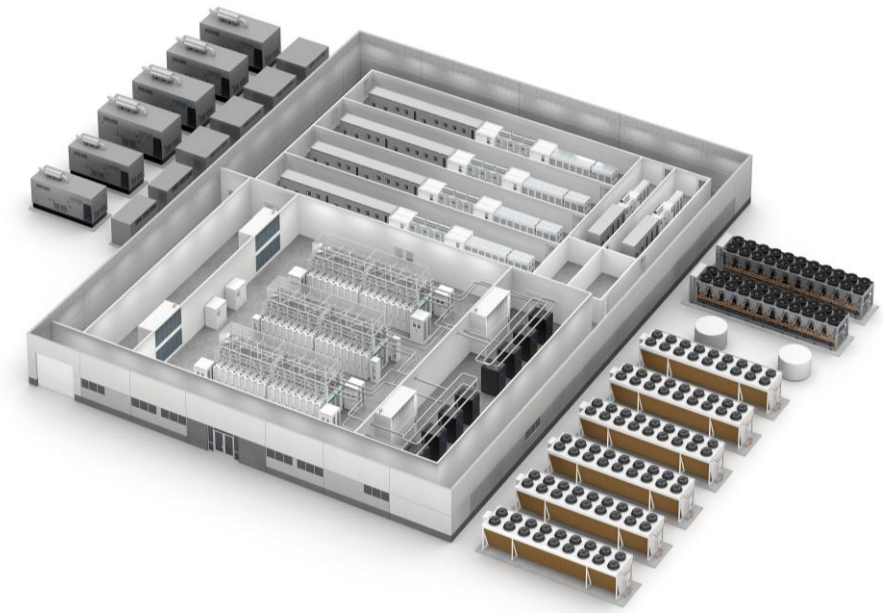
About this Design

- 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

Introduction

AI factories 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 for AI factories, specifically *NVIDIA's Vera Rubin NVL72* racks, in a single data hall. The data hall is purpose-built and optimized for two NVIDIA 2016 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, and reliability.

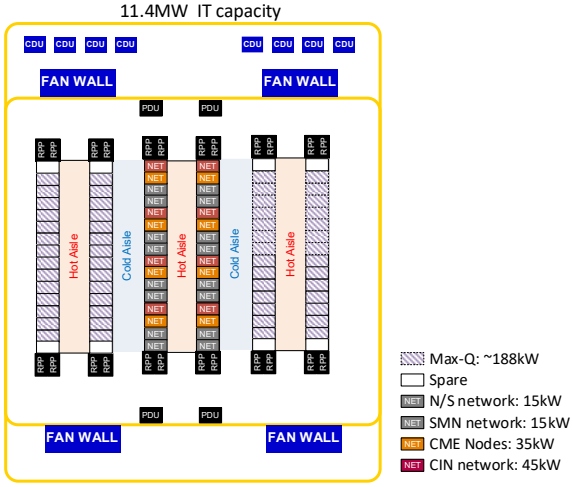
Reference Design 113 includes information on four technical areas: facility power, facility cooling, IT space with power and cooling, and lifecycle software. These areas represent the integrated systems required to meet the design's specifications provided in this overview document.



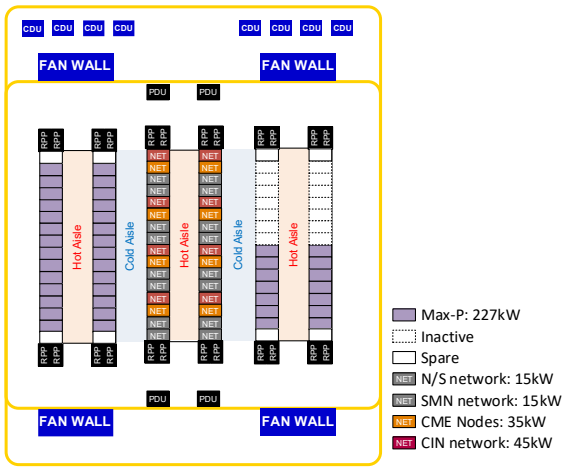
Maximizing Data Hall Output

Data Hall Diagrams

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



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



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, and 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, and for *NVIDIA's Vera Rubin NVL72* it draws around 188 kW at this operating point. 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 operate at Max-P. This operating state fully utilizes the facility assets and maximizes 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 (CDUs) 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.

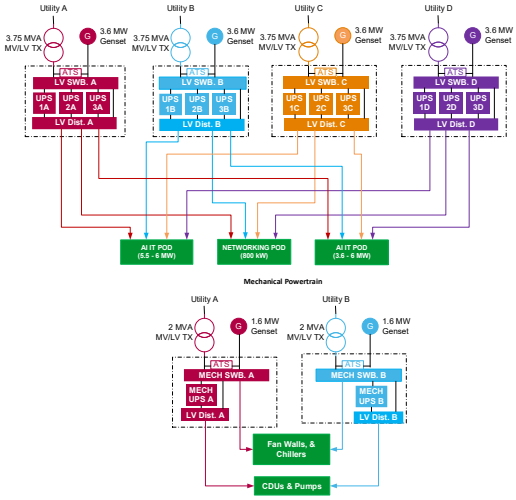
By dynamically operating the AI racks, this design enables the maximum IT capacity on the facility's power equipment. MaxLPS is another key operating point, representing the average rack power across multiple workloads, estimated at ~156 kW per rack. Like Max-Q, it can increase the number of active AI 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



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 3+1 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 *NVIDIA's Vera Rubin NVL72* racks' electrical design point (EDP) power peak. Downstream, these powertrains feed remote power panels that power the AI racks or power distribution units (PDUs) that power the networking racks with 3+1 redundancy. Separately, two 2 MW 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 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

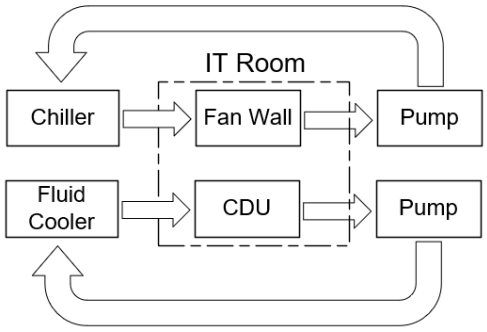
Facility Power Attributes

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

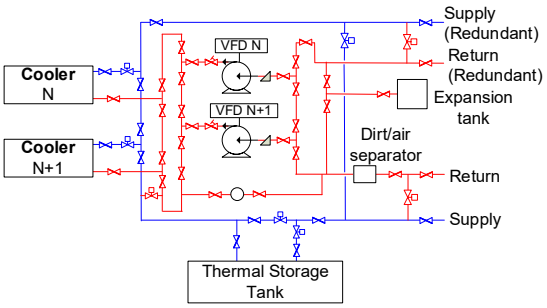
Facility Cooling

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 68°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 104°F water to *Motivair Floor-Mounted 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* fluid coolers. With supply water temperatures approaching 104°F, fluid 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 fluid 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 in the engineering package.

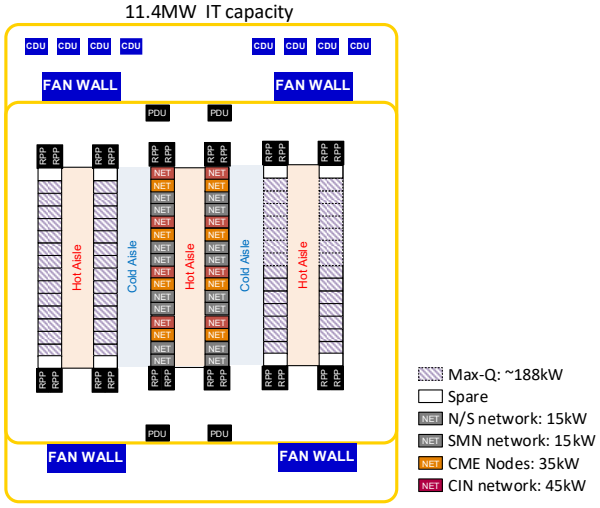
Facility Cooling Attributes

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	68	°F
Chiller CW return temperature	86	°F
Fluid cooler CW supply temperature	104	°F
Fluid cooler CW return temperature	122	°F
Storage tank size	30,000 (high temp. loop) 15,000 (low temp. loop)	gallon
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

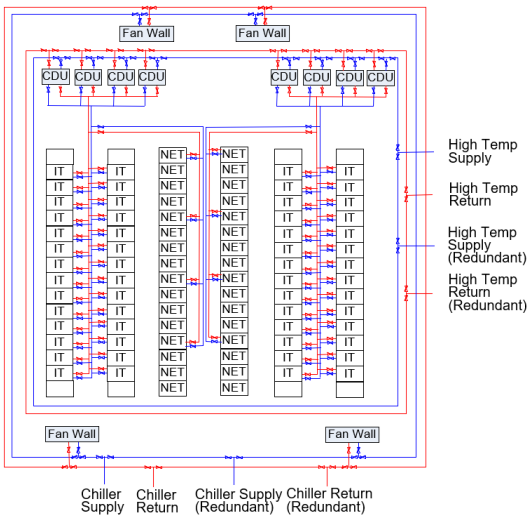
IT Space

IT Space Diagrams

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



Data Hall Piping



The IT space of this design consist 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 network racks 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. The racks are organized to two pods consisting of only AI racks and one central pod of networking racks.

Each AI pod and networking pod is deployed with *Prefabricated Modular EcoStruxure™ Pod Data Center* to provide a 6 ft. wide ducted hot aisle containment for proper airflow, busway 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,

Four *Uniflair FWCV* chilled water fan walls supply conditioned air to the data hall in an N+1 configuration. Four *Motivair Floor-Mounted Liquid to Liquid MCDUs* provide precise liquid cooling to each AI pod and its associated CIN network racks 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.

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 pod of networking racks is powered by four *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 RPPs* 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 RPPs* for 2N redundant feeds.

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

Design Options

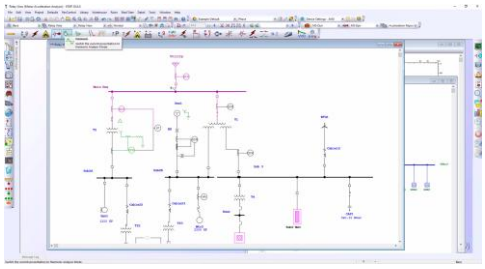
This reference design can be modified as follows without a significant effect on the design's performance attributes:

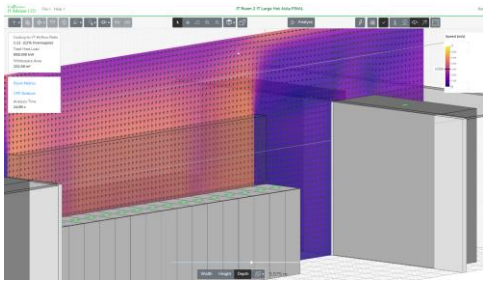
- Use *Uniflair FXCV* fan walls
- Use *NetShelter Aisle Containment*
- Use *NetShelter Pod Containment*
- CRAHs can be selected instead of fan walls
- Variations in AI cluster configuration
- Use TCS loop surge tank to ensure TCS supply temperature stability

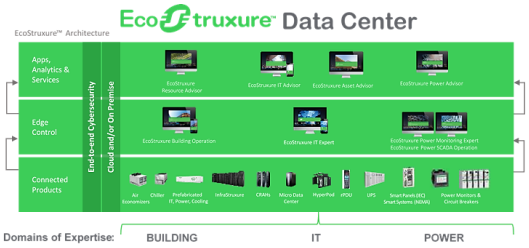
IT Space Attributes

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	79	°F
CW supply temperature	68	°F
CW return temperature	86	°F
Containment type	Ducted hot aisle	
CDU type	L2L	
CDU redundancy	N+1	
CDU CW supply temperature	104	°F
CDU CW return temperature	122	°F
TCS loop supply temperature	113	°F
TCS loop return temperature	131	°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 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
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
Data center overall space	37,391	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	68	°F
Chiller CW return temperature	86	°F
Fluid cooler CW supply temperature	104	°F
Fluid cooler CW return temperature	122	°F
Combined* storage tank size	30,000 (high temp. loop); 15,000 (low temp. loop)	gallon
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	5,486	ft^2
Rack power feed redundancy	Redundant	
CRAC/CRAH redundancy	N+1	
Containment type	Ducted hot aisle	
CDU Type	L2L	
CDU redundancy	N+1	
CRAC/CRAH supply air temperature	79	°F
CW supply temperature	68	°F
CW return temperature	86	°F
CDU CW supply temperature	104	°F
CDU CW return temperature	122	°F
TCS loop supply temperature	113	°F
TCS loop return supply temperature	131	°F

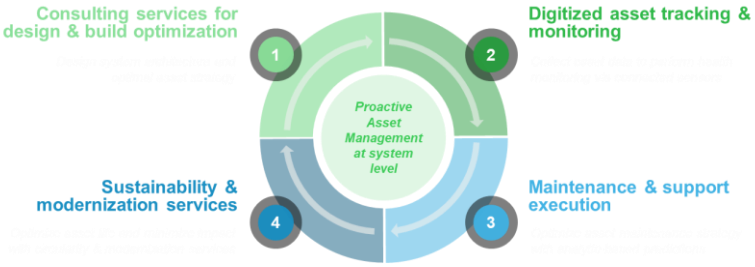
Schneider Electric Lifecycle Services

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.

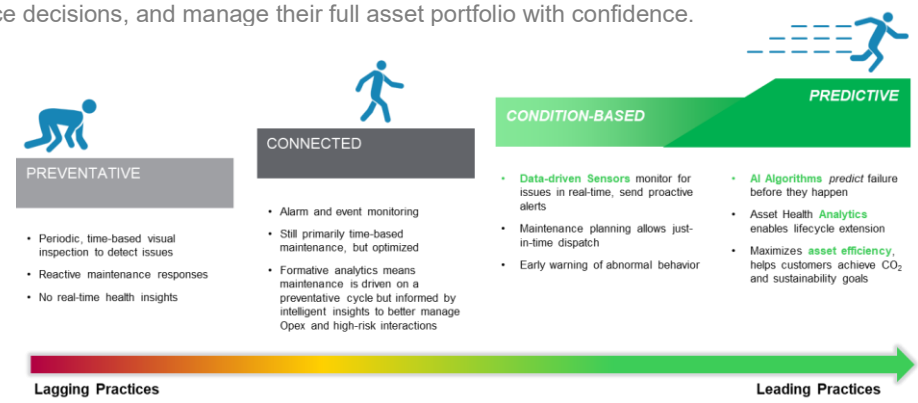
- *EcoFit Essential:* Upgrade your existing equipment with new features.
- *EcoFit Advanced:* Replace the core components with active and natively connectable new components.
- *EcoFit Replacement:* Get a comprehensive replacement with one of our latest equipment products/platforms/solutions.

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?

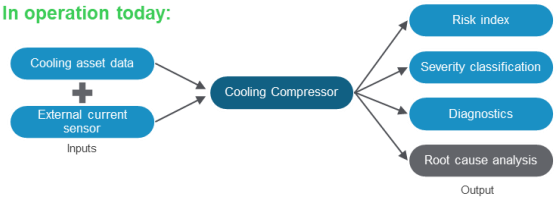
A data center is made up of systems within systems. Any solution monitoring critical assets within a data center needs to not only understand the performance of each individual asset but also the performance of its surrounding systems. Considering the entire system unlocks deeper understanding of critical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

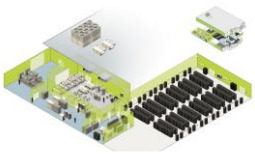
At the asset level, EcoStruxure aggregates data from a variety of intelligent sensors, including, thermal and environmental sensors for switchgear, partial discharge sensors, CDU controllers, oil transformer condition sensing, and leak detection for liquid cooling. Critical asset data can be collected 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

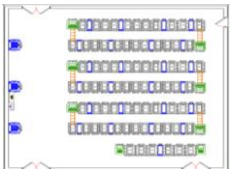
Cybersecurity is embedded by design across Schneider Electric's AI-ready CBM and analytics architecture, ensuring that the massive volumes of operational and performance data generated by high-density AI infrastructure can be securely collected, transported, and analyzed. From intelligent devices and 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.

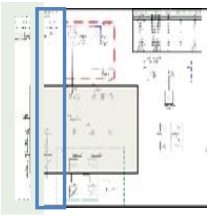
Get more information for this design:



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

This design is complemented by Controls Reference Design 2.

Download Controls Reference Design 2

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