

3818 kW, Tier III, IEC, Chilled Water, Liquid-Cooled & Air-Cooled AI Clusters

Design Overview

Data Center IT Capacity

3818 kW

Adaptable from 1808 kW to 3818 kW

Target Availability

Tier III

Annualized PUE at 100% Load

Paris: 1.15 – 1.16

Singapore: 1.25 – 1.26

(Scenario dependent)

Racks and Density

Total Racks: 128 / 144 (Scenario dependent)

Rack Density:

Max air-cooled: 40 kW

Max liquid-cooled: 73 kW

Data Center Overall Space

3060 m²

Regional Voltage and Frequency

400V, 50Hz

About this Design

- IT space and power distribution designed to accommodate AI clusters with density up to 73 kW per rack
- Various options to support liquid-cooled racks, including liquid-to-air coolant distribution units (CDUs) and liquid-to-liquid CDUs
- Chilled water systems optimized for high water temperatures using *Uniflair FWCV* fan walls and *Uniflair XRAF* air-cooled packaged chillers
- Redundant design for increased availability and concurrent maintainability


Introduction

High-density AI clusters 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 high-density AI clusters with two IT rooms. IT room 1 depicts three retrofit scenarios, where a new, high-density AI cluster is installed alongside existing traditional IT.

- Scenario 1A shows a high-density air-cooled AI cluster.
- Scenario 1B shows a high-density liquid-cooled AI cluster which uses liquid-to-air coolant distribution units (CDUs) for heat rejection. This is ideal for scenarios when you cannot connect to facility water systems.
- Scenario 1C shows a high-density liquid-cooled AI cluster which uses liquid-to-liquid CDUs. This is ideal for scenarios where you can tap into facility water systems.

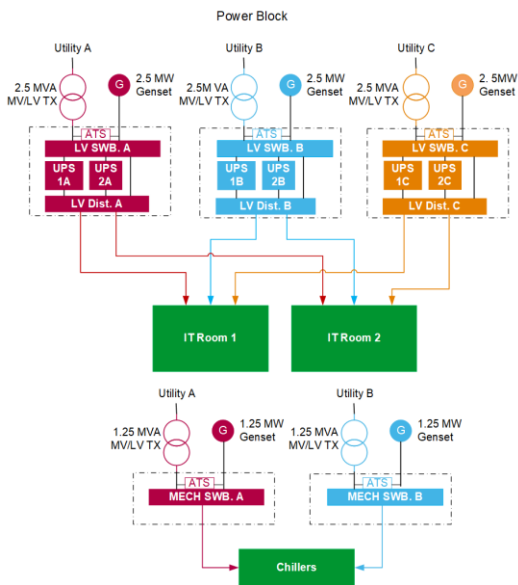
IT room 2 is purpose-built and optimized for a liquid-cooled AI cluster which uses liquid-to-liquid CDUs.

Reference Design 99 includes information for four technical areas: facility power, facility cooling, IT space and lifecycle software. They represent the integrated systems required to meet the design's specifications in this overview document.



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 rooms is supplied through three 2.5 MW powertrains. The three powertrains provide tri-redundant UPS power to the IT space, backed up by diesel generators. Each powertrain consists of a 4000-amp *Okken* main switchboard feeding two 1250 kW *Galaxy VX* UPS with 5 minutes of runtime in parallel and a 4000-amp *Okken* distribution section. The main switchboards also feed the *Uniflair FWCV* fan walls in the two IT rooms. Downstream, these powertrains feed *Canalis* busways that power the IT racks with 2N redundancy. The UPSs also feed the CDUs and chilled water pumps. Separately, two 1.25 MW powertrains feed the chillers with 2N redundant power.

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 IEC or IEEE 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:

- Add *EcoStruxure Power Monitoring Expert*
- Provision for load bank
- Change UPS battery type & runtime
- Add facility cooling UPS
- Add/remove/change standby generators:
 - Location & tank size


Facility Power Attributes

Name	Value	Unit
Total facility peak power (IT and cooling)	6250	kW
Total amps (IT main bus, each)	4000	A
Input voltage (IT main bus)	400	V
Switchboard kAIC (IT main bus)	66	kA
Generator redundancy (IT main bus)	Tri-redundant	
IT power path	Dual	
IT space UPS capacity, per powertrain	2500	kW
IT space UPS redundancy	Tri-redundant	
IT space UPS runtime @ rated load	5	minutes
IT space UPS output voltage	400	V
Total amps (Facility cooling bus, each)	1600	A
Input voltage (Facility cooling bus)	400	V
Switchboard kAIC (Facility cooling bus)	36	kA
Generator redundancy (Facility cooling bus)	2N	
Facility cooling UPS capacity	N/A	kW
Facility cooling UPS redundancy	N/A	
Facility cooling UPS runtime @ rated load	N/A	minutes



Facility Cooling

Facility Cooling Block Diagrams

Retrofit Scenario System Flow



New Build Scenario System Flow



Design Options

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

- Add *EcoStruxure IT Expert*
- Change storage tank size
- Use standard temperature chillers, like *Uniflair XRAF* or *Uniflair BCEF*, chillers for loop with fan walls in IT Room 2

The facility cooling design is based on the specified AI deployment scenarios. For IT Room 1 (retrofit IT room scenario), a chilled water system with dual path piping is implemented. Three *Uniflair BCEF* chillers with free cooling capabilities deliver 20°C chilled water in an N+1 configuration.

The facility cooling design for IT Room 2 (new IT room scenario) is comprised of two separate chilled water loops. A high temperature water loop, with two *Uniflair XRAF* extra high temperature chillers with screw compressors and free cooling capabilities, provides 31°C water to the IT room to cool the IT equipment. A separate chilled water loop, with two *Uniflair XRAF* extra high temperature chillers, provides 20°C water for the air handling units of the IT room. Using the *Uniflair XRAF* extra high temperature chiller for this chilled water loop enables future-readiness for water temperature increase, but *Uniflair BCEF* chillers and standard *Uniflair XRAF* chillers can be used instead.

A thermal storage system provides 5 minutes of continuous cooling after a power outage or chiller restart. The *Uniflair BCEF* and *Uniflair XRAF* chillers can fully restart within 3 minutes.

More information on fan wall and CDU cooling architecture is provided in the IT room section of this document.

This design is instrumented to work with *EcoStruxure IT Expert* and *AVEVA Unified Operations Center*.

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 (chillers)	4993 (Paris) 5522 (Singapore)	kW
Input voltage	400	V
Heat rejection medium	Chilled water	
Chiller redundancy	N+1	
Outdoor heat exchange	Packaged chiller with free cooling	
CW supply temperature	20-21	°C
CW return temperature	30	°C
CW supply temp (Room 2, to CDUs)	31	°C
CW return temp (Room 2, from CDUs)	40	°C
Combined* storage tank size	28	m ³
Ride-through time	5	minutes
Outdoor ambient temperature range	-9.6 to 39.3	°C
Economizer type	Water-side	

*Summation of all three chilled water loops

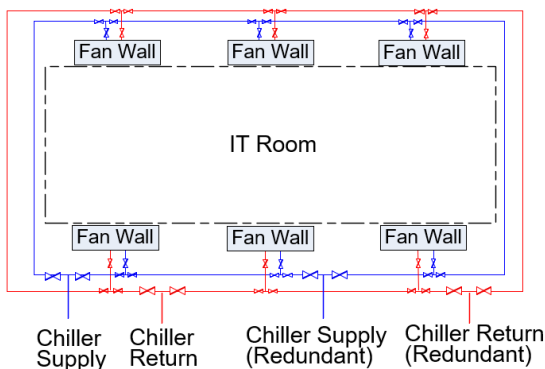
Retrofit IT Room: Scenario 1A

IT Room 1A Diagrams

IT Room 1A Layout



IT Room 1A 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
- CRAHs can be selected instead of fan walls
- Variations in AI cluster configuration

The first retrofit IT room scenario features eighty 12 kW air-cooled IT racks. The load has been expanded with an AI cluster consisting of twenty-four 40 kW air-cooled server racks with six 15 kW air-cooled networking racks (modeled after Nvidia's *DGX SuperPOD*). This scenario demonstrates a 50/50 split in power between low and high-density IT racks. The 12 kW IT racks are configured in pods of 20 racks and share a 1.2 m wide hot aisle. The 40 kW air-cooled AI racks are configured with four racks together and two 15 kW networking racks in the middle of the row. The high-density pod shares a 2.4 m wide hot aisle to allow proper airflow. Ducted hot aisles and a common ceiling plenum return hot air to the fan walls for cooling.

Six *Uniflair FWCV* chilled water fan walls deliver clean and conditioned supply air to the IT room in an N+1 configuration. The redundant piping system across the IT room provides an alternate path for chilled water in case of cooling equipment failure or maintenance.

The 12 kW IT racks and 15 kW networking racks are configured with 1+1 32A *NetShelter* metered rack-mount power distribution units (rPDUs). The 40 kW AI racks are configured with 1+1 63 A *NetShelter Advanced* rPDUs. Each rack is powered by 2N redundant tap-offs from *Canalis KS* busway providing A and B-side power to each rack. Each tap off unit can be configured to house up to two 63 A *NG125* circuit breakers with associated *Acti9 iEM3000* energy meters and auxiliaries. Rows of 12 kW racks are fed by 250 A *Canalis KS* busway, while the air-cooled AI clusters are fed by 630 A *Canalis KS* busway.

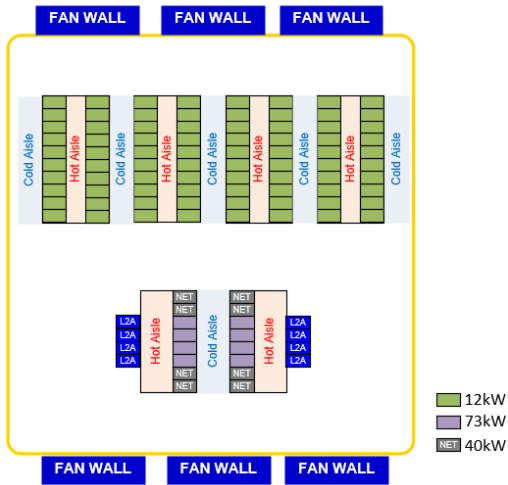
IT Room 1A Attributes

Name	Value	Unit
IT load	2010	kW
Supply voltage to IT	400	V
Single or dual cord	Dual	
Number of 12kW air-cooled racks	80	racks
Number of 40kW air-cooled racks	24	racks
Number of 15kW networking racks	6	racks
IT floor space	415	m ²
CRAC/CRAH type	Fan wall	
CRAC/CRAH redundancy	N+1	
CW supply temperature	20	°C
CW return temperature	30	°C
Containment type	Ducted hot aisle	
CDU type	N/A	
CDU redundancy	N/A	
TCS loop supply temperature	N/A	
TCS loop return temperature	N/A	

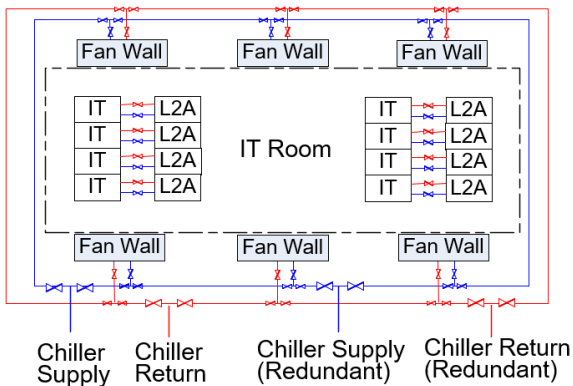
Retrofit IT Room: Scenario 1B

IT Room 1B Diagrams

IT Room 1B Layout



IT Room 1B 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
- CRAHs can be selected instead of fan walls
- Variations in AI cluster configuration

The second retrofit IT room scenario features eighty 12 kW air-cooled IT racks. The load has been expanded with an AI cluster consisting of eight 73 kW liquid-cooled AI racks with eight 40 kW air-cooled networking racks (modeled after Nvidia's *DGX SuperPOD*). The AI cluster is configured with four server racks together in the center and networking racks on each end of the row. For the liquid-cooled racks, *Uniflair ACSX* liquid-to-air (L2A) coolant distribution units (CDUs) are placed on opposite sides of the hot aisle. The liquid cooled servers use direct-to-chip cooling technology. The liquid cooling loop which directly feeds coolant to the racks is known as the Technology Cooling System (TCS). A 2.4 m wide hot aisle is designed for the high-density pods to ensure proper airflow. Ducted hot aisles and a common ceiling plenum return hot air to the fan walls for cooling.

L2A CDUs allow liquid-cooled racks to be deployed in air-only data centers. They supply coolant to the racks, and then reject return coolant heat into the air. In this scenario, the CDUs provide coolant to the racks via piping across the hot aisle. Six *Uniflair FWCV* chilled water fan walls with redundant piping deliver supply air to the IT room in an N+1 configuration.

The 12 kW IT racks are powered by 1+1 32 A *NetShelter* metered rPDUs. The 40 kW networking racks are configured with 1+1 63 A power feeds going to *NetShelter Advanced* rPDUs. The 73 kW liquid-cooled AI racks are configured with three OCP V3 power shelves, fed with 3+3 63 A power feeds. Each rack is powered by 2N redundant tap-offs from *Canalis KS* busway providing A and B-side power to each rack. Each tap off unit can be configured to house up to two 63 A *NG125* circuit breakers with associated *Acti9 iEM3000* energy meters and auxiliaries (e.g., shunt trip for leak detection). Pods of 12 kW racks are fed by 250 A *Canalis KS* busway, while the liquid-cooled AI cluster is fed by 800 A *Canalis KS* busway.


IT Room 1B Attributes

Name	Value	Unit
IT load	1864	kW
Supply voltage to IT	400	V
Single or dual	Dual	
Number of 12kW air-cooled racks	80	racks
Number of 73kW liquid-cooled racks	8	racks
Number of 40kW networking racks	8	racks
IT floor space	415	m ²
CRAC/CRAH type	Fan wall	
CRAC/CRAH redundancy	N+1	
CW supply temperature	20	°C
CW return temperature	30	°C
Containment type	Ducted hot aisle	
CDU type	L2A	
CDU redundancy	N	
TCS loop supply temperature	40	°C
TCS loop return temperature	50	°C

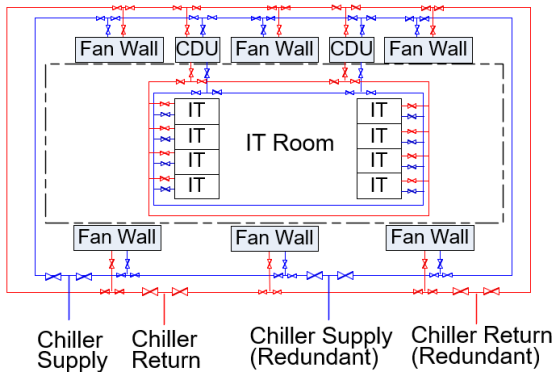
Retrofit IT Room: Scenario 1C

IT Room 1C Diagrams

IT Room 1C Layout



IT Room 1C 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
- CRAHs can be selected instead of fan walls
- Variations in AI cluster configuration

The third retrofit IT room scenario features eighty 12 kW air-cooled IT racks. The load has been expanded with an AI cluster consisting of eight 73 kW liquid-cooled IT racks with eight 40 kW air-cooled networking racks (modeled after Nvidia's *DGX SuperPOD*). The AI cluster is configured with four server racks together in the center and networking racks on each end of the row. For the liquid-cooled racks in the AI cluster, two *Uniflair CPOR* liquid-to-liquid (L2L) CDUs provide coolant to the racks. The L2L CDUs are placed in the service hallway. The liquid cooled servers use direct-to-chip cooling technology. The liquid-cooled pod shares a 2.4m wide hot aisle for proper airflow.

L2L CDUs are the heat exchange interface between liquid-cooled IT racks on the TCS loop and the facility water system (FWS). In this scenario, the CDUs are tied together on a common loop providing N+1 redundancy. The CDUs are fed the same facility supply water as the fan walls. Six *Uniflair FWCV* chilled water fan walls with redundant piping deliver supply air to the IT room in an N+1 configuration.


The 12 kW IT racks are powered by 1+1 32 A *NetShelter* metered rPDUs. The 40 kW networking racks are configured with 1+1 63 A power feeds going to *NetShelter Advanced* rPDUs. The 73 kW liquid-cooled AI racks are configured with three OCP V3 power shelves, fed with 3+3 63 A power feeds. Each rack is powered by 2N redundant tap-offs from *Canalis KS* busway providing A and B-side power to each rack. Each tap off unit can be configured to house up to two 63 A *NG125* circuit breakers with associated *Acti9 iEM3000* energy meters and auxiliaries (e.g., shunt trip for leak detection). Pods of 12 kW racks are fed by 250 A *Canalis KS* busway, while the liquid-cooled AI cluster is fed by 800 A *Canalis KS* busway.

IT Room 1C Attributes

Name	Value	Unit
IT load	1864	kW
Supply voltage to IT	400	V
Single or dual cord	Dual	
Number of 12kW air cooled racks	80	racks
Number of 73kW liquid cooled racks	8	racks
Number of 40kW networking racks	8	racks
IT floor space	415	m ²
CRAC/CRAH type	Fan wall	
CRAC/CRAH redundancy	N+1	
CW supply temperature	20	°C
CW return temperature	30	°C
Containment type	Ducted hot aisle	
CDU type	L2L	
CDU redundancy	2N	
TCS loop supply temperature	40	°C
TCS loop return temperature	50	°C

New Build IT Room 2

IT Room 2 Diagrams

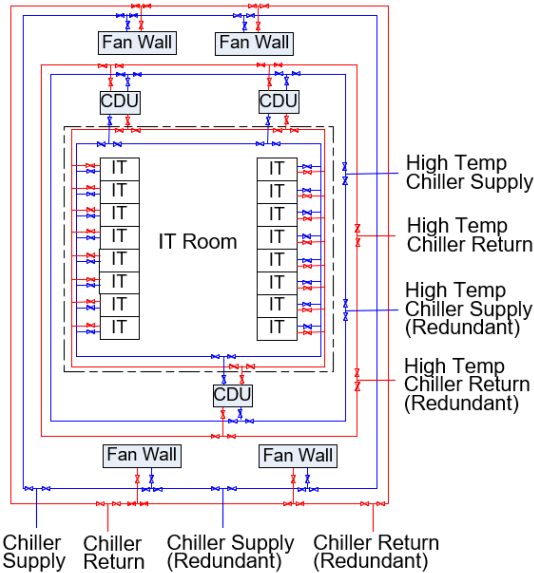


IT Room 2 is dedicated to a new AI cluster and features sixteen 73 kW liquid-cooled IT racks with sixteen 40 kW air-cooled networking racks placed at the ends of the rows. The liquid-cooled and networking racks are configured in one pod and share a 1.8m wide hot aisle. The liquid-cooled servers use direct-to-chip cooling technology. Hot aisle containment is still required to handle the hot air return of the networking racks and remaining heat from the liquid cooled racks.

Four *Uniflair FWCV* chilled water fan walls deliver supply air to the IT room in an N+1 configuration. Three *Uniflair CPORL2L* CDUs are tied together on a common TCS loop with N+1 redundancy to provide coolant to the liquid-cooled racks. The CDUs run on a separate, high-temperature chilled water loop to increase free cooling opportunity. *Uniflair XRAF* extra high temperature chillers make it possible to operate this chiller-based cooling loop at temperatures not seen in the industry today providing unmatched cooling efficiency.

The 40 kW networking racks are configured with 1+1 63 A power feeds going to *NetShelter Advanced* rPDUs. The 73 kW liquid-cooled AI racks are configured with three OCP V3 power shelves, fed with 3+3 63 A power feeds. Each rack is powered by 2N redundant tap-offs from *Canalis KS* busway providing A and B-side power to each tap off unit can be configured to house up to two 63 A *NG125* circuit breakers with associated *Acti9 iEM3000* energy meters and auxiliaries (e.g., shunt trip for leak detection). Each 1 MW row of 73 kW and 40 kW networking racks are fed by four (2N) 800 A *Canalis KS* busway, where each 800 A busway run feeds half of the row. ADD PDU INFO

IT Room 2 Piping



IT Room 2 Attributes



Name	Value	Unit
IT load	1808	kW
Supply voltage to IT	400	V
Single or dual cord	Dual	
Number of 73kW liquid-cooled racks	16	racks
Number of 40kW networking racks	16	racks
IT floor space	159	m ²
CRAC/CRAH type	Fan wall	
CRAC/CRAH redundancy	N+1	
CW supply temperature	21	°C
CW return temperature	30	°C
Containment type	Ducted hot aisle	
CDU type	L2L	
CDU redundancy	N+1	
CDU CW supply temperature	31	°C
CDU CW return temperature	40	°C
TCS loop supply temperature	40	°C
TCS loop return temperature	50	°C

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

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 *IT Advisor 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.



Visit [EcoStruxure for Data Center](#) for more details.

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

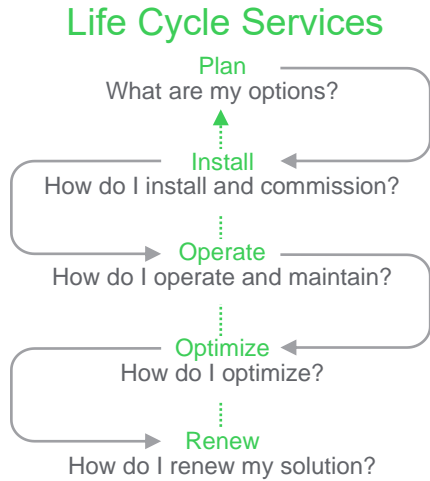
OVERVIEW	Value	Unit
Target availability	III	Tier
Annualized PUE at 100% load (1A & 2 / 1B & 2 / 1C & 2)	1.16 / 1.16 / 1.15	Paris
	1.26 / 1.26 / 1.25	Singapore
Data center IT capacity	3672 – 3818	kW
Data center overall space	3060	m ²
Maximum rack density	73	kW/rack
FACILITY POWER	Value	Unit
Total facility peak power (IT and cooling)	6250	kW
Total amps (IT main bus, each)	4000	A
Input voltage (IT main bus)	400	V
Switchboard kAIC	66	kA
Generator redundancy (IT main bus)	Tri-redundant	
IT Power path	Dual	
IT space UPS capacity, per powertrain	2500	kW
IT space UPS redundancy	Tri-redundant	
IT space UPS runtime @ rated load	5	minutes
IT space UPS output voltage	400	V
Total amps (facility cooling bus, each)	1600	A
Input voltage (facility cooling bus)	400	V
Switchboard kAIC (facility cooling bus)	36	kA
Generator redundancy (facility cooling bus)	2N	
FACILITY COOLING	Value	Unit
Total max cooling capacity (chillers)	4993 (Paris), 5522 (Singapore)	kW
Input voltage	400	V
Heat rejection medium	Chilled water	
Chiller redundancy	N+1	
Outdoor heat exchange	Packaged chiller with free cooling	
CW supply temperature	20	°C
CW return temperature	30	°C
CW supply temp (IT Room 2, to CDUs)	31	°C
CW return temp (IT Room 2, from CDUs)	40	°C
Combined* storage tank size	28	m ³
Ride-through time	5	minutes
Outdoor ambient temperature range	-9.6 to 39.3	°C
Economizer type	Water-side	

*Summation of all three chilled water loops

Design Attributes continued


IT SPACE	Retrofit room			New room	Total	Unit
	1A	1B	1C	2		
IT load	2010	1864	1864	1808	3672 – 3818	kW
Supply voltage to IT	400	400	400	400	400	V
Maximum density	40	73	73	73	73	kW/rack
Number of racks	110	96	96	32	128 – 142	racks
IT floor space	415	415	415	159	574	m ²
Single or dual cord	Dual	Dual	Dual	Dual	Dual	
CRAC/CRAH type	Fan wall	Fan wall	Fan wall	Fan wall	Fan wall	
CRAC/CRAH redundancy	N+1	N+1	N+1	N+1	N+1	
Containment type	Ducted hot aisle	Ducted hot aisle	Ducted hot aisle	Ducted hot aisle	Ducted hot aisle	
CDU Type	N/A	L2A	L2L	L2L		
CDU redundancy	N/A	N	N+1	N+1		
CW supply temperature	20	20	20	21		°C
CW return temperature	30	30	30	30		°C
CDU CW supply temperature	N/A	N/A	20	31		°C
CDU CW return temperature	N/A	N/A	30	40		°C
TCS loop supply temperature	N/A	40	40	40		°C
TCS loop return supply temperature	N/A	50	50	50		°C

Schneider Electric Life-Cycle Services



- 1** Team of **over 7,000 trained specialists** covering every phase and system in the data center
- 2** Standardized, documented, and validated **methodology** leveraging automation tools and repeatable processes **developed over 45 years**
- 3** **Complete portfolio of services** to solve your technical or business challenge, simplify your life, and reduce costs

Get more information for this design:




3D spatial views



Floor layouts

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.



One-line schematics



Bill of materials

Documentation is available in multiple formats to suit the needs of both engineers and managers working on data center projects. For the engineering package of this design please email us at referencedesigns@se.com