HOLISTIC ENERGY MANAGEMENT AND THERMAL WASTE INTEGRATED SYSTEM FOR ENERGY OPTIMIZATION

HEATWISE

Self-assessment framework

Project ID: 101138491

Prepared by: DSTECH

February, 2026

Funded by the European Union

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs,
Education and Research EAER
State Secretariat for Education,
Research and Innovation SERI
EU Framework Programmes

Self-assessment framework

Project acronym	**HEATWISE**
Project title	Holistic Energy Management and Thermal Waste Integrated System for Energy Optimization
GA Number	101138491
Call identifier	HORIZON-CL5-2023-D4-01
Topic identifier	HORIZON-CL5-2023-D4-01-04
Project duration	36 months
Related work package	WP8 Developing self-assessment tool with holistic algorithms
Deliverable no.	D (8.2) Self-assessment framework
Dissemination Level	Public
Deliverable type	Other
Due Date of deliverable	28 February 2026
Completion date of deliverable	26 February 2026
Coordinator	H1 Systems
Website	www.heatwise.eu
Lead beneficiary of deliverable	DSTECH
Author(s)	Kuzay, Mustafa Bayraktar, Basak Ispir, Halil Ibrahim Ethem Demirel, Ender
Contributor(s)	Yilmaz, Cagatay Karnebro, Axel Melgaard, Simon Pommerencke Koirala, Binod Prasad Ciesielski, Sebastian Hatipoglu, Halim
Reviewer(s)	Seres, Tamas Belkin, Shahar
Keywords	Self-assessment, KPI, Data Collection, Sampling and Report

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

Document history:

Revision	Date	Status
V0.1	06.02.2026	Initial draft
V0.8	13.02.2026	Draft for review
V0.9	24.02.2026	Final draft after implementing comments, ready for editing
V0.9.1	25.02.2026	Final draft after editing, ready for approval
V1.0	26.02.2026	Final report

Disclaimer

Copyright Notice

Acknowledgment

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

Table of contents

Executive Summary....4
List of abbreviations....6
1. Introduction....8
    1.1. Why Does It Matter?....9
2. Structure of the Self-Assessment Tool....13
    2.1. Data Collection and Sampling....13
    2.2. Calculation of KPIs....15
    2.3. System architecture of the SAT....18
        2.3.1. Sub-windows of the SAT....22
3. Implementation of SAT at Pilot Data Centres....30
    3.1. Implementation of SAT at Empa....30
    3.2. Implementation of SAT at AAU....33
    3.3. Implementation of SAT at PSNC....34
    3.4. Implementation of SAT at Tofas....36
    3.5. Implementation of SAT at RISE....38
    3.6. Reproreproducibility Guidelines of SAT....44
    3.7. Road map for the implementation of SAT at external pilots....46
4. Conclusion....49
References....51
List of figures....56
List of tables....57
Annexes....58

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

Executive Summary

The rapid digitalization of society, together with the growing adoption of cloud computing, artificial intelligence, and data-intensive services, has led to a continuous increase in global data centre energy consumption. This trend has intensified the need for periodic, transparent, and comparable performance assessments, particularly in the context of emerging regulatory frameworks such as the European Energy Efficiency Directive (EED). Despite the availability of large volumes of operational data, many data centres still face significant challenges in systematically collecting heterogeneous measurements, computing standardized Key Performance Indicators (KPIs), and translating raw monitoring data into actionable and regulation-ready performance assessments within defined reporting periods.

To address these challenges, this deliverable introduces the Self-Assessment Tool (SAT), a unified, modular, and extensible framework designed for the comprehensive evaluation of thermal and energy performance of data centres. SAT integrates IT, cooling, and energy-related data sourced from existing data monitoring systems and supports both real-time and retrospective analysis. In addition to live data acquisition, the framework allows historical datasets collected from external sensors or offline measurement campaigns to be imported and processed using the same analytical workflow, ensuring methodological consistency across online and offline assessments.

At the core of SAT lies a dedicated KPI computation engine that evaluates a comprehensive and standardized set of thermal KPIs and energy-related KPIs. All KPIs are computed using a transparent, deterministic, and reproducible workflow, explicitly documenting data selection, sampling resolution, temporal alignment and aggregation logic. This design not only ensures traceability from raw measurements to final indicators but also enables independent implementation of the KPI methodology, addressing reproducibility requirements increasingly emphasized in both regulatory and scientific contexts.

The framework further combines KPI computation with integrated visualization and reporting capabilities. SAT automatically generates assessment outputs including time-series representations of thermal and energy KPIs, rack-level thermal and power maps, and KPI-based energy efficiency classifications. These outputs are made available through an interactive dashboard environment and are simultaneously embedded into structured, automated PDF self-assessment reports. By decoupling KPI computation and

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

reporting from site-specific dashboards and manual operator interpretation, SAT enables consistent cross-site comparison and reduces subjectivity in performance evaluation.

The applicability and robustness of the proposed framework have been demonstrated and validated through deployments at five pilot data centres (Empa, AAU, PSNC, Tofas and RISE), covering different cooling concepts and monitoring infrastructures. The results confirm that SAT can operate effectively under heterogeneous data availability conditions while maintaining a unified assessment methodology. Owing to its modular architecture, the framework can be deployed either as a standalone desktop application or as a web-based service, allowing flexible integration into existing operational environments.

Overall, SAT provides a practical and regulation-oriented self-assessment framework that bridges the gap between raw monitoring data and standardized performance evaluation. By enabling continuous, KPI-driven assessment of thermal management, energy efficiency and waste heat reuse potential, the proposed framework supports data centre operators in meeting regulatory obligations, improving operational transparency, and identifying efficiency improvement opportunities in a structured and reproducible manner.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

List of abbreviations

Acronym	Meaning
AAU	Aalborg University
AI	Artificial Intelligence
API	Application Programming Interface
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers
CDU	Cooling Distribution Unit
CFD	Computational Fluid Dynamics
COP	Coefficient of Performance
CPU	Central Processing Unit
DC	Data Centre
DCIM	Data Centre Infrastructure Management
DMS	Data Management System
EED	Energy Efficiency Directive
EMPA	Swiss Federal Laboratories for Materials Science and Technology
ERF	Energy Reuse Factor
EU	European Union
GPU	Graphical Processing Unit
GUI	Graphical User Interface
HRU	Heat Recovery Unit
HTE	Heat Transfer Efficiency
IPMI	Intelligent Platform Management Interface
IT	Information Technology
KPI	Key Performance Indicator
LI	Leakage Index

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

MQTT	Message Queuing Telemetry Transport
PES	Primary Energy Savings
PSNC	Poznan Supercomputing and Networking Centre
PUE	Power Usage Effectiveness
RCI	Rack Cooling Index
REF	Renewable Energy Factor
REST	Representational State Transfer
RHI	Return Heat Index
RI	Recirculation Index
RISE	Research Institutes of Sweden Pilot
RTI	Return Temperature Index
SAT	Self-Assessment Tool
SHI	Supply Heat Index
TOFAS	Tofas Türk Otomobil Fabrikası A.S.
WHR	Waste Heat Recovery
WUE	Water Usage Effectiveness

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

1. Introduction

The rapid digitalization of society, growing adoption of Artificial Intelligence (AI) applications and simulation-driven research studies have led to a significant increase in global data centre energy consumption and associated carbon emissions in recent years. The environmental impact of data centres has drawn significant attention from the scientific community and policy makers over the last few years, which resulted in the implementation of new regulations for evaluating data centre energy performance [1-9].

The first phase of the European Union (EU)-wide data centre rating scheme has been established with regulations to define reporting requirements on energy use, efficiency, waste heat use, and related sustainability metrics [6-7]. However, limited transparency in the data provided by operators remains a continuing challenge [10]. This lack of transparency not only reduces the clarity of performance assessments but also makes it difficult to develop strategies for improving energy efficiency, resource utilization, and cooling performance [10-11]. Therefore, the introduction of new regulatory frameworks is essential for periodically assessing energy performance and environmental impact of data centres. These frameworks emphasize the use of widely adopted KPIs for transparent and comparable assessments. However, collecting the necessary data within the corresponding period, calculating the KPIs, and subsequently assessing efficiency with respect to the KPIs remain significant challenges for the majority of data centres. Accordingly, 36% of eligible data centres, around 770 facilities, participated in Europe in the first reporting cycle, and 80% of the submitted data were assessed as accurate and reliable [11]. To address these challenges, a framework focusing on the streamlining of data collection, KPI calculation, and efficiency assessment is essential for the evaluation of data centre performance comprehensively.

Data centres are such complex engineering systems influenced by many factors, such as thermal response, airflow distribution, cooling technology, and energy use and reuse. Therefore, a broader set of performance indicators is needed for a comprehensive and reliable assessment of data centres [10-11]. Meanwhile, monitoring energy efficiency and thermal features in data centres has advanced substantially due to rising computational loads and the increasing need for operational transparency under strict energy and sustainability regulations. Modern practices rely on continuous thermal and electrical measurements to evaluate airflow management, cooling system effectiveness, and overall facility performance. This provides operators with actionable insights to maintain compliance and improve sustainability [1, 4, 6-7, 12-17]. Critical components in this evolution are the Intelligent Platform Management Interface (IPMI) telemetry and external sensor networks, which enable real-time measurements of inlet, outlet, Central/Graphical Processing Unit (CPU and GPU) temperature server power

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

consumption and fan speeds. The IT level measurements capture the direct thermal response of servers to dynamic workloads and have become essential for accurate assessment of airflow efficiency, recirculation behaviour, and cooling performance in operational facilities [18-21]. In parallel, open-source monitoring, database, visualization and analytics systems such as Prometheus and Grafana have become widely adopted due to their modular architecture, scalability, and ability to integrate heterogeneous data sources ranging from server-embedded sensors to environmental monitoring systems. These platforms enable high-resolution time series collection and provide flexible visualization environments for analyzing thermal and energy behaviour at facility, rack, row, and cooling-unit levels [22-27]. Several studies highlight the need for unified, KPI driven assessment platforms capable of harmonizing various data sources and supporting cross site comparisons to address energy efficiency requirements [28-34]. Within this context, the framework proposed in this study aims to introduce three key contributions: (i) Integrating IT, power, and cooling data, external sensor measurements and open-source monitoring pipelines into a unified analytical framework capable of computing standardized KPIs in alignment with emerging regulatory needs, (ii) including a dedicated KPI-calculation module for evaluating KPIs using both real-time and historical datasets and (iii) employing identical visualization and reporting modules used in operational pilot sites to ensure consistency between offline and online analyses and to enable harmonized cross pilot performance assessment.

In this study, a multi-dimensional and KPI-driven assessment framework is developed to evaluate data centre performance in alignment with EU reporting and sustainability requirements. The framework is capable of accessing IT, power, and cooling data through the existing data monitoring system. Additionally, an existing dataset can be imported into the framework, so that the thermal data measured by external sensors can be utilized within the framework. The framework has been developed in both standalone and web-based versions to offer data centre operators a flexible deployment option. The proposed framework was successfully demonstrated and validated on two pilot data centres located in Denmark and Switzerland. Ultimately, the proposed framework provides a functional framework that combines real-time KPI calculation, spatio-temporal visualization, and an automated reporting scheme.

1.1. Why Does It Matter?

The increasing energy demand and environmental impact of data centres have placed them at the centre of both regulatory oversight and operational frameworks. In this context, the development of a dedicated SAT is not only timely but essential for ensuring compliance, transparency, and continuous performance improvement.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

The Energy Efficiency Directive (EED) establishes binding requirements for improving energy efficiency across the European Union. Recent revisions of the directive explicitly identify data centres as high-impact energy consumers and introduce obligations that go beyond conventional energy accounting [35-37].

Key EED expectations for data centres include:

- Mandatory energy performance reporting
- Disclosure of cooling efficiency and thermal management quality
- Identification of waste heat recovery potential
- Promotion of continuous monitoring rather than static audits

The EED shifts the focus from annual or one-off reporting to ongoing performance tracking, requiring operators to demonstrate that efficiency is actively managed throughout the operational lifecycle.

The SAT fulfils these requirements by:

- Automatically calculating standardised thermal and energy KPIs
- Enabling time-resolved analysis of cooling and airflow performance
- Providing traceable and reproducible datasets suitable for regulatory reporting

As a result, the SAT functions as both a compliance-enabling and pre-compliance validation tool, significantly reducing the administrative burden associated with EED reporting while increasing data quality and transparency.

From an operational standpoint, the Uptime Institute has consistently emphasised that availability, reliability, and efficiency must be managed together. Uptime Institute analyses show that [38-40]:

- Thermal management deficiencies are a major contributor to unplanned downtime
- Average temperature compliance often masks local hot spots and airflow inefficiencies
- High-level indicators such as PUE alone are insufficient to characterise real operational risk

Accordingly, Uptime Institute best practices advocate the use of diagnostic thermal KPIs, including temperature distribution, airflow effectiveness, and cooling utilisation metrics.

The SAT aligns with these principles by:

- Providing rack-level and zone-level thermal KPIs
- Detecting hidden inefficiencies that do not appear in aggregate metrics

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

- Supporting risk-aware optimisation, ensuring efficiency gains do not compromise uptime

In this way, the SAT enables a transition from reactive troubleshooting to proactive thermal and operational management, fully aligned with Uptime Institute methodologies. Taken together, EED regulations and Uptime Institute guidance converge on a single conclusion: Modern data centres require continuous, KPI-driven, self-operated performance assessment tools.

Traditional approaches such as periodic audits or manual reporting are limited because they provide only snapshot views of performance, rely on aggregated or estimated data and offer limited insight for day-to-day operational decisions. The SAT overcomes these limitations by embedding assessment directly into operational workflows. It enables operators to monitor compliance-relevant indicators continuously, identify performance degradation at an early stage, quantify the impact of optimisation measures, generate repeatable and structured reports aligned with both regulatory and industry frameworks.

Unlike dashboards and DCIM systems, which mainly focus on data visualization and infrastructure management, SAT provides a dedicated assessment layer for computing and reporting performance KPIs. By separating KPI calculation and reporting from site-specific dashboards and manual operator interpretation, SAT enables consistent cross-site comparison and supports automated, regulation-oriented performance assessment. A structured comparison of SAT with operator-driven assessments, dashboard-based monitoring and DCIM platforms is provided in Table 1.1.

Table 1.1: Comparison of assessment capabilities across operator-driven approaches, dashboards, DCIM platforms and SAT.

Capability	**Operator-driven assessment**	**Dashboard based monitoring tools**	**DCIM platforms**	**SAT**
Primary function	Manual interpretation	Visualization	Infrastructure management	Automated KPI-based assessment
KPI computation	Ad-hoc and/or site-specific	Limited / custom scripts	Vendor dependent	Standardized and automated

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

Cross-site comparability	No	No	Limited	Yes
Regulatory alignment	Manual	Not explicit	Partial	Explicit and embedded
Offline historical analysis	Limited	Limited	Limited	Yes
Automated reporting	No	No	Partial	PDF assessment report

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

2. Structure of the Self-Assessment Tool

2.1. Data Collection and Sampling

The Data Management System (DMS) serves as the central repository for all data collected throughout the pilot data centres in the HEATWISE project. It is designed to aggregate necessary IT, data centre and energy grids data to implement holistic data centre management solutions. The DMS can collect data from IT, cooling units (liquid and air cooling), building, energy storage systems, heat pump and the environment.

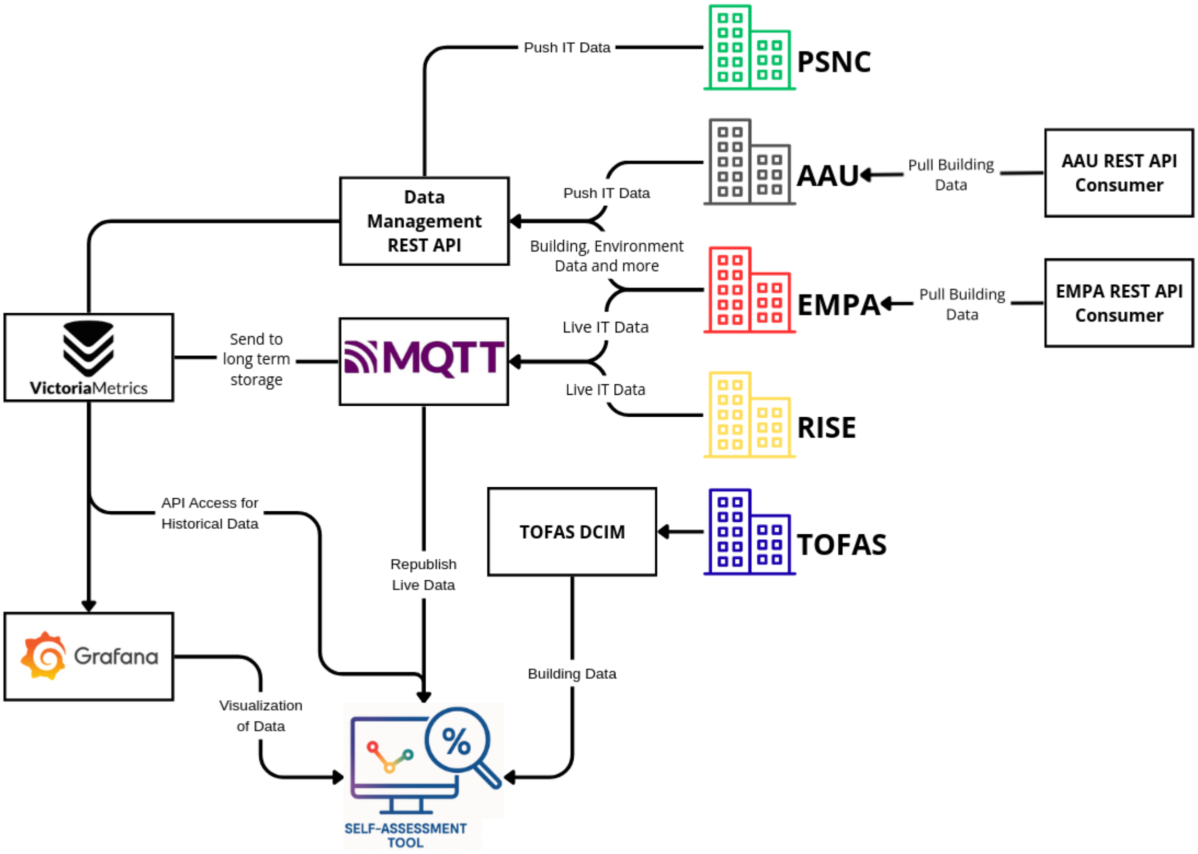

Figure 2.1: Data management system and data pipelines from pilots.

The SAT retrieves operational and thermal data from pilot data centres through a unified data acquisition layer. The primary data source is the DMS, which aggregates real-time and historical measurements from IT and cooling infrastructures. The architecture is designed to aggregate heterogeneous data from pilot data centres and energy grids into a unified structure. The core of this system is Victoria Metrics, a high-performance time-series database that provides synchronized storage and rapid retrieval for both real-time and historical data. Data ingestion is handled through two primary communication pipelines to ensure flexibility: REST API consumers and MQTT Proxy. REST API Consumers are dedicated services that pull building data and environmental metrics from the pilots. Additionally, a Data Management system’s REST API allows for the direct pushing of IT data to the project partners. Real-time live IT data from pilots flows through an MQTT Proxy. This component serves a dual purpose: it routes data to long-term storage (Victoria Metrics) and simultaneously republishes live data streams for immediate consumption by project partners.

As depicted in Fig. 2.2, the SAT employs a comprehensive and adaptable data collection and sampling strategy to handle both real-time data streams and externally submitted datasets.

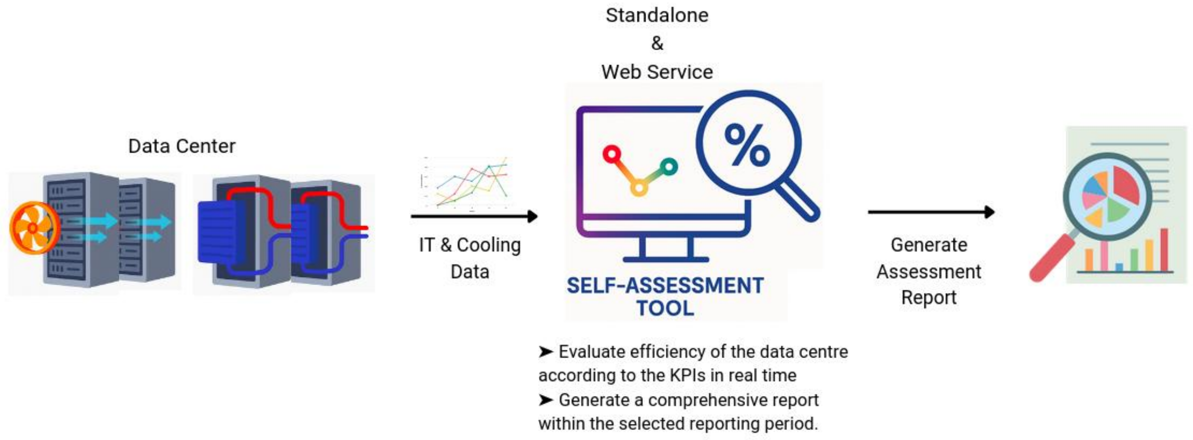

➤ Evaluate efficiency of the data centre according to the KPIs in real time
➤ Generate a comprehensive report within the selected reporting period.

Figure 2.2: End-to-end workflow of the SAT from data acquisition to automated assessment reporting.

The collected data mainly include: (i) server inlet and outlet air temperatures, (ii) server-level thermal measurements, (iii) server power consumptions, (iv) cooling unit data and (v) auxiliary operational parameters required for KPI evaluation. All queries are executed dynamically based on the user-defined configuration selected in the setup interface. SAT supports flexible temporal data selection, allowing users to analyse both short-term operational behaviour and long-term performance trends. Two complementary approaches are implemented: (i) relative time window selection, where users specify the last X hours of operation and (ii) absolute time range selection, where start and end date-time values are explicitly defined. To balance temporal resolution with computational efficiency, SAT introduces a sampling step parameter. This parameter defines the time interval between consecutive data points retrieved from the database. A smaller step size provides higher temporal resolution and captures transient behaviour. A larger step size reduces data volume and is suitable for long-term performance assessment. The step size is defined by the user during the setup phase and

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

is propagated consistently throughout the data collection and KPI computation pipeline. This design avoids implicit resampling and ensures transparency in how temporal aggregation is performed.

After retrieval, raw time-series data undergo a lightweight pre-processing stage before KPI calculation: (i) duplicate sensor entries are removed to avoid bias, (ii) data are grouped based on physical meaning (e.g., server inlet temperatures, server outlet temperatures, chip temperatures) and (iii) for each sensor, representative statistics such as time-averaged values are computed over the selected interval. A key design principle of SAT is reproducible KPI evaluation. All data collection parameters such as time range, sampling step, data centre selection and KPI set are explicitly defined by the user and recorded internally. This guarantees that KPI results can be reproduced for the same configuration and generated reports accurately reflect the underlying data selection and processing logic. This strategy is particularly important for KPI formulations that rely on statistical distributions (e.g., percentage of servers operating outside recommended temperature bounds). The sampled and pre-processed data form the direct input to the KPI evaluation module. Thermal and energy KPIs (see, D8.1) are computed using the same sampled dataset, ensuring internal consistency.

Finally, both the numerical results and the underlying time-series visualisations are seamlessly integrated into the SAT dashboard and exported into automated PDF reports. This end-to-end workflow links data collection and sampling directly to decision-support outputs.

2.2. Calculation of KPIs

The calculation of KPIs in SAT is performed using synchronized thermal and energy-related time series retrieved from the DMS infrastructure. For a given analysis period, inlet and outlet air temperatures at server level, air supply and return temperatures at room or cooling unit level, and power-related measurements are retrieved as time series. To handle potential differences in sampling timestamps between these signals, a nearest-neighbour matching strategy with a defined temporal tolerance is applied. This guarantees that all quantities used in a single KPI evaluation correspond to the same physical operating state. These KPIs directly link operational performance with ASHRAE and EED objectives.

Thermal performance indicators are primarily derived from temperature differences between inlet, outlet, supply, and return air streams. The Return Temperature Index (RTI) is calculated as a normalized ratio comparing the measured return-to-supply temperature rise with the temperature rise across IT equipment. RTI values close to one indicate effective air management, while deviations suggest bypass or short-circuiting

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

effects. The Return Heat Index (RHI) is derived from the Supply Heat Index (SHI), which quantifies the fraction of supplied cooling air that directly contributes to cooling IT equipment. RHI therefore reflects how effectively the heat generated by servers is captured by the return air stream, providing insight into containment quality and thermal mixing behaviour. Temperature compliance with recommended operating limits is evaluated using the Rack Cooling Index (RCI). In the current implementation, RCI is computed using the average inlet temperature over the analysis period, scaled by the total number of servers. Separate high- and low-temperature indices (RCI_{HI} and RCI_{LO}) quantify deviations above and below ASHRAE-recommended temperature ranges, respectively. These metrics provide a normalized measure of overall thermal compliance. Local air management effects are assessed using the Recirculation Index (RI) and the Locality Index (LI). RI estimates the degree of warm air recirculation by comparing inlet and supply temperatures, while LI captures mismatches between outlet and return temperatures, highlighting spatial non-uniformities and localized inefficiencies in airflow distribution.

In parallel with thermal KPIs, SAT computes a comprehensive set of energy-related KPIs. Total IT power consumption, cooling unit power consumption, and heat recovery unit (HRU) power consumption are aggregated for each evaluation timestamp. The Power Usage Effectiveness (PUE) is then calculated as the ratio of total facility power to IT power, explicitly accounting for cooling and heat recovery system energy use. Based on the calculated PUE, a derived Coefficient of Performance (COP) is evaluated to provide an intuitive measure of cooling efficiency. This formulation enables rapid comparison of cooling effectiveness across different operating conditions and pilot data centres. When heat recovery data are available, recovered thermal power is computed from water mass flow rate and inlet–outlet temperature difference of the HRU loop. This recovered heat is used to calculate the Energy Reuse Factor (ERF), which quantifies the fraction of IT power that is effectively reused as useful thermal energy. Finally, an energy-level performance assessment is performed using the Primary Energy Savings (PES) indicator. PES compares the primary energy demand of the actual system configuration with a reference scenario, incorporating electricity and thermal primary energy factors. This metric directly links operational performance with Energy Efficiency Directive (EED) objectives. An energy-level performance assessment is performed using the Primary Energy Savings (PES) indicator. PES compares the primary energy demand of the actual system configuration with a reference scenario, incorporating electricity and thermal primary energy factors. In addition to dimensionless energy efficiency indicators, SAT explicitly computes the recovered thermal power (P_{reuse}) from the heat recovery unit (HRU). This quantity represents the instantaneous amount of useful heat extracted from the data centre cooling loop and is calculated using the measured water mass flow rate and the inlet–outlet temperature difference across the HRU. Unlike normalized KPIs, P_{reuse} is expressed

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

in absolute power units and provides a direct measure of the thermal energy recovery potential of the system.

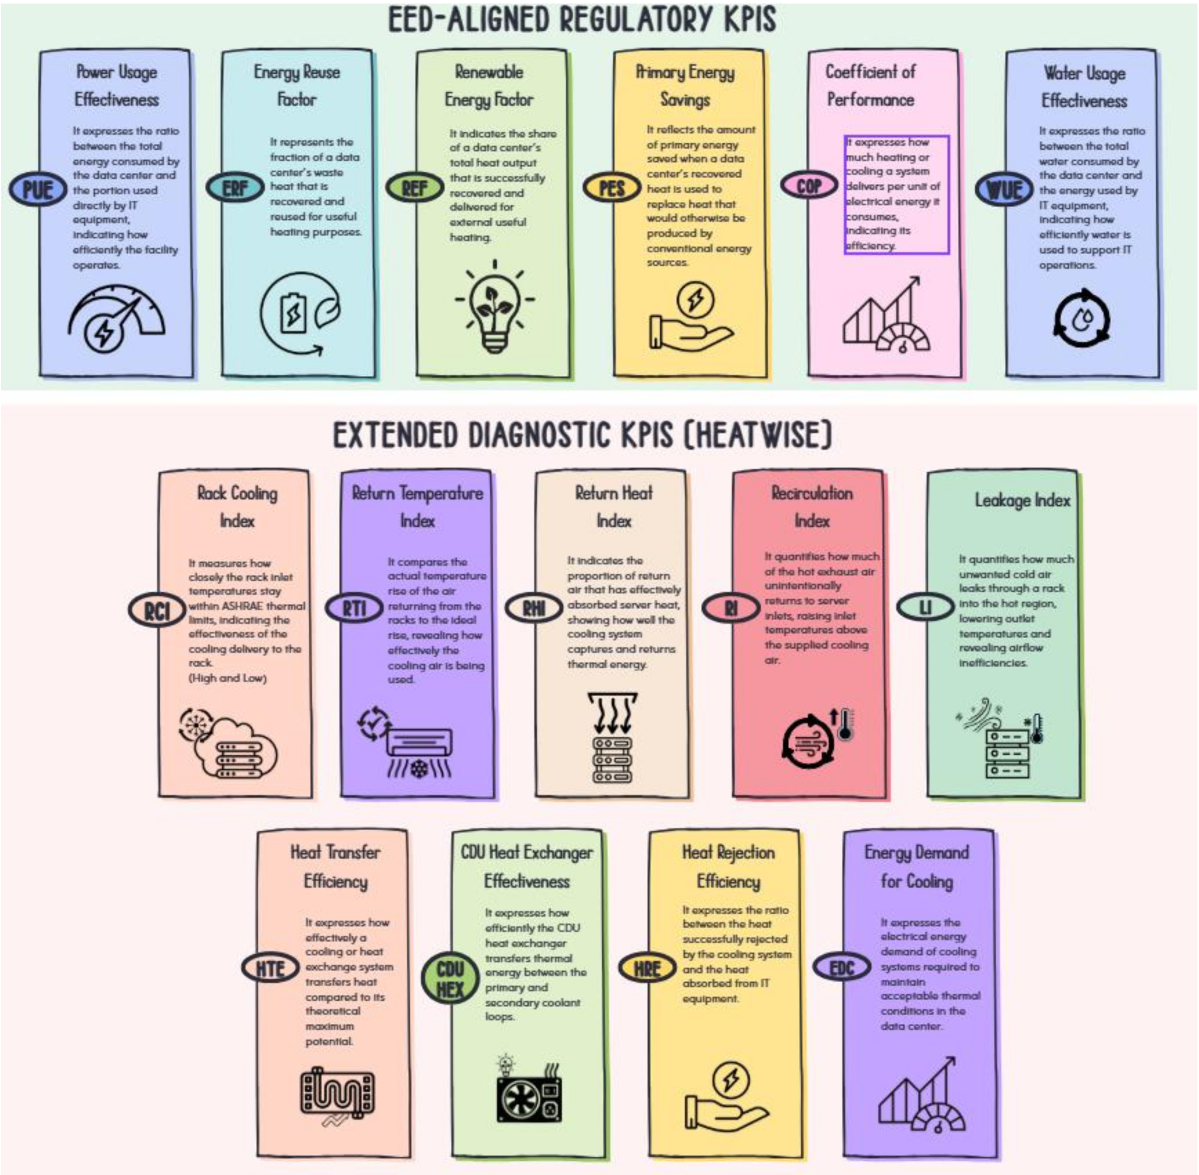

Power Usage Effectiveness
PUE
It expresses the ratio between the total energy consumed by the data center and the portion used directly by IT equipment, indicating how efficiently the facility operates.

Energy Reuse Factor
ERF
It represents the fraction of a data center's waste heat that is recovered and reused for useful heating purposes.

Renewable Energy Factor
REF
It indicates the share of a data center's total heat output that is successfully recovered and delivered for external useful heating.

Primary Energy Savings
PES
It reflects the amount of primary energy saved when a data center's recovered heat is used to replace heat that would otherwise be produced by conventional energy sources.

Coefficient of Performance
COP
It expresses how much heating or cooling a system delivers per unit of electrical energy it consumes, indicating its efficiency.

Water Usage Effectiveness
WUE
It expresses the ratio between the total water consumed by the data center and the energy used by IT equipment, indicating how efficiently water is used to support IT operations.

Rack Cooling Index
RCI
It measures how closely the rack inlet temperatures stay within ASHRAE thermal limits, indicating the effectiveness of the cooling delivery to the rack.
(High and Low)

Return Temperature Index
RTI
It compares the actual temperature rise of the air returning from the racks to the ideal rise, revealing how effectively the cooling air is being used.

Return Heat Index
RHI
It indicates the proportion of return air that has effectively absorbed server heat, showing how well the cooling system captures and returns thermal energy.

Recirculation Index
RI
It quantifies how much of the hot exhaust air unintentionally returns to server inlets, raising inlet temperatures above the supplied cooling air.

Leakage Index
LI
It quantifies how much unwanted cold air leaks through a rack into the hot region, lowering outlet temperatures and revealing airflow inefficiencies.

Heat Transfer Efficiency
HTE
It expresses how effectively a cooling or heat exchange system transfers heat compared to its theoretical maximum potential.

CDU Heat Exchanger Effectiveness
CDU HEX
It expresses how efficiently the CDU heat exchanger transfers thermal energy between the primary and secondary coolant loops.

Heat Rejection Efficiency
HRE
It expresses the ratio between the heat successfully rejected by the cooling system and the heat absorbed from IT equipment.

Energy Demand for Cooling
EDC
It expresses the electrical energy demand of cooling systems required to maintain acceptable thermal conditions in the data center.

Figure 2.3: Set of KPIs computable within the SAT framework.

All KPIs are computed deterministically without additional filtering or resampling beyond the initial time alignment step. The resulting KPI time series are logged and exported in a structured format, ensuring traceability from raw measurements to final performance indicators and enabling seamless integration with visualization and reporting modules. For clarity from a regulatory and client perspective, KPIs are grouped into EED-aligned indicators and extended diagnostic indicators developed within HEATWISE. This separation highlights SAT's role both as a compliance-support tool and as an advanced operational assessment framework.

2.3. System architecture of the SAT

The SAT is designed as a modular and extensible system that integrates data acquisition, KPI computation, visualization, and reporting into a unified workflow. The architecture follows a layered approach, separating data access, processing logic, and user interaction to ensure scalability, transparency, and ease of maintenance. The SAT interfaces with the data centre monitoring infrastructure through a data access layer. This layer retrieves thermal and energy-related time-series data from the Data Management System (DMS), which stores measurements in a time-series database compatible with REST API query semantics. By relying on standardized REST-based queries, SAT remains independent of specific sensor hardware, monitoring vendors, or pilot-specific implementations.

Above the data access layer, a data alignment and preprocessing layer ensures temporal consistency across signals. Temperature measurements, power consumption data, and heat recovery parameters often originate from different subsystems with varying sampling frequencies. SAT resolves this by mapping all signals onto a common evaluation timeline using nearest-neighbour matching with a defined tolerance. This guarantees that thermal and energy KPIs are computed using physically consistent system states. The core analytical layer of SAT is the KPI computation engine, implemented as a standalone and reusable module. This layer processes synchronized input data to calculate both thermal KPIs (such as RTI, RHI, RCI, RI, and LI) and energy-related KPIs (including PUE, COP, REF, ERF, and PES). The computation logic is deterministic and configuration-driven, ensuring that results are reproducible for a given time range, sampling step, and data centre selection. All SAT evaluations are driven by an explicit configuration layer that defines the analysis time window, sampling resolution, data centre selection, and KPI set. These parameters are propagated consistently across data acquisition, preprocessing, KPI computation, and reporting modules. As a result, each assessment can be fully reproduced using the same configuration, enabling transparent comparison between different scenarios, pilot sites, or operational periods.

On top of the analytical layer, SAT provides an interactive user interface that allows users to configure the assessment parameters. Through the setup interface, users select the pilot data centre, define the analysis time window, specify the sampling resolution, and choose the set of KPIs to be evaluated. These configuration parameters are propagated consistently throughout the system, ensuring full traceability between user input and computed results.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

The visualization and reporting layer transforms computed KPIs into actionable insights. KPI time series, summary statistics, and comparative indicators are displayed within the SAT dashboard, enabling immediate interpretation of thermal and energy performance. In parallel, the same results are automatically embedded into structured PDF self-assessment reports, ensuring that visualizations and numerical values remain consistent across interactive and offline outputs. A key architectural feature of SAT is its pilot-agnostic design. Differences between air-cooled, liquid-cooled and hybrid-cooled data centres, availability of heat recovery systems, or variations in sensor coverage are handled through configuration rather than code modification. This allows SAT to be deployed across multiple pilot sites while maintaining a common assessment methodology.

Overall, the SAT architecture enables a seamless end-to-end workflow (Fig 2.4), linking raw operational measurements to high-level performance indicators and decision-support outputs. By combining modular design, standardized data interfaces, and integrated reporting, SAT provides a robust and extensible platform for thermal and energy performance assessment of data centres in line with EED objectives.

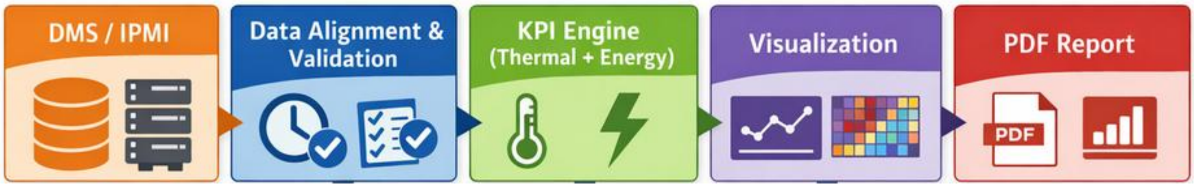

Figure 2.4: End-to-end workflow of SAT from data collection to reporting.

In addition to live data acquisition through the DMS, SAT also supports an offline data ingestion mode based on user-provided IT datasets. This feature enables users to upload historical or campaign-based measurements, such as server inlet temperatures and power consumption extracted from IT logs, in a structured file format. Once uploaded, these datasets are processed using the same data alignment, sampling, and KPI computation pipeline as DMS-sourced data. This dual data ingestion capability significantly increases the flexibility and applicability of SAT. It allows users to perform self-assessments in environments where direct DMS access is not available, during commissioning phases or for post-analysis of experimental campaigns. By ensuring that both online (DMS-based) and offline (IT and cooling unit based) workflows converge into the same analytical process, SAT guarantees methodological consistency and enables fair comparison between real-time monitoring results and historical performance evaluations.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

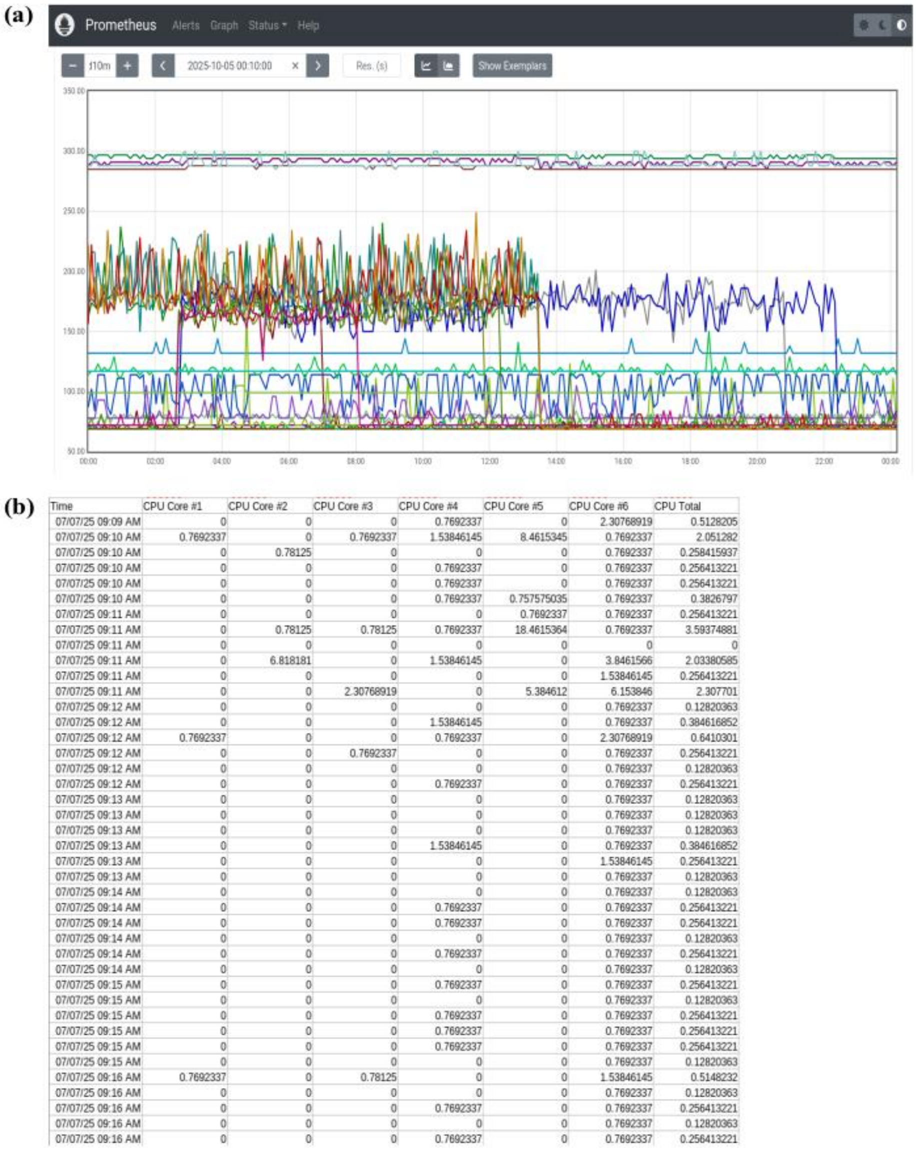

Time	CPU Core #1	CPU Core #2	CPU Core #3	CPU Core #4	CPU Core #5	CPU Core #6	CPU Total
07/07/25 09:09 AM	0	0	0	0.7692337	0	2.30768919	0.5128205
07/07/25 09:10 AM	0.7692337	0	0.7692337	1.53846145	8.4615345	0.7692337	2.051282
07/07/25 09:10 AM	0	0.78125	0	0	0	0.7692337	0.258415937
07/07/25 09:10 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:10 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:10 AM	0	0	0	0.7692337	0.757575035	0.7692337	0.3826797
07/07/25 09:11 AM	0	0	0	0	0.7692337	0.7692337	0.256413221
07/07/25 09:11 AM	0	0.78125	0.78125	0.7692337	18.4615364	0.7692337	3.59374881
07/07/25 09:11 AM	0	0	0	0	0	0	0
07/07/25 09:11 AM	0	6.818181	0	1.53846145	0	3.8461566	2.03380585
07/07/25 09:11 AM	0	0	0	0	0	1.53846145	0.256413221
07/07/25 09:11 AM	0	0	2.30768919	0	5.384612	6.153846	2.307701
07/07/25 09:12 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:12 AM	0	0	0	1.53846145	0	0.7692337	0.384616852
07/07/25 09:12 AM	0.7692337	0	0	0.7692337	0	2.30768919	0.6410301
07/07/25 09:12 AM	0	0	0.7692337	0	0	0.7692337	0.256413221
07/07/25 09:12 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:12 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:13 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:13 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:13 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:13 AM	0	0	0	1.53846145	0	0.7692337	0.384616852
07/07/25 09:13 AM	0	0	0	0	0	1.53846145	0.256413221
07/07/25 09:13 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:14 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:14 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:14 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:14 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:14 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:14 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:15 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:15 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:15 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:15 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:15 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:15 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:16 AM	0.7692337	0	0.78125	0	0	1.53846145	0.5148232
07/07/25 09:16 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:16 AM	0	0	0	0.7692337	0	0.7692337	0.256413221
07/07/25 09:16 AM	0	0	0	0	0	0.7692337	0.12820363
07/07/25 09:16 AM	0	0	0	0.7692337	0	0.7692337	0.256413221

Figure 2.5: Data types used in SAT: (a) real-time data and (b) externally provided dataset.

The SAT architecture is designed to ensure portability, reproducibility, and integration with the broader data infrastructure. The SAT is designed as both a standalone desktop and web-based version. The standalone version is a desktop Graphical User Interface (GUI) application built with Python [41] and Tkinter/ttk components [42]. It provides a fully independent environment for local data retrieval, KPI calculation and visualization. On the other hand, the web version delivers the same functionality through a Representational State Transfer (REST) [43] based Application Programming Interface (API), enabling the frontend to request KPI calculations, data queries and visual outputs from the backend. The frontend communicates with the Flask [44-45] backend via REST endpoints. A Flask backend handles API requests, executes KPI calculations, generates plots and reports, and serves them to an HTML [46-47] frontend. This web-based version enables multiuser access and remote operation.

For computational processes, the SAT standalone version employs pandas [48] and NumPy [49] for tabular data manipulation, numerical operations, and timestamp alignment. Visualization components are implemented using matplotlib [50-51] and seaborn [52], supporting the generation of KPI time-series plots, heatmaps, server-level ASHRAE compliance and other thermal-performance visualizations. Automated reporting is achieved via ReportLab [52] and matplotlib backends, which generate PDF reports with embedded plots, tables and textual summaries. The web interface is built with HTML and with optional use of Bootstrap [53] for styling and Plotly [54-55] for interactive charts. This layered combination of components shown in Fig. 2.6 enables cross-platform deployment, allows scalable integration with all systems and provides consistent workflow across both standalone and web-based SAT environments. For the web-based deployment, SAT supports controlled access to data and computation services through the backend API. User requests for data queries, KPI evaluation, and report generation are handled server-side, allowing integration with authentication mechanisms and secure data handling practices where required by the hosting environment.

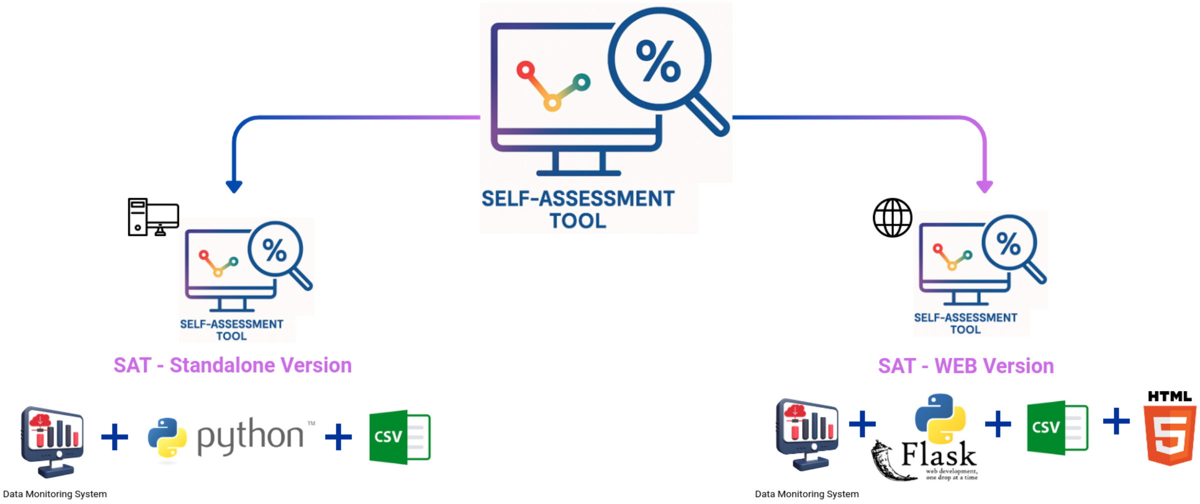

Figure 2.6: Components of SAT in standalone and web-based versions.

The modular architecture of SAT facilitates future extensions without modification of the core system. New KPIs, additional data sources, or alternative visualization components can be integrated by extending the corresponding layer while preserving existing functionality. This design enables SAT to evolve alongside future monitoring technologies, regulatory requirements, and pilot-specific needs.

2.3.1.Sub-windows of the SAT

The SAT user interaction layer is organized around two main components as the setup window and the dashboard, which together guide the user from configuration to performance assessment in a structured and intuitive manner.

The Setup Window serves as the initial interaction point of SAT and is responsible for defining the scope of the assessment. Through this interface, users select the target pilot data centre and specify the temporal boundaries of the analysis, either by defining a relative time window or by selecting explicit start and end timestamps. In addition, users define the sampling resolution used for data retrieval and KPI computation, ensuring that temporal aggregation aligns with the intended analysis scale.

Beyond time and sampling configuration, the setup window allows users to determine the size of the IT population considered in the assessment, such as the number of monitored servers. This parameter is required for normalization in several KPIs and enables SAT to adapt to different data centre scales without structural changes. Once the configuration is confirmed, all selected parameters are passed consistently to downstream components, ensuring traceability between user input and computed results.

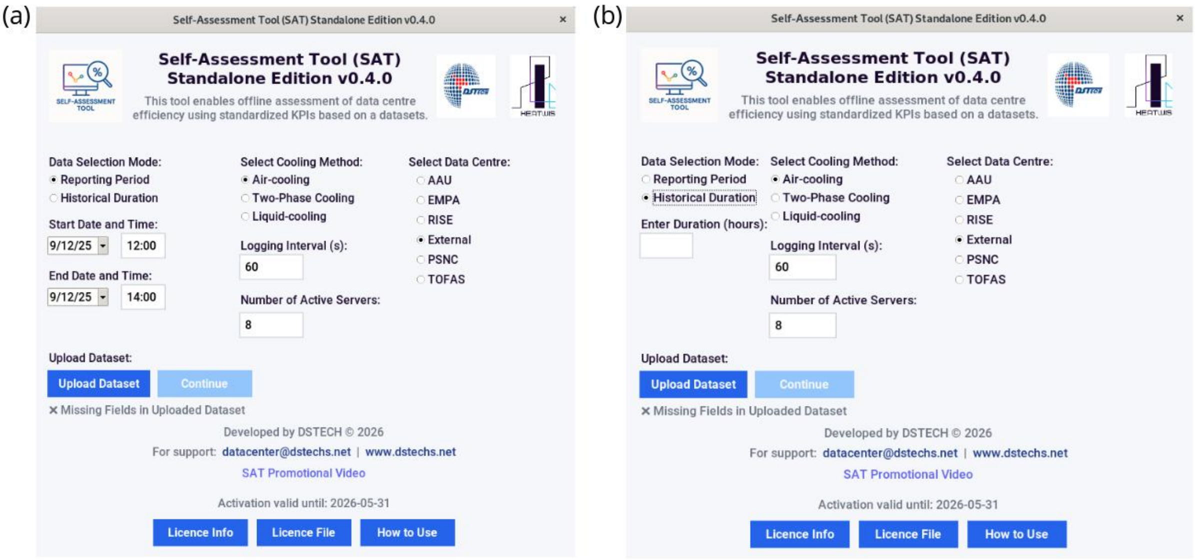
(a)

Self-Assessment Tool (SAT) Standalone Edition v0.4.0

Self-Assessment Tool (SAT)
Standalone Edition v0.4.0
This tool enables offline assessment of data centre efficiency using standardized KPIs based on a datasets.

Data Selection Mode:
● Reporting Period
○ Historical Duration

Start Date and Time: 9/12/25 12:00
End Date and Time: 9/12/25 14:00

Select Cooling Method:
● Air-cooling
○ Two-Phase Cooling
○ Liquid-cooling

Logging Interval (s): 60
Number of Active Servers: 8

Select Data Centre:
○ AAU
○ EMPA
○ RISE
● External
○ PSNC
○ TOFAS

Upload Dataset:
× Missing Fields in Uploaded Dataset

Developed by DSTECH © 2026
For support: datacenter@dstechs.net | www.dstechs.net
SAT Promotional Video
Activation valid until: 2026-05-31

(b)

Self-Assessment Tool (SAT) Standalone Edition v0.4.0

Self-Assessment Tool (SAT)
Standalone Edition v0.4.0
This tool enables offline assessment of data centre efficiency using standardized KPIs based on a datasets.

Data Selection Mode:
○ Reporting Period
● Historical Duration

Enter Duration (hours):

Select Cooling Method:
● Air-cooling
○ Two-Phase Cooling
○ Liquid-cooling

Logging Interval (s): 60
Number of Active Servers: 8

Select Data Centre:
○ AAU
○ EMPA
○ RISE
● External
○ PSNC
○ TOFAS

Upload Dataset:
× Missing Fields in Uploaded Dataset

Developed by DSTECH © 2026
For support: datacenter@dstechs.net | www.dstechs.net
SAT Promotional Video
Activation valid until: 2026-05-31

Figure 2.6: Selection of (a) reporting period and (b) historical duration in the setup window.

Following configuration, SAT transitions to the Dashboard, which functions as the main analysis and visualization environment. The Dashboard presents the computed thermal and energy KPIs in a consolidated view, combining numerical indicators, time-series plots, and distribution-based visualizations. This layout allows users to simultaneously assess instantaneous behaviour, temporal trends, and aggregated performance metrics. Within the Dashboard, KPI values are displayed both as time-resolved curves and as summary statistics, such as average and latest values over the selected period. Thermal KPIs and energy-related KPIs are presented in a coordinated manner, enabling users to directly relate airflow performance and temperature compliance to energy efficiency and heat recovery behaviour.

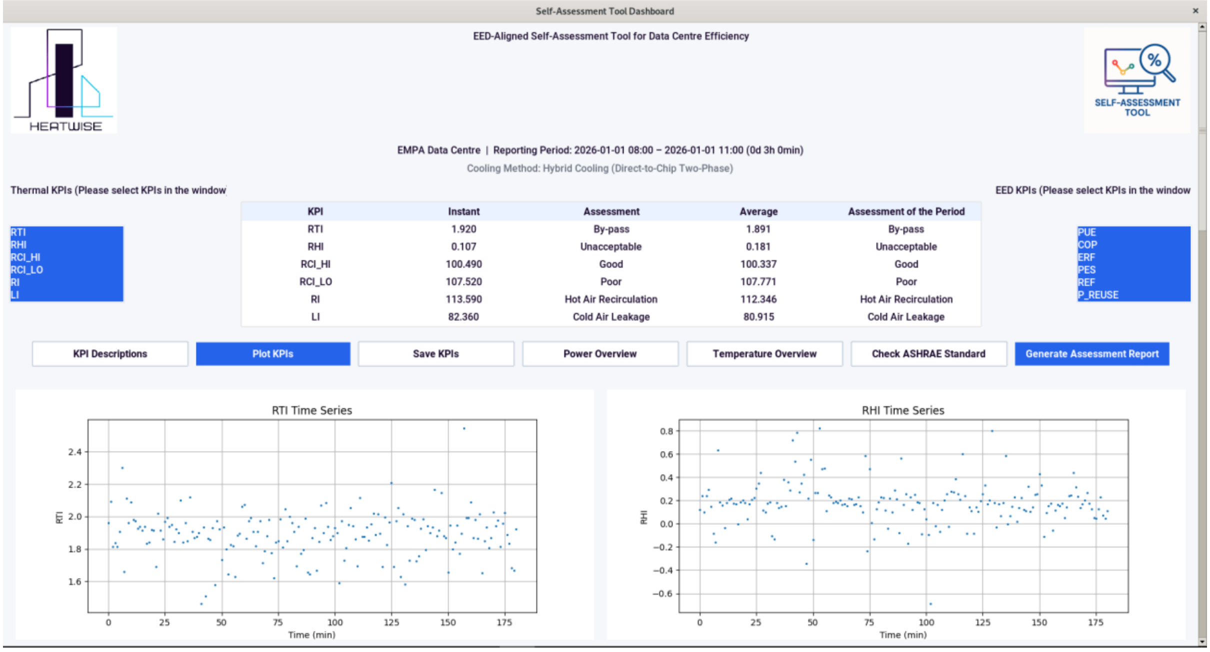

EED-Aligned Self-Assessment Tool for Data Centre Efficiency

EMPA Data Centre | Reporting Period: 2026-01-01 08:00 – 2026-01-01 11:00 (0d 3h 0min)

Cooling Method: Hybrid Cooling (Direct-to-Chip Two-Phase)

KPI	Instant	Assessment	Average	Assessment of the Period
RTI	1.920	By-pass	1.891	By-pass
RHI	0.107	Unacceptable	0.181	Unacceptable
RCI_HI	100.490	Good	100.337	Good
RCI_LO	107.520	Poor	107.771	Poor
RI	113.590	Hot Air Recirculation	112.346	Hot Air Recirculation
LI	82.360	Cold Air Leakage	80.915	Cold Air Leakage

Figure 2.7:. The dashboard of the SAT.

To clearly distinguish between visualization and evaluation, the KPI assessment table shown in Fig. 2.8 is treated as a separate element of the Dashboard. This table represents the assessment layer of SAT, where the computed KPIs are aggregated and interpreted against predefined reference criteria. While the plots provide insight into temporal dynamics and operational variability, the KPI table summarizes the assessed performance over the selected time window, forming the primary basis for performance evaluation and comparison.

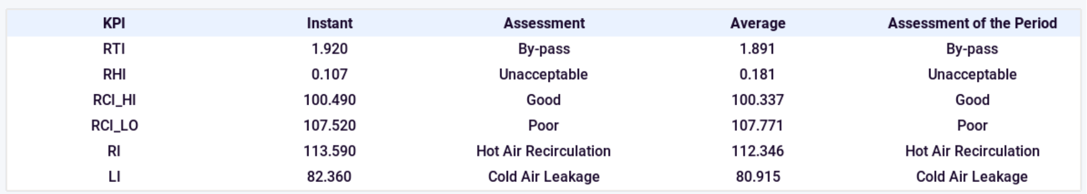

KPI	Instant	Assessment	Average	Assessment of the Period
RTI	1.920	By-pass	1.891	By-pass
RHI	0.107	Unacceptable	0.181	Unacceptable
RCI_HI	100.490	Good	100.337	Good
RCI_LO	107.520	Poor	107.771	Poor
RI	113.590	Hot Air Recirculation	112.346	Hot Air Recirculation
LI	82.360	Cold Air Leakage	80.915	Cold Air Leakage

Figure 2.8: Performance evaluation according to the KPIs.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

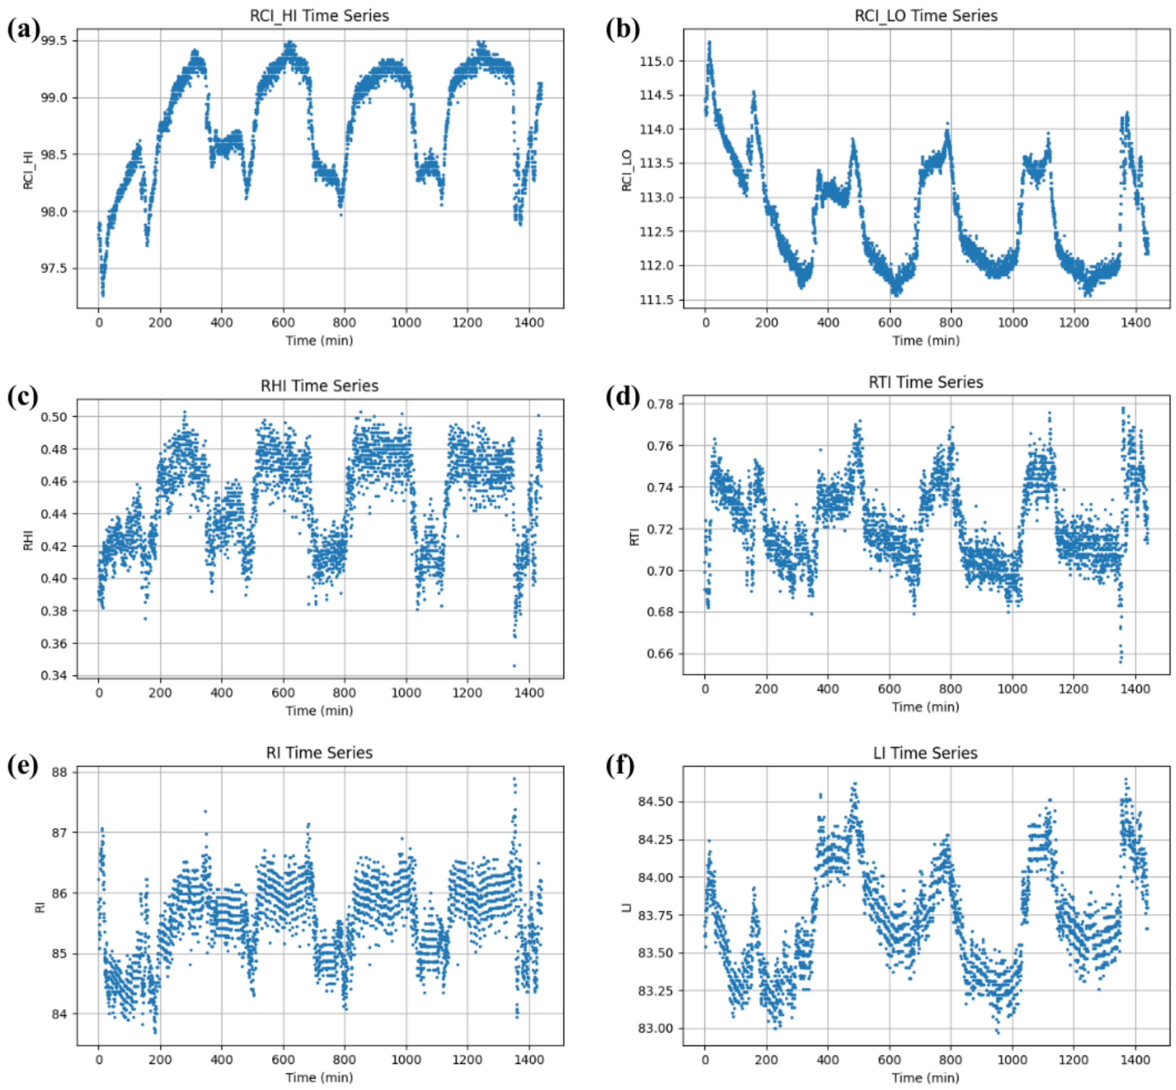

Figure 2.9: Time series of thermal KPIs calculated by SAT.

The dashboard also supports advanced visual diagnostics, including power and temperature bar charts and compliance maps, which highlight spatial and categorical distributions such as ASHRAE temperature compliance at server level. These visual elements facilitate rapid identification of hotspots, recirculation zones, or localized inefficiencies that may not be apparent from global KPIs alone. In addition to interactive analysis, the Dashboard provides direct access to automated reporting functionality. Users can generate self-assessment reports that encapsulate the selected configuration, computed KPIs, and corresponding visualizations in a structured PDF format. This ensures that insights derived during interactive exploration can be preserved, shared, and reused for documentation or compliance purposes.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

In addition to temporal analysis, spatial diagnostics play a key role in interpreting thermal and power distributions within the data centre. Rack-level thermal and power maps enable the identification of hotspots, non-uniform cooling zones, and load concentrations, allowing quantitative KPI deviations to be linked to specific physical locations. Server-level power distributions further highlight clusters of high computational activity or underutilized resources, supporting correlation between power densities and observed thermal behaviour.

Figure 2.10 presents a rack-level spatial analysis, which may correspond to a different selected analysis interval. Rack-based visualization represents an alternative evaluation mode supported by SAT and can be generated independently from server-level time-specific snapshots. Figure 2.10 illustrates an example rack-level heatmap generated by SAT for an air-cooled rack configuration at the EMPA pilot. This visualization does not correspond to the direct-to-chip cooling infrastructure of EMPA but demonstrates the rack-based diagnostic capability supported by SAT.

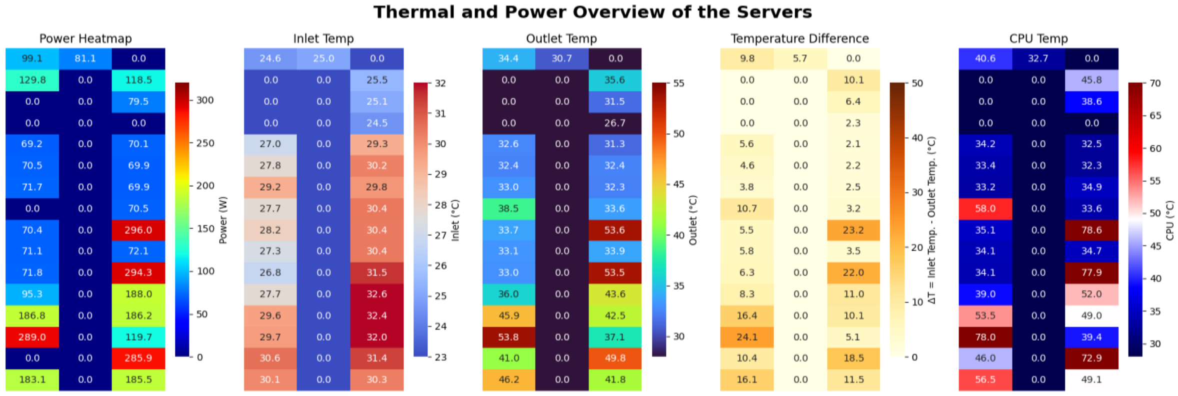

Power Heatmap

99.1	81.1	0.0
129.8	0.0	118.5
0.0	0.0	79.5
0.0	0.0	0.0
69.2	0.0	70.1
70.5	0.0	69.9
71.7	0.0	69.9
0.0	0.0	70.5
70.4	0.0	296.0
71.1	0.0	72.1
71.8	0.0	294.3
95.3	0.0	188.0
186.8	0.0	186.2
289.0	0.0	119.7
0.0	0.0	285.9
183.1	0.0	185.5

Inlet Temp

24.6	25.0	0.0
0.0	0.0	25.5
0.0	0.0	25.1
0.0	0.0	24.5
27.0	0.0	29.3
27.8	0.0	30.2
29.2	0.0	29.8
27.7	0.0	30.4
28.2	0.0	30.4
27.3	0.0	30.4
26.8	0.0	31.5
27.7	0.0	32.6
29.6	0.0	32.4
29.7	0.0	32.0
30.6	0.0	31.4
30.1	0.0	30.3

Outlet Temp

34.4	30.7	0.0
0.0	0.0	35.6
0.0	0.0	31.5
0.0	0.0	26.7
32.6	0.0	31.3
32.4	0.0	32.4
33.0	0.0	32.3
38.5	0.0	33.6
33.7	0.0	53.6
33.1	0.0	33.9
33.0	0.0	53.5
36.0	0.0	43.6
45.9	0.0	42.5
53.8	0.0	37.1
41.0	0.0	49.8
46.2	0.0	41.8

Temperature Difference

9.8	5.7	0.0
0.0	0.0	10.1
0.0	0.0	6.4
0.0	0.0	2.3
5.6	0.0	2.1
4.6	0.0	2.2
3.8	0.0	2.5
10.7	0.0	3.2
5.5	0.0	23.2
5.8	0.0	3.5
6.3	0.0	22.0
8.3	0.0	11.0
16.4	0.0	10.1
24.1	0.0	5.1
10.4	0.0	18.5
16.1	0.0	11.5

CPU Temp

40.6	32.7	0.0
0.0	0.0	45.8
0.0	0.0	38.6
0.0	0.0	0.0
34.2	0.0	32.5
33.4	0.0	32.3
33.2	0.0	34.9
58.0	0.0	33.6
35.1	0.0	78.6
34.1	0.0	34.7
34.1	0.0	77.9
39.0	0.0	52.0
53.5	0.0	49.0
78.0	0.0	39.4
46.0	0.0	72.9
56.5	0.0	49.1

Figure 2.10: Rack-level power and thermal maps.

In rack-level heatmap visualizations, non-populated server slots are assigned a value of 0 for both power consumption and temperature parameters in the internal data structure. However, these empty slots are intentionally excluded from graphical rendering to avoid misinterpretation of zero values as valid operational measurements. Similarly, CPU-level entries may be omitted from plots if corresponding telemetry data are unavailable or flagged as invalid during preprocessing. This approach ensures consistency between declared infrastructure and the set of valid measurements used for visualization and KPI computation.

To improve traceability between the declared infrastructure and the visualized thermal maps, a schematic rack identification diagram is provided in Figure 2.11.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

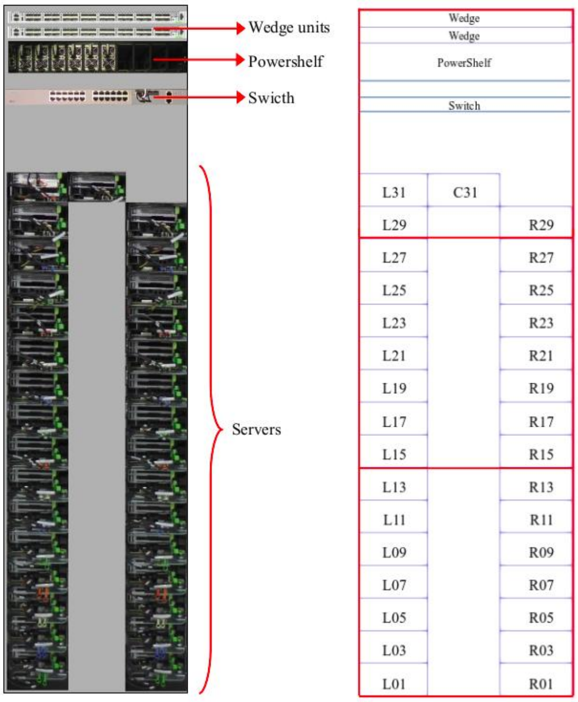

Figure 2.11: Air-cooled rack layout at the EMPA pilot.

Figures 2.12 and 2.13 correspond to the same selected operational timestamp. The ordering along the horizontal axis follows the physical rack/server layout in the data centre rather than numerical sorting.

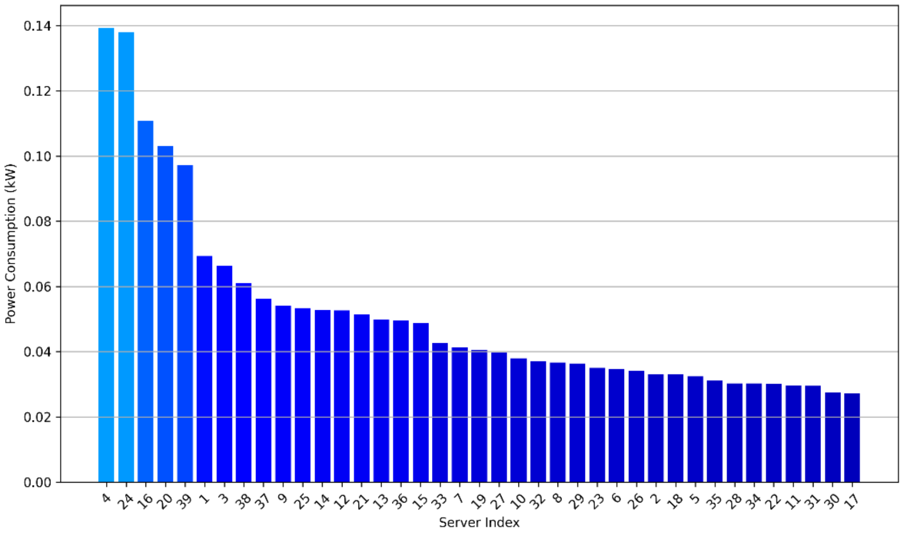

Figure 2.12: Power consumption of servers at the EMPA pilot.

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

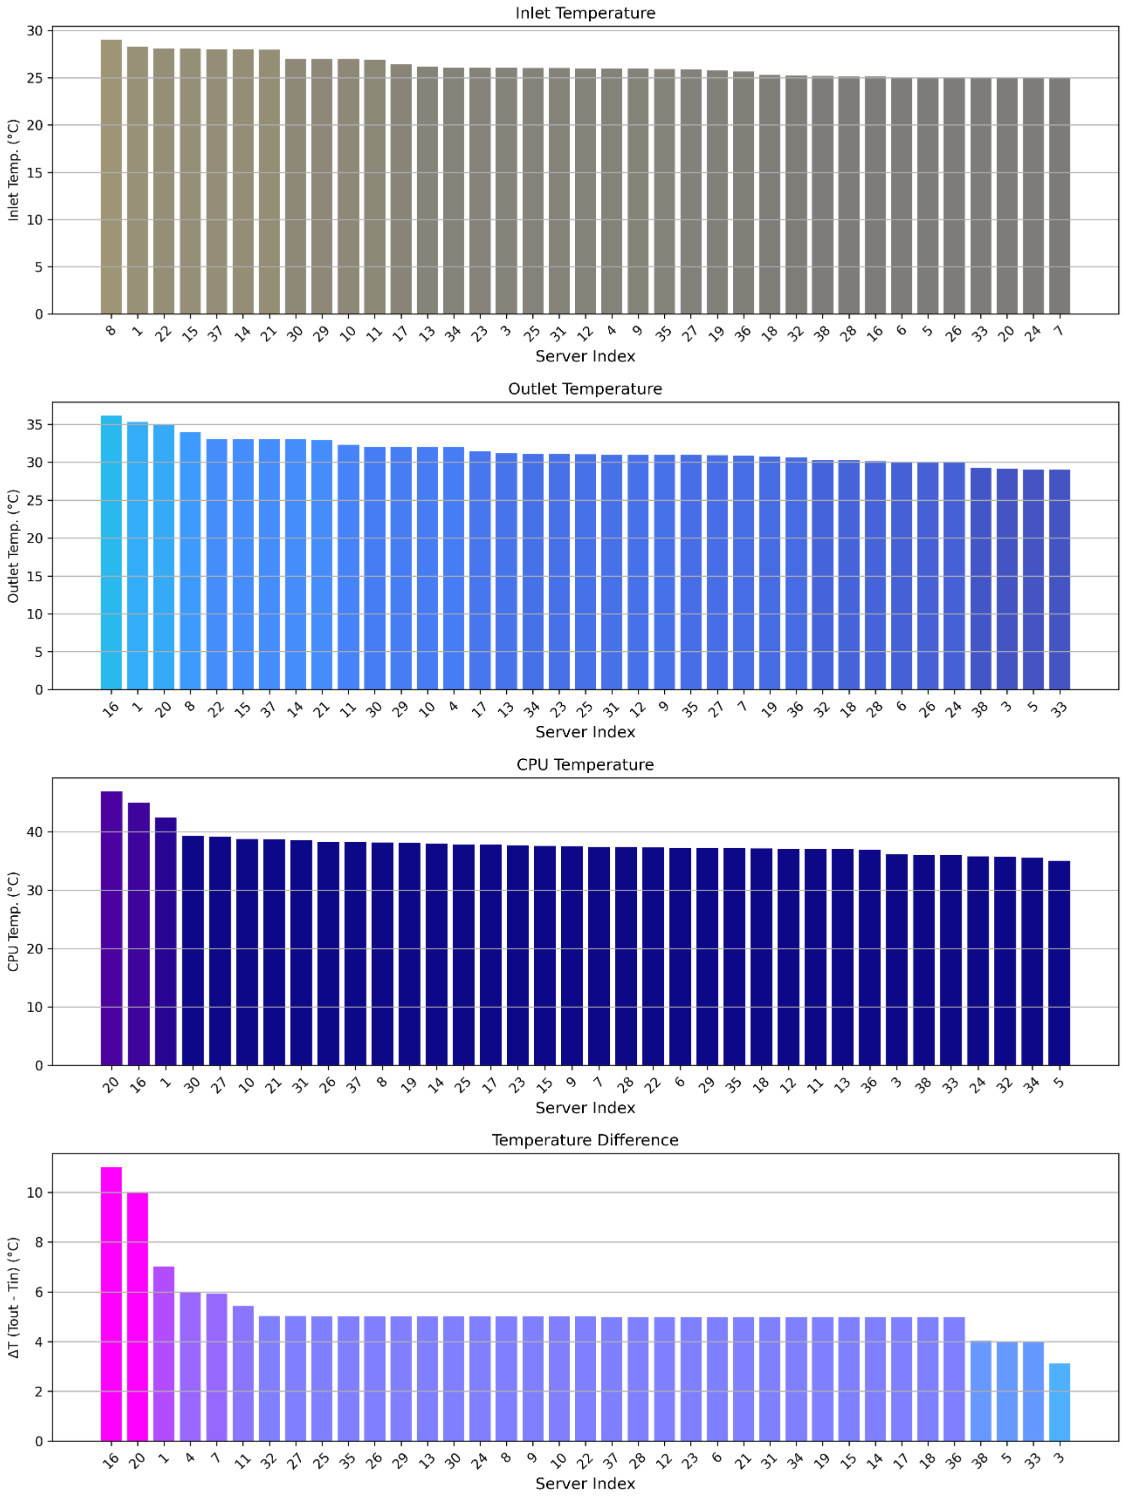

Figure 2.13: Thermal characteristics of servers at the EMPA pilot.

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

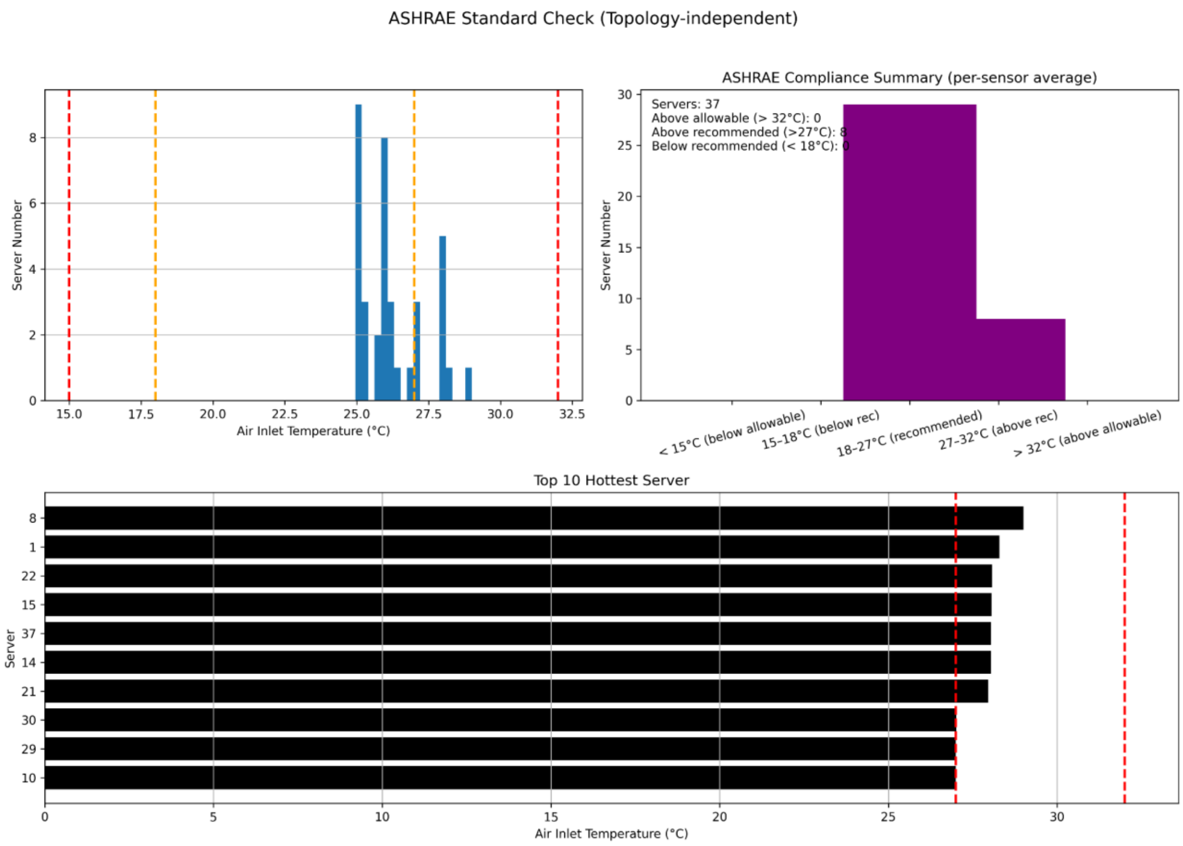

Figure 2.14: ASHRAE compliance assessment for air and hybrid-cooled data centres.

In addition to interactive visualization, SAT provides an integrated reporting functionality through the Generate Report feature available directly within the Dashboard. This functionality enables users to automatically generate a comprehensive PDF self-assessment report based on the current dashboard configuration and computed results. When the Generate Report option is triggered, SAT compiles all relevant assessment parameters, including the selected pilot data centre, analysis time window, sampling resolution, and evaluated KPI set. These configuration details are explicitly documented in the report to ensure transparency and reproducibility of the assessment. The generated PDF report includes a structured summary of both thermal and energy KPIs. For each selected KPI, the report presents representative numerical values, such as time-averaged and final values over the analysis period, allowing stakeholders to quickly assess overall performance without requiring access to the interactive dashboard. In addition to numerical summaries, the report embeds the same visual outputs displayed in the Dashboard. These include KPI time-series plots, distribution-based indicators, and compliance-related visualizations such as ASHRAE temperature classification maps. By reusing the dashboard-generated figures, SAT guarantees consistency between

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

interactive analysis and offline documentation. For pilots equipped with energy monitoring and heat recovery systems, the report also includes energy efficiency indicators such as PUE, COP, REF, Energy Reuse Factor (ERF), and Primary Energy Savings (PES). This ensures that both thermal behaviour and energy performance are jointly documented within a single assessment artifact. The PDF report is generated programmatically and stored locally, enabling easy sharing with stakeholders, project partners, or regulatory bodies. As the report is created using the same analytical pipeline as the dashboard visualizations, it provides a reliable and traceable representation of the system state during the selected evaluation period.

By integrating automated report generation into the dashboard, DC- SAT bridges the gap between interactive performance exploration and formal documentation. This feature supports efficient knowledge transfer, compliance reporting, and long-term performance tracking, particularly in the context of pilot evaluations and EED reporting requirements. Together, the Setup Window and Dashboard establish a clear separation between configuration and analysis while maintaining a seamless workflow. This design reduces user complexity, minimizes configuration errors, and ensures that all assessments performed within SAT follow a consistent and reproducible methodology.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

3. Implementation of SAT at Pilot Data Centres

The SAT has been implemented across the HW pilot data centres following a unified, modular, and scalable system architecture, while allowing pilot-specific adaptations in data acquisition and integration. The core objective of this implementation is to enable a consistent assessment of thermal and energy performance across heterogeneous data centre environments, including different cooling technologies, monitoring infrastructures, and data availability levels.

At some pilot sites, SAT is built around a common data management concept in which raw monitoring data are collected, normalized, and made available through a Data Management System (DMS). Time-series data originating from IT infrastructure, cooling systems, and energy meters are structured using a unified data model and stored in a scalable long-term time-series database (VictoriaMetrics). This approach ensures that SAT can perform KPI calculations over arbitrary time windows, ranging from near-real-time monitoring to long-term historical performance analysis.

SAT supports multiple data ingestion pathways to accommodate the diverse operational constraints of the pilots. Where live connectivity is available, data are retrieved via REST-based APIs and, when required, through MQTT-based streaming for near-real-time analysis. In parallel, SAT provides an offline ingestion mechanism that allows structured CSV datasets to be directly loaded into the platform. This flexibility enables the inclusion of historical datasets, measurement campaigns, or pilot sites with limited API access, without compromising the consistency of KPI evaluation.

On top of the unified data layer, SAT implements standardized thermal and energy KPI calculations, ensuring comparability of results across all pilot data centres. By combining a common architectural backbone with flexible data integration mechanisms, SAT provides a robust and extensible framework for assessing data centre performance across the HW pilot landscape.

3.1. Implementation of SAT at Empa

At the Empa pilot, the SAT implementation is tightly coupled with the existing hybrid cooling monitoring and data acquisition infrastructure, enabling both high-resolution analysis and long-term performance assessment. Hybrid cooling system measurements (including supply/return temperatures, flow rates, and cooling unit operational states) together with IT operational data are continuously transferred to the DMS via REST-based interfaces. Within the DMS, all time-series data are tagged, and persisted in VictoriaMetrics, providing scalable long-term storage and high-frequency monitoring data. SAT accesses these datasets through standardized REST APIs for

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

retrospective KPI evaluation over user-defined time windows, while near-real-time streams are made available through the MQTT proxy for live performance monitoring. This dual access mechanism allows SAT to compute thermal and energy KPIs consistently across different temporal resolutions, correlate cooling system behavior with IT load variations, and identify performance deviations or inefficiencies. The processed results are then visualized via Grafana dashboards, ensuring that SAT outputs are fully interoperable with Empa's existing monitoring ecosystem and can be used both for continuous operational insight and detailed post-analysis of hybrid cooling performance.

The data queries used by SAT for the Empa pilot data centre are summarized in Table 3.1. Server-side queries are used directly on the SAT dashboard for server-level power bar charts, temperature bar charts, and ASHRAE compliance. Additionally, on the KPI calculation module side, total IT power is obtained by summing these series based on timestamps using the get_total_it_power_series() function.

Table 3.1: Server-side and cooling-side data queries used in SAT for the Empa pilot data centre.

Component	Part	Variable	Query
IT Equipments	Server	T_{in}	temperature{pilot="EM",datasource="IT",datatype="TIN"}
		T_{out}	temperature{pilot="EM",datasource="IT",datatype="TEX"}
		$T_{CPU,max}$	temperature{pilot="EM",datasource="IT",datatype="TCPU(2)"}
		P_{IT}	power{pilot="EM",datasource="IT",datatype="PS"}
Cooling Unit	Air-side	T_{supply}	temperature{pilot="EM",datasource="C",datatype="TAS"}
		T_{return}	temperature{pilot="EM",datasource="C",datatype="TAR"}
		P_C	power{pilot="EM",datasource="C",datatype="PC"}
	HRU	P_{HRU}	power{pilot="EM",datasource="HR",datatype="PH"}

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

		Q_{water}	flow{pilot="EM",datasource="HR",datatype="WF"}
		$T_{water,in}$	temperature{pilot="EM",datasource="HR",datatype="WINT"}
		$T_{water,out}$	temperature{pilot="EM",datasource="HR",datatype="WOUT"}
		T_{liquid}	temperature{pilot="EM",datasource="HR",datatype="LT"}
		T_{vapor}	temperature{pilot="EM",datasource="HR",datatype="VT"}

Each data query used by SAT directly contributes to one or more thermal or energy KPIs. Server-side temperature dataset forms the basis of airflow and compliance-related indicators, while power and heat recovery dataset enable energy efficiency, reuse, and primary energy performance evaluation. A detailed mapping between queries and KPIs is provided in Table 3.2 (for detailed information on KPI calculation methodology, please refer to Deliverable D8.1).

Table 3.2: Mapping between SAT data queries and thermal and energy KPIs for the Empa pilot data centre.

Variable	**Description**	**KPI**
T_{in}	Server inlet air temperature	RTI, RHI, RCI$_{HI}$, RCI$_{LO}$, RI
T_{out}	Server outlet air temperature	RTI, LI
P_{IT}	Server IT power consumption	PUE, COP, ERF, PES, REF
T_{supply}	Air supply temperature	RTI, RHI, RI
T_{return}	Air return temperature	RTI, RHI, LI
P_C	Cooling power consumption	PUE, COP,ERF, PES, REF
P_{HRU}	HRU power consumption	PUE, COP,ERF, PES, REF
Q_{water}	HRU water flow rate	ERF, PES, REF
$T_{water,in}$	HRU water inlet (vapor) temperature	ERF, PES, REF

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

$T_{water,out}$	HRU water outlet (liquid) temperature	ERF, PES, REF

At the EMPA pilot, SAT was applied in full scope, including thermal and energy KPI computation, dashboard-based visualization, and automated report generation. The system successfully evaluated RTI, RHI, RCI, RI, LI, PUE, COP, ERF, and PES indicators based on synchronized IT and cooling system measurements. A consolidated standalone SAT report example is provided in the Annexes to demonstrate reporting completeness and traceability. It should be noted that, due to the selected test interval, certain measurements (e.g., HRU water flow rate) were recorded as zero. Consequently, KPIs dependent on these parameters may not reflect representative operational conditions for that specific interval. This confirms that SAT operates without functional limitation at the EMPA pilot.

3.2. Implementation of SAT at AAU

For the AAU pilot, the SAT implementation is tailored to air-cooled data centre operation and is designed to support both continuous monitoring and offline analytical workflows. Temperature measurements together with energy-related data are collected within the local monitoring infrastructure and made available to the Data Management System (DMS). These datasets can be accessed by SAT through REST-based APIs, fully aligned with the common SAT architecture and data model. When retrieved via the DMS, time-series data are stored and managed in VictoriaMetrics, enabling SAT to perform KPI calculations over configurable time intervals, including trend analysis and comparative assessment of air-cooling effectiveness.

In addition to the online, API-driven integration, the AAU pilot supports a flexible offline workflow. Air-cooled temperature measurements and associated energy data can be exported as structured CSV files and directly ingested into SAT using the “load data” functionality. This capability is particularly useful for analyzing historical datasets, experimental measurements or data collected outside the live monitoring environment. Regardless of the data acquisition path, SAT processes the inputs using the same KPI calculation and visualization pipeline, ensuring consistency of results. The computed thermal and energy-related KPIs are subsequently visualized through Grafana dashboards, preserving full interoperability with the VictoriaMetrics-based storage and enabling a unified performance assessment framework across the HW pilots.

At the AAU data centre, only cooling unit data are continuously available through the live data stream. All other measurements are limited to historical datasets and are

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

imported into SAT via the offline data ingestion workflow or DMS. Table 3.3 summarizes the raw variables accessed by SAT at the AAU pilot.

Table 3.3: Cooling-side data queries used in SAT for the AAU pilot.

Level	Part	Variable	Query
Data Centre	General	P_{DC}	power{pilot="AAU", datasource="DC"}
IT Equipments	Server	T_{in}	temperature{pilot="AAU", datasource="IT", datatype="TIN"}
		T_{out}	temperature{pilot="AAU", datasource="IT", datatype="TEX"}
		$T_{CPU,max}$	temperature{pilot="AAU", component="CPU", datasource="IT", datatype="TCPU(7)"}
		P_{IT}	power{pilot="AAU", component="CPU", datasource="IT", datatype="PS"}
Cooling Unit (Fan Coils)	Air-side	T_{supply}	temperature{pilot="AAU", datasource="C", datatype="TAS", id="1"}
		T_{return}	temperature{pilot="AAU", datasource="C", datatype="TAR", id="1"}
	Water Side	$T_{water,in}$	temperature{pilot="AAU", datasource="C", datatype="INT"}
		$T_{water,out}$	temperature{pilot="AAU", datasource="C", datatype="OUT"}

3.3. Implementation of SAT at PSNC

For the PSNC pilot, SAT was deployed in offline mode, enabling KPI computation and reporting without any dependency on a live monitoring API. All required inputs were provided as locally available time-series data and processed directly at runtime. In this configuration, SAT validates the availability of the standard thermal and energy inputs, including IT power, total facility power, CPU temperatures, and cooling loop temperatures. In addition, waste heat–related information was also supplied as part of the offline dataset, allowing energy and waste heat oriented KPIs to be evaluated directly without relying on external heat meters or online interfaces. This fully offline setup

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

ensures consistent and reproducible assessment of both energy efficiency and waste heat reuse performance for the PSNC pilot.

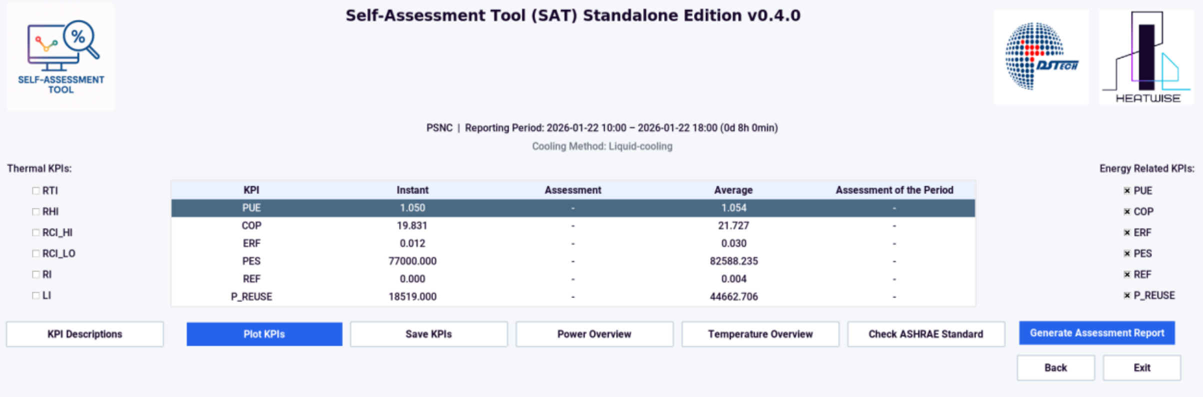

Self-Assessment Tool (SAT) Standalone Edition v0.4.0

PSNC | Reporting Period: 2026-01-22 10:00 – 2026-01-22 18:00 (0d 8h 0min)
Cooling Method: Liquid-cooling

Thermal KPIs:
☐ RTI
☐ RHI
☐ RCI_HI
☐ RCI_LO
☐ RI
☐ LI

KPI	Instant	Assessment	Average	Assessment of the Period
PUE	1.050	-	1.054	-
COP	19.831	-	21.727	-
ERF	0.012	-	0.030	-
PES	77000.000	-	82588.235	-
REF	0.000	-	0.004	-
P_REUSE	18519.000	-	44662.706	-

Energy Related KPIs:
☒ PUE
☒ COP
☒ ERF
☒ PES
☒ REF
☒ P_REUSE

Figure 3.1: Overview of the SAT dashboard for energy KPI evaluation at the PSNC pilot.

The energy KPIs in the PSNC pilot are evaluated exclusively from offline power and heat-reuse measurements, ensuring full consistency with the CSV-based deployment of SAT. The PUE is computed as the ratio of total facility power to total IT power, where the IT power is obtained by aggregating all server-level power measurements at each timestamp, and the facility power is obtained by aggregating the corresponding data centre–level power channels. This approach enables a time-resolved assessment of overall energy efficiency under varying operational conditions. The PSNC pilot relies primarily on liquid-cooled servers. The reported PUE value is calculated based solely on the electrical power associated with the liquid-cooled server environment, using the corresponding facility power boundary and aggregated IT power of these servers.

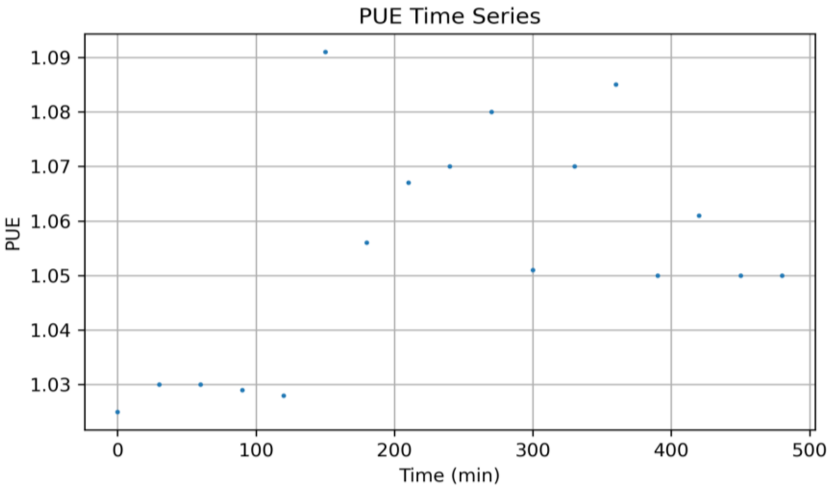

Figure 3.2: Temporal variation of PUE during the selected period at PSNC pilot.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

A distinctive feature of the PSNC offline dataset is that waste heat related information is directly embedded in the available measurements, allowing heat reuse indicators to be evaluated without requiring additional online instrumentation or external heat-meter integration. Based on these data, SAT evaluates a set of complementary energy reuse metrics at matching timestamps. The reused power is taken directly from the measured waste heat reuse signal, while the REF is computed as the ratio between the reusable energy and the instantaneous IT power demand. In addition, the PES is derived as the difference between the total facility power and the IT power, representing the non-IT energy overhead associated with data centre operation.

By combining facility-level power consumption, IT demand, and directly measured waste heat reuse within a single offline processing chain, the PSNC implementation of SAT enables a synchronized and reproducible evaluation of energy efficiency and waste heat reuse performance across the analysed period.

3.4. Implementation of SAT at Tofas

The implementation of the SAT at the TOFAS pilot is primarily focused on energy performance assessment, reflecting both the monitoring priorities of the site and the characteristics of the available data infrastructure. At TOFAS, SAT operates independently of live monitoring platforms while maintaining full methodological consistency with the broader HEATWISE framework.

At the TOFAŞ data centre, energy-related measurements are collected locally from electrical monitoring systems and prepared specifically for SAT analysis. Instead of continuous data streaming via monitoring APIs, the pilot follows an offline, file-based data ingestion strategy, enabling controlled data validation, traceability, and reproducibility. Two core datasets are provided to SAT: (i) total data centre electrical power consumption ($P_{cooling}$, $P_{lightining}$ and P_{IT}) and (ii) IT equipment power consumption, capturing the electrical load directly associated with computing activities.

Based on the available datasets, the KPI scope at the TOFAŞ pilot is intentionally limited to energy-based indicators. SAT computes: PUE, expressing the ratio between total facility power consumption and IT power demand. COP, derived within the SAT framework as an aggregated indicator of energy conversion efficiency. The restriction of KPIs to PUE and COP ensures that all reported metrics are fully supported by measured data, avoiding the use of assumptions or surrogate variables for thermal or airflow-related indicators that are not available at this site. This approach enhances the robustness, transparency, and interpretability of the assessment results.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

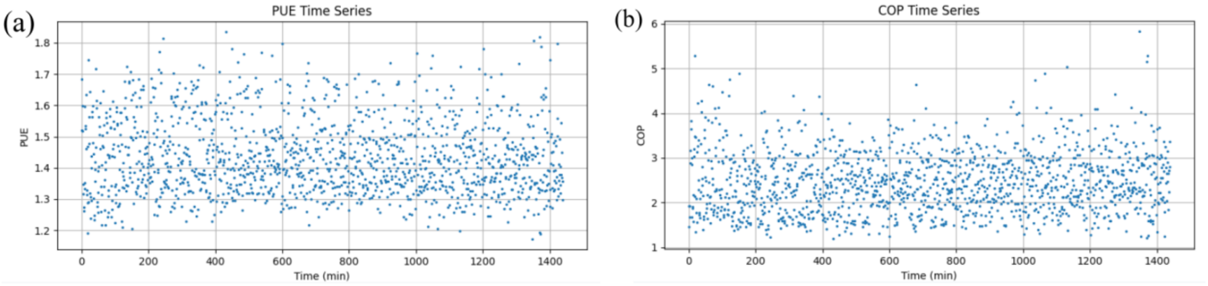

Figure 3.3: Temporal variation of (a) PUE and (b) COP during the selected period at TOFAS pilot.

Despite the use of an offline, CSV-based ingestion approach, the TOFAS implementation is fully integrated into the common SAT processing, visualization and reporting workflow. Once computed, KPI results are stored, visualized, and reported using the same internal structures and presentation logic applied across other HEATWISE pilots. As a result, time-series visualizations, summary statistics, and reporting outputs generated for TOFAS follow the same layout and conventions used throughout the project, enabling consistent comparison of energy performance trends across pilots, even when underlying data acquisition strategies differ.

The TOFAS pilot illustrates SAT's adaptability to heterogeneous data availability scenarios. While some HEATWISE pilots rely on real-time monitoring systems and integrated data management platforms, TOFAS operates under a constrained but well-defined data environment. SAT addresses this diversity by decoupling data ingestion mechanisms from KPI computation and reporting logic. This separation ensures that differences in data access do not affect KPI definitions, calculation procedures, or interpretation principles, thereby preserving methodological consistency across all pilots.

The scope of the SAT implementation at TOFAS is inherently influenced by the availability and resolution of monitored data. Although the provided energy datasets support a robust evaluation of overall electrical efficiency, several limitations should be acknowledged. First, the absence of detailed thermal and airflow measurements, such as rack inlet/outlet temperatures, cooling unit operating conditions, or air and water flow rates, restricts the assessment to energy-based KPIs. Advanced thermal performance indicators and spatial efficiency metrics cannot be evaluated within the current data context. Second, the reliance on offline CSV-based data ingestion limits the assessment to retrospective analyses over predefined time windows. While this approach improves deployment simplicity and data reliability, it prevents real-time monitoring and limits the ability to capture short-term operational dynamics or rapidly assess the impact of control actions. Third, the lack of component-level power breakdowns, particularly for cooling subsystems and auxiliary equipment, means that COP values represent aggregated system-level indicators rather than device-specific efficiencies. As a result, detailed

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

attribution of efficiency changes to individual subsystems is not possible without additional instrumentation.

Finally, while SAT enforces uniform KPI definitions and reporting formats across all HEATWISE pilots, differences in data richness should be considered when performing cross-pilot comparisons. For TOFAS, comparative analyses are most meaningful at the level of overall energy efficiency trends, rather than detailed thermodynamic or spatial performance.

Despite these limitations, the TOFAŞ pilot demonstrates SAT's capability to deliver meaningful and comparable insights under minimal integration conditions. The identified constraints also highlight clear pathways for future enhancements, such as the integration of thermal measurements or subsystem-level power monitoring, should the site's monitoring infrastructure be expanded.

3.5. Implementation of SAT at RISE

In the RISE pilot, the SAT implementation is tightly integrated with the existing hybrid cooling monitoring and data acquisition infrastructure, enabling both high-resolution analysis and long-term performance assessment. Measurements from the hybrid cooling system (including air-side supply and return temperatures, cooling unit power consumption, heat recovery unit (HRU) operation, and liquid-side thermal parameters) are continuously collected alongside detailed IT operational data from the server infrastructure.

All monitoring data are transferred to the Data Management System (DMS) through REST-based interfaces and exposed via standardized query endpoints. SAT accesses these datasets using unified query definitions, allowing consistent retrieval of thermal, electrical, and operational parameters across different subsystems. This approach ensures portability of the assessment framework while preserving site-specific measurement semantics.

SAT supports retrospective KPI evaluation over user-defined time windows, enabling detailed post-analysis of hybrid cooling system performance. In parallel, near-real-time data streams can be accessed for live performance monitoring and diagnostic analysis. At the RISE pilot data centre, this dual capability allows continuous operational insight while also supporting in-depth evaluation of cooling efficiency, thermal management quality, and IT–cooling interaction under varying operating conditions.

Table 3.4 summarizes the raw variables accessed by SAT at the RISE pilot, together with their associated system components and query definitions. These variables form the basis for subsequent KPI calculations within the SAT framework.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

Table 3.4: Server and cooling data queries in SAT for the RISE pilot.

Component	Part	Variable	Query
IT Equipment	Server	T_{in}	temperature{pilot="RI", datasource="IT", datatype="TIN"}
		T_{out}	temperature{pilot="RI", datasource="IT", datatype="TEX"}
		T_{CPU}	temperature{pilot="RI", datasource="IT", datatype="TCPU(1)"}
		P_{IT}	power{pilot="RI", datasource="IT", datatype="PS"}
Cooling Unit	Air-side	T_{supply}	temperature{pilot="RI", datasource="C", datatype="TAS"}
		T_{return}	temperature{pilot="RI", datasource="C", datatype="TAR"}
		P_{C}	power{pilot="RI", datasource="C", datatype="PCO"}
		P_{fan}	power{datasource="DC", datatype="BxMfanP"}
	Water-side	P_{HRU}	power{datasource="HR", datatype="HRUeP"}
		Q_{water}	flow{datasource="HR", datatype="Flw", id="1"}

		$T_{water,in}$	temperature{datasource="HR", datatype="ST1reT"}
		$T_{water,out}$	temperature{datasource="HR", datatype="ST5suT"}
		T_{liquid}	temperature{datasource="HR", datatype="LT"}
		T_{vapor}	temperature{datasource="HR", datatype="VT"}

Each data query used by SAT directly contributes to one or more thermal or energy KPIs. Server-side temperature dataset form the basis of airflow and compliance-related indicators, while power and heat recovery dataset enable energy efficiency, reuse, and primary energy performance evaluation. A detailed mapping between queries and KPIs is provided in Table 3.5.

Table 3.5: Mapping between SAT data queries and KPIs for the RISE pilot.

Variable	Description	KPI
T_{in}	Server inlet air temperature	RTI, RHI, RCI_{HI}, RCI_{LO}, RI
T_{out}	Server outlet air temperature	RTI, LI
P_{IT}	Server IT power consumption	PUE, COP, ERF, PES, REF
T_{supply}	Air supply temperature	RTI, RHI, RI
T_{return}	Air return temperature	RTI, RHI, LI
P_{C}	Cooling power consumption	PUE, COP,ERF, PES, REF
P_{HRU}	HRU power consumption	PUE, COP,ERF, PES, REF
Q_{water}	HRU water flow rate	ERF, PES, REF
$T_{water,in}$	HRU water inlet (vapor) temperature	ERF, PES, REF

$T_{water,out}$	HRU water outlet (liquid) temperature	ERF, PES, REF

At the RISE pilot, certain cooling-side metrics (e.g. BxMfanP) are not available for all time intervals; therefore, SAT treats these variables as optional and automatically excludes them from the calculations when data are missing. For IT power assessment, only power supply measurements are used to prevent double counting of electrical consumption. All variables are synchronized at the timestamp level before being passed to the KPI computation layer, ensuring consistent and temporally aligned performance evaluation.

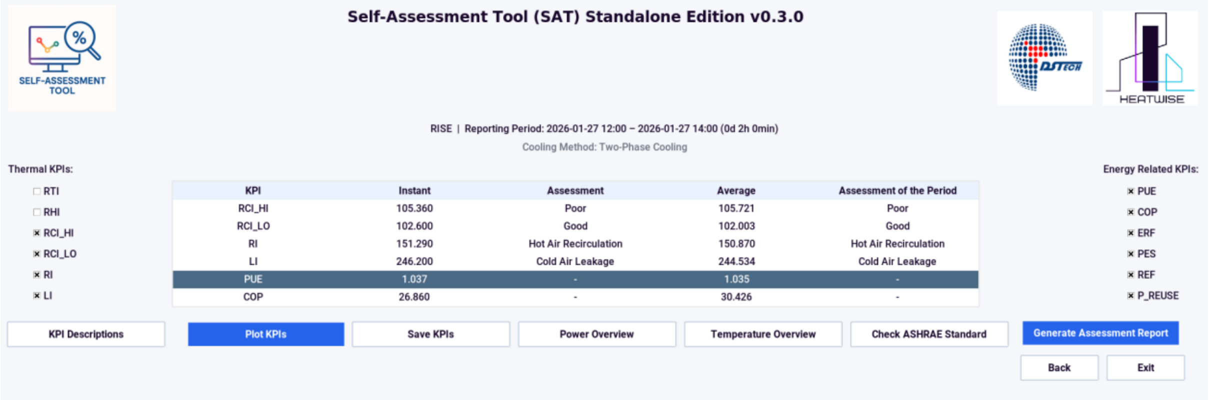

Self-Assessment Tool (SAT) Standalone Edition v0.3.0

RISE | Reporting Period: 2026-01-27 12:00 – 2026-01-27 14:00 (0d 2h 0min)

Cooling Method: Two-Phase Cooling

Thermal KPIs:
☐ RTI
☐ RHI
☒ RCI_HI
☒ RCI_LO
☒ RI
☒ LI

KPI	Instant	Assessment	Average	Assessment of the Period
RCI_HI	105.360	Poor	105.721	Poor
RCI_LO	102.600	Good	102.003	Good
RI	151.290	Hot Air Recirculation	150.870	Hot Air Recirculation
LI	246.200	Cold Air Leakage	244.534	Cold Air Leakage
PUE	1.037	-	1.035	-
COP	26.860	-	30.426	-

Energy Related KPIs:
☒ PUE
☒ COP
☒ ERF
☒ PES
☒ REF
☒ P_REUSE

Figure 3.5: Overview of the SAT dashboard for thermal and energy KPI evaluation at the RISE pilot.

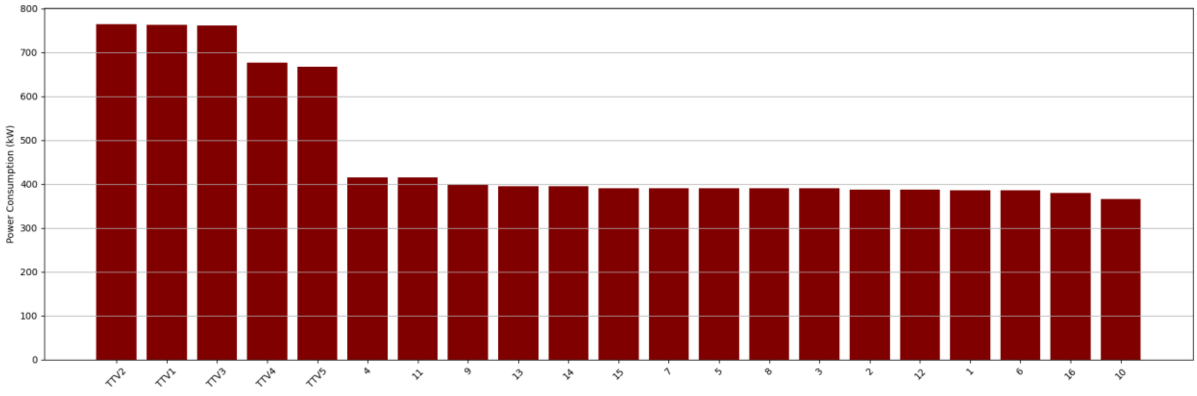

Figure 3.6: Power consumption of servers and TTVs at RISE pilot.

As illustrated in Fig. 3.7, the test conducted at the RISE data centre between 12:00 and 14:00 on 27 January 2026 indicates that the PUE values remain within the target range (PUE<1.05) defined by the EED, demonstrating compliance with the intended energy efficiency objectives during the evaluated operating period.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

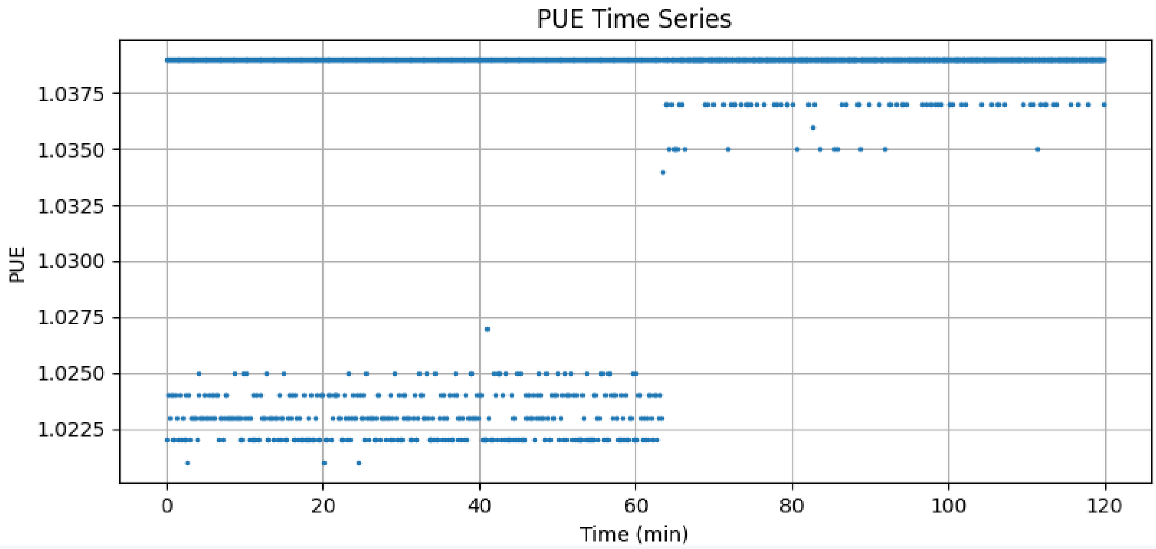

Figure 3.7: Temporal variation of PUE during the selected period at RISE pilot.

To provide a comprehensive overview of KPI coverage and reporting consistency across all pilot sites, a consolidated results table is presented below. Table 3.6 summarizes, for each pilot, the selected reporting interval, computed KPI values, and their corresponding performance classification. This consolidated presentation enhances clarity, traceability and cross-pilot comparability.

Due to data availability constraints, thermal KPIs were not computed for the TOFAŞ pilot. Specifically, inlet and outlet temperature measurements at server level were not available for the air-cooled infrastructure, which prevented the calculation of temperature-based thermal performance indicators (RTI, RHI, RCI, RI, LI). For the AAU pilot, thermal KPIs and core energy efficiency indicators (PUE and COP) were successfully computed based on the available monitoring data. However, since no waste heat recovery system is implemented at this site, reuse-related energy KPIs (such as ERF, PES, and recovered thermal power) were not evaluated. For the EMPA, RISE, and PSNC pilots, the complete set of thermal and energy-related KPIs defined within the SAT framework was computed without limitation.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

Table 3.6: Consolidated KPI results and reporting periods across pilots.

Pilot	Reporting Period	KPI	Value	Status
EMPA	01.01.2026 08.00 01.01.2026 11.00	RCI_{HI}	100.32	Good
		RCI_{LO}	107.36	Poor
		RHI	0.19	Unacceptable
		RTI	1.89	Bypass
		RI	112.27	Hot Air Recirculation
		LI	80.83	Cold Air Leakage
		PUE	1.73	Poor
		COP	2.70	-
		ERF	0.00	-
		PES	0.00	-
		REF	0.00	-
AAU	12.09.2025 10:00 12.09.2025 14:00	RCI_{HI}	122.69	Poor
		RCI_{LO}	99.68	Good
		RHI	1.33	Acceptable
		RTI	0.35	Bypass
		RI	101.66	Hot Air Recirculation
		LI	88.65	Cold Air Leakage
		PUE	4.42	Poor
		COP	0.30	-
PSNC	22.01.2026 10:00 22.01:2026 14:00	RCI_{HI}	100.11	Good
		RCI_{LO}	100.04	Good
		RHI	1.00	Acceptable

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

		RTI	0.04	Bypass
		RI	100.00	Acceptable
		LI	100.95	Acceptable
		PUE	1.05	Good
		COP	19.83	-
RISE	27.01.2026 10:00 27.01.2026 12:00	RCI_{HI}	105.72	Poor
		RCI_{LO}	102.00	Good
		RHI	3.53	Acceptable
		RTI	1.80	Bypass
		RI	150.87	Hot Air Recirculation
		LI	244.53	Cold Air Leakage
		PUE	1.03	Very Good
		COP	30.43	-
		ERF	18.34	Acceptable
		PES	0.51	Poor
		P_{REUSE}	3040	-
TOFAS	16.12.2025 10:00 16.12.2025 18:00	PUE	1.41	Poor
		COP	2.44	-

3.6. Reproreproducibility Guidelines of SAT

To support reproducibility and independent re-implementation, this subsection describes the minimal data requirements and the computational workflow used in the SAT framework. In addition, a synthetic example dataset together with reference outputs is made publicly available to illustrate the complete KPI computation process.

The SAT requires a minimal time-series dataset consisting of synchronized measurements. At minimum, the following columns are required: time stamp, inlet and outlet temperature, and power consumption of servers, power consumption of data

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

centre, supply and return temperatures of cooling units. These variables are sufficient to compute the thermal and energy KPIs reported in this study. A synthetic dataset following this schema has been published to enable independent validation and testing. All input time series are aligned to a common temporal axis prior to KPI computation. If raw measurements are available at different sampling frequencies, the data are resampled to a user-defined time step using time-based aggregation. Temperature values are averaged over each interval, while power measurements may be either averaged or retained as instantaneous values depending on the analysis objective. This step ensures temporal consistency across heterogeneous data sources. KPI calculation follows a fixed and reproducible order. First, synchronized inlet, outlet, supply, and return temperature series are processed. Second, IT and facility-level power measurements are incorporated. Thermal KPIs (RTI, RHI, RCI, RI, LI) are computed before energy-related KPIs (PUE, COP), ensuring that all required intermediate quantities are available. This ordered workflow is identical for both real-time and historical datasets. Time steps with missing or invalid measurements are excluded from KPI computation to avoid introducing artificial bias. No interpolation is applied at the KPI level. This conservative approach ensures that reported KPIs are based solely on valid and physically meaningful measurements. SAT computes both instantaneous KPI values at each time step and aggregated KPI statistics over a selected analysis window. Instantaneous KPIs are used to identify transient effects and short-term deviations, while averaged KPIs support comparative assessment across operational periods. Both representations are demonstrated using the synthetic dataset and corresponding KPI visualizations.

To facilitate reproducibility, the synthetic dataset, reference configuration files and example KPI outputs (including time-series plots and power distribution charts) are publicly available via a GitHub repository (see Data Availability statement). To illustrate reproducibility rather than to present additional performance results, representative KPI visualizations generated from the synthetic dataset are shown in Fig. 3.8.

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

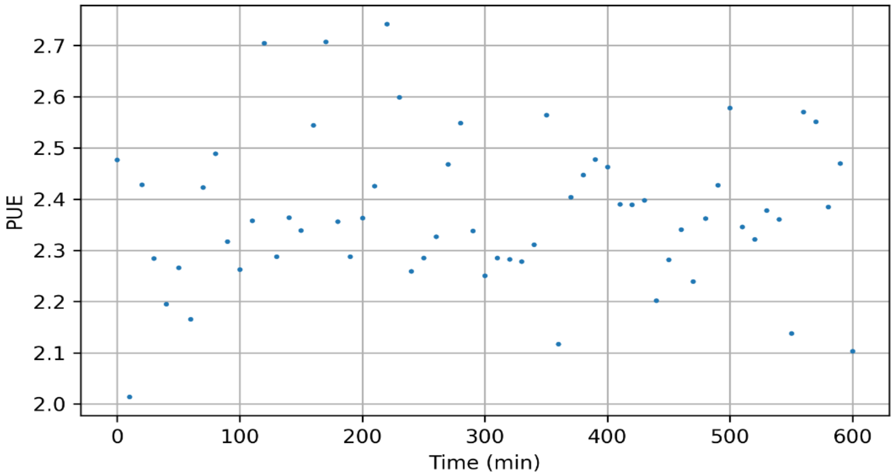

Figure 3.8: Example KPI outputs generated from a dataset to demonstrate reproducibility of the SAT workflow.

3.7. Road map for the implementation of SAT at external pilots

The implementation of SAT at external pilot sites follows a structured and progressive roadmap designed to accommodate different levels of data availability, IT security constraints, and operational maturity. This roadmap enables pilot sites to gradually integrate SAT into their existing monitoring ecosystems while maintaining data confidentiality and minimizing deployment effort. H1 Systems (H1S) supports this process by acting as the primary coordination partner between external pilots and the SAT development team (DS Tech), ensuring alignment with HEATWISE project objectives. Fig. 3.9 illustrates the workflow for the implementation at external pilot sites.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

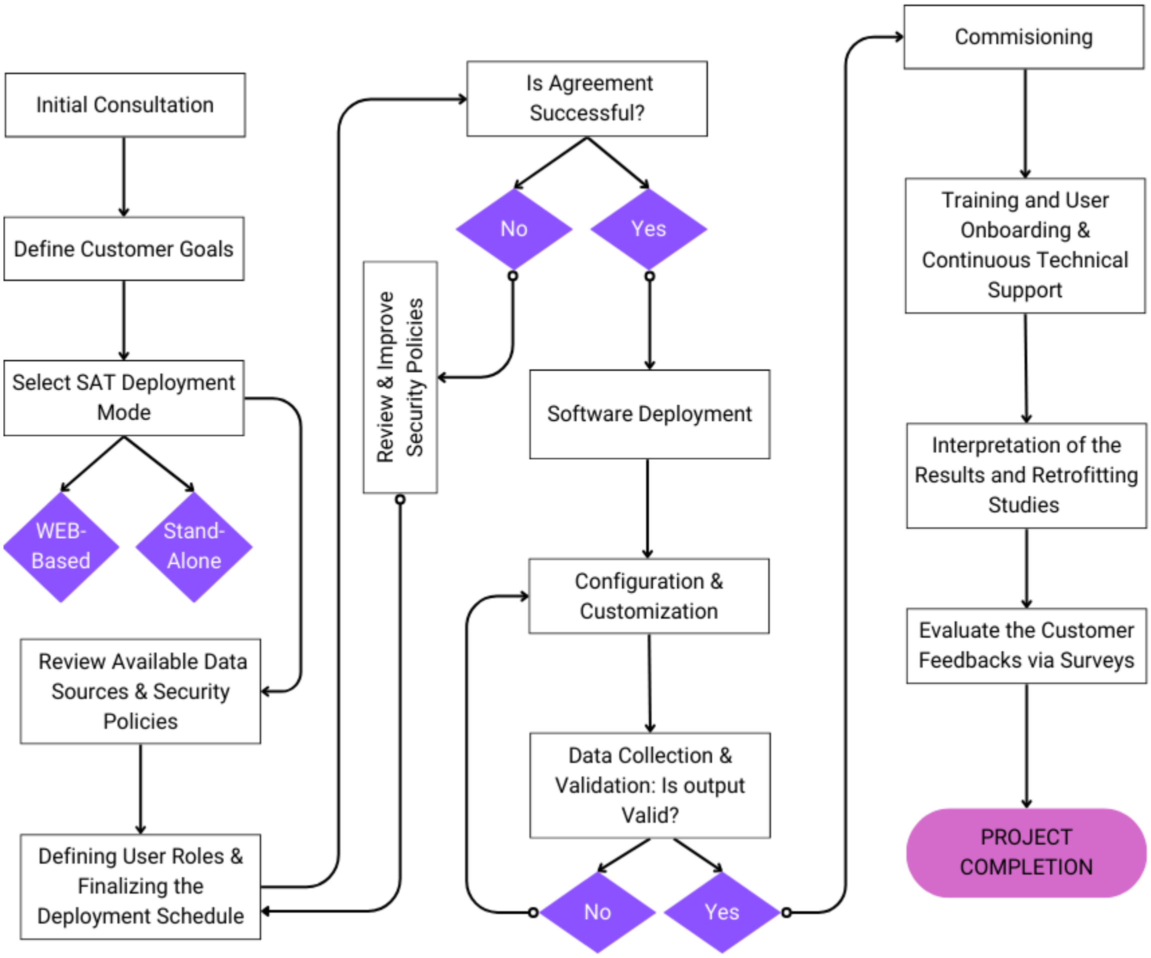

Figure 3.9: Implementation of SAT at external pilot sites.

The first phase of the roadmap focuses on initial engagement and feasibility assessment. During this phase, the operational objectives of the external pilot are identified, and available data sources are reviewed to determine whether live interfaces or offline datasets will be used. H1S leads the interaction with pilot operators at this stage, facilitating requirement definition and supporting the selection of a suitable deployment model.

In the second phase, deployment planning is carried out with external pilots adopting the standalone desktop version of SAT, in line with site-specific IT and security policies. During this phase, data privacy constraints (NDA), access boundaries, and local execution requirements are clarified to ensure that all data processing remains within the pilot's infrastructure. H1S supports coordination between pilot IT teams and DS Tech engineers to ensure that deployment planning remains consistent with site-specific constraints.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

The third phase involves data integration and configuration. Live data streams are connected via REST or MQTT interfaces where available, while historical datasets such as IT and cooling unit-based measurements are uploaded using structured templates. Dataset structure and semantics are jointly validated to ensure compatibility with the SAT KPI computation module.

Once data ingestion is operational, the fourth phase addresses data validation and initial assessment. Imported or streamed data are checked for completeness, temporal consistency, and physical plausibility, followed by a first KPI evaluation cycle. This phase allows early verification of results and alignment of expectations between the pilot site and the project team.

The fifth phase focuses on user onboarding and operational assessment. Training activities support pilot users in understanding the SAT interface, KPI interpretation, and report generation workflow. H1S contributes to the interpretation of assessment outcomes from a system-level perspective, helping translate technical KPIs into operational insights relevant for pilot stakeholders.

Following onboarding, SAT enters regular operation, supporting performance benchmarking, compliance assessment, and automated reporting. Feedback from external pilots is incorporated to refine configurations and workflows, ensuring that SAT remains adaptable and scalable for future deployments.

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

4. Conclusion

This deliverable presented the development, architecture and multi-pilot implementation of the SAT, a unified self-assessment framework designed for comprehensive thermal and energy performance evaluation of data centres. The proposed framework addresses key challenges identified in current regulatory and operational practices, namely the fragmentation of data sources, the complexity of KPI computation, and the lack of continuous, reproducible performance assessment aligned with emerging ASHRAE and EED regulations. By integrating IT, cooling, and energy data into a single analytical workflow, SAT enables systematic and transparent evaluation of data centre performance using standardized thermal and energy KPIs.

The modular architecture of SAT ensures methodological consistency across heterogeneous pilot environments, while allowing flexibility in data ingestion mechanisms. The framework successfully supports both real-time monitoring via data management systems and fully offline, file-based workflows, as demonstrated across the HEATWISE pilot sites. This dual capability enables SAT to operate under diverse infrastructural constraints without compromising the comparability of results. The integration of interactive dashboards with automated PDF reporting further bridges the gap between exploratory analysis and formal documentation, supporting both operational decision-making and regulatory reporting requirements.

In addition, the modular architecture of SAT enables the future integration of additional sustainability, decarbonization, and optimization-oriented KPIs without requiring modification of the core computational system. New indicators can be incorporated as extensions of the KPI computation layer while preserving the existing data acquisition, synchronization, and reporting workflow. This design choice reinforces the forward-looking and scalable nature of the framework, ensuring that SAT can evolve in parallel with emerging regulatory requirements, sustainability metrics, and advanced optimization methodologies.

Pilot implementations at EMPA, AAU, PSNC, TOFAŞ, and RISE demonstrated the robustness and adaptability of the framework across air-cooled, liquid-cooled, and hybrid-cooled data centres with varying levels of data availability. In particular, the PSNC and TOFAŞ pilots highlighted the capability of SAT to deliver meaningful energy efficiency and waste heat reuse insights using offline datasets, while the EMPA and RISE deployments showcased high-resolution, time-resolved performance analysis enabled by live monitoring infrastructures. These case studies confirm that SAT provides a scalable and pilot-agnostic solution for harmonized KPI-driven assessment.

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

Overall, the SAT framework represents a practical step towards continuous, KPI-based self-assessment of data centres in line with the Energy Efficiency Directive and industry best practices. Beyond compliance, the framework supports proactive identification of inefficiencies, evaluation of optimization measures, and long-term performance tracking. Future extensions may include tighter integration with self-optimization algorithms, enhanced subsystem-level diagnostics, and expanded support for emerging sustainability indicators. As such, SAT provides a solid foundation for both current pilot evaluations and wider deployment in operational data centres seeking transparent, reproducible, and regulation-ready performance assessment.

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

References

1. Masanet, E, Shehabi, A, Lei, N, Smith, S, Koomey, J (2020) Recalibrating global data center energy-use estimates. Science, 367(6481) 984-986. https://doi.org/10.1126/science.aba3758
2. Jones, N (2018) How to stop data centres from gobbling up the world's electricity. Nature 561(7722) 163-166. https://doi.org/10.1038/d41586-018-06610-y
3. Hintemann, R, Hinterholzer, S (2019) Energy consumption of data centers worldwide. Business, Computer Science.
4. Shehabi A, Smith S, Sartor D, Brown R, Herrlin M, Koomey J, Masanet E, Horner R, Azevedo I, Lintner W (2016) United states data center energy usage report. Lawrence Berkeley National Laboratory, Berkeley, California. https://doi.org/10.2172/1372902
5. Avgerinou M., Bertoldi P, Castellazzi L (2017) Trends in data centre energy consumption under the European code of conduct for data centre energy efficiency. Energies 10(10) 1470. https://doi.org/10.3390/en10101470
6. European Commission (2023). Energy Efficiency Directive (EED) Document 32023L1791. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ%3AJOL_2023_231_R_0001&qid=1695186598766
7. European Commission (2024). Commission Delegated Regulation 2024/1364. https://eur-lex.europa.eu/eli/reg_del/2024/1364/oj/eng
8. Acton M, Booth J, Paci D (2025) Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency, Publications Office of the European Union, Luxembourg. https://data.europa.eu/doi/10.2760/9449356
9. European Commission (2020). A European Green Deal. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en
10. European Commission (2025) Assessment of next steps to promote the energy performance and sustainability of data centres. Publications Office of the European Union MJ0125123ENN. https://doi.org/10.2833/0828045
11. European Commission (2025) Assessment of the energy performance and sustainability of data centres in EU. https://doi.org/10.2833/3168794
12. Koomey J (2011) Growth in data center electricity use 2005 to 2010. A report by Analytical Press, completed at the request of The New York Times 9 161.
13. Kliazovich D, Bouvry, P, Khan SU (2012) GreenCloud: a packet-level simulator of energy-aware cloud computing data centers. The Journal of Supercomputing, 62(3) 1263-1283. https://doi.org/10.1007/s11227-010-0504-1

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

14. Patel CD, Sharma R, Bash CE, Beitelmal A (2002) Thermal considerations in cooling large scale high compute density data centers. In ITherm 2002. Eighth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (Cat. No. 02CH37258) (pp. 767-776). IEEE. https://doi.org/10.1109/ITHERM.2002.1012532
15. Cho J, Yang J, Park W (2014) Evaluation of air distribution system's airflow performance for cooling energy savings in high-density data centers. Energy and buildings 68 270-279. https://doi.org/10.1016/j.enbuild.2013.09.013
16. Moore J, Chase, JS, Ranganathan P (2006) Weatherman: Automated, online and predictive thermal mapping and management for data centers. In 2006 IEEE international conference on Autonomic Computing (pp. 155-164) IEEE. https://doi.org/10.1109/ICAC.2006.1662394
17. Liu J, Terzis A. (2012) Sensing data centres for energy efficiency. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 370(1958) 136-157. https://doi.org/10.1098/rsta.2011.0245
18. Shor J (2020) Compact thermal sensors for dense CPU thermal monitoring and regulation: A review. IEEE Sensors Journal 21(11) 12774-12788. https://doi.org/10.1109/JSEN.2020.3029852
19. Santos D, Mataloto B, Ferreira JC (2019) Data center environment monitoring system. In Proceedings of the 2019 4th International Conference on Cloud Computing and Internet of Things (pp. 75-81). https://doi.org/10.1145/3361821.3361824
20. Buyya R, Ilager S, Arroba P (2024) Energy-efficiency and sustainability in new generation cloud computing: a vision and directions for integrated management of data centre resources and workloads. Software: Practice and Experience 54(1) 24-38. https://doi.org/10.1002/spe.3248
21. Matko V, Brezovec, B, Milanovič M. (2019). Intelligent monitoring of data center physical infrastructure. Applied Sciences 9(23) 4998. https://doi.org/10.3390/app9234998
22. Turnbull, J (2018) Monitoring with Prometheus. Turnbull Press.
23. Sabharwal, N, Pandey, P (2020) Working with prometheus query language (promql). In Monitoring Microservices and Containerized Applications: Deployment, Configuration, and Best Practices for Prometheus and Alert Manager (pp. 141-167) Berkeley, CA: Apress. https://doi.org/10.1007/978-1-4842-6216-0_5

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

24. Kovac D, Hosek J, Masek, P, Stusek M (2015). Keeping eyes on your home: Open-source network monitoring center for mobile devices. In 2015 38th International Conference on Telecommunications and Signal Processing (TSP) (pp. 612-616) IEEE. https://doi.org/10.1109/TSP.2015.7296336
25. Flores D, Adell C, Vanderaa J (2024) Modern Network Observability: A hands-on approach using open source tools such as Telegraf, Prometheus, and Grafana. Packt Publishing Ltd.
26. Elradi MD (2025) Prometheus & Grafana: A Metrics-focused Monitoring Stack. Journal of Computer Allied Intelligence (JCAI, ISSN: 2584-2676) 3(3) 28-39. https://doi.org/10.69996/jcai.2025015
27. Yu C, Lumezanu C, Sharma A, Xu Q, Jiang G, Madhyastha HV (2015) Software-defined latency monitoring in data center networks. In International Conference on Passive and Active Network Measurement (pp. 360-372). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-15509-8_27
28. Nadjahi, C, Louahlia, H, Lemasson S. (2018) A review of thermal management and innovative cooling strategies for data center. Sustainable Computing: Informatics and Systems 19 14-28. https://doi.org/10.1016/j.suscom.2018.05.002
29. Wang L, Khan SU (2013) Review of performance metrics for green data centers: a taxonomy study. The journal of supercomputing, 63(3) 639-656. https://doi.org/10.1007/s11227-011-0704-3
30. Anthony B, Abdul Majid M, Romli, A (2018) A collaborative agent based green IS practice assessment tool for environmental sustainability attainment in enterprise data centers. Journal of Enterprise Information Management, 31(5), 771-795. https://doi.org/10.1108/JEIM-10-2017-0147
31. Hodgkinson IR, Jackson TW, Al-Busaidi KA (2025) A conceptual tool to self-assess digital decarbonization readiness. Journal of Business Strategy 46(3-4) 90-114. https://doi.org/10.1108/JBS-04-2025-0078
32. Muladi N, Surendro K (2014) The readiness self-assessment model for green IT implementation in organizations. In 2014 International Conference of Advanced Informatics: Concept, Theory and Application (ICAICTA) (pp. 146-151) IEEE. https://doi.org/10.1109/ICAICTA.2014.7005931
33. Banijamali P, Fatima I, Lago P, Heitlager I (2024) Unveiling key performance indicators for the energy efficiency of cloud data storage. In 2024 IEEE 21st International Conference on Software Architecture Companion (ICSA-C) (pp. 222-229) IEEE. https://doi.org/10.1109/ICSA-C63560.2024.00047

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

34. Curry E, Conway G, Donnellan B, Sheridan, C, Ellis K (2012) Measuring energy efficiency practices in mature data center: A maturity model approach. In Computer and Information Sciences III: 27th International Symposium on Computer and Information Sciences (pp. 51-61). London: Springer London. https://doi.org/10.1007/978-1-4471-4594-3_6

35. European Parliament and Council, Directive (EU) 2023/1791 on Energy Efficiency (Recast), Official Journal of the European Union, 2023.

36. European Commission, Data Centres in the EU: Energy Consumption, Energy Efficiency and Policy Options, JRC Science for Policy Report, 2020.

37. European Commission, EU Code of Conduct for Data Centre Energy Efficiency – Best Practices, Version 9.0, 2023.

38. Uptime Institute, Thermal Guidelines for Data Processing Environments, White Paper, 2022.

39. Uptime Institute, Global Data Center Survey, Annual Report, 2023.

40. Uptime Institute, Beyond PUE: Why Efficiency Metrics Must Evolve, Uptime Intelligence Research, 2021.

41. Python Software Foundation (PSF) Python Language Reference, version 3.x. https://docs.python.org/3/

42. Shipman JW (2013) Tkinter 8.5 reference: a GUI for Python. New Mexico Tech Computer Center, 54, 356-359.

43. Costa B, Pires PF, Delicato FC, Merson P (2014) Evaluating a Representational State Transfer (REST) architecture: What is the impact of REST in my architecture?. In 2014 IEEE/IFIP Conference on Software Architecture pp. 105-114). https://doi.org/10.1109/WICSA.2014.29

44. Grinberg M (2018) Flask web development. O'Reilly Media, Inc.

45. Gaspar D, Stouffer J (2018) Mastering Flask Web Development: Build Enterprise-grade, Scalable Python Web Applications. Packt Publishing Ltd.

46. Duckett J (2011) HTML and CSS: Design and Build Websites, John Willey & Sons. Inc. Indianappolis.

47. Robbins J (2012) Learning web design: A beginner's guide to HTML, CSS, JavaScript, and web graphics. "O'Reilly Media, Inc. https://doi.org/10.5860/choice.50-6805

48. McKinney, W. (2010). Data structures for statistical computing in Python. scipy 445(1) 51-56.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

49. Harris CR, Millman KJ, Van Der Walt SJ, Gommers R, Virtanen, P, Cournapeau D, Wieser E, Taylor J, Berg S, Brett M, Haldane A, Del Rio JF, Peterson P, Gerar-Marchant P, Sheppard K, Reddy T, Weckesser W, Abbasi H, Gohlke C, Oliphant TE (2020) Array programming with NumPy. Nature 585(7825) 357-362. https://doi.org/10.1038/s41586-020-2649-2
50. Hunter JD (2007) Matplotlib: A 2D graphics environment. Computing in Science & Engineering 9(03), 90-95. https://doi.org/10.1109/MCSE.2007.55
51. Sial AH, Rashdi SYS, Khan AH (2021) Comparative analysis of data visualization libraries Matplotlib and Seaborn in Python. International Journal, 10(1) 277-281. https://doi.org/10.30534/ijatcse/2021/391012021
52. ReportLab Documentation (2023) ReportLab User Guide. https://www.reportlab.com/documentation/
53. Bootstrap Team (2023) Bootstrap Documentation v5.3.8. https://getbootstrap.com/
54. Plotly Technologies Inc. (2015) Plotly: Collaborative Data Science. https://plot.ly/
55. Sievert C (2020) Interactive web-based data visualization with R, plotly, and shiny. Chapman and Hall/CRC. https://doi.org/10.1201/9780429447273

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

List of figures

Figure 2.1: Data management system and data pipelines from pilots. 13

Figure 2.2: End-to-end workflow of the SAT from data acquisition to automated assessment reporting. ...14

Figure 2.3: Set of KPIs computable within the SAT framework. 17

Figure 2.4: End-to-end workflow of SAT from data collection to reporting. 19

Figure 2.5: Data types used in SAT: (a) real-time data and (b) externally provided dataset. 20

Figure 2.6: Components of SAT in standalone and web-based versions. 21

Figure 2.6: Selection of (a) reporting period and (b) historical duration in the setup window. 22

Figure 2.7:. The dashboard of the SAT. 23

Figure 2.8: Performance evaluation according to the KPIs. 23

Figure 2.9: Time series of thermal KPIs calculated by SAT. 24

Figure 2.10: Rack-level power and thermal maps. 25

Figure 2.11: Air-cooled rack layout at the EMPA pilot. 26

Figure 2.12: Power consumption of servers at the EMPA pilot. 26

Figure 2.13: Thermal characteristics of servers at the EMPA pilot. 27

Figure 2.14: ASHRAE compliance assessment for air and hybrid-cooled data centres. 28

Figure 3.1: Overview of the SAT dashboard for energy KPI evaluation at the PSNC pilot. 35

Figure 3.2: Temporal variation of PUE during the selected period at PSNC pilot. 35

Figure 3.3: Temporal variation of (a) PUE and (b) COP during the selected period at TOFAS pilot. 37

Figure 3.5: Overview of the SAT dashboard for thermal and energy KPI evaluation at the RISE pilot. 41

Figure 3.6: Power consumption of servers and TTVs at RISE pilot. 41

Figure 3.7: Temporal variation of PUE during the selected period at RISE pilot. 42

Figure 3.8: Example KPI outputs generated from a dataset to demonstrate reproducibility of the SAT workflow. 46

Figure 3.9: Implementation of SAT at external pilot sites. 47

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

List of tables

Table 1.1: Comparison of assessment capabilities across operator-driven approaches, dashboards, DCIM platforms and SAT. 11

Table 3.1: Server-side and cooling-side data queries used in SAT for the Empa pilot data centre. 31

Table 3.2: Mapping between SAT data queries and thermal and energy KPIs for the Empa pilot data centre. 32

Table 3.3: Cooling-side data queries used in SAT for the AAU pilot. 34

Table 3.4: Server and cooling data queries in SAT for the RISE pilot. 39

Table 3.5: Mapping between SAT data queries and KPIs for the RISE pilot. 40

Table 3.6: Consolidated KPI results and reporting periods across pilots. 43

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

Annexes

Annex 1. SAT report for EMPA Data Centre.

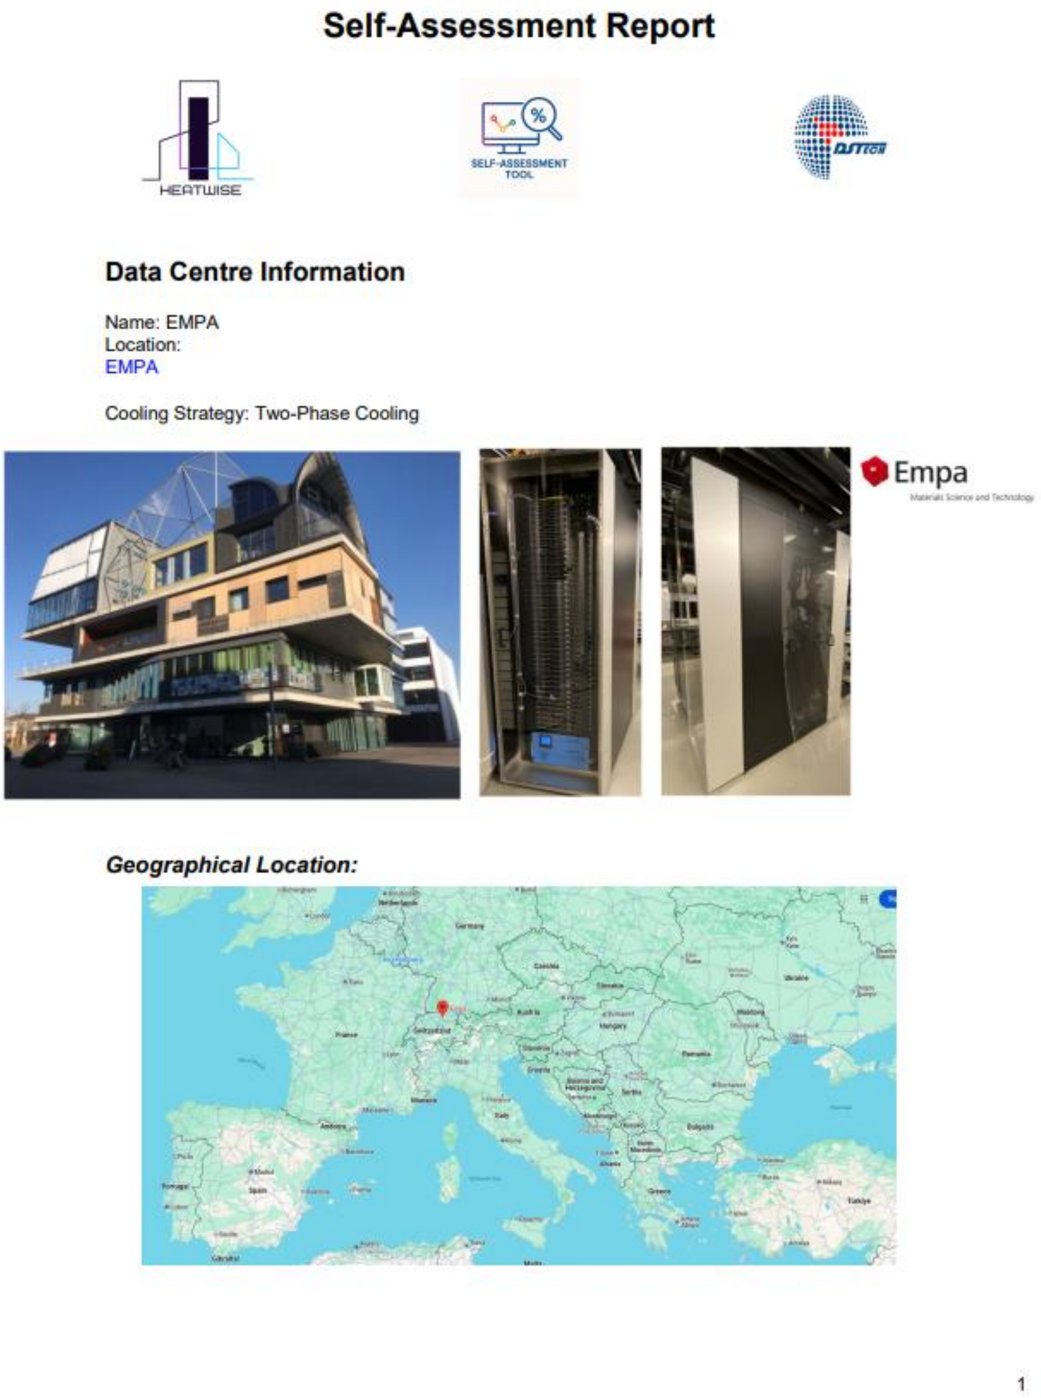

Self-Assessment Report

Data Centre Information

Name: EMPA
Location:
EMPA

Cooling Strategy: Two-Phase Cooling

Geographical Location:

1

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

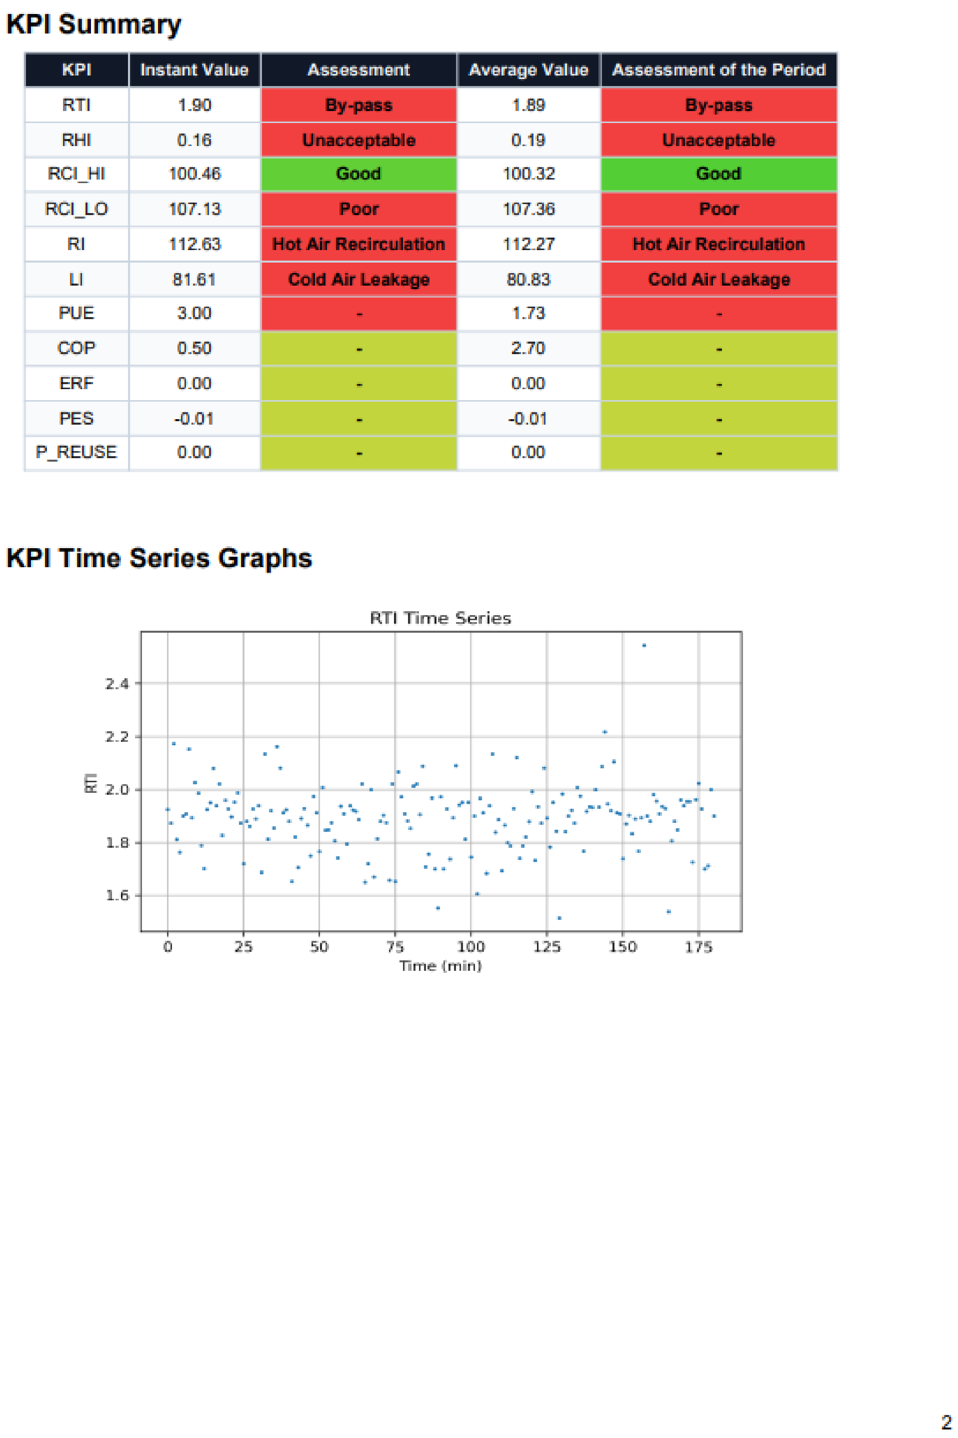

KPI Summary

KPI	Instant Value	Assessment	Average Value	Assessment of the Period
RTI	1.90	By-pass	1.89	By-pass
RHI	0.16	Unacceptable	0.19	Unacceptable
RCI_HI	100.46	Good	100.32	Good
RCI_LO	107.13	Poor	107.36	Poor
RI	112.63	Hot Air Recirculation	112.27	Hot Air Recirculation
LI	81.61	Cold Air Leakage	80.83	Cold Air Leakage
PUE	3.00	-	1.73	-
COP	0.50	-	2.70	-
ERF	0.00	-	0.00	-
PES	-0.01	-	-0.01	-
P_REUSE	0.00	-	0.00	-

KPI Time Series Graphs

2

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

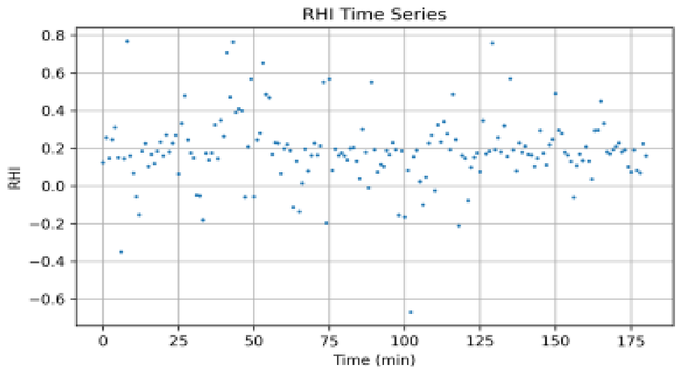

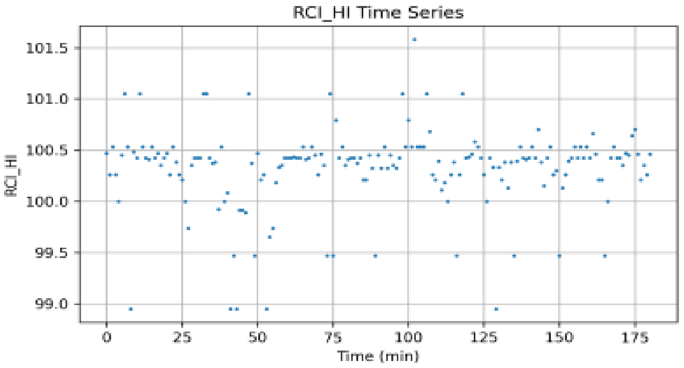

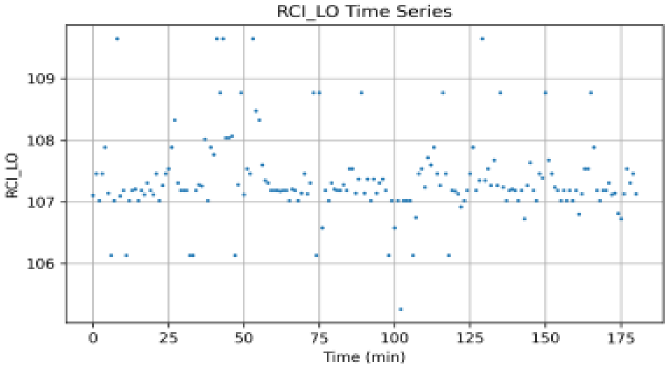

3

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

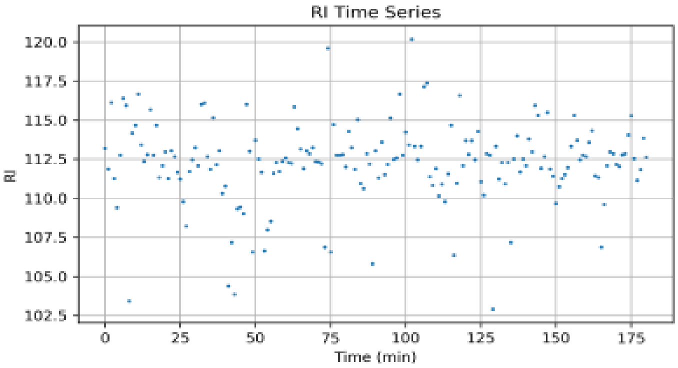

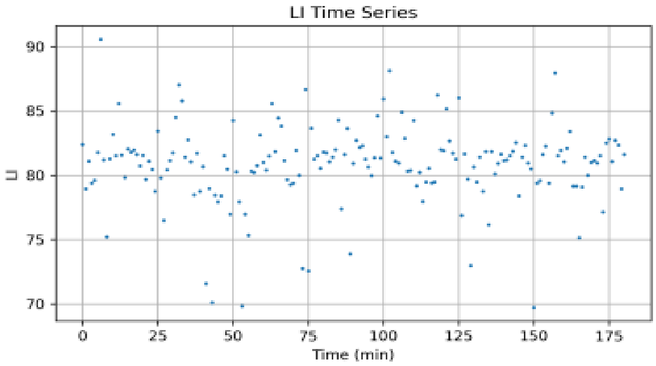

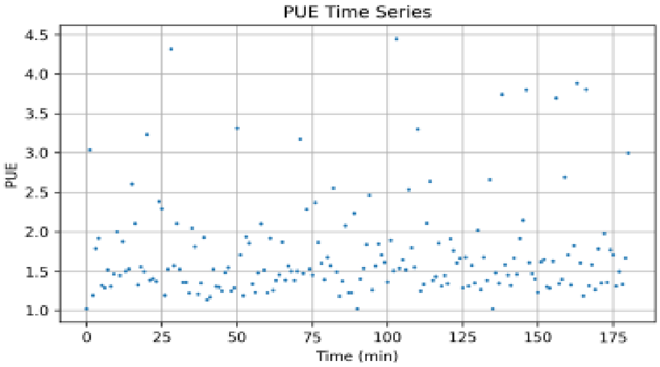

4

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

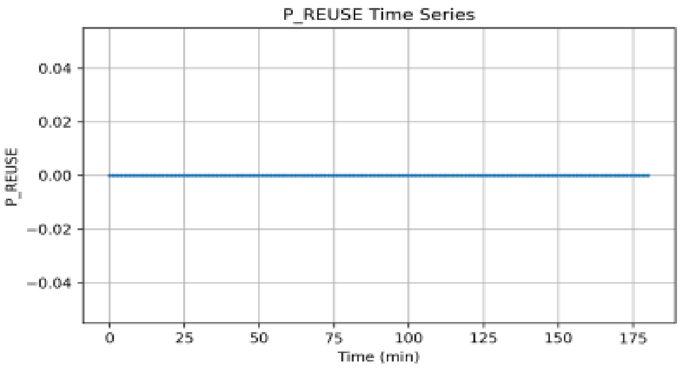

Overview Power

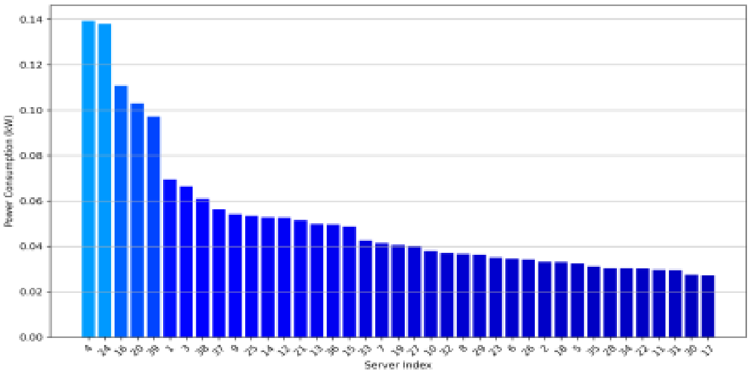

Overview Temperatures

6

This project has received funding from the European Union’s Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.

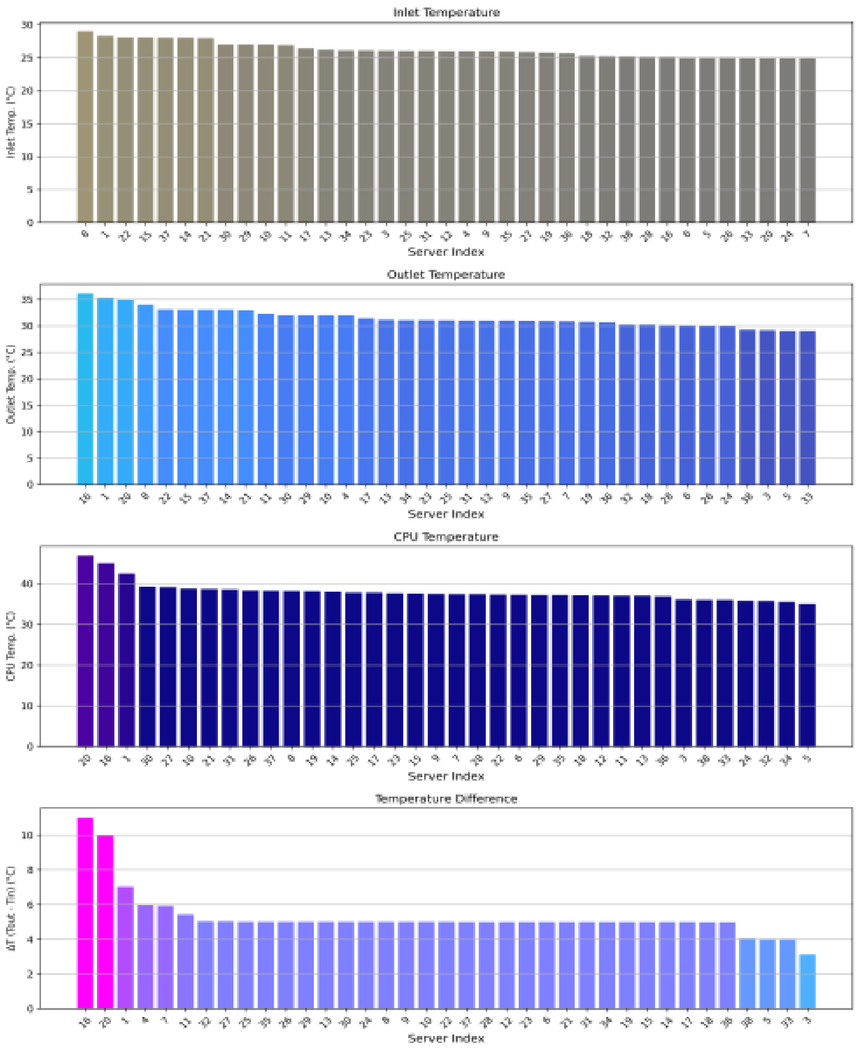

ASHRAE Standards

Category	Minimum (°C)	Maximum (°C)
Recommended	18	27
Allowable	15	32

7

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

ASHRAE Standard Check

This figure shows inlet-temperature compliance vs 27°C (recommended) and 32°C (allowable).

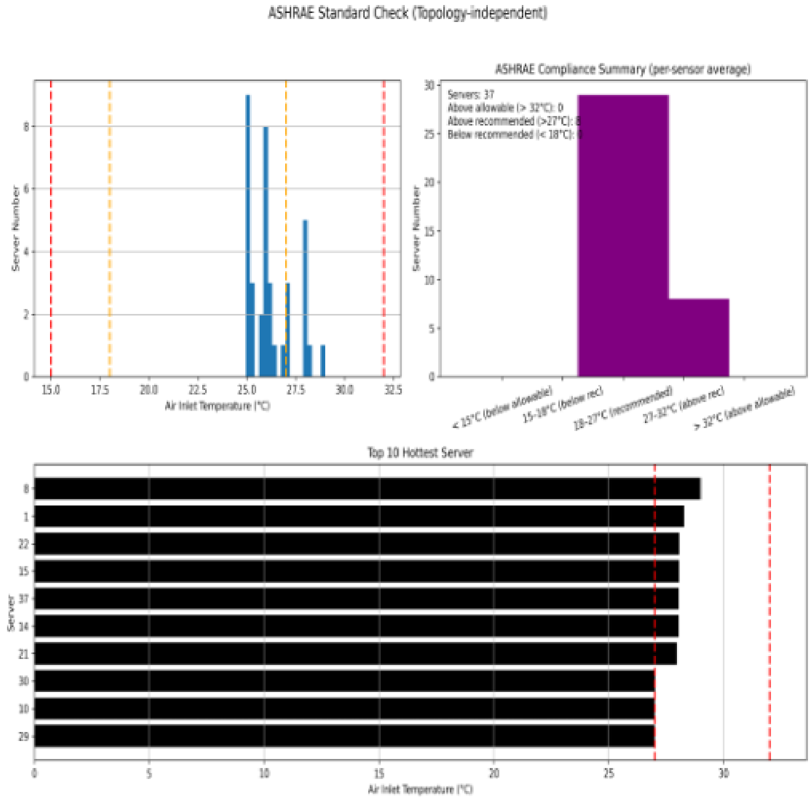

8

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

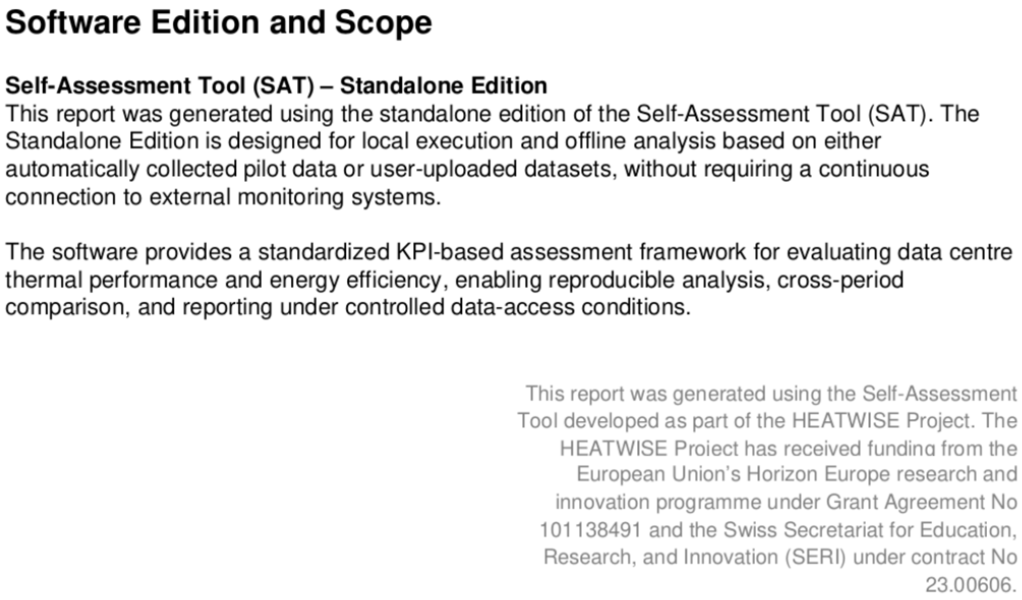

Software Edition and Scope

Self-Assessment Tool (SAT) – Standalone Edition
This report was generated using the standalone edition of the Self-Assessment Tool (SAT). The Standalone Edition is designed for local execution and offline analysis based on either automatically collected pilot data or user-uploaded datasets, without requiring a continuous connection to external monitoring systems.

The software provides a standardized KPI-based assessment framework for evaluating data centre thermal performance and energy efficiency, enabling reproducible analysis, cross-period comparison, and reporting under controlled data-access conditions.

This report was generated using the Self-Assessment Tool developed as part of the HEATWISE Project. The HEATWISE Project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract No 23.00606.

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project Nº 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract Nº 23.00606.