



Sabin Dahal

Leveraging Brick Ontology to Enhance Energy Reporting and Compliance in Data Centers

A Framework for Aligning with EU Energy Efficiency
Directive

Metropolia University of Applied Sciences

Master of Engineering

Computing in Construction

Master's Thesis

02 November 2025

Abstract

Author: Sabin Dahal
Title: Leveraging Brick to Enhance Energy Reporting and Compliance in Data Centers
Number of Pages: 64 pages + 8 appendices
Date: 02 November 2025
Degree: Master of Engineering
Degree Programme: Computing in Construction
Supervisor: Martin Aalto, Energy Expert
Seppo Törmä, Senior Lecturer

This thesis aimed to develop a unified semantic modelling framework utilising Brick and RealEstateCore ontologies to automate energy reporting and ensure compliance with the EU Energy Efficiency Directive (EED) requirements for data centers with an IT capacity exceeding 500 kW. The research followed a four-stage approach: starting with data collection at the data center facility, then moving to semantic modeling with Brick vocabulary for data center infrastructure representation, followed by RDF conversion with SHACL validation for model integrity verification, and finally automated SPARQL-based computation of KPIs. The framework produced a validated RDF model comprising more than 20 entities that describe spatial zones, IT equipment, and sensor data. The system generated KPIs that achieved compliance standards through automated calculations, showing an average PUE of 1.3 and indicating 30% additional energy consumption beyond the needs of IT equipment. The semantic modelling approach reduces manual reporting efforts, improves data accuracy, and enables scalable deployment across multiple data center facilities and was validated using SHACL validation. Overall, the framework establishes a foundation for automated EED compliance reporting while supporting advanced analytics for energy optimization and operational excellence in data center management.

Keywords: Data center, Brick, Interoperability, Sustainability, Energy efficiency, EU Energy Efficiency Directive, Semantic modelling

The originality of this thesis has been checked using the Turnitin Originality Check service.

Acknowledgement

I want to sincerely thank my supervisor at Granlund Oy, Martin Aalto, for his steady guidance, helpful feedback, and practical insights throughout this project. I am also grateful to Seppo Törmä at Metropolia University of Applied Sciences for his academic support. Additionally, I appreciate Taija Salminen for her thoughtful advice and meticulous review of the grammar, which enhanced the clarity and readability of this thesis. I also thank my team manager, Teemu Hyvärinen, for encouraging me to explore this topic and for supporting me in balancing project duties with research. I am grateful to our department manager, Timo Ranne, for trusting me with this opportunity and for his ongoing confidence in my professional development. Lastly, I am thankful to Gabe Fierro, founder of Brick; his generous guidance during the initial phase, especially in model choices and key resources, helped shape a clear technical direction.

Finally, my deepest gratitude is extended to my wife, Diksha Dhakal, whose unwavering encouragement, patience, and understanding carried me through long nights of studying and writing.

Sabin Dahal
02 November 2025
Helsinki

Contents

List of Abbreviations

1	Introduction	1
1.1	Background	1
1.2	Problem Statement	2
1.3	Research Questions	3
1.4	Thesis Structure	3
2	Granlund Oy	4
3	Literature Review	4
3.1	Data Center and Its Energy Challenges	5
3.2	Performance metrics	9
3.3	Energy Efficiency Directive of EU	12
3.4	Brick for Building and Infrastructure Modelling	13
3.5	Semantic Technologies in Energy Management	17
3.6	Core Semantic Models	18
3.6.1	BOT	19
3.6.2	SAREF	20
3.6.3	Project Haystack	21
3.6.4	RealEstateCore	22
3.6.5	SSN/SOSA T	23
3.7	Research Gaps	24
4	Methodology	27
4.1	Overview	27
4.2	Research Approach	28
4.3	Ontology Setup	31
4.4	Data Collection	34
4.5	Data Mapping & RDF Conversion	36
4.6	Metric Computation	45
4.7	Compliance Reporting	45
5	Results	46
5.1	Case Study Description	47
5.2	Results and Analysis	47

6	Discussion	54
6.1	Implications and Recommendations	54
6.2	Limitations	55
7	Conclusion and Future Work	56
7.1	Summary	56
7.2	Future Work	58
	References	60

Appendices

Appendix 1: Turtle file of the Data Center

Appendix 2: SHACL Validation

Appendix 3: EED Compliance Reporting

Appendix 4:PUE, ERE & ERF Line Chart

Appendix 5:Raw energy consumption data from the data Center

Appendix 6:Energy Consumption dataset

Appendix 7:EED Compliance Report of a data Center

Appendix 8: RDF generation script using Python and Brick

List of Abbreviations

- AEC:** Architecture, Engineering, and Construction. Architects and engineers are employed for the design and execution of commercial and residential projects during the planning and construction phase.
- AHU:** Air Handling Unit. An AHU consists of elements that are installed in large, accessible box-shaped units called modules, which provide the necessary ventilation requirements for purifying and air conditioning the indoor air in a building
- AI:** Artificial Intelligence. Artificial intelligence represents a developing technology that seeks to duplicate human intelligence through machine systems.
- AGI:** Artificial General Intelligence. An AGI is a theoretical concept that represents an artificial intelligence capable of self-control and performing various basic tasks, much like a person.
- API:** Application Programming Interface. APIs define an application as any software program that performs a specific task. The interface functions as a service agreement between two applications.
- ASHRAE:** American Society of Heating, Refrigerating, and Air-Conditioning Engineers. An American professional association seeking to advance heating, ventilation, air conditioning, and refrigeration systems design and construction.
- BACnet:** Building Automation Controls Network. Based on ISO 16484-5, this network protocol specifically facilitates communication across buildings for HVAC, lighting, and fire detection.

- BAS:** Building Automation System. BAS functions as an intelligent system that streamlines building management operations. The system encompasses multiple functions that enable HVAC control, management of lighting and security systems, and energy management capabilities. BAS operates within the BMS (Building Management System) framework because BMS describes the complete system, which includes BAS along with other integrated components for building-wide control and optimisation.
- BIM:** Building Information Modelling. This involves creating a virtual 3D model and managing the information related to a construction project throughout its lifecycle.
- BMS:** Building Management System. The system operates through computer technology to monitor and control building electrical equipment, including ventilation systems, lighting systems, energy systems, fire systems, and security systems.
- BOT:** Building Topology Ontology. The Building Topology Ontology (BOT) is a minimal ontology for describing the core topological concepts of a building.
- CRAC:** Computer Room Air Conditioners. A cooling system for data centers uses refrigerant to regulate temperature and humidity.
- CRAH:** Computer Room Air Handlers. These systems use chilled water to regulate temperature and humidity in data centers.
- DC:** Data Center. Organisations typically use a large group of networked computer servers to store, process, or distribute large amounts of data remotely.

- DCIM:** Data Center Infrastructure Management. DCIM represents the integration of IT functions and building facility management within an organisation.
- DT:** Digital Twin. A digital twin contextualises a reliable and up-to-date digital representation of a physical object, person, or process in a digital version of its environment.
- EED:** Energy Efficiency Directive. The revised Energy Efficiency Directive (EU/2023/1791) establishes "energy efficiency first" as a legal principle, requiring EU countries to prioritise energy efficiency in their policies and investments.
- EU:** European Union: The EU is a unique partnership between 27 European countries, known as Member States, or EU countries.
- FM:** Facility Management. A professional discipline focuses on delivering efficient logistics and support services that relate to real estate and buildings.
- GHG:** Greenhouse Gas. GHGs are the gases in the Earth's atmosphere that trap heat.
- HVAC:** Heating, Ventilation, and Air Conditioning. These technologies are used to regulate the temperature, humidity, and air quality in an enclosed space.
- IFC:** Industry Foundation Class. A 3D design data standard for the AEC (Architecture, Engineering, and Construction) industry. It allows the description of building structures and semantic properties, promoting interoperability and collaboration.
- IoT:** Internet of Things. Sensors, processing power, software, and additional technologies equip these devices. These components

enable the devices to communicate and exchange data with other devices and systems via the Internet or alternative communication networks.

- IT: Information Technology. Refers to the science and practice of utilising computers and other electronic devices to store and transmit data.
- JSON: JavaScript Object Notation. It is commonly used for data storage and transfer. JSON is a popular choice for applications that benefit from a simple and easy-to-use data format.
- KPIs: Key Performance Indicators are the key targets that a company should track to make the most impact on the company's strategic business outcomes.
- KNX: Konnex. The fieldbus protocol operates as a specific communication standard for building automation systems. The system allows for unbroken communication and control of building systems, which include lighting, HVAC, security, and other applications.
- kWh: Kilowatt-hour. A non-SI unit of energy; it is the energy delivered by one kilowatt of power for one hour.
- LOD: Level of Development. It refers to the degree of commitment to and maturity of the design of a BIM model element at a specific project stage.
- MEP: Mechanical, Electrical, and Plumbing systems responsible for managing a building's mechanical, electrical, and plumbing infrastructure.

- MQTT:** Message Queuing Telemetry Transport. The protocol works as a simple messaging system for devices that have limited abilities and for networks that have low bandwidth, high delays, and unreliable connections.
- OWL:** Web Ontology Language. A semantic web language designed to represent rich and complex knowledge about things, groups of things, and relations between things.
- PAAS:** Platform as a Service. PaaS, or Platform as a Service, is a cloud computing platform that includes hardware, software, and infrastructure for development and deployment
- PUE:** Power Usage Effectiveness. PUE is a metric used to determine the energy efficiency of a data center.
- RDF:** Resource Description Framework. The Resource Description Framework (RDF) is a general framework for representing interconnected data on the web.
- SHACL:** Shapes Constraint Language. W3C standard language for describing Resource Description Framework graphs. SHACL has been designed to enhance the semantics and technical interoperability of ontologies, expressed as RDF graphs.
- SAAS:** Software as a Service. A method of software delivery and licensing in which software is accessed online via subscription, rather than bought and installed on individual computers.
- TCP:** Transmission Control Protocol. TCP is a protocol that sits on the top of Internet Protocol (IP) to provide reliable, ordered, and error-checked delivery of a stream of bytes between hosts communicating via IP.

- UPS: Uninterruptible Power Supply. A UPS is a device that maintains power in the event of a failure. One of its components is a battery that is always ready and charged.
- VAV: Variable Air Volume. Variable-volume systems enable energy-efficient HVAC system distribution by optimising the amount and temperature of the distributed air.
- W3C: World Wide Web Consortium. An international, public-interest non-profit organisation in which member organisations, full-time staff, and the public collaborate to develop web standards.

1 Introduction

The digital economy fundamentally relies on data centers, which serve as its operational backbone. Organizations are increasingly dependent on cloud computing, artificial intelligence, and data-driven services, elevating the importance of efficiency and sustainability in data center operations. The following section outlines the research background that underpins the investigation of power usage in data centers and its resulting implications.

1.1 Background

Data centers support various digital services, including cloud platforms, social media, e-commerce, and artificial intelligence workloads. They are among the most energy-intensive building types today. A single facility can use hundreds of times more electricity per square meter than a conventional office, and global data center electricity consumption is increasing due to the rapid growth of digital services. The International Energy Agency estimated that data center consumption could reach approximately 945 TWh by 2030 [1]. Cooling and power distribution infrastructure add significant overhead to the IT load, pushing overall energy use well beyond what the servers consume. As intelligence models expand, their energy footprints also vary greatly: some large language models consume orders of magnitude more power than more efficient alternatives.[2.] These trends have substantial environmental impacts and increase operating costs.

Data centers in the US used about 1.7 billion liters of water between 2023 and 2024. Reports from the Houston Advanced Research Center suggest that if the capacity of these data centers increases, water consumption could reach approximately 185 billion liters by the end of 2025 and 1.51 trillion liters by 2030. Evaporative cooling uses significantly more water in hot climates than dry coolers or closed-loop liquid systems. Therefore, data centers should be built with water-efficient cooling, seasonal free cooling, and waste heat recovery in mind. They should be equipped with advanced technology to continuously monitor and report

Water Usage Effectiveness (WUE). [3.] To address rising energy and water use, the EU's Energy Efficiency Directive (EED) establishes an “energy efficiency first” principle and mandates reporting for data center operators with an installed IT capacity above 500 kW. [4.] Starting in 2024, operators must submit annual reports to a European database with KPIs such as Power Usage Effectiveness(PUE), Water Usage Effectiveness(WUE), Energy Re-use Factor(ERF), Renewable Energy Factor (REF), carbon emissions, and renewable energy use details. The delegated regulation adopted in 2024 specifies the data points and templates required for these reports. However, current reporting processes are largely manual, and existing Building Management Systems (BMS) and Data Center Infrastructure Management (DCIM) systems utilise proprietary formats with inconsistent naming, which makes it challenging to aggregate sensor data and automate metric calculations.[4.]

1.2 Problem Statement

The central question addressed in this thesis is whether Brick, an open and extensible ontology for buildings, can be applied to model a data center's physical and logical structure, enabling the automated generation of EED-compliant energy reports. The specific objectives are semantic modelling focusing in developing a Brick ontology for a data center that covers spaces, IT equipment, cooling and power systems, and their sensors. Use Brick relationships—such as `brick:hasPart` for composition and `brick:hasPoint/brick:isPointOf` to link equipment to telemetry points and supplement with RealEstateCore classes included in the BrickSchema for zone definitions.[14.]Data mapping and RDF conversion needs to done by connecting different CSV files from BMS/DCIM systems to Brick instances; link time-series identifiers to sensors and change the data into RDF triples for storage and analysis.

Automated KPI computation is done by creating SPARQL queries to compute PUE, WUE, and other metrics from the RDF model. The integration of time-series

data into the semantic graph enables the automatic aggregation of data, producing more accurate results.

For Validation and reporting the system needs to implement a multi-stage validation pipeline that includes syntax checking, vocabulary compliance, and SHACL constraints to verify model quality. A prototype reporting application should be developed to produce CSV reports which follow the format specified in the delegated regulation. For Evaluation the framework is implemented with a case study of a Nordic data center to assess semantic completeness, computational correctness, and practical benefits.

1.3 Research Questions

This research is guided by three questions. First, how can Brick ontology be applied to model the physical and logical and metering infrastructure of a data center in alignment with EU-EED reporting requirements?

Second, what are the technical approaches and challenges in linking time-series measurement data to a Brick ontology to enable automated computation of EED performance indicators (PUE, WUE, ERF, REF, etc.), and what impact does this automation have on reporting efficiency and accuracy?

Third, which validation methods and tools ensure that a Brick-based RDF model remains reliable, complete and compliant with EU regulations before KPI computation and submission?

1.4 Thesis Structure

The thesis begins with a literature review covering data center architecture, energy performance metrics PUE, WUE, ERF, REF, and CUE, the EU EED requirements, and existing semantic ontologies Brick, BOT, SAREF, Project Haystack, RealEstateCore, and SSN/SOSA. The methodology chapter outlines the research approach, details the ontology setup, explains data

collection methods, describes mapping and RDF conversion processes, and covers validation procedures, metric computation, and compliance reporting.

The prototype chapter applies these methods to simulated data. The results chapter analyses the semantic graph and presents visualisations and KPIs, while the discussion chapter addresses implications, recommendations, and limitations. Finally, the conclusion summarises the findings and outlines future research directions.

2 Granlund Oy

This thesis is carried out in collaboration with Granlund Oy's Data Center Department, which plays a key role in planning, designing, and delivering mission-critical facilities in Finland and abroad. The department prioritises energy efficiency, sustainability, and adherence to international and EU regulations, ensuring that data centers meet the highest standards of operational reliability. Founded in 1960, Granlund Oy is a prominent organisation in the real estate and construction sectors, employing around 1,500 specialists across various fields. In addition to its activities in Finland, Granlund Oy has expanded to Sweden and the United Kingdom. [5.]

The company functions as a consortium of enterprises specialising in various fields, including Mechanical, Electrical, and Plumbing (MEP) design; property management; software development; energy and environmental consulting; real estate consulting; construction management and supervision; and building management. Granlund Oy reports that its annual net sales exceed EUR 139 million, with six per cent of its turnover allocated to innovation and research, demonstrating its commitment to improving industry standards. [5.]

3 Literature Review

EED requires data centers to meet energy efficiency standards, as these facilities consume a significant amount of energy. The process of achieving compliance

demands both precise metric calculations and standardised machine-readable reporting methods.[5.] The following section examines the technical, regulatory, and semantic modelling requirements that enable Brick ontology to implement this challenge.

3.1 Data Center and Its Energy Challenges

Data centers form the backbone of a wide array of Internet-based services, including social networking, e-commerce, web hosting, and more specialised services such as platform as a service (PaaS), software as a service (SaaS), and grid/cloud computing. For example, Microsoft's Azure platform, Google App Engine, Amazon's EC2 platform, and Sun's Grid Engine illustrate these fundamental service platforms. Many of these services are enabled by virtualisation, which is increasingly adopted in data centers to boost server utilisation and enable more flexible resource allocation. However, virtualisation also adds complexity to various aspects of data center management. [5.] A data center comprises different systems that operate efficiently. These systems can be categorised into power Systems, cooling systems, cabling and connectivity, networking equipment and servers.

The operation of a data center is dependent on having a power infrastructure that is both consistent and reliable. UPS, PDU, and generators are some of the devices that contribute to the achievement of this objective.

Server cooling is needed to prevent heat from reducing device efficiency and lifespan. There are several methods for managing the heat generated by the servers. The regulation of temperature and humidity in the data center is achieved through a combination of computer room air conditioners (CRACs) and computer room air handlers (CRAHs). To cool the air, CRAC units utilise refrigerants, whereas CRAH units employ chilled water. Prudnikov says that Brick is helpful in improving energy use for cooling systems, which are the largest consumers of power.[6.]

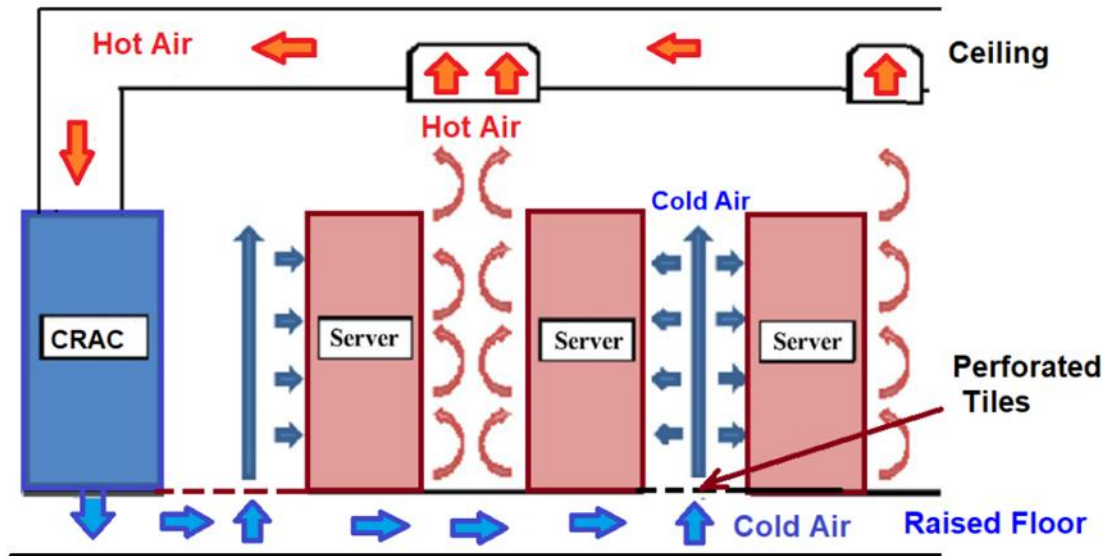


Figure 1. Cooling in a data center. [7.]

The heated air generated by the servers is directed upwards and then to the CRAC units, which cool the air and deliver it back to the server via perforated tiles, as shown in Figure 1. In terms of temperature, the cold air is 21 degrees Celsius, while the hot air ranges from 35 to 41 degrees Celsius. [7.]

All IT equipment in a data center must be interconnected via a combination of copper and fibre-optic connections. Effective cable management eliminates ventilation and cooling issues, reduces downtime, and ensures easy access to equipment. To establish an organised and efficient layout, techniques such as structured cabling are used to consolidate and direct the lines systematically. Networking equipment comprises connectivity devices such as switches, routers, load balancers, and firewalls, which facilitate the efficient transfer of data.

High-performance computers housed in data centers are commonly referred to as servers. They provide applications and users with the storage, resources, and services needed to operate over a network. These servers, built with durable hardware such as multi-core CPUs, ample RAM, and redundant power supplies, ensure reliability and speed. They come in various types, including

web, application, database, and virtualisation servers. Each type serves a specific purpose. They are often arranged in rack-mounted or blade-mounted setups to optimise space and efficiency. In contrast, virtualisation technology allows multiple virtual machines to run on a single server, thereby enhancing resource utilisation. Data centers feature advanced cooling systems, robust security measures, and redundancy to ensure operational continuity and protect against failures, while gradually adopting sustainable practices to reduce energy consumption and minimise environmental impact. [8.]

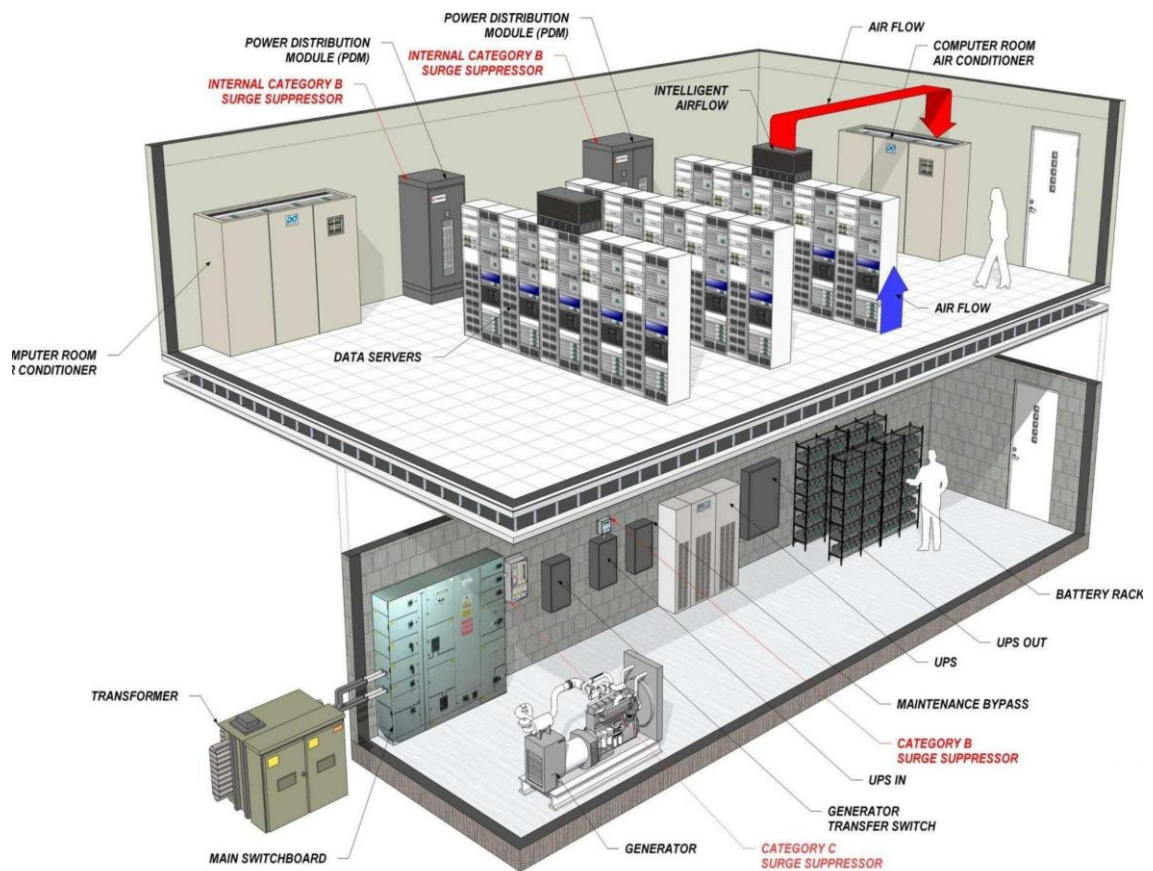


Figure 2. Fundamentals of a Data Center. [8.]

The main components and basic structure of a data center appear in figure 2. The upper section displays the IT infrastructure which consists of data servers and network racks and computer room air conditioners that use intelligent airflow management to maintain proper temperature. The blue and red arrows in the airflow pattern show the cooling system which delivers cold air to server racks while removing hot air for temperature control.

The lower section shows the electrical and backup systems which maintain ongoing operation. It includes transformers, switchboards, uninterruptible power supply (UPS) units, battery racks, and a generator for emergency power. The system components function as one unit to deliver reliable power distribution and surge protection which maintains uninterrupted data center operations.

The figure 2 demonstrates how physical infrastructure and power systems operate together to achieve energy efficiency and reliability which are central themes of this thesis about energy reporting and compliance.

The energy efficiency of data centers has lately emerged as a prominent issue because to its substantial economic, environmental, and performance implications for society. Data centers exert a significant financial influence on society for various reasons. An average data center uses the same amount of power as many as 25000 homes. The space in the data center can consume up to 200 times as much electricity as a standard office. Additionally, the energy expenses of operating a regular data center double every five years. Thus, for today's data centers, electricity bills have become a significant outlay, given the sharp increase in electricity use and growing electricity costs. Occasionally, the power costs exceed the cost of purchasing hardware.

Energy usage in the data center generates several environmental issues. For instance, the total data center power consumption in 2005 accounted for one per cent of the US power consumption overall. It produced emissions comparable to those of a mid-sized nation, such as Argentina. The total electricity consumed by data centers in the US was 1.7% to 2.2% of the country's total electrical consumption [9]. The global electricity usage by data centers was estimated to be between 1.1% and 1.5% of the total worldwide electricity usage [9.] Data centers worldwide consumed 270 TWh in 2012 [9] and 300-380 TWh in 2023 [8], and it is expected to reach 945 TWh by 2030 [1]. These numbers demonstrate the importance of data centers reducing energy consumption, implementing precise energy monitoring, providing transparent reporting, and aligning with EED

requirements. Therefore, data centers are under more scrutiny than ever to address rising energy costs, mitigate their environmental impact, and implement robust energy management plans to ensure continued operation in a manner that is both compliant and sustainable.

Due to this concern, this thesis examines how Brick, a standardised metadata schema for building systems, can be utilised to optimise energy consumption and enhance energy reporting and compliance frameworks in data centers. These improvements will facilitate more intelligent energy management, improved data interoperability, and the achievement of sustainability goals. As artificial intelligence (AI) and the pursuit of artificial general intelligence (AGI) progress rapidly, the energy requirements of these systems are increasing at an unprecedented rate.

Recent benchmarks indicate that the energy consumption of large language models can differ significantly, with models such as O3 and DeepSeek-R1 utilising almost 70 times more electricity per task compared to energy-efficient alternatives like GPT-4.1 nano [10.] This trend underscores a significant challenge: future data centers must support increasingly intricate AI workloads while reducing their environmental footprint. If neglected, the extensive implementation of energy-intensive AI models may severely strain electrical networks, increase carbon emissions, and hinder sustainability objectives. Consequently, the data center needs to implement sophisticated energy monitoring and reporting systems. Utilising semantic models, such as Brick, provides a practical approach, allowing operators to outline and oversee energy use based on specific workloads. This not only enhances energy efficiency and operational optimisation but also bolsters compliance with regulatory frameworks and environmental norms as AI expands globally.

3.2 Performance metrics

In recent years, numerous performance metrics have been established to evaluate and compare the efficacy and efficiency of data centers. These metrics

can assess individual components (cooling, IT, power supply) of the data center or the data center in its entirety. Performance measures encompass total energy use, water usage, and carbon emissions, along with subsystem efficiency, including temperature distribution.[11.] The following are the primary metrics used to evaluate the performance of an energy-efficient data center.

PUE, which shows the ratio between the total power used by the building and the power used by IT equipment, is the most common performance measure: [12.]

The ideal PUE value is 1.0, but the upper limit is infinite.

$$PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}; 1 \leq PUE \quad (1)$$

CUE is the main energy impact and carbon emissions of a data center are measured by Carbon Usage Effectiveness (CUE). [11.]

$$CUE = \frac{CEF * \text{Total energy}}{\text{IT energy}} \left[\frac{\text{kgCO}_2}{\text{kWh}} \right]; 0 \leq CUE \quad (2)$$

where CEF means Carbon Emission Factor

Data center generate extra heat by turning electrical energy into thermal energy in their IT infrastructure. This extra heat can be utilised in various ways, depending on local heating needs. Although quantifying energy reuse is challenging, it presents a considerable opportunity to enhance energy efficiency. The Energy Reuse Factor (ERF) measures the extent to which energy is reused. [12.]

$$ERF = \frac{\text{Reuse energy}}{\text{Total energy}}; 0 \leq ERF \leq 1 \quad (3)$$

Even though PUE doesn't account for water usage, about 4% of the electricity demand in the US goes toward water treatment and transportation. The Water Usage Effectiveness (WUE) is available to monitor the ecological impact of this. [12.] It is defined as:

$$WUE = \frac{\text{Water Usage}}{\text{IT energy}} \left[\frac{\text{L}}{\text{kWh}} \right]; 0 \leq WUE \quad (4)$$

The Renewable Energy Factor (REF) measures the proportion of renewable energy used by data centers relative to their overall energy usage. [12.]

$$REF = \frac{E_{REN}}{E_{DC}}, 0 \leq REF \leq 1 \quad (5)$$

E_{REN} : Total amount of energy coming from renewables in kWh

E_{DC} : Total amount of energy consumed by the data center in kWh

The cooling effectiveness (CER) ratio has been one of the most popular efficiency metrics recently because of the extreme heat produced during the operations within a data center. It determines the cooling system's efficiency by dividing the amount of heat extracted by the energy consumed by the system.

$$CER = \frac{Q(\text{removed})}{E(\text{cooling})} \quad (6)$$

The quantity of heat evacuated from the data center in kWh per year is denoted as $Q(\text{removed})$.

ERE is calculated as:

$$ERE = \frac{\text{Cooling} + \text{Power} + \text{Lighting} + \text{IT} - \text{Reuse}}{IT} \quad (7)$$

The cooling system consumes energy measured in kWh per year, denoted as $E(\text{cooling})$.

According to the European Union's Energy Efficiency Directive (EED), data centers with an installed IT capacity exceeding 500 kW are required to submit annual energy usage statistics, with compliance commencing in 2023 and the first submissions due by May 2024. [4.] The incorporation of Brick—a standardised metadata schema intended to support portable innovative building applications, encompassing detailed energy modelling and analytics—will be crucial in enabling precise generation and reporting of energy consumption data, thus improving operational efficiency and sustainability in data centers. This development will influence future views of data centers, categorising them either as major global energy consumers, resembling an environmental liability, or as vital infrastructure in a digitised society that fosters sustainability, exemplified by

waste heat recovery for residential heating and greenhouse agriculture support. [13.]

3.3 Energy Efficiency Directive of EU

Adopted on September 13, 2023, the EU Energy Efficiency Directive (EED) (Directive EU 2023/1791) reforms Directive 2012/27/EU and provides a comprehensive framework to promote energy efficiency throughout the Union in line with the 'Fit for 55' package and the European Green Deal. This directive emphasises the 'energy efficiency first' principle, which means focusing on reducing energy use in areas such as buildings, transportation, and industry. It mandates that final energy consumption must be cut to no more than 763 million tonnes of oil equivalent by 2030. In this context, the directive addresses the growing energy needs of the ICT industry, where data centers are a crucial sector due to their significant contribution to electricity consumption (roughly 3% of the EU total, according to recent estimates, with projections indicating further growth driven by digitalisation and advancements in artificial intelligence). [4.]

The directive's Recitals 86 and 87 emphasise the importance of increasing transparency in data center operations to support sustainability. These clauses acknowledge that the demand for ICT energy is expected to increase dramatically by 2030 and support mandatory reporting to raise awareness among operators and stakeholders. This will foster improvements in energy efficiency, the integration of renewable energy, waste heat utilisation, cooling effectiveness, carbon management, and water conservation. In order to achieve the Union's overarching goals, which comprise a 55 percent drop in GHG emissions from 1990 levels by 2030, these steps are being taken to ensure that relevant laws, like the Renewable Energy Directive (EU) 2018/2001 and the Industrial Emissions Directive 2010/75/EU, work together.[4.]

The particular requirements for data centers with an installed information technology (IT) power demand of at least 500 kW are outlined in Article 12. Key performance indicators (KPIs), including power usage effectiveness (PUE),

temperature set points, data usage, energy reuse factor (ERF), renewable energy factor(REF), water usage effectiveness (WUE), IT power demand, and storage capacity, must be publicly reported by owners and operators starting on May 15, 2024, and then every year after that. Delegated Regulation (EU) 2024/1364, which outlines the initial stage of a Union-wide sustainability rating system, was adopted by the European Commission on 14 March 2024 to operationalise these requirements. [4.]

3.4 Brick for Building and Infrastructure Modelling

Brick is an open-source ontology designed to standardise semantic descriptions of physical, logical, and virtual assets within buildings, as well as their interrelationships. Brick has two parts: a flexible class structure that represents the physical and logical elements in buildings, and a basic set of connections that illustrate the relationships between these elements. A labelled, directed graph with nodes representing entities and edges representing relationships is a Brick model of a building. Brick is described using the Resource Description Framework (RDF). RDF represents graph-based knowledge as tuples of (subject, predicate, object) termed 'triples'. A triple is defined in such a manner that a subject entity is related to an object entity through a predicate. [14.]

```
1 # Example of sensor instantiation for a Light Sensor in a building automation system,
  demonstrating various instantiation methods using the BRICK ontology
2 # instantiate class explicitly
3 :lightSensor1 a brick:Light_Sensor .
4
5 # instantiate a class implicitly through application of tags
6 :lightSensor1 brick:hasTag tag:Light .
7 :lightSensor1 brick:hasTag tag:Illuminance .
8 :lightSensor1 brick:hasTag tag:Sensor .
9
10 # combination of explicit class and tags
11 :lightSensor1 a brick:Illuminance_Sensor .
12 :lightSensor1 brick:hasTag tag:Light .
13
14 # instantiation from behavior
15 :lightSensor1 a brick:Sensor .
16 :lightSensor1 brick:measures brick:Light .
17 :lightSensor1 brick:measures brick:Illuminance .
18
19 # alternative instantiation from behavior
20 :lightSensor1 a brick:Illuminance_Sensor .
21 :lightSensor1 brick:measures brick:Light .
```

Listing 1. An example of sensor instantiation for a light sensor in a building automation system using Brick. [15.]

Listing 1 illustrates the application of the tag system within Brick. Line 3 of listing 1 is a triple which consists of `:lightSensor1` as the subject, the predicate (relationship) is 'a', and the object is `brick:Light_Sensor`. Brick is defined using two languages called Resource Description Framework Schema (RDFS) and Web Ontology Language (OWL). The implementation of Brick to establish a structured and semantic representation of a building's systems, sensors, and equipment facilitates efficient control, monitoring, and automation of energy-intensive activities, hence optimising energy consumption in data centres.[14].

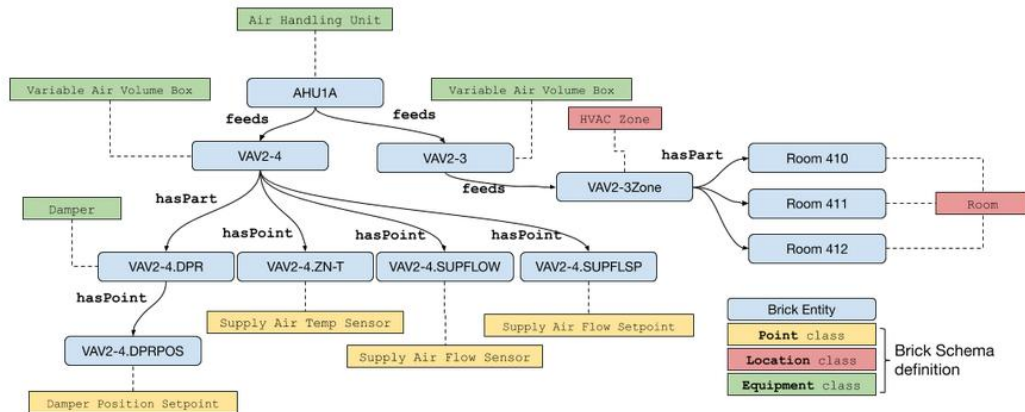


Figure 3. Graph representation of a building in Brick. [14.]

Figure 3 represents the Brick of a building, where the central Air Handling Unit (AHU) supplies air to the system. An example shown in figure 4, it is defined as AHU1A. It supplies Variable Air Volume (VAV) boxes, which regulate airflow to the various HVAC zones. The VAV2-3Zone consists of three rooms: Room 410, Room 411, and Room 412. VAV2-4 has control points known as VAV2-4.DPR as a damper, and that damper has another point, VAV2-4.DPRPOS, known as the damper position setpoint. Brick comprises five core concepts, which are explained as **Entity**, any physical, logical, or digital object—the real "things" in a building—can be abstracted into an entity. In this context, anything with a tangible form is considered a physical entity. AHUs, VAV boxes, luminaires, lighting systems, electric meters, thermostats, and chargers for electric vehicles are all examples of mechanical equipment. Rooms and floors are examples of spatial elements. Everything whose representation is rooted in software is considered a virtual entity. At actuation points, software can write values; at sensing and status points, it can read the world as it is right now. [14.]

Tag is a property of an object is its tag. Examples of tags include sensors, setpoints, air, and water. Project Haystack created tags, and Brick has adopted the concept from Project Haystack. However, Brick does not entirely rely on tags to determine an entity's type. [14.]

A **class** is a category of entities that is defined by a particular shared purpose and properties that belong to the same group. Classes are structured in a hierarchy, and entities represent instances of one or several classes. The hierarchy also includes additional tags that help locate annotations. [14.]

Relationship is the character of the connection between two related things is defined by a relationship. *Encapsulation, sequencing, and instantiation* are examples of relationships that exist between objects. There are other perspectives on the philosophy of Brick relationships, like composition, topology and telemetry. [14.]

The **graph** is defined as an organised data structure, which represents a set of entities called "nodes" and relationships called "edges." A directed and labelled graph functions as a representation of Brick. In Brick, the graph is represented using the RDF data model [14].

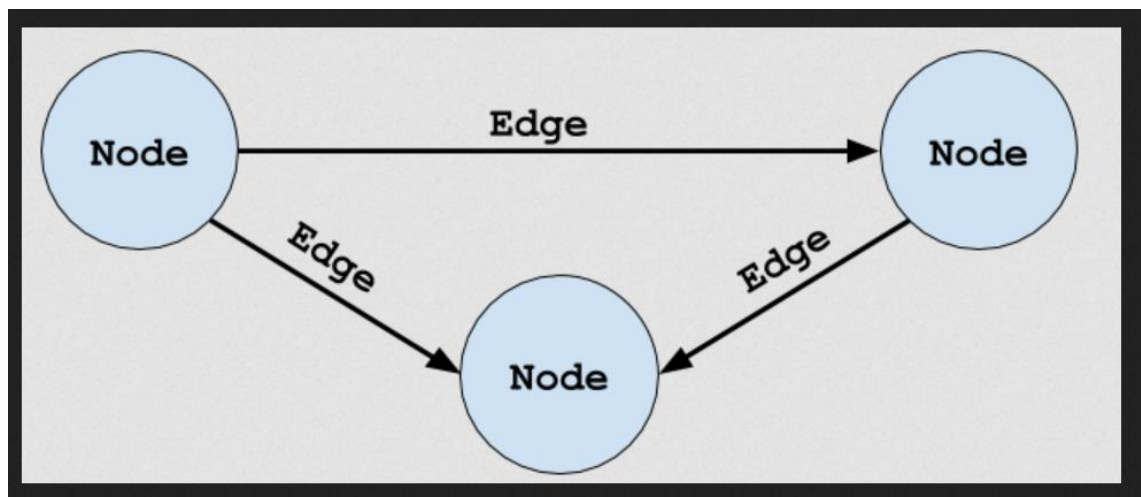


Figure 4: Data structure representation in Brick using nodes and edges. [15.]

Figure 4 is a graph illustrates a directed graph, also known as a digraph. The circles designated as "Node" represent the vertices of the graph. The arrows marked "Edge" denote directed connections between the nodes, illustrating the direction of relationships. Each arrow directs from a source node to a target node, signifying the potential flow of data or control inside a system.[15.]

3.5 Semantic Technologies in Energy Management

Structured data representation and interoperability are enabled in energy management systems through the use of semantic technologies, including knowledge graphs, linked data frameworks, and ontologies. For better building energy management (BEM) decisions, these technologies let computers understand and process data from various sources, including sensors installed in buildings, weather reports, and metrics measuring energy usage.[16.]

Buildings account for roughly 30–40% of worldwide energy use, indicating an urgent need for enhanced energy efficiency, operational effectiveness, and sustainability within the built environment. Achieving these objectives requires the synthesis and astute use of diverse data generated by various building systems, sensors, and apparatus from multiple suppliers. This diversity presents considerable issues for data interoperability, as hardware and software systems frequently employ proprietary or incompatible data formats. [16.]

Semantic modelling has emerged as an effective way to overcome these difficulties. Semantic models utilise formal ontologies to define the relationships between concepts in a machine-readable format, enabling consistent data interpretation, enhanced interoperability, and automated data processing across platforms. Semantic modelling integrates meaning into data structures, enabling sophisticated functions like automated reasoning, which aids in fault identification, predictive maintenance, demand response optimisation, and energy forecasting. Moreover, these models incorporate intrinsic scalability, allowing for the effortless integration of additional systems and the augmentation of applications without requiring significant reconfiguration. [16.]

The evolution of digital modelling in the construction sector demonstrates a distinct path from conventional building information models (BIMs) to more nuanced semantic representations. The Industry Foundation Classes (IFC) schema initially served as the primary standard for BIM data sharing, offering a comprehensive yet rigid framework. Although IFC is essential during the design

and construction phases, it lacks the flexibility and inferencing capabilities required for dynamic building operations. Recent initiatives have aimed to convert IFC models into OWL (Web Ontology Language)-based representations to leverage the benefits of semantic web technologies, thereby addressing these limitations. This transition facilitates enhanced integration of BIM data with real-time building management systems and other data sources, establishing a basis for genuinely intelligent and adaptable building operations. [16.]

These advantages are particularly significant for data centers, which are among the most energy-intensive building types. As data centers become increasingly complex and extensive, effective energy management is essential—not only for cost containment but also to mitigate carbon emissions and achieve sustainability objectives. Semantic modelling enhances visibility and integration among diverse data sources, such as IT load measurements, cooling systems, power distribution units (PDUs), and environmental sensors. Ontology-driven frameworks can integrate these varied data streams to facilitate automated analytics, dynamic task allocation, and intelligent cooling methodologies. This semantic methodology is especially beneficial for the implementation of energy performance measures, such as Power Usage Effectiveness (PUE) and Carbon Usage Effectiveness (CUE), which facilitate ongoing monitoring and optimisation. [16.]

Furthermore, semantic integration facilitates regulatory compliance, reporting, and engagement in innovative grid initiatives by enhancing communication with external energy systems and stakeholders. The integration of BIM, building operations, and semantic technologies positions data centers as beneficiaries of advanced digital infrastructure, enabling them to address real-time demands while reducing their environmental impact.

3.6 Core Semantic Models

The foundational ontologies that have gained popularity in building operations due to their practicality and conformity to semantic web standards are highlighted

in this section. Among these, BOT (Building Topology Ontology), SAREF (Smart Applications REFerence ontology), SSN/SOSA (Semantic Sensor Network / Sensor, Observation, Sample, and Actuator), and Brick (see 2.4 above) stand out because they each address a different layer of building representation, ranging from devices, sensors, and metadata to spatial topology. These ontologies enable semantic interoperability, automated reasoning, and system integration. Although other models, such as Project Haystack, offer novel features like tagging, their limited use prevents a more thorough analysis. Schemas based on XML and UML are excluded from this core set due to their poor alignment with linked data principles and their lack of native semantic support. [16.]

3.6.1 BOT

In 2019, the Linked Building Data Community Group, under the W3C Consortium, produced the Building Topology Ontology (BOT), a lightweight yet extensible semantic framework for documenting the structural and spatial composition of buildings. BOT is notable for employing semantic web technologies and linked data to facilitate AECO interoperability despite not being a W3C standard. [17.]

BOT focuses on building topology, covering zones (sites, buildings, storeys, and spaces), elements (systems, equipment, and sensors), and their interactions. BOT's core class `bot:Zone` models spatial volumes and hierarchies. The links between these zones, established by `bot:containsZone` and `bot:adjacentZone`, enable coherent spatial modelling. [17.]

Another essential idea is `bot:Element`, which represents the building's structural, control, and other physical components. For a networked and integrated picture of the architecture, the ontology includes the class `bot:Interface` to model spatial intersections or connection points between zones and elements. [17.]

Despite its simplicity and lack of domain coverage (such as product specifications, sensor data, and project management), BOT is compatible with

Brick, SAREF, and SSN/SOSA. Its modular design allows it to be customised for energy management, building operations, and facility monitoring. [17.]

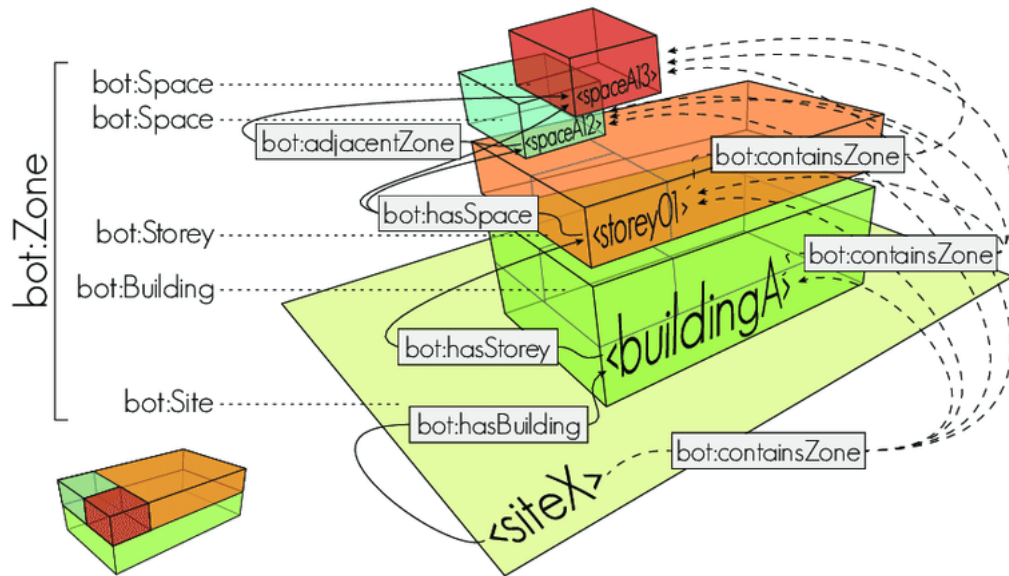


Figure 5. BOT Zone Connectivity [17.]

Figure 5 shows the spatial relationships and hierarchy of the Building Topology Ontology (BOT). It displays a location (siteX) with a building (buildingA1), a storey (storey01), and several spaces (spaceA1, spaceA2, and spaceA3). In the BOT framework, relationships such as bot:hasBuilding, bot:hasStorey, bot:hasSpace, bot:containsZone, and bot:adjacentZone demonstrate how zones and spatial elements are semantically connected.[17.]

3.6.2 SAREF

A core model and several domain-specific extensions, including SAREF4BLDG for buildings, SAREF4ENER for energy, and SAREF4SYST for system hierarchies, are features of SAREF, a semantic ontology created by ETSI and the European Commission to facilitate interoperability among IoT devices. Reliable data exchange in SAREF-based semantic models is achieved through the use of TCP protocols for networked device communication.[17.]

To improve building management, SAREF4BLDG focuses on building devices and spaces, partially aligning with ontologies such as BOT and IFC. Smart grids

can use energy more efficiently thanks to SAREF4ENER's scheduling and power profiles, which facilitate demand response and optimisation. SAREF4SYST helps set up systems by illustrating how devices in innovative environments interact and connect. All things considered, the ontology places a strong emphasis on modularity and interoperability, making it easy to integrate with smart building and energy systems. [17.]

3.6.3 Project Haystack

Project Haystack, a 2014 non-profit, develops and maintains the open-source Haystack Ontology. Siemens, Intel, SkyFoundry, Legrand, and others make up its industry board. Haystack's semantic data model for the building automation industry uses dictionaries of name–value pairs to identify building concepts and define their meanings uniquely. Combining *Defs*, or definitions, in libraries can express complex concepts. Hierarchical concepts are linked using parent–child structures, while instance relationships are stated via *Refs* (pointers) or child protos. In project haystack Zinc format is used to serialise Haystack's data into structured text-based representation known as Trio format, and the Filter language is lightweight query syntax within haystack that allows users to retrieve or match specific data points based on tags and their values. Haystack 4 added formal hierarchies and relationships to the ontology, which still doesn't support semantic web standards like RDF or OWL. This deficiency often causes incorrect or inaccessible data representations. The Haystack Tagging Ontology (HTO) formalises tag usage and expands vocabulary for interoperability, utilising semantic web standards such as RDF, OWL, and SPARQL. [17.]

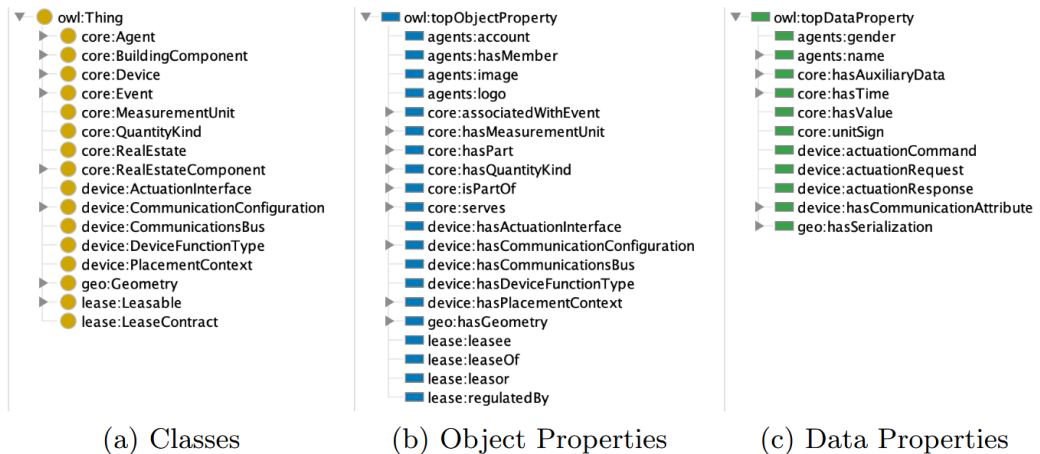
```
id: @whitehouse
dis: "White House"
site
area: 55000ft²
geoAddr: "1600 Pennsylvania Avenue NW, Washington, DC"
tz: "New_York"
weatherStationRef: @weather.washington
```

Listing 2. Example of a Haystack Ontology Instance in Trio Format. [18.]

Listing 2 shows a Haystack Trio site entity. The site's URL is identified as *@whitehouse* and named "White House" (dis). With a geographical address of "1600 Pennsylvania Avenue NW, Washington, DC" and 55,000 square feet, it is a site. "New_York" is the time zone, and *weatherStationRef* points to *@weather.washington*. It shows how to semantically model a physical building item in Haystack using name–value pairs (tags and definitions).[18.]

3.6.4 RealEstateCore

RealEstateCore ontology is a modular OWL 2 framework designed specifically for real estate companies that work with smart buildings and need to integrate data from various sources. It enables other building systems to work together semantically by providing standard representations for essential elements, such as buildings, rooms, sensors, actuators, leases, and organisations. RealEstateCore helps with crucial tasks, such as optimising energy use and analysing presence, by utilising concepts like *QuantityKind* (what is measured), *PlacementContext* (where it is measured), and *DeviceFunctionType* (what the device does). Its design focuses on being straightforward, easy to use, and easy to maintain over time, and it adheres to standards such as Brick, SSN/SOSA, and QUDT when applicable. RealEstateCore standardises data from various building automation systems, enabling extensive analysis and optimisation of real estate portfolios. Major Nordic real estate firms are already using it. [19.]



Listing 3. RealEstateCore top-level constructs. [19.]














Listing 3 shows the real estate core structure, which is divided into three sections: classes, object properties, and data properties.[19.]

3.6.5 SSN/SOSA T

The Sensor, Observation, Sample, and Actuator (SOSA) ontology was combined with the original Semantic Sensor Network (SSN) ontology to make the present SSN/SOSA ontology. This new version makes it easier to model various aspects, including the properties of devices, the sensors' observations of the environment, the processes involved, the observations made, the actuators (which trigger actions), and the sensors themselves. It also provides a complete framework for describing sensor-based systems.[16.]

SOSA makes it easier to explain sensors and actuators that perform actions or respond to stimuli. SSN, on the other hand, examines entire systems and how their components work together, much like an HVAC system. SOSA serves as a basic framework for real-world interaction modelling because it demonstrates how sensors detect physical events including temperature and window position and how actuators execute specific actions based on these observations. Table 1 summarises the six most essential ontologies used in a built environment and their primary areas of focus.[16.]

Table 1. An overview of the key ontologies discussed above. [16.]

Ontology	Year	Uses/Extends	Classes	Formats	Scope	Main Focus
BOT	2017	None	7	T	 	Building topology
SAREF	2015	None	28	T	 	Energy efficiency
SSN / SOSA	2019	None	19	T		Sensors and measurements
Brick	2016	PH	5	T	  	Building entities
PH	2014	None	479*	N/A	  	Building info. from devices
RealEstateCore	2019	SSN / SOSA	141	T	 	Buildings

Scopes:  Home,  Smart Device,  Building

Formats:  Turtle

Table 1 summarises the six most essential ontologies used in a built environment and their primary areas of focus.

3.7 Research Gaps

Buildings in the United States comprise 35% of energy usage and thirty-eight per cent of energy-related greenhouse gas (GHG) emissions [20], so a significant step is needed for the transformation in the Architecture, Engineering, and Construction (AEC) sector, where standardised models like Brick can streamline design and operations. Brick was developed to standardise building data and enhance energy efficiency and operations management, ultimately saving energy while also contributing to the occupants' health, comfort, and productivity. Brick has modelled hundreds of buildings and received validation from various buildings worldwide. It covers most of the HVAC and lighting concepts included in any current BMS, exactly as specified initially. Therefore, it is not too difficult to import data from preexisting BMS systems. [21.]

Nevertheless, Brick lacks comprehensive classes and terminology suitable for data center infrastructure. PDUs, liquid cooling, high density (e.g., GPU racks), UPS, and detailed power/cooling flows are examples of elements that Brick does not include. For the accurate representation and analysis of energy systems in a data center, Brick needs to be extended to domain-specific vocabulary. *The terms brick:Generator_Room and brick:Transformer_Room have been deprecated, and their replacements have not yet been included in the Brick or REC ontology.* [22.] The addition of classes and relationships related to specialised data center equipment would enhance Brick's ability to model energy systems. Data centers generate high-frequency, high-volume telemetry from thousands of sensors across IT power and cooling systems. Brick was designed primarily for building metadata and does not support real-time time-series data. [23.] As a result, there is a need for research into scalable methods for linking rapid telemetry to ontology instances, storage-efficient representation for large-volume data, and real-time mapping from building management systems to brick-based knowledge graphs. [24.] Brick is currently unable to model waste heat flow or connections to external networks explicitly, nor can it support renewable energy sources in direct relation to EED reporting. This is because the most recent laws from the European Union focus on waste heat recovery and the integration of renewable energy. [25.]

Although Brick is recognised as a robust semantic framework for modelling building systems and assets, it has yet to see widespread use in practice. Its ability to unify disparate data sources and provide machine-readable energy reporting remains untouched mainly in this sector, leaving an opportunity to explore various ontologies.

This thesis aims to close inefficiency and inaccuracy gap by proposing and testing a straightforward method for applying Brick in a data center. It will outline a scalable, user-oriented approach for representing essential data-center elements—such as servers, power distribution, cooling systems, and environmental sensors—and linking their telemetry data to gain deeper insights into energy use and operational performance.

Prior research has examined data-center performance indicators and created semantic ontologies for building systems; however, none have integrated these strands to meet the EU Energy Efficiency Directive reporting criteria. Ontologies lack full classes for data-center-specific equipment, inadequate support for automated KPI computation, or clear pathways for aligning heterogeneous sensor data with regulatory templates. Many contributions only describe ontological models without demonstrating how they minimise manual effort or enhance reporting accuracy. Developing a framework that (1) models data center infrastructure in accordance with EED requirements, (2) integrates time-series data for automated PUE, WUE, ERF, and REF calculation, and (3) validates for regulatory compliance is lacking. This thesis tackles gaps by using Brick and RealEstateCore to develop a scalable semantic model and implement an end-to-end pipeline for data integration, KPI computation, and EED-compliant reporting.

By adapting and implementing Brick in this manner, this research aims to demonstrate how it can facilitate standardised energy reporting and promote regulatory compliance for data centers, ultimately contributing to the transition toward more sustainable and efficient digital infrastructure. Table 2 below maps EED data center energy efficiency metrics to Brick support. Calculating PUE and WUE is possible with existing meters, but they lack dedicated computation classes. For calculating ERF, classes such as `Brick:District_Heating` and waste heat flow properties are needed to model recovered heat pathways. For CUE, linkage to carbon-intensity datasets is missing, and for REF, generation classes such as `brick:Hydropower_Turbine` and `brick:Wind_Turbine` extensions are required.

Table 2. Mapping EED data center efficiency metrics to Brick support.

<i>EED Requirement</i>	<i>What Does EED Need?</i>	<i>Brick Current Support</i>	<i>Gap / Limitation</i>
<i>PUE</i>	IT load power/Total facility power	<i>brick: Energy_Sensor</i> , <i>brick: Meter</i> and <i>brick: isPointOfEquipment</i> . <i>brick: Power_Sensor</i> gives total power.	Missing PUE calculation classes: manual computation required
<i>CUE</i>	Total CO ₂ emissions (kg) / Total IT energy (kWh)	There is no built-in linkage between energy sensors and carbon intensity datasets.	Cannot directly calculate CUE from the model alone
<i>ERF</i>	Total reused energy / total energy consumed	<i>brick: Heat_Exchanger</i> and <i>brick: Heat_Recovery_Hot_Water_System</i>	Waste heat flow properties & <i>brick: District_Heating</i> are missing
<i>REF</i>	Renewable share of total consumption	<i>brick: Electrical_Meter</i> <i>brick: PV_Panel</i> <i>brick: Solar_Panel</i> <i>brick: Thermal_Power_Meter</i>	Missing classes <i>brick: Hydropower_Turbine</i> <i>brick: Wind_Turbine</i>
<i>WUE</i>	Water use / IT energy use	<i>brick: Water_Meter</i> <i>brick: Water_System</i> <i>brick: Water_Level_Sensor</i> <i>brick: Water_Flow_Sensor</i> <i>brick: Electrical_Meter</i>	Missing WUE calculation classes: manual computation required

4 Methodology

4.1 Overview

This methodology chapter assesses whether Brick can serve as a comprehensive framework for modelling data center infrastructure and facilitating EED-compliant reporting. It comprises four stages—data collection, semantic modelling, RDF conversion, and compliance reporting—chosen to directly address the objectives of accuracy, automation, and regulatory alignment. This sequence follows the EED report process shown in figure 6 below:

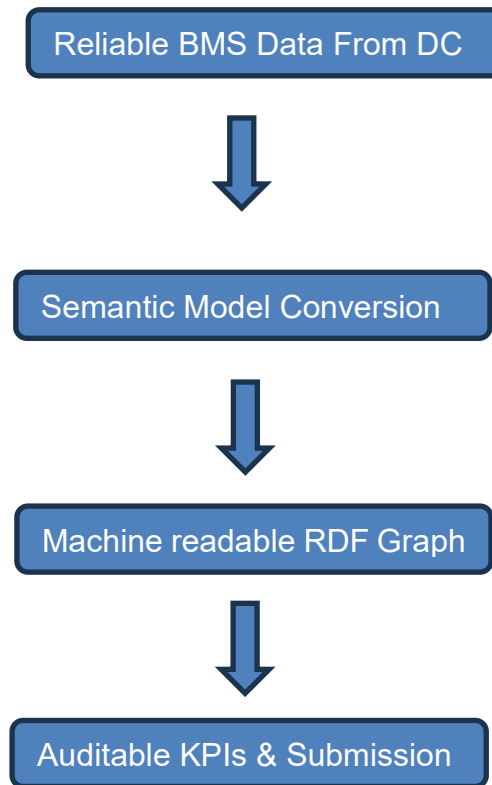


Figure 6: The EED reporting workflow illustrating the four-stage methodology.

Figure 6 illustrates the process of collecting data, converting the semantic model, generating RDF graphs, and reporting KPIs, all based on the integration of the Brick ontology for energy compliance in data centers.

4.2 Research Approach

The research approach section explains the reasoning behind selecting Brick as the backbone for the semantic framework in this work, due to its mature class hierarchy for building systems while mitigating current gaps in data-center-specific concepts. The goal is to demonstrate that utilising Brick's existing features, while addressing its shortcomings, leads to the most feasible approach for developing a working compliance reporting prototype without creating a new ontology layer. The data center energy consumption dataset has been linked to EED reporting standards by extending the ontology's existing schema in this thesis. I have utilised the energy consumption records from the Data center,

located in Finland, as the primary method for producing EED-compliant reports. This thesis uses an energy consumption report from the Data center, generated at five-minute intervals, to deliver the EED report. The reason for adopting Brick is outlined in table 3. Brick, accompanied by Real Estate Core (REC), offers the most comprehensive coverage among other ontologies and schemas. [26.]

Table 3: Comparison of Brick with other schemas. [26.]

Modelling Support	Brick	Project Haystack	IFC	BOT	SAREF
HVAC Systems	Full	Full	Full	None	None
Lighting Systems	Full	Limited	Full	None	None
Electrical Systems	Full	Full	Full	None	None
Spatial Information	Full	None	Full	Full	None
Sensor Systems	Full	Full	Basic	None	Full
Control Relationships	Full	None	Basic	None	None
Operational Relationships	Full	None	Basic	None	None
Formal Definitions	Full	None	Full	Full	Full

Table 3 shows that Brick has clearly taken the lead in terms of modelling support for smart buildings.[26.] The methodology that I will be implementing is described in Figure 8 below. Furthermore, this process is designed to assess whether modelling a data center using Brick can simplify the generation of regulatory reports and improve compliance with the European Union's new energy-efficiency data and data-center sustainability rules. It illustrates the end-to-end workflow, with sections 3.3 to 3.8 explains in detail.

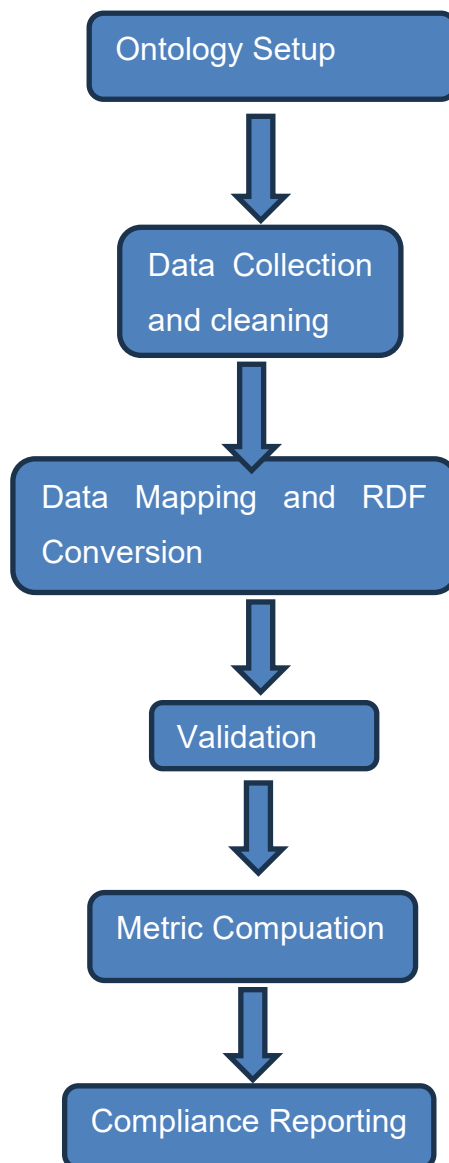


Figure 7. The workflow process shows how Brick links raw IoT data to EED KPI reporting.

4.3 Ontology Setup

The process of accurate compliance reporting requires integrating various inconsistent metadata from data center systems. The ontology setup phase standardises heterogeneous BMS/IoT metadata in Brick, which improves interoperability, reduces manual data preparation, and supports automated queries. The standardisation of semantic layers in this step is essential because it underpins the computation of automated EED reporting KPIs. [10]. Brick is a formal metadata representation designed to bridge this gap by focusing on completeness encompassing the necessary concepts for applications, expressiveness encompassing the requisite relationships and properties for applications, and extensibility the ability to define new concepts as needed. This amalgamation of attributes renders Brick an enticing metadata repository for data-centric applications. Applications conduct queries on a Brick model to obtain the necessary metadata for configuration and operation. [27.]

Building Management Systems (BMS) are the most common source of digitised sensor and actuator metadata in operational data centers. However, the point labels in these systems, which are used to denote sensors, alarms, controllers, and actuators, are typically inconsistent, unstructured, and vendor-specific, making automated reasoning and analysis difficult. Moreover, these point labels contain only partial semantic information, such as the equipment's name, zone, or measurement type, without the standardised definitions that would enable interoperability with analytics platforms. This is especially true when attempting to scale up such operations across multiple sites or vendors.[27.]

This thesis utilises Brick version 1.4.4, with namespaces as shown in table 4 below:

Table 4: Ontology namespaces and their corresponding Uniform Resource Identifiers(URIs)

Prefix	Namespace URI
Brick:	https://brickschema.org/schema/1.4/Brick#
Bf:	https://brickschema.org/schema/1.4/BrickFrame#
Rec:	https://w3id.org/rec#
Qudt:	http://qudt.org/schema/qudt/
Unit:	http://qudt.org/vocab/unit/
Time:	http://www.w3.org/2006/time#
Ex:	http://example.com/project#

Project namespace ex: as a semantic standard layer to unify heterogeneous, unstructured BMS metadata. To overcome the challenge, Brick will serve as a standard semantic layer for managing inconsistent BMS metadata. The following diagram, Figure 8, is adapted from the OpenBuildingControl and Mortar frameworks, illustrating how, in this architecture, Brick serves as a broker to link design models (e.g., IFC from BIM), simulation tools (e.g., EnergyPlus, Modelica), and open-building systems (e.g., BACnet-based Building Management Systems, BMS). In the whole model, metadata from points in BMS is first normalised using rule-based translation or machine learning methods into standardised Brick classes (e.g., *brick:Power_Sensor*), linked to time series streams of data in Mortar-compatible databases. Figure 9 illustrates how Brick is used to normalise building metadata and timeseries data from design, modelling, control, and BMS systems, and put it into a shared data store (Mortar), which enables structured analytics. [28.]

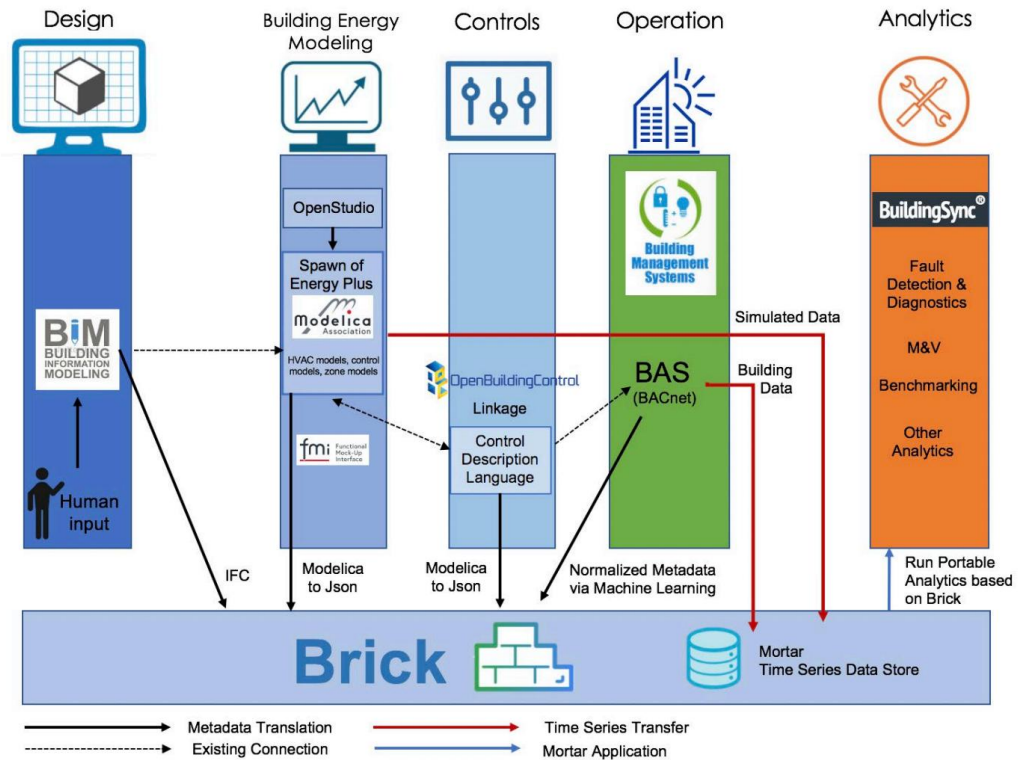


Figure 8. Integration of BIM, Building Energy Modeling, Controls, and Analytics using Brick and Mortar.

For this thesis, the Excel datasets extracted from the data center server racks and panels, as shown in figure 8, represent heterogeneous, unstructured BMS exports. By parsing sensor IDs, assuming it to be 3A07_ES401010007, as in the example below, the corresponding equipment and measurement types were identified for Brick class mapping:

```
ex:IT_Power_Sensor_3A07 a brick:Active_Power_Sensor;
brick:isPointOf ex:IT_Power_Meter;
brick:hasUnit unit:kiloW .
```

Listing 4. Brick class mapping method.

By associating the data with their corresponding panels and functions (e.g., IT load, cooling infrastructure), we constructed a semantically enriched model using Brick classes and units. This data is then linked to high-resolution time series

data (one-hour interval), allowing compliance with the EU Energy Efficiency Directive's structured queries and energy analysis indicators.

4.4 Data Collection

The EED requires data centers to collect precise data about energy usage and other relevant metrics at a high frequency for accurate energy reporting. The original data collected at a fifteen-minute interval to ensure accurate energy performance metrics for precise energy reporting. However, because the values remain relatively steady for an extended period, the dataset is then aggregated into one-hour time intervals to better illustrate the highs and lows over a 24-hour cycle.

The production of valid compliance metrics through semantic models depends on reliable source data. To evaluate and calculate key performance metrics related to its operational efficiency, an energy consumption dataset from the data center's server racks (`rec:ITRack`) was retrieved. Nevertheless, the dataset provided by the data center lacked meter and sensor IDs, which made it impossible to calculate all the metrics mentioned in the Energy Efficiency Directive (EED). Additionally, mapping the dataset to the correct Brick classes presented challenges, and certain assumptions were made.

As a result, some classes may require future extensions to accurately map the datasets into Brick. Because the available dataset did not contain detailed identifiers for all meters and sensors, it was not feasible to calculate every EED-defined metric. Therefore, this study focuses explicitly on calculating Power Usage Effectiveness (PUE), Energy Reuse Factor (ERF), Energy Reuse Efficiency (ERE), Water Usage Effectiveness (WUE), IT energy and the total energy consumed within the data center over a one-hour time interval. Figure 11 below displays the energy consumption data directly from the BAS system of the data center.

The current dataset lacks sufficient information for technical analysis and reporting because the absence of sensor and meter IDs hinders data traceability, consistency verification, and energy use comparison between different systems or facilities, which are essential for data center regulatory compliance and effective energy management. Because of this, assumptions regarding energy readings corresponds to `brick:Power_Meter`, non-IT electrical consumption was mapped to `brick:Electrical_Meter` and IT load readings were assigned to `brick:Energy_Sensor`. The raw dataset contained major structural and formatting issues for example some power values are recorded as numbers(1250) and others as strings with units(1.25MW). Another example was readings are at 15-minute intervals but some rows were duplicated or skipped and timestamps were formatted inconsistently, which made systematic analysis and semantic modelling difficult. The energy consumption measurements received complete data cleaning and standardised treatment through OpenRefine, as this tool specialises in handling disorganised datasets. The advanced clustering algorithms and faceting features of OpenRefine enabled me to identify and correct mistakes, duplicate entries, and formatting issues that frequently occur in large datasets related to energy monitoring.

The data transformation process utilised OpenRefine's automated reconciliation tools to standardise time formats and normalise measurement values, creating uniform data structures that adhered to semantic web standards. The systematic cleaning process transformed the unorganised dataset into a well-structured format that supported both statistical analysis and Brick semantic modelling, as shown in table 4. The preprocessing step became vital for maintaining data quality standards in energy efficiency reporting and for reliable integration with building automation systems' ontologies.

4.5 Data Mapping & RDF Conversion

After standardising the raw Excel data from meters and sensors, it is converted into a machine-readable, standardised format, which enables automated querying and reporting. Mapping raw data into Brick classes and converting it to RDF creates links between data points and their corresponding assets, units, and timestamps, which enables automated metric generation. This process directly supports the objective of reducing manual effort and eliminating errors in EED report preparation.

In this phase, the collected energy performance data from the data center—including PUEs, ERFs, WUEs, and EREs—recorded at one-hour intervals, was translated into a structured, machine-readable format using the Brick. This approach models the energy infrastructure and sensor measurements of the facility in a way that can automatically generate reports and increase data interoperability, while also making other processes easier for workflow.

Initially, physical assets critical for data center's operations were described using the appropriate Brick classes. This included concepts such as IT. Energy is measured by `brick:Energy_Sensor`, is related to equipment through `brick:isPointOf` equipment property. The property `brick:hasSubMeter` indicates a submeter of `brick:Meter`, `brick:isMeteredBy` indicates the meter measuring the consumption from `brick:Meter`. All these classes falls into the category `brick:Electrical_System` as they are part of the electrical subsystem in the building which is a subclass of `brick:System` and connects its equipment, points, space (`rec:space`) using `brick:hasPart` relationship. Furthermore, `brick:Water_Meter` is used for measuring water consumption in the data center, which utilises `brick:hasPoint` followed by `brick:Water_Flow_Sensor` point to calculate the value from `brick:Water_Meter`. `brick:ICT_Equipment` and it has property `brick:hasLocation` `rec:Space` for describing the location of the IT equipment. In the server room of the data center, server racks can be mapped

to `brick:ICT_Rack`. The object properties used to describe the relationship between these entities were chosen by utilising a few standard Brick object properties (`brick:hasPoint`, `brick:hasPart`, `brick:feeds`, and `brick:isFedBy`) provided by the ontology.

For calculating KPIs, Brick should be extended to include the following Brick concepts `brick:Power_Usage_Effectiveness`, `brick:Energy_Reuse_Effectiveness`, `brick:Energy_Reuse_Factor`, and `brick:Water_Usage_Effectiveness`, to align with the EED requirements. In addition to these, all meters and sensors were also annotated with their measurement units using QUDT unit definitions through *the* `brick:hasUnit` property. Timeseries data references were linked to each point using the `ref:TimeseriesReference`, which has the `ref:hasTimeSeriesId` property.

The Brickschema is written as shown in listing 5.

```
class ref:TimeseriesReference extends ref:ExternalReference
{
  + ref:hasTimeseriesId xsd:string;
  + ref:storedAt rdfs:Resource;
}
```

[22]

Listing 5:Timeseries class

During this phase, several Python libraries were installed to support data transformation and semantic modelling. These included:

- `Brick[all]` - for working with the Brick
- `rdflib` - for working with RDF and creating RDF triples
- `pandas`-for processing and handling CSV data
- `matplotlib` for generating visualisations
- `networkx` - for visualising RDF graphs

Once the required environment was set up in Visual Studio Code (VS Code), the raw energy data, which was formatted as an Excel file, was then transformed into RDF triples using Python and RDFLib. Next, the dataset column containing IT

power was converted into RDF triples for `brick:Energy_Sensor` and `brick:Electric_Power_Sensor`, which represent a timestamped metric reading. This involved connecting the measurement values to the correct sensor entities using triples such as `brick:hasValue` for the recorded metric, `brick:hasTime` for the timestamp, and `brick:hasUnit` for the unit of measurement. The model was written in Turtle (.ttl) format to create a complete RDF-based documentation of the energy performance within this data center.

Converting this data to RDF enables us to run semantic queries and generate compliance-based reports. For example, stakeholders have the option of creating extracts using SPARQL e.g., an average PUE over a specific period or identifying unusual system behaviours to detect potential inefficiencies. Brick also helps maintain a consistent structure of information, ensuring that data can be added and modelled separately. It also provides an opportunity to extend the model together if new data becomes available.

Furthermore, this approach enables the system to be easily integrated with Building Management Systems (BMS) that communicate via protocols such as BACnet, MQTT brokers, or KNX devices. APIs enable seamless integration and efficient data exchange between external systems. Sensors for temperature, humidity, carbon dioxide, and other variables can be integrated into a single semantic framework, allowing for access to a comprehensive view of building operations and enhancing fault detection and real-time monitoring. [26.] The data extracted in the Excel format is from the Internet of Things (IoT) sensors. The turtle file is also checked using GraphDB software to verify the relationships between different Brick and REC ontology entities, as shown in figure 12.

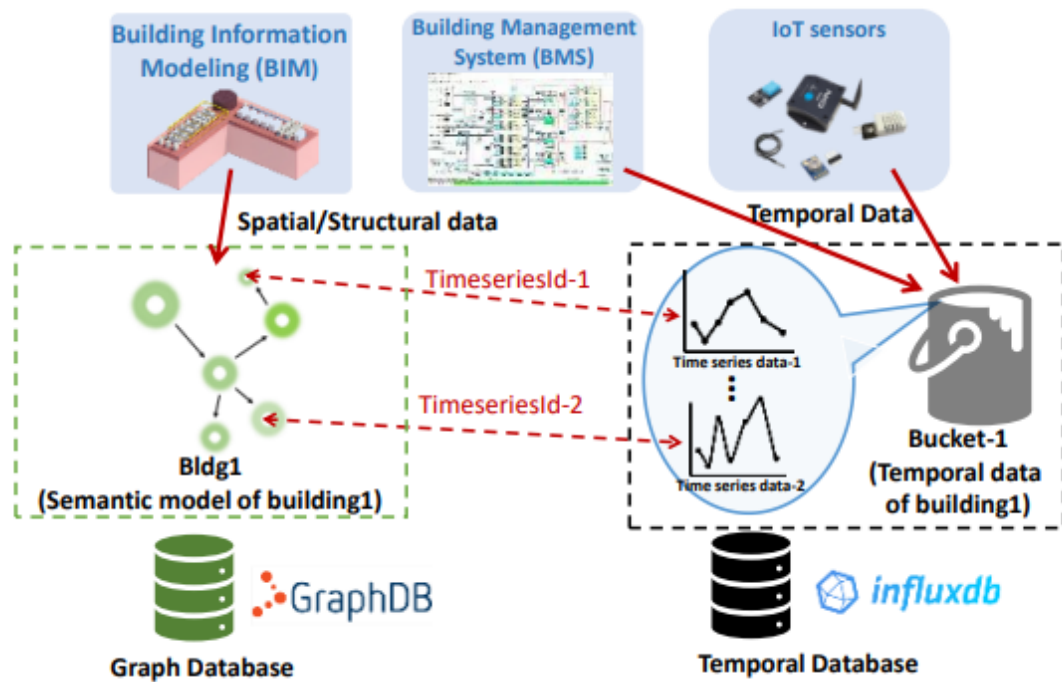


Figure 9. Multiple sources of information management using Brick [29].

The schematic in figure 9 represents a Brick-oriented multi-source data integration effort. These integration activities include creating a Brick model based on the spatial/structural information extracted from the BIM model, obtaining temporal data from BMS and IoT sensors, and mapping spatial and structural data with data in the Brick model through TimeseriesId. The data was generated and mapped to the Brick using a similar method, as shown in the above figure.

For accurate energy reporting, real-world raw metadata from the building must be transformed into a standard, structured graph that represents the building's physical spaces, equipment, sensors, and their relationships according to Brick's classes and properties, as shown in table 7.

Table 5: Converting a real-world energy consumption report to Brick Point and Relationship.

Data Column	Brick Point/Relationship	Description
IT Power	brick:hasPoint -> brick:Energy_Sensor brick:hasUnit unit:KiloW-HR	Measures peak hourly real power consumption of IT equipment.
Total Power	brick:hasPoint->brick:Power_Sensor ; brick:hasUnit unit:KiloW-HR	Measures instantaneous total facility power.
Water Consumption	brick:hasPoint -> brick:Water_Flow_Sensor brick:hasUnit unit:L-PER-HR	Measures water flow rate consumption

The relationships among IT power, total power, and water consumption are represented in the listing shown below.

```

:DC_IT_Hourly_Peak_Energy_Sensor a brick:Energy_Sensor ;
  brick:hasUnit unit:KiloW-HR ;
  brick:aggregate [
    brick:aggregationFunction "max" ;
    brick:aggregationInterval "RPT1H"
  ] ;
  brick:powerComplexity [
    brick:value "real"
  ] ;
  brick:isPointOf :IT_equipment .

:DC_Hourly_Total_Power_Meter a brick:Power_Sensor ;
brick:hasUnit unit:KiloW-HR ;
brick:aggregate [
brick:aggregateFucntion "max" ;
brick: aggregationInterval "RPT1H" ;
];
brick:powerComplexity [
    brick:value "real" ;
]
.

:DC_Water_Flow_Sensor a brick:Water_Flow_Sensor ;
  brick:hasUnit unit:L-PER-HR ;
  brick:isPointOf :water_meter .

:water_meter a brick:Water_Meter ;
  brick:hasPoint :water_flow_sensor ;
  brick:meters :DC .

```

Listing 6: Relationships among IT, total power and water consumptions.

Each existing real-world item, such as a room, chiller, electricity meter, or temperature sensor, becomes a node, which is an instance of a Brick Class (e.g., `brick:Room`, `brick:Chiller`, `brick:Electric_Meter`, `brick:Temperature_Sensor`). Connections between nodes represent physical or logical relationships, using Brick's defined properties, such as `brick:hasLocation`, `brick:hasPoint`, and `brick:feeds`, among others. After this mapping, the schema forms a graph where:

- Equipment (e.g., a chiller) has sensors and setpoints as points.
- Location relationships link spaces (zone, floor, room).
- Logical flows (e.g., an electric meter feeding specific equipment) are encoded.

The snippets of Python code for generating a turtle file of a data center are shown below in listing 6, and the complete code is listed in Appendix List 1. After that, the RDF is visualised using Brick Studio (<https://Brick.github.io/brick-studio/>).

As shown in appendix 8, this Python script generates a semantic model of a data center utilising Brick and Real Estate Core (REC) ontologies. It initialises an RDF graph, defines namespaces (Brick, Rec, and Ex), and sets up ontology imports. Then, it creates a building representing the data center and assigns it both a `brick:building` and `rec:DataCenter` type. It then creates three major subsystems `brick:Electrical_System`, `brick:Cooling_System` and `brick:Water_System` and links them to the building via the `brick:isPartOf` relationship.

By highlighting important elements and their linkages, the RDF graph in page 56 illustrates how the data center building can be modelled using semantic modelling. `IT_Zone` has `IT_Rack` and `Server_Rack`. Infrastructure components, such as `UPS1` and `PDU1`, are linked through power distribution relationships. The clustering system represents both physical and functional aspects of the facility by accurately mapping equipment to spaces and zones. The visualisation functions as both a validation and communication tool within the methodology to verify that RDF triples match Brick classes and relationships, while providing stakeholders with clarity on the model structure. The graph allows stakeholders to perform specific semantic queries, which are necessary for EED-compliant energy reporting, such as retrieving all meters in the IT zone and following the energy flow from the UPS to individual racks as shown in figure 18. The design is scalable; additional sensors, equipment, or spatial zones can be integrated by simply extending the RDF graph with new triples. [30.] This allows for the approach to remain consistent with Brick's ontology while simultaneously responding to the ever-changing requirements of the data center.

A multi-stage validation process was employed to ensure the syntactic and semantic integrity of the RDF model. This process combined RDF syntax verification, vocabulary compliance checks, and optional SHACL-based conformance testing [30]. This method helped to detect problems at an early stage while maintaining the standards of the Brick vocabulary and testing the SHACL shape-based formal constraints. The Python validation script developed

using Brick, rdflib, and pySHACL [31] executes validation through three distinct phases.

The RDF syntax verification process handles the Turtle file data center model. Still, it stops execution immediately upon detecting any syntax errors, including improper triples or incorrect delimiters, with an exact error notification.

The detection of undefined classes occurs through a SPARQL query which extracts distinct model classes and then verifies them against the Brick namespace (<https://Brick.org/schema/Brick#>). The system flags all unrecognised or misspelt classes, which helps maintain the use of valid Brick terminology.

The validation process becomes optional when an SHACL shapes file is provided to pySHACL with RDFS inference enabled for model validation. The validation process utilises the Brick SHACL specification to verify class-property compatibility, along with required property cardinalities and data type constraints.

The sequential validation procedure exceeds standard SHACL validation to detect vocabulary-level errors, which would otherwise escape detection. The script execution, as displayed in figure 11 above, shows valid RDF syntax along with Brick vocabulary-matching classes and SHACL validation success when enabled.

Validation runs as a pre-commit step; if there are failures, they block graph export and include the offending triple, the violated shape, and a remediation hint. For example, a *brick:Power_Meter* without any *brick:hasPoint* is considered a violation. The output displays three pieces of information when validation fails, including the problematic triple or class identification, along with SHACL constraint violation details and a text description for resolving the issue. The automated process maintains query reliability and enables EED-compliant reporting from a properly defined standards-compliant model. The code is shown in appendix 2.

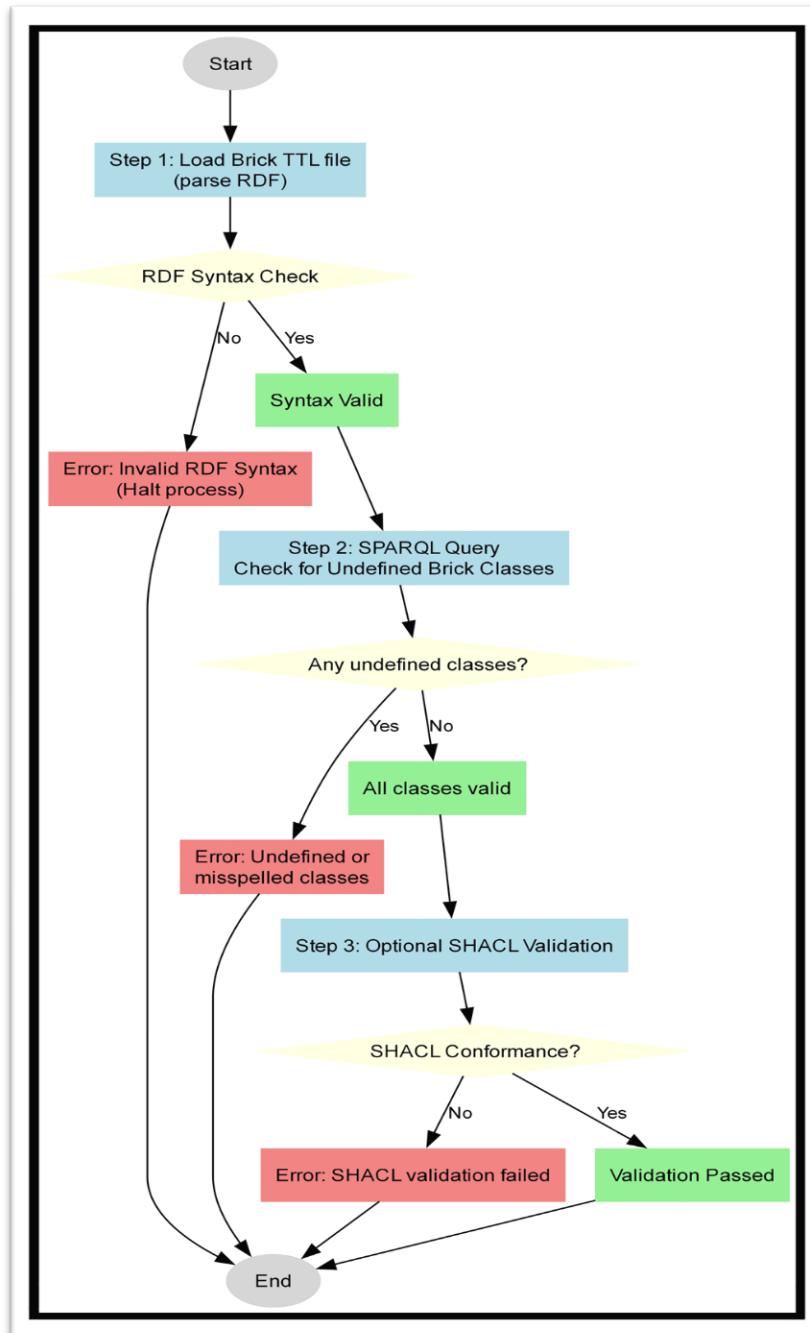


Figure 10. Multi-stage validation workflow for the Brick RDF model.

As shown in figure 10, the process begins by loading the model while checking the RDF/Turtle syntax for validity. The following query displays all classes used in the system and checks their membership in the Brick namespace to detect both spelling errors and terms outside the vocabulary. The pySHACL system verifies class–property compatibility and required cardinalities and datatypes through RDFS inference when a SHACL shapes file is provided. The “Validation

Passed” outcome indicates `sh:conforms = true`, but the report shows offending nodes and constraints for fixes when this outcome is not achieved.

4.6 Metric Computation

During this stage, the Brick model receives direct SPARQL [32] queries to calculate Power Usage Effectiveness (PUE). The `brick:Electric_Meter` `brick:Power_Meter` instances provided facility and IT electrical measurement data. The query shown below in Figure 18 starts at two reference nodes: `FacilityRoot` and `:ITRoot` before using `(brick:isPartOf|brick:hasPart) / brick:isMeteredBy` to find all relevant meters.

Three sub-queries calculate the necessary values for a specified time period to obtain total facility energy (E_{fac}) and total IT energy (E_{it}). The outer SELECT statement calculates the metrics through the standard formula $PUE = E_{fac} / E_{it}$ after converting all values to `xsd:decimal` for proper numeric division. The measurements underwent uniform processing with a 5-minute sampling interval, followed by unit normalisation (kWh for energy). The declarative approach enables site-independent and adaptable metric computation for building topology changes and metering configuration modifications, the code is shown in appendix 3.

4.7 Compliance Reporting

A Python script, shown in appendix 3, generates compliance-grade reports by mapping raw meter streams to Brick entities, harmonising units/intervals, and computing KPIs. A mapping table connects each Excel column to a `brick:Meter`, allowing every data point to be traced back to its corresponding typed asset. The workflow accepts time-stamped measurements while standardising units and time intervals before calculating total values for the reporting period. The system utilises SPARQL to calculate key indicators, including PUE (power usage effectiveness, or facility/IT energy) and WUE (water usage effectiveness, or water/IT energy), before presenting them in a lightweight

Flask dashboard for review. The application generates flat files (CSV/JSON), which can be used for submission and archiving purposes alongside site metadata, including the site name, owner/operator, floor area, installed power, and data lineage. The same workflow becomes reusable across sites through Brick classes and relationships, rather than hard-coded meter IDs, which maintains auditability and reproducibility.

The generated web interface creates an EED Compliance Report for the site in Finland as illustrated in appendix 7 chapter 5.2) and presents three key sections in a single view. The first section of the report contains data center information, which includes site metadata, such as name, owner, operator, municipality, floor area, installed power, and annual data traffic or storage figures, with “N/A” placeholders for missing data. The second section of the report presents calculated KPIs essential for EU Energy Efficiency Directive compliance, including PUE, ERF, REF, CER, CUE, and WUE. In this case, placeholder values are displayed because the dataset contains readings from the Excel file. A download link for the annual energy report in Excel format is also included to support archival and regulatory submission. The interface is configurable via environment variables and a YAML configuration file, ensuring compliance with EU regulations.

5 Results

The Brick framework is applied to the data center case study, utilising simulated data based on real-world patterns. Data is synthesised from typical data center metrics, such as those in theory [e.g., average PUE 1.2-1.5], as the data center’s dataset was limited. The simulation assumes five-minute intervals for energy readings from IT racks, cooling systems, and facility totals. Outputs demonstrate automated KPI computations, visualisations, and compliance reports, validating the framework’s accuracy and efficiency for EED reporting.

5.1 Case Study Description

The data center located in Finland, operates as a high-performance computing facility managed by the one of the prestigious company in Finland. The facility enables scientific research through its powerful computing capabilities, while utilising substantial amounts of energy (10-20 MW peak) across its IT subsystems (servers and GPU racks), cooling systems (liquid-cooled systems with district heating waste heat recovery), power subsystems (UPS and PDUs), and networking subsystems. This research focuses on modelling a specific system consisting of 10 IT racks (brick:IT_Rack) with servers (brick:Server) and two cooling units (brick:Chiller and brick:Heat_Exchange), and power and water measurement devices (brick:Electric_Meter for IT/facility power and brick:Water_Meter for cooling water).

The simulated CSV exports replicate BMS/IoT sensor data at five-minute intervals throughout 24 hours (IT power ranges between 8 and 12 MW while following load patterns, and facility power ranges between 10 and 15 MW, including overhead). The RDF model was generated through a Python script (Appendix 1) which assigned sensors to Brick classes (e.g., :IT_Power_Sensor a brick:Electric_Power_Sensor; brick:hasTimeseriesId "sim_IT_meter"). The setup evaluates Energy Efficiency Directive (EED) compliance for Power Usage Effectiveness (PUE) by simulating an energy consumption of 250 MWh per day, with 200 MWh attributed to IT and 50 MWh to non-IT.

5.2 Results and Analysis

The RDF graph shown in figure 16 prevailed a detailed verification to ensure the semantic accuracy and completeness of the data center model. The graph contains more than 20 entities, which include spatial zones, as well as IT and mechanical equipment, and their associated sensors. UPS units, PDUs, and server racks exhibit direct connections, effectively representing all essential energy pathways required for PUE calculation. The visual connections between

elements are vital in obtaining precise and complete results from SPARQL queries.




Figure 11. Chord diagram illustrating class-level relationships among various Brick and REC ontology entities within the data center model.

The chord diagram illustrates energy flow relationships, which help us understand how much each subsystem contributes to the total facility load. The diagram shows that cooling systems account for a significant portion of non-IT energy consumption. These flows form the analytical basis for identifying opportunities for efficiency improvement.

Figure 11 visually represents how entities such as `brick:Site`, `rec:Zone` and `brick:IT_Rack` are interconnected based on `isPartOf`, `hasPart`, and other semantic relationships. The thickness of the arcs reflects the number of relationships between classes, while the direction of the curves represents the flow of relationships, such as `brick:hasPart` from higher-level spatial entities (`brick:Site`) to their components (`rec:Zone`, `brick:ICT_Rack`) and the inverse `rec:isPartOf`. This visualisation is particularly valuable during the modelling phase, as it verifies that spatial, equipment, and functional classes are correctly linked in accordance with both the Brick and REC ontologies. By clarifying these high-level relationships before populating the model with instance data, it ensures that subsequent RDF triples align with standardised semantics, enabling more precise SPARQL queries and facilitating energy performance reporting in compliance with the EU Energy Efficiency Directive.

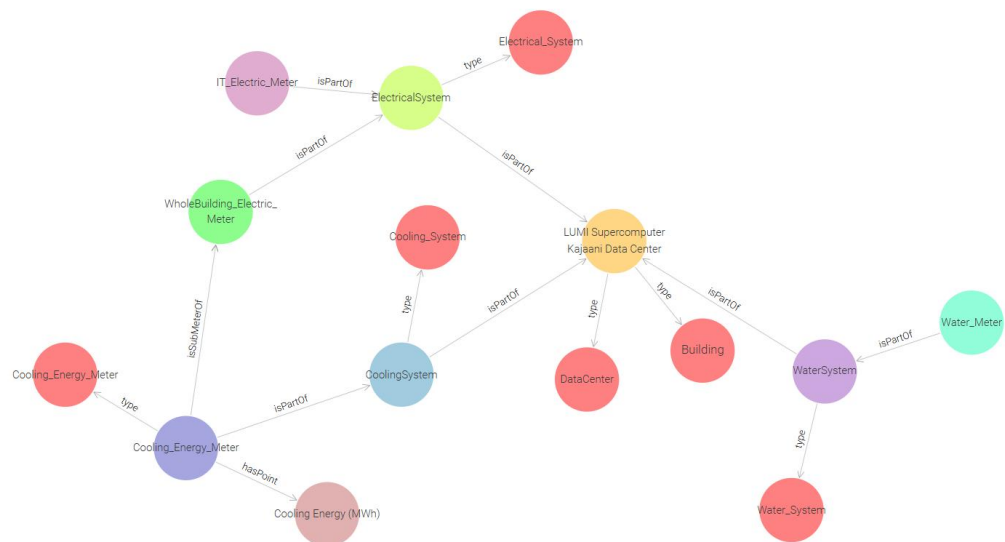


Figure 12. Graph-based visualisation of the data center RDF model showing individual entities and their semantic relationships.

Figure 12 shows a labelled directed graph where each node represents an entity from the Brick (e.g., equipment, spaces, and locations), and each edge is a brick relationship (predicate), encoding how those entities connect as RDF triples. The model follows Brick's principle that a data center is expressed as a

graph of triples (subject–predicate–object), enabling standardised querying and reasoning across heterogeneous systems. [33.]

The workflow diagram figure 10 shows that validation results proved the RDF model met 100% of the SHACL shapes' requirements. The model adheres to Brick schema definitions, enabling reliable querying without semantic mismatches or missing data. The modelling process is robust because it does not contain syntax or class definition errors.

The energy reporting web interface, which meets EED criteria, is shown in appendix 7. It is also possible to retrieve all the data in a CSV file through the online interface; this file can then be sent to the EU database for energy reporting purposes.

The EED-compliant reporting output, shown in appendix 7, demonstrates the framework's capabilities to produce standardised compliance documents in CSV format. The example report included site metadata (e.g., floor area, installed power capacity) along with calculated performance metrics, such as PUE and WUE. This automated output format aligns directly with EU Energy Efficiency Directive requirements and is suitable for direct submission without additional manual formatting.




Figure 13. A graph that illustrates data center's PUE over a specific time.

The PUE time-series analysis figure 13 showed an average daily PUE of 1.33 for the period between January 1, 2024, and January 27, 2024, with fluctuations ranging from 1 to 1.4 based on one-hour interval measurements. There are several sharp drops, ranging from around 1.05 to 1.1, between 5 January 2024 and 17 January 2024, indicating periods of optimal efficiency. The latter part of January exhibits an upward trend, with PUE climbing toward 1.5 to 1.55, indicating decreased efficiency. The facility operates more efficiently during periods when IT load values are high relative to total facility power consumption. [34.]The facility exhibits higher values during periods of increased cooling demand and reduced IT load, when overhead systems consume a larger share of total energy use. These values indicate that the data center is performing efficiently, as they are lower than 1.3, according to the Uptime Institute [35]. The adoption of Brick enables data centers to establish a unified semantic layer that connects IT load, facility load, and cooling load, as well as heat reuse, into a machine-readable model. [21.] The system will then be able to automatically

calculate PUE at any time interval, eliminating the need for manual spreadsheet input. The enhanced transparency reveals inefficient periods along with their fundamental causes, which enables operators to implement corrective actions through cooling setpoint adjustments and load distribution. The accumulated insights from these efforts lead to flattened PUE variations, which maintain a lower average PUE value that meets EED benchmarks. As shown in this graph, Brick enables operators to establish semantic infrastructures, allowing them to understand why these patterns occur and take proactive steps to increase data center efficiency. Furthermore, Brick would also enable sophisticated predictive models to anticipate efficiency degradation before it happens. [36.]




Figure 14. Data center's ERE for a time period.

As shown in Figure 14, the graph illustrates the energy reuse effectiveness (ERE) in the data center for a specific time period in January 2024. ERE is a sustainability metric that measures the effectiveness of a data center in capturing and reusing waste heat from IT equipment, with lower ERE values indicates better performance, meaning greater portion of waste heat is captured and reused. The graph shows exceptional performance peaks between 4 and 6

January 2024, with spikes reaching 1 ERE, meaning 100% heat reuse. On average, the values remain steady between 0.2 and 0.4, indicating moderate and consistent heat recovery capabilities. With the help of Brick, standardised queries would enable real-time optimisation algorithms to maximise heat recovery. As shown in the graph, when the ERE drops in late January, automated systems can adjust heat exchange operations, IT equipment placement strategies, and other relevant factors. [21.]

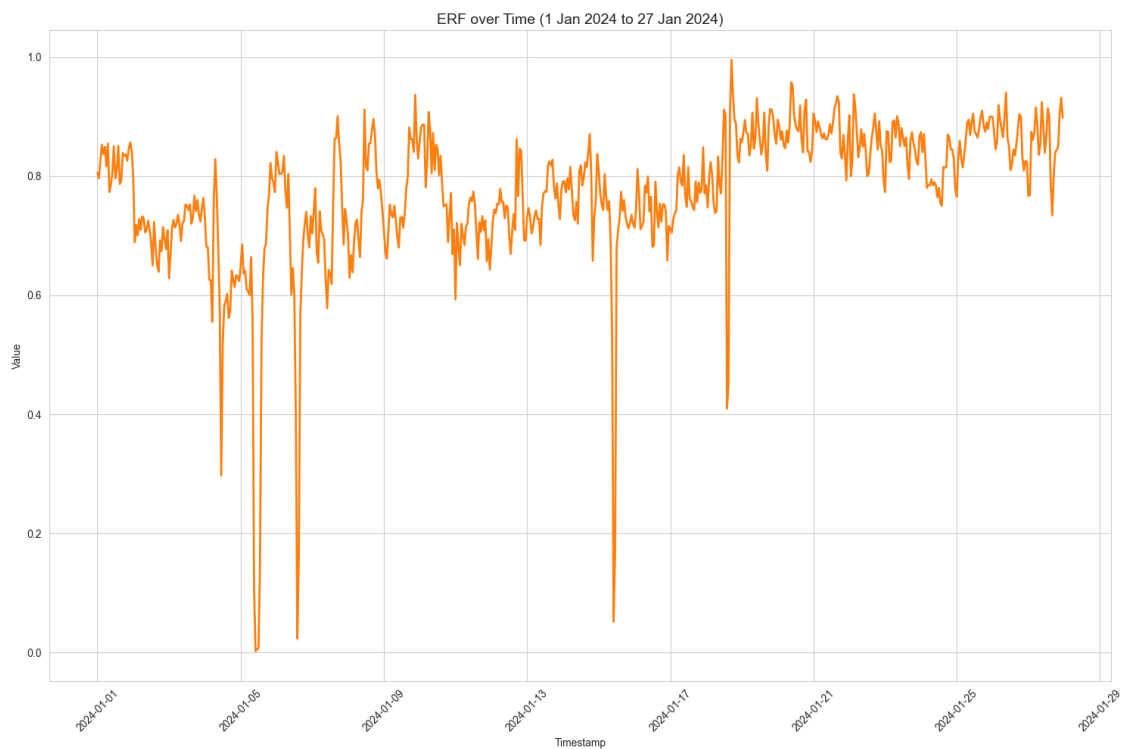


Figure 15. Data center's ERF for a specific time in January 2024.

Figure 15 shows the energy reuse factor (ERF) over the time period in January 2024. ERF is also another sustainability metric that shows how effectively a data center utilises its total energy input, which includes both direct consumption and energy recovery systems. Most of the monitoring period showed ERF values between 0.7 and 0.9, indicating strong overall energy utilisation efficiency. The two significant drops to 0 ERF on January 5th and January 16th show major operational incidents or planned maintenance activities that impacted multiple facility systems, while PUE and ERE values decreased at the same time. When

these types of sudden drops occur, Brick queries could instantly identify the root cause from:

- Power distribution failures affecting PUE
- Heat recovery system outages impacting ERE
- External grid conditions
- Equipment maintenance activities [14].

The above evaluation results demonstrate that for semantic modelling value Brick offers a systematic and flexible way to represent both physical and logical connections within the data center, allowing direct interaction with analytical tools. For automation of the report the SPARQL-based computational method eliminates human aggregation processes, ensuring consistency and scalability across multi-site deployments.

The average PUE value of around 1.4 shows that non-IT infrastructure uses about 40% more energy than IT equipment. This suggests that major changes in cooling operations could save a lot of energy without affecting performance. The automatically generated EED-compliant reports enable immediate EU reporting with minimal operational effort.

The above results confirm that the framework can seamlessly incorporate semantic modelling into regulatory energy reporting for all data centers, thereby reducing the time and effort needed for these data centers, which range from small-scale to hyperscale.

6 Discussion

6.1 Implications and Recommendations

The proposed Brick framework offers several advantages in generating energy consumption reports for data centers, and some of the main points are Brick provides a uniform semantic model for subsystem components, including

sensors, places, and equipment. This model connects BIM, BMS, DCIM, and IoT data to enable consistent analysis. The framework captures relationships between sensors and equipment, enabling SPARQL queries to automatically detect key time series and compute metrics such as PUE, WUE, ERF, and REF without manual mapping. This eliminates the chance of reporting faults and minimises the administrative cost.

The reporting and analytics system can be deployed across multiple data centers because Brick is an open, vendor-neutral standard. Applications created using the Brick model can be reused, enabling cost savings. [21.] The semantic model enables the implementation of standards, including temperature set-point control and waste-heat recovery, while monitoring the integration of fluctuating renewable energy sources. Operators will have the capability to evaluate their performance and identify areas for improvement. An integrated graph and time-series database develop a digital twin (DT) of the facility to conduct advanced analytics, such as fault detection and predictive maintenance. [37.]

6.2 Limitations

Despite Brick's extensive range of classes and relationships for describing a building, specific data center-specific concepts, such as high-density server enclosures, advanced cooling technologies like liquid cooling, modular UPS systems, and connections to district heating, may not be comprehensively represented. To address this, it is essential to extend Brick with domain-specific classes and features. [38.]

Mapping sensors, spaces, floors, and equipment onto the ontology to appropriately represent the building's infrastructure is required to create a Brick model, which requires time and a skilled workforce. Although manual validation is still necessary to ensure utmost accuracy, automated technologies that extract and arrange this data could significantly accelerate the adoption of this ontology. [39.]

The accuracy of KPIs depends on the reliability of sensor data. The accuracy of computations depends on obtaining complete and accurate readings, as missing or incorrect readings can lead to inaccurate results. Maintaining data consistency is an essential requirement. The implementation of validation and redundancy measures helps organisations maintain high data quality. [40.] Data centers are complex facilities that do not remain static during their entire lifespan. Equipment and devices require regular addition, removal, or upgrades. The Brick concept demands ongoing revisions and close collaboration with facilities management teams. Maintaining and modifying the automated reporting code requires significant time and effort. [41.]

7 Conclusion and Future Work

7.1 Summary

This thesis establishes a semantic modelling system to automate energy reporting and compliance for EU Energy Efficiency Directive (EED) data centers. The digital economy based on cloud computing, artificial intelligence, and data services requires continuous growth in electrical power and cooling infrastructure to support its operations. The traditional building-management and data-center infrastructure management systems generate large amounts of sensor data, which is stored in proprietary formats until manually aggregated before being entered into compliance templates. The current reporting system under the EED faces multiple problems, as it requires extensive manual work, produces frequent errors, and fails to meet the directive's detailed multi-metric requirements. The research investigated four main questions regarding data center infrastructure representation through Brick and the integration of time-series measurements for KPI calculation, semantic model validation, and the practical implementation of the framework to enhance reporting accuracy and reduce manual work as discussed in chapter 6 above.

The research began by conducting a comprehensive assessment of the current data center energy problems, along with existing regulatory requirements. The literature study confirmed that PUE serves as a standard benchmark for data center efficiency. Yet, it only measures electrical sustainability while disregarding water usage, renewable energy utilisation, and waste heat recovery. The EU Energy Efficiency Directive underwent updates through its recast version and delegated regulation, which established new reporting requirements for data centers with an IT power capacity exceeding 500 kW. These requirements include measuring and submitting PUE, WUE, ERF, and REF, along with facility metadata. The multiple performance metrics in this system require better data integration and contextual understanding, which current diverse systems fail to deliver. The evaluation of numerous semantic ontologies led to the conclusion that Brick, combined with RealEstateCore, provides the most comprehensive and adaptable framework for representing building and data center infrastructure for regulatory compliance, data integration, and inference purposes.

The research established a fundamental prototype system which contained five essential components. The first part of the research created a data center ontology through Brick and RealEstateCore classes to represent physical areas (sites, zones, rooms) and equipment (servers, racks, power distribution units, uninterruptible power supplies, chillers, pumps and meters) and their connections (hasPart, isPartOf, hasPoint, isPointOf). The model links electric power sensors to their respective racks, rooms, zones, and sites through unambiguous connections which maintain spatial context for flexible query operations. The system received measurement data from BMS and DCIM systems via a data integration channel, operating at five-minute intervals. The system used unique identifiers to connect sensor readings to ontology before converting them into RDF triples, which were stored in a triple store. The validation system, based on Shapes Constraint Language (SHACL), checked the semantic graph for domain and range constraints to verify correct sensor-equipment type associations and ensure the presence of all mandatory attributes. The validation process achieved 100% compliance, demonstrating that the model adhered to ontological rules and prevented potential semantic errors that could impact KPI calculations.

The research team created SPARQL queries and Python scripts to calculate the EED metrics. The calculation of PUE required dividing facility energy consumption by IT energy consumption, while WUE required dividing cooling water volume by IT energy. ERF and REF involved the reuse of energy needed or the comparison of energy generation to total energy input. The RDF graph processing system executed these calculations to produce results that became new properties, linking them to their respective equipment or facility locations. The framework produced reports in CSV and JSON formats, which adhered to the EED template structure, demonstrating that all operations, from data collection to KPI calculation and report creation, could be performed automatically.

7.2 Future Work

In the future, an extension of Brick, backed by the community, that standardises classifications for racks, PDUs, modular data center microgrids, immersion cooling systems, and other data center-specific equipment will facilitate industry-wide adoption. An automated modelling workflow could convert data from DCIM configuration files, BMS point lists, and BIM models straight into Brick representations, thus reducing or eradicating manual effort. Machine learning techniques could further improve these models by detecting previously hidden relationships and identifying correlations within sensor networks.[42.] Simultaneously, dynamic alerting for performance metrics such as PUE, WUE, ERF, and REF—derived from continuous monitoring against established benchmarks—would provide actionable reference signals, enabling operators to adjust cooling methods and participate in demand response activities in real-time.

Furthermore, integrating real-time, automated fault detection and diagnosis (FDD) capabilities within the brick-based energy modelling framework is a key future direction. This aligns with the objectives of ASHRAE Guideline 223P, which supports standardised semantic data models to assist intelligent, automated fault detection in building systems, including data centers.[43.] The traditional

approach of monitoring system performance and troubleshooting relies on automated log interpretation over extended periods of time without human intervention. An increasing number of ontologies and semantic web technologies, such as Brick, indicates a shift toward machine-readable, interoperable models enabling continuous, real-time performance analysis. By embedding semantic definitions of expected operational behaviours and thresholds, automated systems could proactively detect anomalies, such as unexpected PUE spikes, cooling inefficiencies, or energy reuse discrepancies. [19.]

Adding these innovative FDD systems to the current Brick model will improve not only compliance reporting but also strategies for predictive maintenance, resilience, and energy optimisation. Further research may consider integrating reasoning engines based on SPARQL or SHACL constraints to enable real-time alerts for facility operators regarding value discrepancies and operational deviations from specific values.

References

- 1 World Energy Outlook.2024.Online.iea.org.
<<https://www.iea.org/reports/world-energy-outlook-2024>>. Accessed 16 May 2025.
- 2 Zschache, Johannes & Hartwich, Tilman.2025.Comparing energy consumption and accuracy in text classification inference.arXiv.
- 3 Unquenchable Thirst – Texas Data Centers Consume 50 Billion Gallons Of Water As State Grapples With Historic Drought.2025.Online.Yahoo Finance.<<https://finance.yahoo.com/news/unquenchable-thirst-texas-data-centers-010128313.html>>.Updated August 20 2025. Accessed 16 September 2025.
- 4 European Parliament and Council of the European Union.2023. Online.Europa.eu.<<http://data.europa.eu/eli/dir/2023/1791/oj>>. Updated September 20, 2023. Accessed 16 September 2025.
- 5 Granlund.2025.Online.<Granlund.fi>. Accessed 15 May 2025.
- 6 Kant, Krishna.2009.Data center evolution: A tutorial on state of the art, issues, and challenges.Online.Intel Corporation.
<<https://doi.org/10.1016/j.comnet.2009.10.004>>. Updated November 18 2009. Accessed 16 September 2025.
- 7 Arregoces, Mauricio; & Portolani, Maurizio.2004.Data center Fundamentals. Online. Cisco Systems, Inc.<https://ptgmedia.pearsoncmg.com/images/9781587050237/samplepages/1587050234_sample.pdf>. Updated August 2009. Accessed 16 September 2025.
- 8 Heddeghem, Van Ward; Lambert, Sofie; Lannoo, Bart; Colle, Didier; Pickavert, Mario; & Demeester, Piet.2014.Trends in worldwide ICT electricity consumption from 2007 to 2012.Computer Communications, Volume 50,2014, pp. 64-76.
- 9 Kamiya, George & Coroamă, Vlad C.2025.Online. Data Center Energy Use:Critical Review of Models and Results.<<https://www.iea-4e.org/wp-content/uploads/2025/05/Data-Centre-Energy-Use-Critical-Review-of-Models-and-Results.pdf>>. Updated March 2025. Accessed 10 September 2025.
- 10 Jegham, Nader; Abdelatti, Mohamed; Elmoubarki, Laila & Hendawi, Ahmed.2025.Online. How Hungry is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference.arXiv.<<https://arxiv.org/abs/2505.09598>>. Updated May 14, 2025. Accessed 30 July 2025.

- 11 Lijun, Wang & Khan, Samee Ullah.2011. Online. Review of performance metrics for green data centers: A taxonomy study.The journal of Supercomputing. 2011;63:639-56.<<https://doi.org/10.1007/s11227-011-0704-3>>. Updated March, 2013. Accessed 20 September 2025.
- 12 Van De Voort, Thijs; Zavrel, Valentijn; Torrens Galdiz, Itziar & Hensen, Jan.2017. Analysis of performance metrics for data center efficiency - should the Power Utilization Effectiveness (PUE) still be used as the leading indicator?.Eindhoven University of Technology.Volume 54, Issue 1, February 2017.
- 13 Sandberg, Magnus; Risberg, Maria; Ljung, Ann-Louise; Varagnolo, Damiano, Xiong, Duming & Nilsson, Mikael.2017.Online. A modelling methodology for assessing use of datacenter waste heat in greenhouses.<<https://urn.kb.se/resolve?urn=urn:nbn:se:ltu:diva-66466>>. Accessed 15 September 2025.
- 14 BrickSchema.2025.Online.<brickschema.org>. Accessed 15 May 2025.
- 15 Fierro, Gabe; Koh, Jason; Agarwal, Yuvraj; Gupta, Rajesh & Culler, David E.2019.Beyond a House of Sticks: Formalizing Metadata Tags with Brick. In: Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '19);2019 Nov 13–14; New York, NY, USA. New York (NY): ACM; 2019. p. 1–10.
- 16 Aniakor, Michael; Cogo, Vitor & Ferreira, Pedro M.2024. A survey on semantic modeling for building energy management. Lisbon (Portugal): LASIGE, Informatics, Faculty of Sciences, University of Lisbon.
- 17 Rasmussen, Mads Holten; Pauwels, Pieter; Lefrançois, Maxime; Schneider, Gerhard; Hviid, Christian A & Karlshøj, Jan.2017.Online.Recent changes in the Building Topology Ontology.<<https://doi.org/10.13140/RG.2.2.32365.28647>>. Accessed 25 September 2025.
- 18 Project Haystack.2025.Online. <<https://project-haystack.org/doc/docHaystack/Intro>>. Accessed 24 July 2025
- 19 Hammar, Karl; Wallin, Erik; Karlberg, Pontus & Hälleberg, David.2019.Online.The RealEstateCore Ontology. In: The Semantic Web – ISWC 2019 : 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part I. Cham: Springer; 2019. p. 130–45. <<https://urn.kb.se/resolve?urn=urn:nbn:se:hj:diva-46606>>. Accessed 25 September 2025.
- 20 Balaji, Bharathan; Bhattacharya, Arka; Fierro, Gabriel; Gao, Jie; Gluck, Jonathan & Hong, Dong.2018.Brick: Metadata schema for portable smart building applications. Appl Energy. 2018;226:1273–92.

- 21 Foroushani, Sepehr; Bernhardt, Rob & Bernhardt, Mark.2022.On the use of the reference building approach in modern building energy codes. *Energy and Buildings*, Volume 256:111726.
- 22 Balaji, Bharathan; Bhattacharya, Arka; Fierro, Gabe; Gao, Jie; Gluck, Jonathan & Hong, Dong.2016.Online. Brick: Towards a unified metadata schema for buildings. White Paper. <https://www.memoori.com/wp-content/uploads/2016/06/Brick_Schema_Whitepaper.pdf>. Accessed 1 September 2025.
- 23 BrickSchema Entity.2025. Online.<<https://ontology.BrickSchema.org/brick/Entity.html>>.Accessed 29 July 2025.
- 24 BrickSchema Timeseries.2025. Online.<<https://docs.BrickSchema.org/brick/timeseries.html>>. Accessed 29 July 2025.
- 25 Luo, Nan; Fierro, Gabe; Liu, Yu; Dong, Bing & Hong, Tianzhen.2022. Extending the Brick schema to represent metadata of occupants. *Automat Constr.* 2022;139:104307.
- 26 Pritoni, Marco; Paine, Drew; Fierro, Gabe; Mosiman, Chris; Poplawski, Mark & Saha, Arjun.2024. Metadata schemas and ontologies for building energy applications: A critical review and use case analysis. *Energies*. 2021;14(7).
- 27 Peffer, Therese; Fierro,Gabe & Raftery, Paul.2023.Online.Skewering the silos: using Brick to enable portable analytics, modeling and controls in buildings. Berkeley (CA): UC Berkeley, Controls and Information Technology.< <https://escholarship.org/uc/item/04w0b9n2>>. Accessed 16 September 2025.
- 28 Fierro, Gabe; Prakash, Abhinav Kumar; Mosiman, Chris; Pritoni, Marco; Raftery, Paul & Wetter, Michael.2020.Online. Shepherding Metadata Through the Building Lifecycle. In: *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation [Internet]*. Virtual Event Japan: ACM; 2020 [cited 2025 July 30]. p. 70–9. <<https://dl.acm.org/doi/10.1145/3408308.3427627>>.Accessed 10 September 2025.
- 29 Fierro, Gabe; Raftery, Paul; Roa, Christian; Pritoni, Marco & Wetter, Michael.2023. Online. Skewering the silos: using Brick to enable portable analytics, modeling and controls in buildings. p. DOE-CIEE--08681, 2229614. Report No.: DOE-CIEE--08681, 2229614. <https://www.osti.gov/servlets/purl/2229614/>. Accessed 30 July 2025
- 30 Xu, Kang & Xiao,Fei. 2023. Online.Development and application of Brick schema based multi-sources data integration for smart building energy

- management. The Hong Kong Polytechnic University.
<https://publications.ibpsa.org/proceedings/bs/2023/papers/bs2023_1726.pdf>. Accessed 29 september 2025.
- 31 Knublauch, Holger & Kontokostas, Dimitris.2017.Online. Shapes Constraint Language (SHACL). W3C Recommendation. <<https://www.w3.org/TR/shacl/>>. Updated. July 20,2017. Accessed 10 August 2025.
 - 32 Brain, Nick.2025.Online. pySHACL: A Python SHACL Validator. GitHub repository. <<https://github.com/RDFLib/pySHACL>>. Accessed 10 August 2025.
 - 33 Harris, Steve & Seaborne, Andy.2013.Online. SPARQL 1.1 Query Language. W3C.<<http://www.w3.org/TR/2013/REC-sparql11-query-20130321/>>. Updated March 21. Accessed 11 August 2025.
 - 34 Fierro, Gabe.2019.Online.Design of an Effective Ontology and Query Processor Enabling Portable Building Applications [Internet]. Berkeley (CA): University of California, Berkeley, Electrical Engineering and Computer Sciences.Report No.: UCB/EECS-2019-106. <<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-106.html>>. Accessed 1 August 2025.
 - 35 Van de Voort, Thijs; Zavrel, Valentijn; Torrens, Galdiz; Jose Ignacio, Hensen & Jan, LM.2017.Analysis of performance metrics for data center efficiency: should the Power Utilization Effectiveness (PUE) still be used as the main indicator? Part 1. REHVA Eur HVAC Journal 54(1):5-11
 - 36 Lawrence, Andy.2013.Online. Data center PUEs flat since 2013. Uptime Institute. <<https://journal.uptimeinstitute.com/data-center-pues-flat-since-2013/>>. Updated May 7,2020. Accessed 10 August 2025.
 - 37 Building better: Building data, brick by brick.2018.Online. Trellis; <<https://trellis.net/article/building-better-building-data-brick-brick/>>. Accessed 1 September 2025.
 - 38 Merino, Jorge; Xie, Xia; Moretti, Nicola; Chang, Jan Y & Parlikad, Ajith.2023.Online. Data integration for digital twins in the built environment based on federated data models. Proc Inst Civ Eng Smart Infrastruct Constr. 2023 Dec 1;176(4):194 211. <10.1680/jsmic.23.00002>. Accessed 21 August 2025.
 - 39 BrickSchema. Issue #619.2025.Online. GitHub; <<https://github.com/BrickSchema/Brick/issues/619>>. Accessed 20 August 2025.
 - 40 Duarte Roa, Cristian; Raftery, Paul; Prakash, Arka & Peffer Therese.2023.Online. Field Demonstration of the BrickSchema to Scale up the Deployment of ASHRAE Guideline 36 Control Sequences. UC Berkeley: Center for the Built

Environment.<<https://escholarship.org/uc/item/5zt2d66r>>. Accessed 1 September 2025.

- 41 How Does Technology Ensure Data Validation? 2025.Online. Sustainability Directory.<<https://energy.sustainability-directory.com/question/how-does-technology-ensure-data-validation/>>.Updated March 18, 2025. Accessed 2 September 2025.
- 42 Matterport Editorial Team. 2025. Online. Data Center Facility Management: Investing In Digital Resilience. Matterport; 22 April 2025 <<https://matterport.com/blog/data-center-facility-management>>. Accessed 7 August 2025.
- 43 Teymourzadeh, Reza & Nakazawa, Yuki.2025.Online. A scalable and interoperable platform for transforming building information with Brick ontology. arXiv:2509.16259 [cs.CY]. <10.48550/arXiv.2509.16259>. Accessed 18 September 2025.
- 44 American Society of Heating Refrigerating and Air-Conditioning Engineers.2022. ASHRAE Guideline 223P – Standardized Semantic Modeling of Building Data. Atlanta, GA: American Society of Heating, Refrigerating and Air-Conditioning Engineers.

Appendix

1. Turtle file of the Data Center

The content of this appendix is confidential.

2. SHACL Validation

The content of this appendix is confidential.

3. EED Compliance Reporting

The content of this appendix is confidential.

4. PUE, ERE & ERF Chart

The content of this appendix is confidential.

5. Raw energy consumption data from the data center in 15-minute intervals

The content of this appendix is confidential.

6. Energy Consumption Dataset

The content of this appendix is confidential.

7. EED Compliance Report of a data center.

The content of this appendix is confidential.

8. RDF generation script using Python and Brick to define the semantic model of a data center using both Brick and REC ontologies.

The content of this appendix is confidential.