The Innovation Hub

for Affordable Heating and Cooling

DCH Product features

Akram Hameed, John McCulloch, Subbu Sethuvenkatraman

May, 2022

CSIRO

About i-Hub

The Innovation Hub for Affordable Heating and Cooling (i-Hub) is an initiative led by the Australian Institute of Refrigeration, Air Conditioning and Heating (AIRAH) in conjunction with CSIRO, Queensland University of Technology (QUT), the University of Melbourne and the University of Wollongong and supported by Australian Renewable Energy Agency (ARENA) to facilitate the heating, ventilation, air conditioning and refrigeration (HVAC&R) industry's transition to a low emissions future, stimulate jobs growth, and showcase HVAC&R innovation in buildings.

The objective of i-Hub is to support the broader HVAC&R industry with knowledge dissemination, skills-development and capacity-building. By facilitating a collaborative approach to innovation, i-Hub brings together leading universities, researchers, consultants, building owners and equipment manufacturers to create a connected research and development community in Australia.

This Project received funding from ARENA as part of ARENA's Advancing Renewables Program.
The views expressed herein are not necessarily the views of the Australian Government, and the Australian Government does not accept responsibility for any information or advice contained herein.

Primary Project Partner

The information or advice contained in this document is intended for use only by persons who have had adequate technical training in the field to which the Report relates. The information or advice should be verified before it is put to use by any person. Reasonable efforts have been taken to ensure that the information or advice is accurate, reliable and accords with current standards as at the date of publication. To maximum extent permitted by law, the Australian Institute of Refrigeration, Air Conditioning and Heating Inc. (AIRAH), its officers, employees and agents:

a) disclaim all responsibility and all liability (including without limitation, liability in negligence) for all expenses, losses, damages and costs, whether direct, indirect, consequential or special you might incur as a result of the information in this publication being inaccurate or incomplete in any way, and for any reason; and

b) exclude any warranty, condition, guarantee, description or representation in relation to this publication, whether express or implied.

In all cases, the user should be able to establish the accuracy, currency and applicability of the information or advice in relation to any specific circumstances and must rely on his or her professional judgment at all times.

Project title: CSIRO Senaps data platform demonstration and development of the Data Clearing House

This project will oversee development of the Data Clearing House (DCH), a cloud based building data management and application enablement platform. The DCH connects Internet of Things (IoT) systems from buildings and supports complex data analytics. The DCH will underpin the development of applications that improve renewable energy integration in buildings and unlock new opportunities for delivering Buildings to Grid (B2G) services.

This project will investigate features of the CSIRO Senaps data platform and their suitability for the DCH. It will combine these findings with results from the DCH2 Switch data platform subproject to develop the Data Clearing House.

Lead organisation

CSIRO

Project commencement date

July, 2019

Completion date

June 2022

Date published

27 May 2022

Contact name

Subbu Sethuvenkatraman

Email

Subbu.sethuvenkatraman@csiro.au

Project website

List of figures

1 SUMMARY

This document provides details of DCH platform features and functionalities available through the release version May 2022. The document starts with a summary of DCH users and their expectations of the platform. Details of unique functionalities of the platform have been described. A summary of DCH features is provided for easy reference.

2 Vision

CSIRO and the Australian Institute of Refrigeration Airconditioning and Heating (AIRAH) are leading a whole-of-industry initiative to overcome the barriers to digitalization in the built environment sector and to make sourcing of smart-buildings energy advice easy by implementing a new digital sharing platform (the ***i***-Hub Data Clearing House) for building services.

The new ***i***-Hub Data Clearing House platform provides IT infrastructure for Australian building owners to easily connect their buildings to a wide range of commercial providers of energy productivity and building management services. This creates a means for competitive sourcing of energy services, and for stimulating building services product innovation across Australian industry. The ***i***-Hub Data Clearing House platform avoids the need for each service provider to replicate data infrastructure for buildings.

It also aims to provide a platform for innovators to create new data-driven 'Software as a Service' (SaaS) offerings. DCH will have a thriving marketplace for building applications. This will trigger additional energy productivity innovation in the buildings space, allowing new players to enter the market by lowering the barriers around existing data silos.

The DCH underpins a broad vision to support innovation that will accelerate transition to net zero energy buildings, self-controlled buildings and seamless integration with smart cities infrastructure.

3 DCH platform users & benefits

A detailed stakeholder engagement survey carried out in 2020, 2021 has identified the following categories of main users of the DCH platform.

Building management related users Example roles: Facility manager, sustainability team, Maintenance contractors	Building management users will be the primary, frequent users of the platform. They will be benefitting from the data monitoring, display features of the platform, use the data for reporting, auditing purposes, utilise available applications in the platform to deliver services. They will be authorised to act on behalf of the building/portfolio owner. They will be interacting with system integrators to facilitate onboarding of buildings.
Service providers Example roles: Application developer IOT sensor manufactures	DCH enables service providers to offer their services without having to worry about the data access and management aspects. DCH also provides them a path to market for their services. Data driven algorithm developers (app developers) and IOT sensor manufacturers will greatly benefit from the access to data and tools available in the data platform for offering their services. Application developers will have access to rich set of data from different sources including the BOM data for offering deploying data driven algorithms. They will benefit from DCH features such as semantic model based querying, data source management.
System integrators	System integrators integrate various building systems such as Building Management Systems (BMS), Energy Management Systems (EMS), facility management, security systems in a building. As a result, they have intimate knowledge of the field devices, automation, communication capabilities of these sub systems in a building, and also access to relevant metadata about these systems (e.g. point labels, drawings). System integrators play the critical role of being the interface between the buildings and DCH. Their involvement can reduce the cost of onboarding a building and enable easy deployment of applications and controls. The nature of their involvement will vary depending upon the expectations from the building owner. For large buildings, system integrators are likely to facilitate interfacing of building hardware to DCH and support provisioning of access to various building metadata. System integrators also help the service providers by facilitating deployment of applications such as FDD in a building. DCH's onboarding tools and easy deployment of applications capability will benefit system integrators.
Building owners/portfolio owners Other similar roles: CIO/CFO of large organisations	Building owners are the economical buyers of the DCH service. Their motivation is primarily economical, i.e. operational cost saving and productivity improvements that can be achieved through DCH. DCH aims to lowers the cost of delivery of data driven services by sharing the cost of data acquisition, hosting across multiple providers. They are the data owners and will authorise access to data that belongs to their building /portfolio. As a result, building owners may have login access to the platform to set /grant permissions. However, majority of building owners are not normally expected to use the platform and their nominated representatives will be using the platform.

Other stakeholders who will interact with DCH and benefit from the in-built features of DCH:

- DCH aims to serve the building industry through delivery of various services. It is envisaged that a third-party platform operator will operate and manage the platform operations. As a result, DCH platform will have features (e.g. multi tenancy) that help delivery of managed services to the stakeholders identified above.
- Other data users: Building contractors, consultants, researchers will benefit from data available in DCH. Although authorised by data owner (building owner), data users will have specific expectations around how the data is provided to them.
- Semantic building models are currently created off the platform and used in the platform. DCH is currently working on creation of tools for supporting semantic model creation. Semantic modellers will use some features of the platform for carrying out tasks such as validation and verification of building models. It is envisaged that in the future semantic models will be created by system integrators utilising model creation tools. Semantic modellers will use some features of the platform for carrying out tasks such as validation and verification of building models.

3.1 Selection of DCH platform features

Following are the list of typical questions users will be asking while deciding to engage with or use DCH as the platform :

Key questions	Who is asking
Can I connect my legacy building with various data sources?	Building owner, system integrator
What is the onboarding cost, how do I onboard a building with minimal time & effort?	Building owner, system integrator
Is it useful as a monitoring and display platform?	Asset manager, facility manager
Is it useful as a data management and sharing platform?	Facility manager, building owner
Does it have a rich set of applications that I can make use of?	Building owner, facility manager
Does it have features to facilitate deployment of my applications?	Service provider
Is it satisfactory from cyber security, user management perspective?	All

In response to these expectations, the project team have developed a platform that aims to deliver the following features

1. Simplify onboarding, support data collection, data sharing without vendor lock ins
2. Facilitate easy deployment of application-based services
3. Facilitate access to data, building model-based capabilities
4. Serve as a secure platform that can be used by various partners

4 DCH Platform summary

Data Clearing House (DCH1) Platform is a cloud-based building data management and application enablement platform. The conceptual framework of the DCH1 platform is shown in Figure 1 Conceptual framework of the DCH1 platform

. A summary of these features is provided below. Further details are available in other sections of the document.

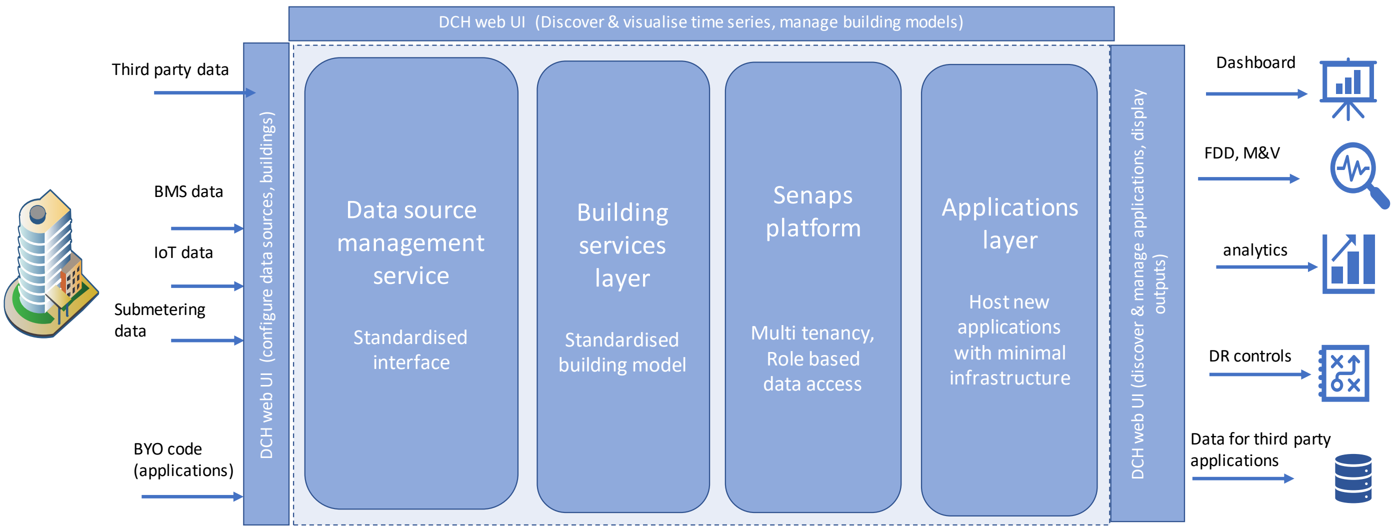

Figure 1 Conceptual framework of the DCH1 platform

The **Data Source Management Service (DSMS)** is a service used by the DCH platform to facilitate the collection and storage of heterogenous near real-time operational building data from BMS systems and IoT services. It is provided to DCH by the Senaps platform (see below).

Building Services Layer (BSL) has features to host semantic models of buildings. Semantic building models are used to query various building data and their relationships. This layer underpins the portability and extensibility features of the data platform.

The data platform (**Senaps**) is a mature platform for storing and handling building data (either geospatial gridded or point timeseries) and automating analysis workflows. DCH leverages several capabilities provided by Senaps. Senaps' user management framework has been adopted for managing role-based access to data in DCH1. Further DCH uses

Senaps' multi tenancy framework to manage multiple users and application hosting in DCH, as well as the DSMS mentioned above.

The **applications layer** supports hosting of independent, certified applications/algorithms. Application algorithms will use data available in the platform to deliver various services such as (i) Data Analytics (with resulting charts and dashboards that identify trends about historic and forecast on-site generation and load information), (ii) Fault Detection and Diagnosis (FDD), (iii) Measurement and Verification (M&V), (iv) Demand Response (DR), to name a few. The applications can be designed in a building agnostic manner, allow a write-once but deploy-many model for scaling application adoption.

DCH also provides features for hosting applications off the platform by making the data available for external software developers through APIs (Application Programming Interfaces).

In addition to the Application Programming Interfaces (Data Source Management Service API, Building Services API, third party data API), DCH provides **web user interfaces** (web UI) that enable authorised users to interact with the platform. The DCH web UI allows users to onboard new building data to DCH and supports construction and exploring of semantic models of buildings and building systems. Time series data visualisation, new applications hosting in DCH is enabled through the web UI.

To illustrate the platform's capabilities, let us consider the smart battery management application that is operating in one of the pilot sites. The application (Figure 2) uses load forecast, weather forecast and time dependent electricity pricing (or other pricing signals) to estimate future PV generation and operational costs of procuring electricity from the grid. The application will then decide if storage needs to be charged or discharged in order to optimise financial or emissions results and will deliver a control signal to the battery management system hardware.

As part of site onboarding, DSMS was used to connect the related site hardware (e.g. Building management system, PV, battery systems, energy meters). A hardware gateway was used to communicate to DCH using MQTT protocol. A Semantic model of the site was developed that included points that are required for operation of the battery algorithm. These two steps are independent of the application and are useful for delivery of various other applications or digital services. The control algorithm developed in Matlab operated off the DCH platform and used historic data from DCH and the relevant third-party data (weather data, electricity price forecasts) already available in DCH to deliver charging and discharging recommendations. The Semantic model of the site was used to query for relevant time series data points. Control algorithm can also be hosted on the platform using the application management capabilities of the platform. Control recommendations are dispatched via DCH MQTT broker to the site hardware.
This algorithm can be deployed across multiple sites that are connected to DCH (and have a semantic model associated with it) without undergoing labour intensive process of mapping the application relevant points.

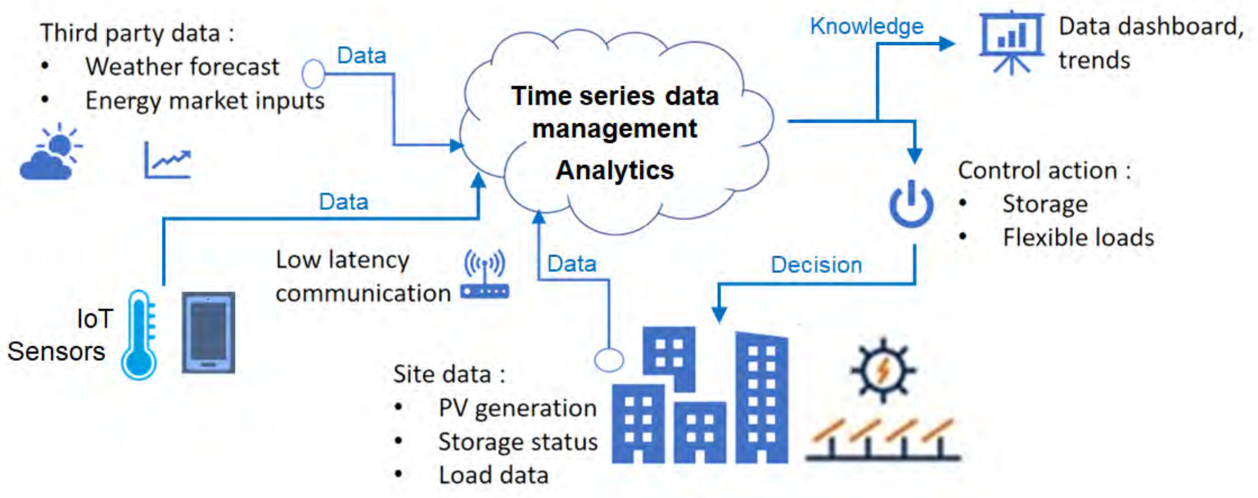

Figure 2: Example of smart battery management application

5 Unique functionalities & capabilities of DCH platform

Key technologies and tools that underpin development of DCH1 data platform are listed here.

5.1 Data Source Management Service

DCH's core objective of lowering barriers to onboarding smart building data into cloud systems is supported by the Data Source Management Service (DSMS) component. The DSMS is an extensible service with a goal of supporting new and existing building systems and services such as BMS, IoT and more. The service's core purpose is to put management of data pipelines back in the hands of technical users.

The DCH team currently support a variety of methods for ingesting data into the platform; via DSMS and other means. These can be categorised into self-service and managed ingestion methods.

Self-Service Data Ingestion

Self-service data ingestion is recommended by the DCH team in most cases. The platform provides a variety of interfaces on common internet protocols that system integrators and developers can use to interact with the platform. Some key protocols are captured below:

- MQTT/MQTTS: a message-based machine-to-machine network protocol. MQTT is both lightweight and mature and has good industry uptake. It is the recommended way to ingest data into DCH. To use MQTT, DCH provides comprehensive instructions and examples for system integrators and developers. The DCH team recommends using secure MQTT (MQTTS), so data is encrypted in transit to prevent snooping by third parties. BMS systems that have MQTT capability or MQTT enabled gateways can use this method of data ingestion.

- DCH supports the use-case where user's equipment connects directly to DCH's MQTT message broker. Examples here are modern BMS systems that have MQTT capabilities, hardware gateways such as Tridium JACE, EdgeX or EasyIO. Additionally, software gateways/middleware such as those provided by Tridium or NodeRED can be used. Users who have access to developers can create their own DCH-compatible custom software to run on their own servers.
- DCH also supports the use-case where a user has an existing message broker. In this case, DCH software will connect to the user's broker and subscribe to data for ingestion. This has been used in practice to capture data from The Things Stack (TTS) formerly The Things Network (TTN) as part of a CSIRO research project.
- AMQP: another common message-based network protocol, AMQP is commonly used for business-to-business (B2B) scenarios where integration into existing cloud services might be necessary. The DCH team strongly encourages the use of secure AMQP.
- HTTPS REST: DCH provides via its underlying data platform (Senaps), a set of RESTful APIs to let users send and receive data from the platform. Whilst MQTT is the recommended approach for ingesting data into DCH, HTTP is an alternative that can be used to good effect. HTTP also enables the user to update and delete existing data in the platform. All API endpoints for REST are fully documented with live documentation. HTTP REST connections can be made from a variety of systems and devices.

- FTP/FTPS: the File Transfer Protocol (FTP) is an older but pervasive data transfer protocol for sending small to large files over the Internet. DCH provides methods for the user to set up scheduled FTP ingestions of Comma Separated Values (CSV) files. The DCH team strongly encourages the use of FTP with Transport Layer Security (TLS) encryption applied (FTPS) to prevent snooping by third parties.

Managed Data Ingestion

Managed data ingestion is an approach where the user makes a request to the DCH team to create a DCH-managed ingestion pipeline. There are two instances where this makes sense to pursue:

1. The user wishes to take advantage of one of the existing managed ingestion methods that the underlying Senaps platform provides
2. The user wishes to request a new or novel managed ingestion method to be supported.

For the first case, there are a variety of managed ingestion pipeline options available to users on request:

- SFTP: Secure File Transfer Protocol (SFTP) is a 'privacy first' evolution from the older FTP approach of sending files over the Internet. There are plans to support this protocol in DCH's self-service pipelines, but for the moment, an interested user can contact the DCH team to set up a managed SFTP ingestion pipeline. As with the FTP self-service system, DCH supports to ingest data from CSV files using SFTP managed pipelines.
- Third party REST API-based services: the DCH team has managed ingestion pipelines for a number of popular built environment IoT services that expose their data via REST API, including Wattwatchers.

In the second case, the DCH team can work with users to support their more bespoke needs where required and strategically relevant. A summary of DCH supported data ingestion options are provided in **Error! Reference source not found.**.

Table 1 : Summary of data ingestion options in DCH

Data source	Ingestion options	Example hardware currently supported
BMS	Direct BMS interface	Modern BMS systems with MQTT/AQMP
	Hardware gateway	Tridum JACE, EdgeX, EasyIO
	Software gateway/middleware	Edgex, Tridium, NodeRED
Other data sources	DCH web API	
	DCH managed third party IoT provider API	Wattwatchers,
	Third party MQTT broker	The Things Stack (TTS) (formerly The Things Network)
	FTP	

PARSING of Data for Self-Service Data Sources

For users who employ the MQTT or AMQP message-based ingestion pipelines, a choice of data parsing method is available. To facilitate uptake of the platform, the DCH team has created a simple 'JSON Schema' that enables a BMS-first approach to data collection. This can be found at the following URL: https://bitbucket.csiro.au/projects/SBDCH/repos/bms-json/browse

Using the DCH-published *bms-json* has the advantage of being supported by documentation and equipment set up guides created by the DCH team.

While the name of the schema in the aforementioned repository includes the term BMS, it should be noted that this schema can be used by any system capable of producing JSON data and sending it via one of the messaging protocols supported by DCH. An example of this is the NodeRED workflow software, which can be described as 'IoT glue' for solving problems in the IoT domain.

Data Security

The DCH team take a serious but pragmatic approach to security. Wherever possible, data ingestion methods can take advantage of encryption protocols like TLS. This means data is encrypted 'in-flight' between the source of data and DCH. In many cases, however, it is impractical or impossible to deploy encryption to existing field devices. Due to this fact, DCH provides unencrypted communications options for protocols that support them. This approach is discouraged, but available as an option of last resort.

In addition to data being encrypted 'in flight' whenever possible, DCH also stores data encrypted 'at rest'. This means, any user data, timeseries or otherwise that is provided to the system is stored encrypted on 'disk' in the cloud.

5.2 The Building Services Layer (BSL)

Interoperability in DCH is achieved by providing APIs to edit and query semantic models of buildings. During onboarding, information about a site and its buildings is composed into semantic models which comply with the Brick Schema. These models describe buildings in terms of structure (wings, floors, rooms), equipment, and associated instrumentation and controls points (which are linked to timeseries data streams) as detailed in 0.

The Brick Schema is a large ontology of classes, including structural features (floors, rooms etc), equipment (heating and ventilation, metering etc), and 'points' (sensors, set points, limits etc). Most importantly, Brick also defines possible relationships between building components, such as where they are located, what sub-components they have, and the ways they can connect to each other. This makes Brick far more expressive than simple tagging systems. An example of Brick schema representation of a building air conditioning system is shown in Figure 3 . This system uses a chilled water based cooling system and the cooling is distributed using an air distribution system of air handling units (AHU's), dampers (vents) and fans.

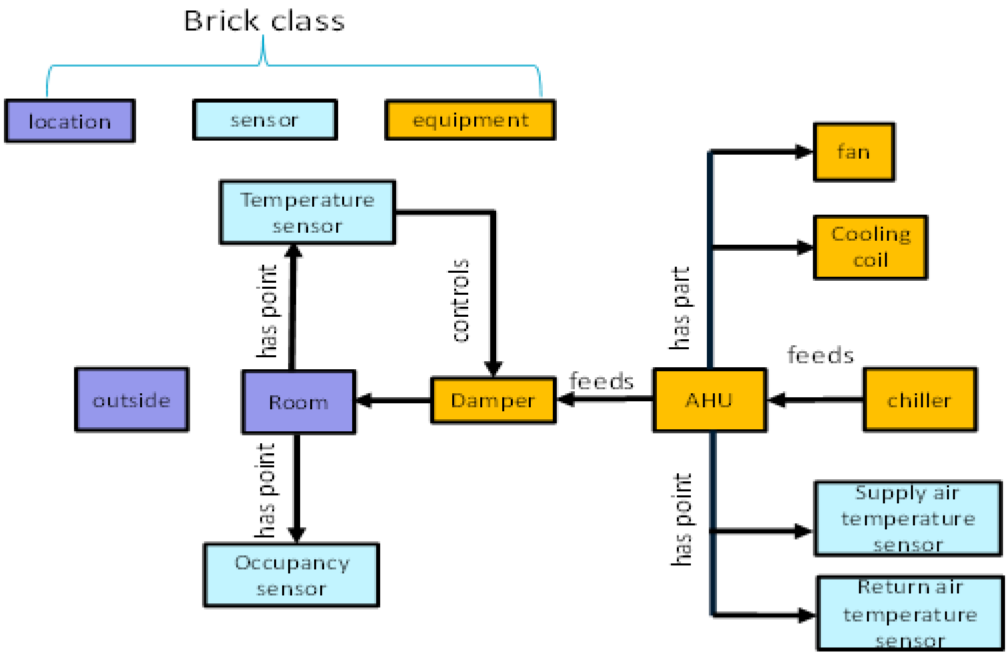

Figure 3: Example of Brick schema capturing relationship between different HVAC entities in a building

The use of Brick models makes building data discoverable by software, using an open query language. This allows applications to interact with buildings without prior knowledge about them. This in turn allows an 'app' ecosystem to exist by enabling the deployment of one application across many building models, and for one building model to be used by many applications. Semantic models are extensible, allowing modifications and changes to the building and adding more data streams over time.

Semantic model of buildings

Semantic models express real-world systems as 'knowledge graphs' (i.e. networks) of entities and the connections between them, using schemas to define valid classes and relationships. The value of this is that an application can express the related things it needs to discover and then let a graph database do the hard work of finding entities of interest, rather than having to download a whole model and explore it programmatically. Using the topology of Figure 3 as an example, an application might submit queries that effectively ask questions such as:

- "which AHU feeds Room 101?"
- "which Rooms are fed by AHU 2?"
- "list all Room temperature sensors and their associated upstream AHU supply air temperature sensors."

DCH's modelling system additionally supports a concept called *Entity Properties*, which are typically a key, a value, and an optional unit. These properties allow a modeller to describe entities inside their model with attributes that may or may not change over time. Importantly, the query language can seamlessly work across these entity properties as well as other model relationships. Entity properties can describe the area of a room, annotate entity serial numbers, asset classification tags (e.g., VBIS[1]), manufacturing year or orientation of solar panels. Existing systems allow users to capture information of this nature, but when it is integrated into DCH, it becomes possible for users and applications to create rich queries to discover relationships that would traditionally require significant data discovery and pre-processing activities.

As a more detailed example, consider a Fault Detection and Diagnosis (FDD) application that is used to identify the root cause of offices that are too warm. While locating the Air Handler Unit (AHU) that delivers air to this office area is straight forward, the FDD application needs to search for anomalies across all the potential related equipment and sensors, including details of temperature sensors (air inlet, cooling coil inlet and outlet, etc) and fan speed inside the AHU, as well as related temperature sensors in offices fed by the AHU. All this is possible with Brick based models and the related querying ability.

Another application might want to discover hierarchies of electrical meters and their connected sub-meters, which might be part of switchboards in other parts of the building. Brick semantic models can fulfill these needs because of two key things: The first is that complex connections between systems can be mapped (simplistic name-based metadata is limited as it indicates the function of individual sensors in isolation, but has difficulty describing the rich set of possible connections); The second is the existence of a query engine that can answer questions about properties and relationships. Armed with these, a query expressing a set of connections becomes portable across many different buildings. In the FDD example described, the application's query can be re-used for any building with AHUs and offices with temperature sensors.

DCH features to support Semantic model development

The DCH Building Service Layer fulfils the need for a semantic building model service. It supports handling of semantic models, linking of semantic models with third-party applications and performing powerful queries over buildings using query interfaces. The workflow for creating building models allows for incremental editing and validation, followed by publishing authoritative versions of buildings to a queryable database.

The DCH BSL is implemented as a RESTful API, backed by a graph database which stores building models as Resource Description Framework (RDF) graphs and performs queries submitted via the web user interface. User queries in the Brick Query Language (BRIQL) are translated to SPARQL Protocol and RDF Query Language (SPARQL), a standard language for interrogating RDF databases. Queries can be composed in terms of node class, the tags associated with node classes, entity properties, and transitive connections between nodes.

Where Brick Schema v1.2 doesn't comprehensively define domain classes, DCH has introduced a set of extensions features. It is the intension of DCH to feed these extensions back into the Brick community discussions, and to continually work towards a merging rather than deviation from the Brick standard. Examples include the addition of a VBIS_tag and serial number entity properties, and several sub-classes around electrical metering, inverters and battery systems.

While RDF is a well understood format in the Semantic Web community, it is not typically familiar to web and app developers. Offering only an RDF-based interface could be a blocker to adoption of the DCH. For this reason, DCH also provides interfaces to a JSON-based interchange schema, using Brick's classes and relationships, for use with the Building Services Layer API. This allows a user to supply building model data as either JSON or RDF, eliminating RDF

[1] https://vbis.com.au/

skills and tools as barriers to entry for developers. DCH users may also configure assets using a web frontend, which interacts with the Building Service Layer. Users of this interface can define or modify buildings, floors, rooms, equipment, and sensor points.

Model Generation Tooling

To facilitate bulk creation of building models outside the platform, DCH team have developed various tools to use building metadata to create building models. CSV based templates and Python based scripts were created that enable scalable, repeatable model creation, often a laborious process. Although this is an offline feature in the current version of DCH, these model creation tools will be made available via DCH UI in the upcoming versions. As part of the on-boarding process, BMS Point metadata and IoT sensor metadata is being collected and stored in the data platform (see **Error! Reference source not found.**). This metadata can then be used to inform the construction of a semantic model. For example, relevant sensor metadata from BACnet and Modbus devices and existing tags such as Haystack Tags are forwarded to DCH platform in a standardised JSON schema. This enables consistent a message payload format while transferring data to DCH and will allow the DCH team to support building owners and system integrators wishing to connect existing data sources. Metadata available from these points will be used to support guided (semi-automated) generation of building models.

Validation of Models

RDF-based schemas, including Brick, can express what can exist in a valid model, but struggle to express what is forbidden. This means that it is possible to create a technically Brick-compliant model that contains conceptual inconsistencies. For example, it makes sense for rooms to be part of HVAC zones, but not for HVAC zones to be part of rooms. For this reason, DCH has augmented the Brick schema by adding a model validation mechanism to the DCH Building Service and has raised model validation as a priority for future work within the Brick Schema development community. The output of a validation operation can be used to warn a user about minor problems (such as possible incompleteness) during the editing phase, or to enforce hard constraints that prevent invalid models from being published into the database.

Database selection for Building Service Layer

The project team compared several graph database systems, including the Resource Description Framework (RDF) and property graph (PG) data paradigms, for their suitability for use in the DCH Building Service Layer based on cost, concurrency, segregation and other features. Candidate database systems used for this study were Blazegraph, GraphDB, HodDB, Neo4J, Neptune, and Virtuoso. A subset of the RDF systems was systematically benchmarked. To do this, datasets of building models were synthesised, templated on the topologies of available Brick models of real-world buildings. Queries for building components adapted from literature[2] were evaluated in terms of execution time. These experiments give us confidence that a responsive building service can be deployed and will scale to large datasets. They also revealed important effects in some systems, such as nonlinear ratio of query time to database scale, and time penalties for requesting nodes with unspecified degrees of separation (e.g. using n-or-more property path quantifiers). However, no single graph database product consistently outperformed its peers across all query types. After balancing cost, performance and required features, a suitable database has been chosen for the Building Service Layer implementation. Since graph access is abstracted, it is possible to switch to alternative databases (including cloud-based SaaS options) in the future with minimal development effort depending on the future operational needs of the DCH platform.

[2] Fierro et al, 2018, " Design and analysis of a query processor for Brick", ACM transactions on sensor networks, Nov, 2018.

5.3 Utilisation of Senaps platform features for DCH

Senaps is a mature platform for storing and accessing sensor and model data (ether geospatial gridded or point timeseries) and automating analysis workflows. Senaps platform is already used by industry partners in agriculture and environmental sensing domain for over five years. DCH uses Senaps for its timeseries sensor data service, its data source management service, its user authentication, and its analysis service. All historic data is stored in the Senaps database. As a result, all DCH applications will access the database to perform analytic functions.

DCH will benefit from underlying data handling capabilities that exist in Senaps. Beyond the building models and related building data streams, DCH applications will have access to a range of other time series, geo spatial and third-party data. For example, by signing up to DCH, users will have access to BoM data services such as observation data and gridded weather forecast data.

Data authentication and authorisation

Security and trust are paramount when developing a DCH style ecosystem where many different organisations or actors interact. Inherently individual people or organisations involved have no way to trust other actors in the system, so the DCH system itself must be the trust provider. A useful analogy is phone “app stores”. Phone users generally do not verify installed apps, because they know that when installing the app the phone’s operating system will ask them to confirm the individual sub-systems that the app will access. This gives basic sanity check and combined with the measures the app stores have in place for approving apps, this provides high level of trust. DCH provides inbuilt authentication and authorisation systems to deliver a trusted data platform for various users. DCH leverages the Senaps platforms’ authentication and authorisation systems to provide these mechanisms for trust, while introducing its own novel semantic model based authorisation system to extend upon the role-based system provided by Senaps.

Authentication

On the authentication side, DCH has extended Senaps’ capabilities to act as an identity provider to the wider DCH system. For example, The Building Services Layer (and applications that are built on top of it) will use Senaps in a similar way to the familiar ‘Sign in with Google’ or ‘Sign in with Facebook’ techniques. This is implemented using the Oauth2 standard (similar to that used by Google and Facebook). As a result, signing in to DCH will allow the users to access various apps based on their credentials. DCH also provides a secure invitation system to enable new users to access components of the ecosystem, based on their privileges. A “You have been invited to join DCH” style email allows new users to sign in for the first time. Combining the existing Senaps authentication tools with the new Oauth2 functionality provides the basis for delivering a secure environment where building managers and application developers can interact with a high degree of trust.

Authorisation

DCH provides a flexible authorisation system for Identity and Access Management (IAM). At the core of this system is the concept of an entity. An entity in DCH parlance can be any object inside a semantic model, but the concept is extended to include the encapsulating concepts of sites, buildings, and organisations. Each entity has an inherent set of traits, such as ‘metadata’, ‘timeseries data’, or ‘semantic models’, among others. Combined with a set of verbs that define actions that can be taken over an entity’s traits (read, write etc), and a set of reaches that determine which related entities should be included, DCH permits authorised users to compose sets of permissions rules defining access to data and metadata in the platform as needed. Higher-level constructs (role sets and user groups) allow users and roles to be managed on a many-to-many basis.

By leveraging the DCH authorisation system, applications can request access to specific buildings and their data. In the future, an application could request access to individual building subsystems using a finer grained permissions scheme.

Multi tenancy

DCH is a multi-tenanted ecosystem using a Software as a Service (SaaS) model. This means all DCH customers interact with a single instance of the DCH software platform. In addition, all building models and time series building data and application analysis are managed in the single instance of the platform. The advantages of this approach include straight forward (with explicit permission) sharing of common datasets between users, efficient use of cloud computing resources and the ease of deploying applications across multiple buildings. DCH uses the following Senaps' features in managing multi tenancy.

1) Application Isolation: Each application running in DCH uses Containerised Computing technologies to isolate the application's execution environment from core platform services and other applications.
2) Data isolation: Data is stored in shared instances of redundant and scalable data storage systems, business logic implemented in DCH services is used to segregate customer data.
3) Role-based Authorisation: A building level granularity authorisation framework implements DCH authorisation rules to enforce isolation and enable sharing in a secure way.
4) Performance and efficiency: DCH's single instance multi-tenancy approach allows for efficient use of compute resources for applications and reduced redundant storage of data. Additionally, the query service in DCH can perform cross-building and cross-customer data queries without time and complexity of distributed federated queries.

Complex data handling and data preparedness

The Senaps platform's analysis Service allows arbitrary data analysis workflows to be developed, containerised, and deployed on dynamically allocated cloud computing resources. Through Senaps' analysis workflow, DCH offers inbuilt models and tools, that applications can access, that provide data manipulation and cleaning capability, reducing the need for applications to build these functionalities internally. For example, one common challenge with building data acquisition is that input data are acquired from different systems at different data rates. Further, these systems most likely will have different timestamps. Algorithms are usually not designed to operate on such data, so a common first step is to interpolate/resample the relevant data streams to a common sample time base. Along similar lines most algorithms find it challenging to work with time series data that has some missing data points, which is another common occurrence in the real world. Senaps provides the functionalities to deal with these challenges on the application's behalf. Applications can directly request gap filled, aligned and resampled versions of multiple data streams directly. The system can also provide aggregated data on request, which is useful, and commonly required, when an application uses daily averaged values, but the underlying data source may be providing raw data to DCH at 1 minute or 5-minute intervals.

DCH has been adding capabilities specifically for building data management. DCH has built tools to carry out diagnostics on sensor data flows. This includes data health notifications via email and dashboard style tools to help building managers ensure all the data points linked to DCH are available in the data platform and are operating as expected. This 'health check' tool performs abnormality/outlier and missing data detection.

5.4 Applications layer

DCH supports third party innovators and application developers through mechanisms for hosting their solutions and allowing users to choose and apply third party applications to their building(s). The supporting Senaps analysis service allows applications to be discovered and used without the need for application developers to manage any hosting infrastructure. The system allows developers to share their application without risking IP leakage by distributing code beyond DCH. It also provides a way for application developers to update/upgrade applications without having to distribute updates to all users. DCH is the first data platform to offer such services that enables third party applications to be deployed as an app and facilitates hosting of the application across multiple buildings using the semantic models.

DCH combines its functionalities to deliver a unique experience to both application developers and users. DCH supports application developers through:

1. Giving application developers access to semantic building models and data without needing to go through the complex process of acquiring data and linking their application to that data for every building.
2. Provision of a powerful querying capability, available through the Building Services Layer, that allows discovery of data and associated components/sensors from building models.
3. Providing the necessary security for application hosting using a containerised environment. Applications can be provided to the platform in various forms including the "bring your own code" that could be written in one of many commonly used languages (e.g. R/Python/MATLAB/C++).
4. Enabling users to take advantage of third party developed software without compromising data sovereignty.

DCH also supports externally hosted applications. Externally hosted applications will interact with the Building Services Layer via API to discover sensor data streams and the sensor data API to provide data services.

It is a goal of DCH to provide in the future a "marketplace" or an "app store" to provide a way for application developers or service providers to reach customers easily.

Currently, a reference Measurement and Verification (M&V) application is available in DCH as a hosted application along with Grafana as a dashboard for visualising results.

5.5 DCH User interface

DCH DASHBOARD

DCH provides a contemporary web dashboard for users to interact with data and models, with the following capabilities:

- User management and authorisation services. The DCH authorisation system is organisation-based.
  - Creation of groups of users for easier user management
  - Definition of new roles to control access to resources
  - Invitation of new users to the inviter's organisation
- Site and building management:
  - Creation of sites/buildings models
  - Interactive map to show sites' locations
- Search & explore
  - Search for sites and buildings
  - Search for data points using types of location, equipment and points

- Find and view all available data streams and assign them to the appropriate models.
  - Graph-based building model visualisation (in-development)
- Installation of applications to gain further insights
- Data source management
  - Adding DCH bms-json compatible devices or services via an MQTT client connection is supported.
  - Data sources can be connected to the points in a building model that correspond to the sensors or other devices in that building.
  - Interactive charts to load data in a chosen period
  - Downloading data as CSV for off-system analysis
- Data health analytics
  - Display varied metrics of all data streams from a set of buildings (including Missing per cent, outlier per cent, data mean, data mode, data std., number of samples, std. Sampling period(s), time to detect missing, number of outlier samples, …)
  - Data stream outage detection
  - Search for data streams and compare their metrics with others

DCH APIs

DCH provides open Application Programming Interfaces (APIs) that enable authorised users to construct and explore semantic models of buildings and building systems. In addition, DCH also provides fully open APIs for users to access their time series data, and to manage their self-managed data sources.

All DCH HTTP REST APIs are documented with so-called 'live' OpenAPI documentation. This allows developers to rapidly become productive with the APIs. For Non-HTTP interfaces such as MQTT, payload schema documentation is available, as well as comprehensive examples.

6 Detailed list of DCH platform features

Functionality	Description
Rapid onboarding of buildings : Simplify onboarding, support data collection via various sources, building model creation	
Add new data source	Authenticated BMS gateways or software services can be connected to DCH and data streams can be linked to building model. DSMS can provide users with credentials for linking data sources to URLs
Add new site	New site and building details can be entered through the web UI
Add asset details of a building or create a building model (building model wizard)	Users can input information regarding each asset they have in their building, grouped by zone, location etc according to Brick schema

Functionality	Description
	and create a building model. No prior experience with Brick Schema or semantic modelling is required
Ingest data from BMS, IOT devices, third party systems via MQTT	Facilitate data ingestion from various devices or services using the external MQTT brokers. e.g. BMS systems through gateways or IoT devices through interfaces like The Things Stack
Ingest data from a BMS and Modbus devices via common gateways	Driver for ingesting data using Tridium JACE along with JSON schema. Data is sent direct to DCH's MQTT broker.
Ingest data from third party providers through Web API	e.g. Wattwatcher energy metering
Ingest data from CSV files	Can be used to ingest historic data, batches of data
Upload a building model to DCH	Already existing building models in RDF format can be directly imported and linked to existing building. This enables thirdparties to provide/create their own model creation tooling and become part of the DCH ecosystem.
Asset metadata available	Display asset metadata wherever possible. In addition, when captured using entity properties, this metadata becomes queryable
Model creation tools	Offline tools to support creation of building models
Facilitate access to data, building model based capabilities : Data prep, data quality, thirdparty data, data explore, search	
Ability to view buildings in a site, find buildings	Use a drop down menu to search for buildings in UI front display grouped according to location
View and download timeseries data linked to a building model	Depending upon permissions, building data can be accessed, visualised as time series chart over a select period, downloaded as CSV
View and download timeseries data in DCH, including thirdparty data	Time series data related to building whether linked to a building model or not can be found using the data explore feature
Find building model specific relationships through a search function	Brick schema specific building model explorer – carry out search based on Brick classes, Brick point names etc. Permits to help understand building subsystems and the data available from them.
Data query capability through building model	Utilise SPARQL, BRIQL features of search. Key feature for app portability.
Data and building model management	users have full abilities to create, update, and destroy their own data
Data source health monitoring	Identify data sources that may be faulty and report any data outlier issues (only with data collected at regular intervals) Explore faulty data sources via UI and as CSV download
Create alerts when there is a data issue	Email alerts to targeted users to help identify and resolve issues.
Data sharing with third parties	API for sharing of data to third party users. Non DCH Service providers application developers can use this feature to access DCH data

Functionality	Description
Data export from DCH	Export data to third parties in a suitable format Currently available for Envizi
ingest data from the BOM and other third-party cloud data sources via APIs	BOM data (historic ambient temperature forecasts, ambient temperature forecasts, solar irradiance forecast, historic irradiance forecast), AEMO data (historic spot price forecast),
Display contextual relationships while navigating a building	Ability to view/display Brick relationships (location, part of, feed) while navigating a building and its components. Contractors or building managers can now see an extra dimension to their data.
Ability to edit, update, delete building asset details (building model)	Ability to make changes to a building model through the UI without having to access the RDF model. Changes made are accessible via BSL API and UI.
Building model available in DCH has undergone some form of validation	Use of SHACL and other verification methods to ensure the building models have been built according to current Brick and DCH modelling norms and conventions.
Building model sharing	API for sharing building models without accessing the UI
Offer app-based value added services : Facilitate easy deployment of applications from third-party providers and offer reference applications	
Applications can be hosted in the platform	Applications written in Python, R, and Matlab can be hosted and executed in a containerised (docker) environment on the platform. This enables a Bring Your Own Code (BYOC) approach where only interfaces to existing software must be adapted/managed to participate in DCH.
Configure applications for deployment across buildings through linking relevant datastreams	Application providers can use building model features (SPARQL for querying) to find relevant data for application deployment
Available applications can be discovered/found easily	Applications available for users including reference application can be easily found, if relevant permissions exist
Available applications can be deployed easily	Applications available for users including reference application can be easily found and deployed, if relevant permissions exist
Application inputs can be configured and outputs can be viewed, downloaded	Provision of user inputs for setting up inputs for the application Provision of displaying application results using Graffana dashboard, download results
Applications with control features can be deployed in DCH	Application outputs can use native MQTT broker to facilitate sending of control signals back to the site/building
DCH can support off platform hosted applications	Use data sharing, export features to support off platform hosted applications
Reference M&V app available for use	Whole of site energy use comparison using weather normalised M&V algorithm User configurable settings – period of M&V, tariff details to compute cost saving
Serve as a secure data platform that can be used by various partners : user management, data management, governance	
Login screen, registration process	Users can get access to DCH through an online registration process. Only approved /registered users can login to DCH

Functionality	Description
Authorised users can see (visualise) only buildings, data source when they access DCH	Role based user management, access control with minimal involvement of DCH team
Building owners can assign view and edit permissions to access data by role and by data type	OAuth2 authentication service for Senaps developed and deployed
Remove data and building models, applications when needed	Data clients must be able to choose to remove their data from the DCH at will and export it in a useful format that preserves the structure. This ensures that their ownership of the data is complete. Similarly, it should be possible to remove applications from DCH
Third-party application evaluation for security	Process to approve third party app hosting in DCH
Data storage, backup	Data stored in AWS (Australia) and backed up at regular intervals
Cyber security features of platform can be demonstrated/proven	Platform has undergone penetration testing Platform uses encrypted communication (all public APIs are encrypted)

7 Glossary of terms

AHU	Air handling unit: a HVAC component
AIRAH	Australian Institute of Refrigeration Airconditioning and Heating
API	Application programming interface
BMS	Building management system
Brick	a semantic schema for modelling buildings and building systems
CSIRO	Commonwealth Scientific and Industrial Research Organisation
DCH	Data Clearinghouse
DR	Demand response
FDD	Fault detection & diagnosis
Haystack	an ontology for tagging building components
HVAC	Heating, ventilation and air conditioning
IoT	Internet of Things: internet-capable sensors & actuators
JSON	JavaScript Object Notation: a data serialisation format
Oauth2	An authorisation framework used to grant permissions, access
Onboarding	A process of linking buildings to DCH
M&V	Measurement & verification
MQTT	MQ Telemetry Transport: a standard protocol for message transport (ISO/IEC 20922)
PV	Photovoltaic: solar panel
RDF	Resource Description Framework: a specification family for semantic modelling
REST	Representational State Transfer: a software methodology for the exchange of data commonly using web-based APIs
SaaS	Software as a service: a business model for software products
Senaps	a sensor data storage and analysis platform developed by CSIRO
SHACL	Shape Constraint Language: a language for semantic model validation rules

SPARQL	SPARQL Protocol and RDF Query Language: a language for querying semantic models