

REPORT

D4.2 Second year research and development report

WP4 Sustainable energy strategy with artificial intelligence: Dynamic energy efficiency control

Title of the work package: **AI-Driven buildings**

Responsible partner: **IT4Innovations National Supercomputing Center (VSB-TU Ostrava)**

Partners: STU BA, Asseco CE, DFKI, TSK

July 2026

www.innovaite.sk



Financované
Európskou úniou
NextGenerationEU

PLÁN [OBNOVY]



ÚRAD PODPRESEDU VLÁDY
SLOVENSKEJ REPUBLIKY
PRE PLÁN OBNOVY
A ZNALOSTNÚ EKONOMIKU



REPORT

D4.2 Second year research and development report

Report summarising the research and development activities of the 2nd year of the project implementation. Description of all key elements of the developed solution, the workflow for the creation of the digital image of the building, and the strategy for automated building energy efficiency management.

Content

Content	2
1 Executive summary	3
2 Project Context and Status of WP4	5
3 Refined scope and research priorities	8
4 Implementation status overview	12
Topic chapters	19
5 (Topic 1) — Communication, visualisation interface and BIM	19
6 (Topic 2) — Data preparation pipeline	52
7 (Topic 3) — Predictive and digital twin models	65
8 (Topic 4) — Research EMOS platform layer	72
9 Optimisation and Dynamic Control	85
10 Pilot deployment, verification and outlook	86
11 Risks and mitigation	92
12 Publications, dissemination and community building	93
Annex	97

1 Executive summary

This deliverable, D4.2 Second-year research and development report, summarises the activities of Work Package 4 (AI-Driven Buildings) during the second year of the InnovAite project. It follows D4.1 First requirements and implementation plan (M8), which collected user requirements and defined the research priorities and the implementation plan, and it reports on how the consortium is progressing along that plan. Not all modules are complete at this stage of the project, but WP4 is on track overall, and we expect the remaining work to be finished according to plan and reported in the subsequent deliverables D4.3 and D4.4.

In Year 2, WP4 specified the data formats and designed the architecture of the end-to-end energy-efficiency management chain: data input → validation → data modification → correlation analytics → model building → inefficiency detection → what-if simulation → implementation of measures and control adjustment → impact monitoring.

All parts of this chain are being progressively implemented as working modules; some modules are split into several implementation steps and are being completed incrementally:

- **Data preparation:** tools for format conversion of incoming data.
- **Validation:** removal of errors, conversion of recording logic, interpolation of missing values.
- **Data modification:** a configurable tool for adjusting the content of the data.
- **Models:** a model of ideal consumption, a model predicting real consumption (capable of predicting inefficiencies), and a specialised model for building thermal management.
- **Inefficiency detection and simulation:** calculation of inefficiency vectors for different types of inefficiency and simulation of the impact of their removal.
- **Impact monitoring:** calculation of running and cumulative impact values of implemented measures.
- **In preparation:** preparation of measures (return-on-investment calculation) and control adjustment (computation of the ideal operating profile).

The individual modules are developed by the members of the WP4 team (IT4I, STU, Asseco CE). To enable joint testing of the complete chain, a dedicated research software, BackBone, was

created, which connects the individual workflow steps in the form of microservices and serves as the shared integration and validation environment for the WP4 research outputs.

On the data side, NDAs were concluded with five pilot partners, whose datasets are being gradually acquired and used for testing and validation of the modules. This effort also includes the deployment of metering devices and the configuration of EMOS metrics at the pilot sites (currently in preparation), supported by the procurement tenders VO1 and VO2 carried out within the project.

With these results, WP4 declares the fulfilment of milestone M4.2: *“Digital building image structures defined and pipelines for the data ingestion created. AI tools for data preprocessing integrated.”* The milestone is fulfilled in the form of modules integrated in the BackBone software; a dedicated milestone report is being prepared in parallel with the submission of this deliverable. Tasks T4.1 and T4.2 have been completed, and tasks T4.3–T4.6 are progressing in line with the time plan (see chapter 2 for details).

The key R&D outputs of Year 2 include conference contributions, the BackBone research software, completed analyses of the Hobra and Lasselsberger data-block datasets, and testing on the Buildings – Střešovičky dataset. The structure of this report follows the WP4 topics: communication, visualisation interface and BIM (chapter 5), the data preparation pipeline (chapter 6), predictive and digital twin models (chapter 7), the EMOS platform layer (chapter 8), cross-cutting optimisation and dynamic control activities (chapter 9), and reference verification at pilot sites (chapter 10).

2 Project Context and Status of WP4

The *InnovAite Slovakia: Illuminating Pathways for AI-Driven Breakthroughs* project (TIC call 09I02-03-V2), funded by the Recovery and Resilience Plan – NextGenerationEU, lasts 30 months (M1–M30) and brings together 11 partners from academia and the private and public sector, led by Gratex International, a.s. Its main objective is to advance the state of knowledge in the field of artificial intelligence applications across multiple domains and to deliver measurable innovations supporting economic growth, sustainability, and quality of life.

Within this project, Work Package 4 (AI-Driven Buildings) applies artificial intelligence to building energy management. WP4 translates stakeholder requirements into research, development, and validation activities that support energy efficiency, comfort, and the digital transformation of buildings. The work package part D4.2. is led by IT4Innovations National Supercomputing Center, VSB – Technical University of Ostrava (IT4I) in cooperation with the Slovak University of Technology in Bratislava (STU), Asseco Central Europe (Asseco CE), DFKI, and TSK.

This chapter briefly recapitulates the WP4 objectives defined in D4.1, presents the updated status of the WP4 tasks at the end of Year 2, and declares the fulfilment of milestone M4.2.

2.1 Objectives O4.1 – O4.3 (recap)

The objectives of WP4, as defined in D4.1, remain unchanged. Their wording and the progress achieved in Year 2 are summarised below:

O4.1 – Enable AI-driven dynamic (flexible) management of energy supply and consumption with regard to economics, energy efficiency, and carbon footprint minimisation, by creating an intelligent AI superstructure on top of existing EMS systems. *Year 2 progress:* On track. The data foundation (T4.1, T4.2) is complete, and the monitoring layer (T4.3, ~50 %) together with the optimisation and deployment work (T4.4, T4.5, ~30 %) that form this AI superstructure are progressing on schedule; full dynamic management will be demonstrated as these tasks complete by M24.

O4.2 – Introduce modern AI/ML algorithms to improve current EMS solutions and optimise building management, in connection with O4.1. *Year 2 progress:* On track. AI/ML methods (LSTM, correlation and sensitivity analysis, reinforcement learning) are being implemented and

evaluated within T4.3 and T4.4, with the optimisation modules now at an early but progressing stage.

O4.3 – Create a functional multidisciplinary connection of the addressed domains (AI and energy), with every step communicated to potential users and EMS developers to ensure the highest quality and relevance. *Year 2 progress:* Largely achieved on the foundational side and continuing. The multidisciplinary AI–energy connection and user involvement were established in T4.1 (completed) and are being maintained through ongoing partner engagement in T4.5 and T4.6 (five pilot partners, NDAs concluded, data being collected).

2.2 Status of WP4 tasks

At the end of Year 2, all WP4 tasks are proceeding in line with the time plan. Tasks T4.1 and T4.2 have been completed, and the remaining tasks T4.3–T4.6 are in progress and on track for completion at M24. The table below summarises the status of the individual tasks; a detailed description of the activities carried out within each task is provided in section 4.3.

Task	Description	Timing	Status
T4.1	User requirements collection	M1–M8	Completed ☒
T4.2	Digital building image data collection tools	M3–M16	Completed ☒
T4.3	Tools for advanced automatic smart building monitoring	M9–M24	In progress ☐ (~50 %)
T4.4	Module for building operation efficiency optimisation	M11–M24	In progress ☐ (~30 %)
T4.5	Building management and solution deployment	M14–M24	In progress ☐ (~30 %)
T4.6	Research the solution on pilot use case	M2–M24	Ongoing — on track

2.3 Milestone M4.2 — declaration of fulfilment

Milestone M4.2 is defined as: *“Digital building image structures defined and pipelines for the data ingestion created. AI tools for data preprocessing integrated.”* WP4 hereby declares this milestone as fulfilled at M16, in line with the time plan.

The digital building image structures have been defined and the data-ingestion pipelines built, covering third-party Information Systems, the BIM model, and sensor data, with all inputs converted into a common structured format. The AI preprocessing tools have been integrated into the pipeline, alongside the automatic detection and removal of invalid data segments and methods for data annotation. The individual components are implemented as modules integrated in the BackBone research software, which serves as the shared platform for testing the complete processing chain.

The table below maps the M4.2 acceptance criteria to the sections of this report that provide the evidence of fulfilment. A dedicated milestone report is being prepared and will be submitted in parallel with this deliverable.

Criterion (M4.2)	Evidence / section	Status
Digital building image structures defined	Topic 1 — chapter 5 (BIM)	Fulfilled ☒
Pipelines for data ingestion created	Topic 2 — chapter 6 (Data preparation)	Fulfilled ☒
AI tools for data preprocessing integrated	Topic 2, sec. 6.5 (AI/SLM components)	Fulfilled ☒

3 Refined scope and research priorities

The scope of WP4 defined in D4.1 remains valid. The practical experience gained during Year 2 — above all the work with real operational data from the pilot partners — confirmed the research priorities defined in D4.1 (section 4.5) and led to several refinements described in this chapter. None of these refinements changes the WP4 objectives or the task structure; they sharpen the way the objectives are pursued.

3.1 Priorities confirmed from D4.1

All four research priorities defined in D4.1, section 4.5, were confirmed by the Year 2 work and are actively being developed:

- **Building digital twins** integrating BIM and sensor data — confirmed and substantially developed through the energy-performance and thermal-management models (chapter 7) and the BIM modelling approaches (chapter 5).
- **Predictive and optimisation algorithms** — confirmed; predictive models were selected, trained and validated (chapter 7), and the optimisation modules are under development in T4.4 (chapter 9).
- **What-if scenario simulations** — confirmed; the estimation of the potential impact of changes through what-if analysis across target areas and inefficiency types is the main focus of the Year 2 work in T4.4.
- **Validation frameworks** — confirmed and strengthened; validation runs continuously on the datasets of the pilot partners (chapter 10), with the BackBone research software (section 3.4) becoming the main instrument for it.

3.2 Lessons learned from practical work with data sources

The most significant refinements of the WP4 approach result from the practical experience with real data sources, which proved more challenging than anticipated in D4.1:

- **Data sparsity and absence of reference vectors.** Some acquired datasets exhibited insufficient coverage and lacked the reference (ground-truth) vectors needed for training and validation of the monitoring and digital-twin models. This is being addressed iteratively through the progressive acquisition of operational data from the pilot partners

and through the preprocessing and data-conditioning methods developed in T4.2 (interpolation, outlier detection, elimination of invalid data segments).

- **Insufficient spatial and temporal granularity of measurements.** In several cases the resolution of the available measurements did not allow reliable localisation and attribution of inefficiencies to individual functional units. This led to supplementary instrumentation at the sites and, more importantly, to a stronger emphasis on hierarchical aggregation structures of end-use consumption: for the analysis to produce meaningful results, a building (or an industrial site) must be decomposed as well as possible into functional consumption blocks. The hierarchical block structure has therefore become a core element of the digital building image and a precondition for the inefficiency analytics.
- **Attrition of pilot partners.** Several originally intended partners withdrew from the data-provision arrangement. They were replaced by alternative partners; data-provision NDAs are now concluded with all five currently participating sites, restoring the required data base.
- **Unavailability of BIM models.** A usable BIM model was not available for all pilot buildings. WP4 therefore formalised two parallel implementation scenarios — with and without a pre-existing BIM model — and broadened its research towards methods for the efficient creation of a BIM representation from minimal prior information. This resulted in the portfolio of five complementary BIM creation approaches described in section 5.4, ranging from a simple box model up to detailed modelling from geodetic point clouds.

3.3 Generalisation of the energy-performance model

The energy consumption of a site is never determined by the building structure alone. A substantial part of consumption is driven by what happens inside — by the operational, technological, or production processes taking place there. In office buildings consumption follows occupancy and working schedules, in public buildings it follows visitor numbers, and in industrial facilities it is largely determined by production. Detecting inefficient consumption always means the same thing methodologically: comparing the measured consumption against a reference of what the consumption should be under the given conditions. Only the reference quantity differs from site to site — weather, occupancy, visitor numbers, or production output.

For this reason, the WP4 energy-performance model was designed as a general module built on the relation between consumption and a configurable reference quantity, rather than as a set of separate, site-type-specific models. This design brings a single coherent methodology, transferability of the trained approaches between site types, and applicability of the developed solution to the whole class of energy consumers addressed by the project — buildings, industry, and municipal units.

For industrial sites in particular, this generalisation is a necessity rather than an option: the consumption of the building itself and the consumption of the production processes are measured together and cannot be meaningfully separated. An analysis restricted to the building alone would leave the dominant part of consumption unexplained, and inefficiencies could not be reliably attributed. The analysis of industrial consumption therefore deliberately includes the general operational and production processes of the site. Buildings benefit from the same design as well, since a significant share of building consumption is likewise linked to operational processes. The WP4 objectives, tasks, and deliverables remain unchanged by this refinement.

3.4 BackBone EMOS — research software for integration and validation

The project application foresees the transformation of the WP4 research outputs into a functional prototype within WP8 (deliverable D8.4 Prototype of System Design, M25–M30). To develop and test the WP4 modules already during the R&D phase, however, the team needed a shared environment in which the individual modules — developed in parallel by IT4I, STU, and Asseco CE — can be connected into the complete processing chain. For this purpose, the BackBone research software was created, with its first version delivered as part of milestone M4.1.

BackBone is not the end-user EMOS software; it is a tool for the research team. It connects the developed modules in the form of microservices, allows reference data to be entered and exported, and enables the complete workflow to be executed and debugged end to end. During Year 2 it became the key instrument for the validation of the WP4 research outputs, and it will provide a natural foundation for the prototyping work in WP8.

3.5 Updated research priorities per topic

Based on the above, the research priorities for the remaining project period are updated per topic as follows:

- **Topic 1 — Communication, visualisation and BIM:** finalise the portfolio of BIM creation approaches and the energy BIM standard; align the visualisation with the needs of operators and the EMOS frontend.
- **Topic 2 — Data preparation:** pipeline orchestration, AI/SLM components, and performance and scalability of the pipeline on larger datasets.
- **Topic 3 — Models:** finalisation of the digital model, broadening of the recognised inefficiency types, and validation of the thermal-management model on pilot data.
- **Topic 4 — EMOS platform layer:** inefficiency analytics, the end-user frontend, and the continued integration of all modules in BackBone.
- **Cross-cutting (Optimisation and Dynamic Control):** the optimisation module (return-on-investment calculation) and the dynamic control (computation of the ideal operating profile) are in preparation and represent the main priority for Year 3.

4 Implementation status overview

This chapter gives an overview of how the WP4 implementation is organised and where it stands at M16. It presents the updated topic and partner structure of the team (section 4.1), the updated timeline across the WP4 tasks (section 4.2), the detailed status of the individual tasks (section 4.3), and a summary of the key technical decisions taken in Year 2 that shaped the design of the solution across topics (section 4.4).


4.1 Updated topic / partner matrix

The WP4 work is organised into four topics led by the individual partners, complemented by the cross-cutting optimisation and dynamic-control activities and by the pilot verification. The figure below shows the team structure (partner → area → activities); the table summarises the topic and partner matrix together with the area owners and the corresponding chapters of this report.

Lead partner	Topic	Chapter						
Asseco CE	Topic 2 Data preparation pipeline	Data Input	Standardization	Roman Rigo	6			
			Validation	Roman Rigo				
			Modification	Roman Rigo				
		Data Pre-processing	Correlation - preparation for model - > numbers + report	Petra Hlavinová				
			Correlation - automatic division into blocks	Petra Hlavinová				
			Correlation - data shift	Petra Hlavinová				
IT4I@VSB	Topic 3 Predictive and digital twin models	Industry Model	Industry - consumption model (historical data)	Tomáš Karásek (Kristina Motyčková)	7			
			Industry - consumption model (prediction)	Tomáš Karásek (Marcel Smolík)				
		Model Building	Buildings - thermal management model	Tomáš Karásek (Tomáš Blejchaf)				
	Topic 4 EMOS platform layer	BackBone	Technological architecture	Tomáš Jungman	8			
			BackBone	Robert Šimon				
		Analytics	Microservices integration	Robert Šimon				
			Inefficiency calculation	Robert Šimon				
		Frontend	Frontend for users	Radek Fiala				
SVF STU BA	Topic 1 Communication, visualization interface and BIM	Interface	Communication interface	Michal Krajčík	5			
			Energy audit from existing data	Angelika Václavová				
			Module describing technologies and their changes - buildings	Angelika V / Michal K.				
		BIM Visualization	BIM frontend model (custom-made boxes)	Richard Varga				
			BIM created from address	Ján Erdélyi				
			BIM from minimal documentation (dwg, project)	Tomáš Funtík				
			Automatic creation of box model (from point cloud)	Ján Erdélyi				
			Modeling from surveyors' point cloud (detailed BIM)	Tomáš Funtík				
		BIM Energy Model	3D data energy BIM model STANDARD	Tomáš Funtík				
		DFKI	validace	Data Strategy		Data strategy validation	Ingo Zinnikus	8; 6
		TSK	validace	Functional strategy		Functional strategy validation	team	5;8

4.2 Updated timeline (Gantt) across topics

The figure below shows the updated WP4 timeline at M16. The plan follows the four project phases; green bars indicate completed work, the vertical line marks the current point of the project (M16, submission of D4.2), and the markers on the left summarise the status of the individual tasks (completed / percentage of completion with the planned end date). Tasks T4.1 and T4.2 are completed, tasks T4.3–T4.6 continue in line with the original plan towards M24; no shift of activities against the planned timeline has been necessary.



4.3 Detailed status of WP4 tasks

The following overview summarises, for each WP4 task, its leader and participants, the current status, and the activities carried out up to M16.

T4.1 User requirements collection [M1-M8] — Leader: STU; Participants: IT4I, Asseco CE, DFKI
Status: completed in line with the plan

Activities: Task T4.1 was completed in line with the time plan. The consortium established a network of user contacts beyond the project partners and collected requirements from the wider user community through several iterative feedback rounds. Building on this, we refined the WP4 solution structure, divided the path to the objectives into individual modules (distinct workflow stages), defined functional requirements for each, and distributed the module-creation tasks among WP4 members. *Related deliverable: D4.1 First requirements and implementation plan — delivered M8 (report provided).*

T4.2 Digital building image data collection tools [M3-M16] — Leader: IT4I; Participants: STU, DFKI, Asseco CE **Status: completed in line with the plan**

Activities: Task T4.2 was completed. We developed tools that ingest data from third-party IS, the BIM model, and sensor data, and preprocess them into a structured digital building image used as the shared data foundation for T4.3–T4.6, applying both standard techniques and AI methods

(LSTM, CNNs). The task also delivered tools for hierarchical structures (aggregated end-use consumption blocks) and a standardised data-record structure for the EMOS processing and inefficiency-evaluation cycle; a specific focus was integrating the BIM model and establishing a standard for the energy information needed for energy analytics. By providing this unified, validated data layer, T4.2 directly underpins O4.1 and O4.2 and forms the foundation referenced in D4.2.

T4.3 Tools for advanced automatic smart building monitoring [M9-M24] — Leader: IT4I; Participants: STU, DFKI, Asseco CE **Status: in progress in line with the plan (M9–M24), ~50 % complete**

Activities: In line with the time plan, T4.3 has reached its midpoint and is progressing successfully. We have implemented the data-preparation stage, the correlation analytics, recognition of several fundamental types of suboptimal operation, and visualisation tailored to both research and operator use, working across several efficiency dimensions (performance, procurement, investment, technological) — with heat-management analysis as a specific case of performance efficiency. A key element is building models (digital twins) that establish optimal/ideal consumption and thus locate the moments, places, and mechanisms of inefficient consumption. Remaining work covers finalizing the digital model, broadening the recognised issue types, and validating at pilot partner sites; by turning the data layer into actionable findings, T4.3 advances O4.1 and O4.2 and feeds the project's publication goals. *Related deliverable: D4.2 Second-year R&D report (M16) — on track (report in preparation).*

T4.4 Module for building operation efficiency optimisation [M11-M24] — Leader: Asseco CE; Participants: STU, DFKI, IT4I **Status: in progress in line with the plan (M11–M24), ~30 % complete**

Activities: In line with the time plan, T4.4 is underway at roughly 30 %. So far we have focused mainly on tools for estimating the potential impact of changes via what-if analysis across various target areas and inefficiency types, drawing both on data sets (predicting changes in consumption profiles) and on knowledge blocks (typical standardised estimates of structured changes and similar reference cases). The task builds on T4.3 inefficiencies and T4.2 digital objects, combined with historical component data, and explores reinforcement learning for automation; by letting

users simulate process/technology adjustments and quantify savings and carbon-footprint reductions, T4.4 directly advances O4.1 and O4.2.

T4.5 Building management and solution deployment [M14-M24] — Leader: STU; Participants: IT4I, Asseco CE, DFKI **Status: in progress in line with the plan (M14–M24), ~30 % complete**

Activities: In line with the time plan, T4.5 is progressing. We work to identify, in the data received, the controllable parts of energy management through which consumption efficiency can be continuously influenced, and on that basis propose a management strategy to the user, including a methodology for calculating and quantifying the impact of the proposed measures. The task consolidates the T4.2–T4.4 algorithms and provides simplified strategies for programmable operation-critical components; by converting monitoring and optimisation outputs into deployable, self-optimising control of real buildings, T4.5 directly advances O4.1 and O4.2.

T4.6 Research the solution on pilot use case [M2-M24] — Leader: STU; Participants: IT4I, Asseco CE, DFKI **Status: in progress in line with the plan (M2–M24)**

Activities: In line with the time plan, T4.6 is progressing. We have established contact with five partners, completed the NDAs covering data provision, and are gradually obtaining their data sets, which serve as validator projects; using this data we identify energy-saving opportunities at the sites, which feeds back into refining the relevant modules. By deploying and testing modules under real-like conditions and collecting site feedback, T4.6 verifies the applicability and modularity of the solution and strengthens the project's path toward practical validation and sustainability, contributing to the benefit quantification underpinning D4.3 and D4.4.

4.4 Key technical decisions in Year 2

Several technical decisions taken in Year 2 shaped the design of the solution across all topics:

- **Standardised data records (SeriesSchema).** A standardised data-record structure was adopted as the single exchange format for transferring data content between all modules and through the EMOS processing and inefficiency-evaluation cycle. Every module consumes and produces data in this format, which makes the modules freely combinable (section 8.2.3).
- **Hierarchical block decomposition.** Aggregated end-use consumption blocks organised in a hierarchical structure were established as a core element of the digital building image and a precondition for the inefficiency analytics (sections 3.2 and 6.3).
- **Microservices and API-first integration.** The individual modules are developed as independent microservices with API-first interfaces, allowing the partners to develop and test them in parallel and to connect them into the complete processing chain in the BackBone research software (sections 3.4 and 8.3).
- **AI/SLM orchestration (LangGraph / MCP).** For the AI and SLM components of the data pipeline, an orchestration approach based on LangGraph and MCP was selected (section 6.5).
- **Explainability-first model selection.** For the historical inefficiency model, explainability was prioritised over raw predictive accuracy, leading to the selection of quantile regression; for consumption prediction, a sequence-to-sequence transformer architecture was selected as the best-performing approach (chapter 7).
- **Physics-based thermal model with open data sources.** The building thermal-management model combines BIM-derived geometry and material parameters with meteorological data from Open-Meteo and shading derived from OpenStreetMap geometry computed in Blender, solved as a transient problem (section 7.3).
- **Portfolio of BIM creation approaches.** Instead of relying on a single method, five complementary approaches to BIM model creation were developed, covering the whole range from minimal input data to detailed geodetic point clouds, together with a 3D energy BIM data standard (sections 5.4 and 5.5).

Topic chapters

The following chapters report the second-year research and development work organised by WP4 topic. Each topic chapter follows a common structure, so that the reader and the reviewers can navigate the deliverable consistently and compare progress across topics. The structure of every topic chapter is the same and comprises: (a) an introduction and scope, defining the role of the topic within WP4; (b) the activities and methodology carried out in Year 2; (c) the implementation outputs, including prototypes, modules and models; (d) the research findings and results; (e) the validation status; and (f) the open issues and next steps planned for Year 3.

The topic chapters cover communication, the visualisation interface and BIM (chapter 5), the data preparation pipeline (chapter 6), predictive and digital twin models (chapter 7), and the EMOS platform layer (chapter 8). They are followed by the cross-cutting optimisation and dynamic control activities (chapter 9) and by the pilot deployment and verification at the reference sites (chapter 10).

5 (Topic 1) — Communication, visualisation interface and BIM

5.1 Introduction and scope

Topic 1 focuses on the communication, visualisation and BIM-related components of the WP4 solution. Its role is to ensure that the digital building image developed in WP4 is not only technically usable by analytical modules, but also understandable and actionable for building owners, facility managers, operators, tenants and other stakeholders involved in energy management.

The work carried out in Year 2 addressed four closely connected areas. First, the team defined the communication and visualisation interface through which building data, energy performance indicators, alerts and optimisation recommendations can be presented to different user roles. Second, several complementary approaches to BIM model creation were developed and evaluated, covering a broad range of practical situations: from buildings with no available documentation, through buildings with partial 2D documentation, to buildings surveyed by point clouds and reconstructed as detailed BIM models. Third, a data standard for an energy-oriented

BIM model was specified, defining the non-graphical information required for energy calculation, BIM visualisation and integration with the EMOS platform. Fourth, an energy-assessment strand was developed — an energy audit derived from existing operational data, together with a module of changes for evaluating energy-saving measures — which links the building data layer to the energy BIM model standard (Section 5.5) and the energy models in Topic 3 (chapter 7).

The outputs described in this chapter form the BIM and visualisation component of the digital building image developed within WP4. They provide the spatial, semantic and user-interface layer required for linking building geometry, energy data, sensor measurements, digital-twin models and EMOS visualisation workflows. This work therefore directly supports the fulfilment of milestone M4.2, particularly the definition of digital building image structures and their integration into the data-ingestion and preprocessing chain.

The chapter is structured as follows. Section 5.2 describes the communication interface and role-based visualisation concept. Section 5.3 presents the energy audit module and the module of changes. Section 5.4 summarises the portfolio of BIM model creation approaches. Section 5.5 specifies the energy BIM model standard. Section 5.6 explains how BIM outputs connect to the EMOS platform and the WP4 data pipeline. Section 5.7 summarises open issues and next steps for Year 3.

5.2 Communication interface

The communication interface is the layer through which the digital building image and the related energy data are presented in a form understandable to users. Activities in Year 2 covered the definition of user roles, the development of a role-based basket of functions, the first design of the visualisation interface and the preparation of a validation survey for building professionals.

The objective was not to create a final graphical user interface at this stage, but to define the functional and informational basis for the end-user frontend. The interface concept is therefore understood as an input to the EMOS platform layer, where the production implementation of the frontend is further developed.

5.2.1 Methodology and approach

Three user roles were distinguished based on practical experience and the expected information needs of different stakeholders:

- **Owner / investor**, who requires a fast, high-level overview of savings, costs, asset performance and investment priorities;
- **Operator / facility manager**, who works with detailed operational parameters, alarms, faults, equipment status and optimisation actions;
- **User / tenant**, who is mainly interested in personal or unit-level consumption, indoor comfort and simple actionable advice.

The design of the communication interface followed four steps:

1. definition of user roles;
2. assembly of a basket of candidate functions;
3. mapping of prioritised functions onto an initial interface concept;
4. preparation of a validation survey to test preferred forms of visualisation with practitioners.

Several cross-cutting design principles were established and apply across all roles:

- navigation across the building hierarchy: portfolio → building → floor → zone → device;
- selectable time granularity: day, week, month and year, with a distinction between workdays and weekends;
- progressive disclosure, where a single headline indicator can be expanded into more detailed information on demand;
- traffic-light status coding, distinguishing good, acceptable and poor states;
- contextualisation of values through benchmarking per m² and weather normalisation using degree-days.

These principles ensure that the same data foundation can be used for different user groups without creating separate and incompatible applications.

5.2.2 User roles and the basket of functions

For each role, a basket of candidate functions and information items was defined. The main functional groups are consolidated in Figure 5.1.

The **Owner / investor** view focuses on the building passport, energy consumption and costs, financial indicators, document archive, energy performance and ESG-related information. This view should support fast strategic decisions, investment prioritisation and asset-performance monitoring.

The **Operator / facility manager** view extends the common data base with a digital-twin monitoring layer, zone heat maps, live BAS/smart-meter data, an alarm centre with priority triage, an AI optimisation assistant, equipment and inspection records, open points and photovoltaic or microgrid utilisation. This view supports day-to-day operational management and corrective action.

The **User / tenant** view prioritises indoor-environment parameters, personal consumption and costs, AI-based behavioural advice, feedback and helpdesk functions. For commercial tenants, the view can also include indicators linking operating costs to revenue or profit.

Owner	Operator / manager	Tenant / occupant
Building passport — name, address, category, year built/renovated, floor area, floors, zones, units. Energy consumption — live use (electricity, gas, district heat, water); split by carrier; weather normalisation (degree-days); comparison vs previous period and baseline; peak loads; benchmark per m²; end-of-month prediction; CO₂. Energy costs — month-to-date; split by carrier; period comparison; weather-normalised; benchmark per m²; end-of-month prediction. Finance — commodity prices (daily/monthly/annual); invoices to pay; cash-flow; defaulters list. Documents — invoice archive; supplier and tenant contracts. Data platform & EPC — energy-performance indicators; classification class; ESG report.	Twin & monitoring — digital twin; zone heat maps; live BAS/smart-meter data — environment (°C, %, ppm), HVAC, energy, control. Alarm centre — anomaly and leak notifications; setpoint exceedances; fault reports. AI optimisation — consumption optimisation (analysis + monitoring + feedback); priority triage of issues. Equipment & revisions — device overview; service interventions; meter administration; documentation archive; planned maintenance. Open points / to-do — tasks by urgency; change reports; status (open / in progress / overdue / done). PV / microgrid — production-consumption balance; PV vs grid; overflows; estimated savings.	Indoor environment — temperature, humidity, CO₂ vs setpoints; IEQ indicator. Consumption & costs — monthly and annual analysis with predictions and period comparison; split by carrier. AI advice — behavioural recommendations (e.g. -1 °C heating ≈ 6% monthly saving; CO₂ > 1000 ppm — ventilate). Feedback & helpdesk — comfort reporting; requests and incident reports; outage notices; contacts. Finance & documents — invoices to pay; contract with owner; invoice archive. Commercial tenant — operating-cost-to-profit and energy-cost-to-revenue indicators; profit vs energy-cost trend; benchmark.

Fig. 5.1 Basket of functions per role, consolidated from the function catalogue

Several functions recur across roles, including consumption, costs, benchmarking and indoor-environment monitoring. The roles are therefore not separate systems, but filtered and differently weighted views of a single catalogue. Figure 5.2 maps each function group to the three roles and distinguishes between detailed, summary and hidden views. This coverage view

provides the design argument for a single role-aware interface rather than three independent applications, and it directly informs the end-user frontend implemented in the EMOS platform layer (section 8.5).

Function group	Owner	Operator / manager	Tenant
Energy & performance			
Building passport	•	•	⦿
Consumption & costs	•	•	•
Benchmarking	•	•	⦿
Prediction (energy / cost)	•	•	⦿
Weather normalisation	•	•	–
Carbon / CO ₂	•	•	⦿
Comfort & operations			
Indoor environment (IEQ)	–	•	•
Alarms & fault reports	⦿	•	⦿
Equipment & maintenance	–	•	–
Open points / tasks	–	•	–
PV / microgrid	⦿	•	–
Intelligence			
AI optimisation	–	•	–
AI advice / recommendations	⦿	⦿	•
What-if / measures (ROI)	•	⦿	–
Helpdesk / communication	–	•	•

• detailed view ⦿ summary view – not shown

Fig. 5.2 Coverage of the function catalogue across the three roles.

5.2.3 Initial design of the visualisation interface

Building on the function catalogue, an initial top-layer interface concept was drafted. The concept structures each role's screen into a limited number of sections to avoid overloading the user with technical detail.

A reference layout, informed by an existing industrial dashboard, organises the view into four main areas:

1. **Header section**, containing building identification, role-specific context and a headline efficiency indicator;
2. **Realised-costs section**, showing the largest sources of inefficiency and cost split by category;
3. **Cost-effectiveness section**, presenting an inefficiency tree across operations and zones, including status coding and trends;
4. **Supplementary-data section**, containing events, system status, metering status and related contextual information.

Role-specific variants were then drafted. They share a common header and a consumption-and-cost block, but differ in their additional sections. The Owner variant emphasises what-if measures, ranking of measures, investment cost, savings, emission impacts, scenario combinations and finance. The Manager variant adds indoor-environment quality, a richer alarm and notification centre and an equipment-and-inspections section. The Tenant variant focuses on personal consumption, indoor comfort, AI advice and a helpdesk.

The design follows the progressive-disclosure principle, so the same underlying data can be shown at the level of detail required by each user role. These designs represent first-iteration concepts. The production of the end-user frontend is carried out within the EMOS platform layer, with which the communication-interface design is being co-developed.

5.2.4 Validation survey of building professionals

To validate the basket of functions and the preferred forms of data visualisation, a short questionnaire was prepared for distribution to several dozen building professionals. The questionnaire captures:

- the respondent's professional role;
- which functions and information items they consider necessary;
- which visualisation or interface types they would most likely use;

- which single function they consider most useful;
- open feedback on missing functionality or usability expectations.

At the time of writing, the questionnaire has been finalised and data collection is in preparation. The results will be used to prioritise the basket of functions, refine role-specific views and select default visualisations for the EMOS frontend.

5.3 Energy audit from existing data and Module of changes

As part of WP4, an energy audit module based on existing operational, technical and economic building data was developed. The proposed approach was tested and validated on a pilot residential building for which energy and water consumption data, selected technical information and a digital building model were available. The energy audit and the module of changes link the building data layer with the energy BIM model standard (Section 5.5) and with the predictive and thermal models developed in Topic 3 (chapter 7); the module of changes also provides the building-level basis for the what-if evaluation of measures pursued in T4.4 (chapter 9).

The main goal of the module is to support a fast digital assessment of a building's energy performance without the need for time-consuming manual processing of all input data. The work focused on verifying whether an energy audit can be prepared from available data sources and on identifying the requirements for automated collection and processing of energy-related information from a digital building model.

5.3.1 Approach

The proposed approach is based on the use of available data on energy and water consumption, technical building systems, building envelope characteristics and operational information. These data are organised into a unified data structure that enables the creation of an energy balance, identification of major energy flows and evaluation of the building's energy performance.

The energy audit module is complemented by a **module of changes**, which supports the assessment of proposed energy-saving measures. The module enables comparison of the current and proposed building state and evaluates the technical and economic benefits of individual measures or their combinations.

One of the main challenges was working with incomplete input data, which is common in existing buildings. The pilot building represented a typical example of such a situation. While operational data on energy and water consumption were available, the digital model contained mainly graphical information and lacked detailed non-graphical data about building constructions and

technical systems. As a result, part of the required input information had to be derived from operational data, engineering assumptions and information provided by the building owner.

This situation provided a valuable opportunity to validate the proposed methodology and identify the minimum dataset required for automated generation of an energy audit and building energy balance.

5.3.2 R&D progress

During the development process, a methodology was created for transforming operational data into the structure of an energy audit. A data model was also developed to support automated processing of building energy inputs and outputs. In parallel, requirements for input data and procedures for energy-balance calculations were defined.

A detailed calculation tool was developed in MS Excel, primarily focused on building technical systems, including HVAC and domestic hot water systems. The tool includes calculations of energy flows from heating demand through heat generation and distribution to domestic hot water preparation.

An important part of the development was the creation of the module of changes, which enables simulation of the impact of energy-saving measures on energy consumption, operating costs and environmental indicators. The module also considers the synergy effects of multiple measures. The objective was to verify the accuracy of the proposed methodology and identify areas requiring further refinement.

Another activity addressed the integration of operational data with the digital building model. The project identified missing non-graphical information required for energy assessment and defined requirements for its structure and storage. The result is a proposed data model that will support future automated extraction of technical building parameters directly from a digital twin.

In addition, a sample energy audit was prepared as a reference document defining the required structure, mandatory chapters, calculation procedures and expected outputs. This document serves as a methodological basis for future automation of the energy-audit process.

5.3.3 Outputs

The main outputs achieved so far include:

- methodology for preparing an energy audit from existing data;
- energy-balance calculation tool focused on technical building systems;
- module for evaluation of energy-saving measures;
- specification of non-graphical data required in a digital building model;

- sample energy audit defining the structure and content of an automatically generated report.

These outputs provide the foundation for integrating energy audits, energy management systems, digital twins and artificial intelligence into a single digital solution.

5.3.4 Future goal

The next step is to transfer the developed calculation procedures into a software environment and integrate them with digital tools using artificial intelligence. The long-term objective is to gradually replace Excel-based calculations with an automated software solution.

The overall goal is to create a system capable of automatically generating an energy audit based on data from energy meters, utility bills and a digital twin of the building. Such a system should automatically create an energy balance, evaluate the technical condition of the building, propose energy-saving measures and support decision-making by building owners and facility managers.

The final result should be an intelligent tool that significantly reduces manual work, increases the efficiency of energy specialists and supports wider adoption of data-driven building energy management. The proposed solution represents an important step towards the digitalisation of energy audits and creates the foundation for automated assessment of building energy performance based on operational data and digital building models.

5.4 BIM visualisation and BIM Model Creation Approaches

The BIM-related work in Topic 1 addresses a practical challenge identified in WP4: energy-management workflows require a spatial and semantic representation of the building, but the availability and quality of BIM models differ significantly across real buildings. Some buildings already have detailed models, others only have partial project documentation, and many have no usable documentation at all.

For this reason, WP4 does not rely on a single BIM creation method. Instead, a portfolio of complementary approaches was developed and evaluated. These approaches cover the full range from minimal input data to detailed geodetic point clouds:

1. a custom-built IFC4 box model for the EMOS frontend;
2. simplified BIM generation without documentation, using airborne laser scanning or OpenStreetMap data;
3. BIM reconstruction from minimal 2D documentation;
4. semi-automatic box model creation from indoor point clouds;

5. detailed BIM modelling from terrestrial laser-scanning point clouds.

Together, these methods make it possible to select the modelling workflow according to the available documentation, required geometric accuracy, available resources and intended use in energy management.

5.4.1 BIM frontend model - custom-built "box" model

This methodology originated from a specific requirement of the EMOS platform team. EMOS collects and processes a wide range of building data and time series and needs to visualise individual measurement and consumption points inside a building. For this purpose, a detailed BIM model is not necessary. What is required is a deliberately simple, schematic model delivered as a strictly schema-conformant IFC file, which can serve as a stable spatial and identity framework for the data handled by the EMOS processing microservices.

The work presented here provides such a framework: a custom-built application that produces valid IFC4 box models and assigns every element a persistent identity.

The result is a lightweight web application rather than a full BIM authoring environment. Its geometry is deliberately schematic, with building footprints extruded into box-like prisms. The primary deliverable is the IFC file itself, while the built-in 3D viewer is used mainly for visual checking of the generated model.

All input workflows feed into a single generation engine. This engine is a Python backend built on FastAPI and IfcOpenShell, accessed by the user through a browser interface that combines a form, a 2D drawing canvas, a map and an IFC viewer based on ThatOpen Components with the web-ifc engine. The available ways of working differ only in how the footprint and storey heights are obtained.




Fig. 5.3 Input options feeding the common generation engine, the resulting IFC4 box model, and its use in the EMOS platform

Three input options were developed step by step.

The first option is a simple form. The user enters approximate building dimensions, and the service arranges the elements into a schematic layout. This already satisfied the original EMOS request and confirmed that a lightweight spatial representation can be sufficient for data visualisation.

The second option is used when at least a scanned floor plan, such as a PDF or JPG, is available. The user sets the scale from one known dimension and traces the outline, zones and machines over the underlay. This enables zones and equipment to be positioned more realistically inside the building footprint.

The third option starts from a map. The building is found by address or selected on an OpenStreetMap background, and its outline is taken directly from the map. This option is particularly useful when no building documentation is available.

The floor-plan and map-based workflows can also be combined. A floor plan can be loaded on top of the map and aligned to the building outline using a two-point alignment, so that a real plan can be traced in its correct geographic position. The map-based way of obtaining footprints for buildings without documentation is discussed from the data-source perspective in Section 5.4.2 (Methodology 2).




Fig. 5.4 The web application: the parametric form on the left, and the resulting IFC4 box model (a translucent building envelope with zones and machine boxes inside) shown in the built-in 3D viewer on the right.




Fig. 5.5 Combining the map with a floor plan: an uploaded plan (lower panel) is fitted onto the building outline taken from the map by a two-point alignment, so that the traced zones land in the correct geographic position.

Regardless of which input option is used, the model is written with the correct IFC spatial hierarchy, from *IfcProject* through *IfcSite*, *IfcBuilding* and *IfcBuildingStorey* to *IfcSpace*. Spatial levels are linked by aggregation, and physical elements are linked by containment. Each footprint, whether a rectangle or a free polygon, is extruded by the storey height into an *IfcExtrudedAreaSolid*. Zones and the building envelope are made translucent so that their contents remain visible while the whole volume can still be highlighted.

In IFC terms, a storey becomes an *IfcBuildingStorey*, a zone or local area becomes an *IfcSpace*, a machine becomes an *IfcBuildingElementProxy* and a production line becomes a logical *IfcSystem* group without its own geometry.




Fig. 5.6 The IFC4 spatial hierarchy, the entity mapping, and the deterministic identity given to every element.

The main distinguishing feature of the tool is not its geometry, but the identity it assigns to every element. Each element receives a deterministic GlobalId derived from its own identifier through a fixed-namespace hash. As a result, the same element always keeps the same IFC GUID even after the model is regenerated. In parallel, a Pset_EMOS property set records the element identifier, level, parent and name. This enables EMOS to find and colour any element at any level of the hierarchy, either by its GUID or by a property query. This behaviour is essential for visualising live measurement and consumption data.

Models created from the map are also georeferenced into the Slovak national coordinate system S-JTSK (JTSK03 / Krovak East North, EPSG:8353) using IfcProjectedCRS and IfcMapConversion. The model records the easting and northing of the site origin and the orientation of the model relative to the grid, including the meridian convergence of the Krovak projection. The geometry remains close to a local origin to preserve numerical precision in viewers, while the real-world position and rotation are stored as metadata readable by GIS tools and the EMOS platform.

Every generated model is opened again with IfcOpenShell and checked automatically. The checks cover the IFC4 schema, entity counts, spatial relationships, containment and aggregation, deterministic GlobalIds, Pset_EMOS properties and, for georeferenced models, the agreement of a reprojected control point with a direct coordinate transformation. For georeferenced models, the control-point difference is expected to remain below 0.1 m.

At the current stage, all three input options are available and are being evaluated side by side. The next steps are to validate schema conformance and georeferencing more thoroughly and to assess the accuracy of dimensions and positions against reference data. The final selection of the approach to be integrated into EMOS will be based on the results of this evaluation.

Inputs: approximate building dimensions entered through a form, a scanned floor plan (PDF or JPG) traced over a scaled underlay, or a footprint selected from an OpenStreetMap background by address; optional storey heights and zone and machine positions.

Outputs: a schema-conformant IFC4 box model with the full spatial hierarchy (IfcProject → IfcSite → IfcBuilding → IfcBuildingStorey → IfcSpace), a deterministic GlobalId and a Pset_EMOS property set on every element and, for map-based models, georeferencing to S-JTSK (JTSK03).

Validation: every generated model is re-opened with IfcOpenShell and checked automatically for IFC4 schema conformance, spatial relationships, deterministic GlobalIds and Pset_EMOS properties; for georeferenced models the reprojected control-point difference is expected to remain below 0.1 m. Dimensional and positional accuracy against reference data is still to be assessed.

5.4.2 Simplified BIM Model Generation without Building Documentation

At the current stage of the R&D activities, two complementary methodologies for automated generation of simplified BIM models without building documentation are being developed and evaluated. While Section 5.4.1 describes the common IFC generation engine, this subsection focuses on external data sources and methodological assumptions used when no building documentation is available

Methodology 1 – Automatic BIM Box Model Generation from Airborne Laser Scanning Data




Fig. 5.7 Workflow of the developed methodology for automatic BIM box model generation from classified airborne laser scanning (ALS) data and export to IFC format.

The process starts with loading classified point clouds in LAZ format. A Digital Terrain Model (DTM) is generated from terrain points classified as Class 2, while points classified as buildings (Class 6) are used to determine building heights above the terrain surface.

Individual buildings are identified through point-cloud segmentation using the DBSCAN clustering algorithm. Building footprints are then generated from the segmented clusters and extruded according to the estimated building height. The result is a simplified 3D box representation of each building, exported into the Industry Foundation Classes (IFC) format.

Inputs: classified ALS point clouds in LAZ format, specifically terrain points (Classification 2) and building points (Classification 6).

Outputs: Digital Terrain Model and simplified BIM box models in IFC format.

Validation: The methodology is currently under development. Initial testing indicates that positional accuracy, height estimation and overall building dimensions are not yet sufficiently accurate for reliable deployment in production workflows. Further research is therefore focused on improving geometric reconstruction, segmentation quality and validation against reference datasets.

Methodology 2 – BIM Box Model Generation from Address and OpenStreetMap Data

The second methodology addresses buildings for which no documentation is available by deriving the building geometry from address information and OpenStreetMap data. Through a graphical user interface with an OpenStreetMap background, the user searches for the building by address or selects it directly on the map.

The building footprint is obtained from the building geometry already mapped in OpenStreetMap. The outer outline of the way or relation that contains, or lies nearest to, the selected point is retrieved and, for multi-part buildings, stitched into a single closed polygon. The number of storeys is read from the OpenStreetMap building:levels attribute when present, defaulting to one when absent. This value can be adjusted by the user, who also specifies the floor-to-floor height. The total building height is then calculated as the product of the number of storeys and the floor height.

The retrieved footprint and parameters are passed to the box model generator described in Section 5.4.1. The generator creates the IFC4 model, assigns the spatial hierarchy and element identities and georeferences the result. Local areas or zones can be drawn within the footprint and are exported as IfcSpace entities.




Fig. 5.8 Map-based selection: the building is located by address and its OpenStreetMap footprint is retrieved as the model footprint (example: the Faculty of Civil Engineering, STU in Bratislava)

Inputs: building address, building footprint obtained from OpenStreetMap, number of storeys from OpenStreetMap or user input and user-defined floor-to-floor height.

Outputs: simplified IFC4 box model of the selected building, with individual storeys represented as IfcBuildingStorey entities and user-defined local areas represented as IfcSpace entities, georeferenced to S-JTSK (JTSK03).

Validation: The footprint and storey count depend on the completeness of OpenStreetMap data and are not yet validated against reference building data. IFC schema conformance and georeferencing are validated as part of the generation engine. Future work will assess

dimensional accuracy, orientation, positional correctness and comparison with reference building data.

5.4.3 BIM from minimal documentation

A BIM model of an existing building can be reconstructed from 2D documentation either manually or through automated 2D-to-3D generation. The decisive factors are the modelling time required and the geometric accuracy achieved. To evaluate this trade-off in measurable terms, both routes were applied to the same pilot building and the resulting IFC models were compared.

The main research question was: how much time does automation save, and at what cost in accuracy and completeness?

Pilot building and inputs

The pilot building is a two-storey detached family house with 230 m² floor space and no basement. The input was DWG project documentation, including floor plans and sections. A terrestrial 3D scan served only as an independent as-built reference.

Two models were produced from the same documentation and exported to IFC:

- a manually authored model in Autodesk Revit, used as the reference;
- a model created by an automated 2D-to-3D generation service.

Capture-based routes starting from point clouds are addressed separately in Sections 5.4.4 and 5.4.5.

The 3D scan showed that the DWG documentation did not fully match the as-built state. For this reason, the manually authored model, corrected against the scan, was used as the reference throughout the comparison rather than the documentation itself.




Fig. 5.9 Manually authored IFC model (architecture), used as the reference as-built state

Achievable model level depends on the input

The achievable level of detail depends directly on the available input. Documentation alone, such as PDF or DWG drawings, reliably supports L1 and L2 models, while L3 additionally requires a 3D scan for as-built verification and 360° photographs for the MEP layer. The documentation-only route therefore delivers verified architecture but stops short of a complete L3 model.

Tab. 5.1 Model levels and the documentation typically required for each.

Level	Model content	Typical input	Accuracy
L1 — Box	Volumes per storey; external envelope only; no openings	Hand survey or higher	±500 mm
L2 — Mid-detail	Walls, windows, doors, slabs, basic roof; topology not exact	PDF / DWG plans + sections	±100–200 mm
L3 — Full BIM	Complete as-built architecture + parameters; optional MEP	Docs + 3D scan + 360° photos	±50 mm

With full project documentation, an independent scan and 360° photographs, the pilot building qualifies for L3, which is why both modelling routes were compared at this level.

Manual modelling versus automated generation

The manual route provides the highest accuracy and full parameter control, but it is the most time-consuming option. For the pilot building, the total time was approximately 27 hours, consisting of approximately 20 hours for architecture and 7 hours for MEP modelling.

The automated route generates a structured IFC model in seconds. Elements are correctly classified as walls, windows, doors, slabs and spaces, but the output contains geometry only, without parameters, and includes systematic deviations. Stepped or recessed floors first require a calibration step to align storeys.




Fig. 5.10 Model produced by the automated 2D-to-3D generation service from the same floor plans.

Tab. 5.2 Manual modelling versus automated generation, including modelling time and estimated time saving for the pilot building.

Aspect	Manual (reference)	Automated 2D→3D
Maturity	Production	Production, with limitations
Time to L2	~10 h	~3–4 h (~60% saving)
Time to L3 (architecture)	~27 h	~16–19 h (~20–30% saving)
Input requirement	Any documentation quality	Clean plans; stepped plans need calibration
Operator skill	High (BIM specialist)	Medium
Geometric accuracy	Highest (verified vs. scan)	Medium; systematic deviations
Physical parameters	Fully assigned	None — added manually
MEP capability	Yes (with 360° photos)	No
Scalability	Low	High for L1/L2

Geometry is generated in seconds; however, the times above include input preparation, calibration and manual refinement to each level. Time savings are largest at L1 and L2, where the generated geometry is directly reusable. At L3, the savings decrease because manual parameter enrichment, as-built correction and MEP modelling dominate and are not automated.

Accuracy of the generated model


Measured against the manual reference, the generated model matches it to within about 90% of external dimensions. The remaining deviations are not primarily a scale problem but a data-extraction problem. The generation service reads only the opening width visible in the 2D floor plan and has no access to height dimensions provided as text annotations. As a result, each window and door is assigned a default height instead of the actual value. For doors, only the width is recovered, while height is defaulted in the same way.

Five of the 20 doors are missing entirely, and a similar default-and-omission pattern affects the staircase and several wall junctions. Most of these remaining errors could be reduced by more careful preparation of the input drawing, but the additional preparation time would likely exceed

the time saved by automation. Physical parameters such as U-values and materials are absent from the output and must be added manually regardless of the modelling route.


MANUAL BIM MODEL

Ground Floor




AUTOMATED BIM MODEL

Ground Floor



MANUAL BIM MODEL

Second Floor



AUTOMATED BIM MODEL

Second Floor



Fig. 5.11 Ground-floor and second-floor plans of the manual reference and the generated model, showing geometric differences in wall layout and openings

Key findings

Automated 2D-to-3D generation is best used to create L1/L2 models or to accelerate the early stage of an L3 model. It should not be considered a full replacement for manual modelling. Every generated element still requires verification, and physical parameters must be added manually.

The next steps include testing alternative 2D-to-3D generation tools against the same documentation to determine whether the observed limitations, such as default heights and missing elements, are tool-specific or inherent to the approach. Further work will also explore AI-

assisted population of the generated scaffold with physical parameters to reduce the manual enrichment effort that currently limits its use for L3 models.

5.4.4 Automatic box model creation from point cloud

The developed software implements a semi-automatic Scan-to-BIM workflow designed to transform a 3D point cloud into a structured IFC representation of the indoor environment of buildings. The software is organised as a sequential geometric pipeline consisting of point-cloud import, preprocessing, plane segmentation, semantic classification of structural surfaces, space reconstruction, manual correction and export to IFC format.

Approach

The input point cloud can be imported in several formats, including *.e57, *.ply, *.pcd, *.pts, *.txt and *.csv. To improve processing stability, the data can be voxelised and cleaned using Statistical Outlier Removal or Radius Outlier Removal.

Plane segmentation is based on the estimation of local normals using PCA over the k-nearest neighbours. For better computational efficiency, the calculations are performed in chunks with the support of a KD Tree. Planes are extracted iteratively through a region-growing approach, where a local region is formed from a seed point, the plane model is refined using SVD and inliers are selected according to their distance from the plane and normal deviation.




Fig. 5.12 Segmented and classified point cloud (Green – floor, Blue – walls, Yellow – Ceilings)

The detected planes are then classified as floors, walls, ceilings, ambiguous elements or discarded elements. Room reconstruction is performed floor by floor. Ceiling planes are projected onto the XY plane and converted into an occupancy raster mask, from which room contours are extracted, refined and aligned with wall lines derived from the classified wall planes. The resulting polygons are extruded between the floor and ceiling levels to create 3D room volumes, i.e. a space-based box model.




Fig. 5.13 Created IfcSpace with adjustable vertexes (blue dots)

A key feature of the software is the possibility of interactive manual correction. The user can create, delete and edit rooms, modify room contours, move, add or remove vertices and manually adjust room height. This allows the user to inspect intermediate and final results and correct the reconstructed geometry when necessary.

Inputs

The methodology uses indoor 3D point clouds captured by terrestrial laser scanning or similar surveying methods. Supported import formats include *.e57, *.ply, *.pcd, *.pts, *.txt and *.csv.

Outputs

The final output is exported to IFC4 as IfcSpace entities with metric units, local placement, closed room profiles and extrusion-based geometry. The model provides a lightweight spatial representation of interior spaces suitable for visualisation and linkage with energy or sensor data.

Validation

The method has been tested on point-cloud data with the aim of verifying the feasibility of semi-automatic room reconstruction. The current implementation successfully supports plane segmentation, classification, room contour generation and manual correction. Further validation is required to quantify geometric accuracy, robustness across different point-cloud densities and sensitivity to incomplete scans or occluded areas.

5.4.5 Detailed BIM modelling from point clouds

This methodology addresses situations in which the energy-oriented BIM workflow has to be applied to an existing building that has no available BIM model. Many buildings are modelled during design as part of project documentation, but for a large part of the existing building stock such a model is missing and must be created retrospectively.

Detailed manual modelling from a geodetic point cloud is the most accurate way to obtain such a model. A surveyor captures the building by terrestrial laser scanning and delivers a georeferenced point cloud that serves as the underlay for modelling the structure. Only surfaces that are visible and capturable by the TLS scanner can be reconstructed, while concealed layers and hidden building services cannot. For example, the total thickness of a slab can be measured, but its individual layers and material composition remain unknown unless this information is supplied from other sources. Despite this limitation, visible geometry can be reconstructed with high accuracy.

Description of the Approach

Modelling is performed in a BIM authoring tool, such as Autodesk Revit, into which the georeferenced point cloud is linked. The first and decisive step is aligning the coordinate systems of the BIM project and the point cloud, which is referenced in the Slovak national coordinate system S-JTSK. This ensures that the final model is located and oriented exactly as built.

The point cloud is then sectioned in plan and elevation. The resulting slices expose the edges of walls and other elements, from which walls, windows, doors, slabs, roofs and stairs are modelled. The resulting detailed geometry is then enriched with non-graphical information, such as element and material parameters where known, building name and address, and names and functions of individual spaces.

The workflow is demonstrated on the RD Zobor case study, a two-storey detached house. The georeferenced point cloud serves as the underlay from which the detailed BIM model is reconstructed.



Fig. 5.14 Input georeferenced point cloud (left) and the resulting detailed BIM model (right), RD Zobor case study

Inputs

The input is a georeferenced terrestrial laser-scanning point cloud of the building, delivered by a surveyor in the S-JTSK coordinate system.

Outputs

The output is a detailed, georeferenced IFC model that reproduces the as-built geometry and is organised as a spatial hierarchy from IfcBuilding through IfcBuildingStorey to individual IfcSpace entities. Each room is represented as a separate space named according to its function, such as entrance hall, corridor, bathroom or living room. Components such as sensors and heating elements are assigned to the relevant storey and space.

Beyond geometry, the model provides spatial data required for energy management, including room volumes, orientation of spaces relative to cardinal directions and distinction between heated and unheated spaces. Structural elements carry thermal and material parameters such as thermal conductivity, density, specific heat capacity and emissivity, which are used by the energy model.




Fig. 5.15 Structural elements enriched with material and thermal parameters used by the energy model, RD Zobor case study.

Validation

Compared with the automated approaches in Sections 5.4.1 to 5.4.4, detailed modelling from a point cloud is the most time-consuming option, but it provides the highest geometric fidelity and the richest information. Because the point cloud captures the real as-built shape, which can differ significantly from the original design documentation, the resulting model is the closest achievable approximation of reality and the most reliable basis for the BIM-based energy workflow.

The main limitation is the inability to capture concealed layers and hidden services. These must be supplied from documentation, expert assumptions, destructive inspection or other sources where required.

5.5 BIM energy model

5.5.1 3D data energy BIM model STANDARD

While Section 5.4.5 describes how a detailed BIM model can be obtained, this section specifies the data standard that such a model must meet in order to serve energy management. The standard defines how the model is structured, which non-graphical parameters it must contain, to which elements these parameters are assigned and how the model is delivered. This ensures that the model can be reliably reused for energy calculation, BIM visualisation and integration within the EMOS layer.

Open Data Format

The model is delivered in the open IFC format, an internationally recognised, machine-readable standard that supports interoperability between different software tools. Every entity is defined by its class, such as IfcWall, IfcMaterial, IfcSpace or IfcBuilding, and information is attached through objectified relationships, such as IfcRelDefinesByProperties.

A property can be assigned either to an individual object or to a superior spatial class, such as IfcBuildingStorey or IfcBuilding, from which other objects can inherit it. This allows values defined once to be referenced rather than repeated for every element, keeping the IFC file compact and unambiguous.

Required Information

For energy management, the model must contain a defined set of non-graphical parameters grouped by object type. Each parameter is stored with an explicit data type, such as length for thickness or real number for thermal conductivity, so that the value remains machine-readable.

Tab. 5.3 Required non-graphical parameters of the energy BIM model standard, by object group.

Object group	Required non-graphical parameters [unit]
Windows and doors	Thermal transmittance $[W/(m^2 \cdot K)]$; light transmittance [-]
Walls, slabs, floors and internal partitions (per layer)	Layer thickness [m]; thermal conductivity $[W/(m \cdot K)]$; specific heat capacity $[J/(kg \cdot K)]$; density $[kg/m^3]$
External walls (surface)	Surface emissivity or albedo [-]
Room equipment (solid bodies)	Mass [kg]; specific heat capacity $[J/(kg \cdot K)]$

Implementation

Several ways of writing the required parameters into the model were evaluated. Exporting the material property sets offered directly by the authoring tool proved unsuitable. In Autodesk Revit, the built-in exporter can export material property sets only as a whole, which adds many more parameters than needed and unnecessarily increases the size and complexity of the model.

For this reason, a dedicated property set was defined for each object group, such as walls, windows, doors, room equipment and spaces. Each property set is exported as a separate IFC set containing only the required parameters.

Construction layers, which are usually modelled as a single element rather than as individual layers, are exported as parts. For example, the screed, acoustic insulation and walking layer of a floor are exported so that each layer becomes an IfcBuildingElementPart with its own property set.

Covering	
Summary	Location
Material	PartOf
Conflicts	IFC_Parts
Property	
Value	
Hrubka	0.03 m
Hustota	500
Kategoria	Floors
Material	DŘEVO - podlaha
Merna_tepelna_kapacita	2510
Miestnost	2.4.04
Podlažie	2NP
Poradie	1
Tepelna_vodivost	0.13




Fig. 5.16 A standardised property set (IFC_Parts) assigned to an element, with its parameters and values; apartment building in Střešovičky.

```
#1740=IFCPROPERTYSINGLEVALUE('Merna_tepelna_kapacita',$,IFCREAL(900.),$);
#1741=IFCPROPERTYSINGLEVALUE('Tepelna_vodivost',$,IFCREAL(0.7000000000000007),$);
```

Fig. 5.17 Excerpt of the IFC file, where each parameter is stored as a property value with an explicit data type (here IFCREAL).

The outcome is a fixed data structure that must always be followed so that the model can be used reliably as an input and information container for the energy calculation developed by the IT4I team from the supplied BIM model.

The standard has been applied to two case studies: the RD Zobor family house used in the previous sections and an apartment building in Střešovičky, for which an energy audit was carried out and a BIM model was created for energy calculation.

Visualisation and EMOS

Because the model represents the real geometry together with the spatial and material structure of the building, it is also suitable for visualisation. Spaces and other elements are separate data objects, each with a unique identifier, placement and orientation. Therefore, any of them can be individually addressed in the EMOS user environment and coloured, highlighted, filtered or hidden, with every piece of information linked to the correct object.

The energy BIM model standard therefore serves two purposes. It provides structured technical input for energy calculations, and at the same time it provides a consistent spatial and semantic layer for user-facing visualisation.

5.6 Integration with EMOS / data pipeline

The BIM outputs described in Sections 5.4 and 5.5 are integrated into the broader WP4 data pipeline and EMOS platform layer. Their role is to provide the spatial and semantic structure required to connect measurements, energy data, equipment, zones, rooms, building systems and user-facing visualisations.

From the perspective of the data pipeline, the BIM model provides:

- spatial hierarchy, including building, storey, space, zone and equipment relationships;
- stable identifiers for linking time-series data to physical or spatial objects;
- geometric information required for visualisation and selected energy calculations;
- non-graphical properties required for energy modelling;
- object classification and grouping needed for analytics and reporting.

From the perspective of EMOS, the BIM model provides the visual and identity framework for displaying live and historical data. Sensor readings, consumption values, detected inefficiencies and optimisation recommendations can be mapped to the corresponding building spaces or components. This allows the user to interpret data spatially rather than only as tables or charts.

Two levels of BIM integration are therefore supported:

1. **Lightweight visualisation models**, such as IFC4 box models, which are sufficient for displaying measurement points, zones and consumption blocks in the EMOS frontend;
2. **Detailed energy BIM models**, which contain enriched material and thermal parameters required for energy calculation and more advanced digital-twin workflows.

The selection of the appropriate model type depends on the use case. For rapid deployment and visualisation of sensor data, a lightweight box model may be sufficient. For thermal simulation, energy audits or detailed digital-twin modelling, a more detailed model with non-graphical parameters is required.

The integration with the EMOS platform is based on stable identifiers, particularly IFC GUIDs and the Pset_EMOS property structure. These identifiers ensure that data streams remain linked to

the correct spaces, equipment or building elements even when models are regenerated or updated. This is essential for long-term operation, historical analysis and validation of optimisation measures.

The BIM integration layer therefore acts as a bridge between the technical data pipeline and the user-facing EMOS environment. It connects data ingestion, preprocessing, analytics, digital-twin modelling and visualisation into a coherent spatial framework.

5.7 Open issues and next steps

The work carried out in Year 2 established the conceptual and technical basis for the BIM and visualisation components of WP4. Several outputs are already available, including the role-based interface concept, the energy audit methodology, the module of changes, the custom IFC4 box-model generator, simplified BIM generation approaches, point-cloud-based modelling workflows and the energy BIM model standard.

However, several issues remain open and will be addressed in Year 3.

First, the role-based communication interface must be validated with building professionals. The questionnaire prepared in Year 2 will be distributed and evaluated, and its results will be used to refine the basket of functions and the default visualisations for each user role.

Second, the different BIM creation approaches must be evaluated against reference data. This includes validation of dimensional accuracy, positional accuracy, georeferencing, schema conformance and suitability for EMOS integration. Particular attention will be paid to the comparison between automated and manual workflows, including their time requirements, achievable model quality and required level of manual correction.

Third, the custom IFC4 box-model generator requires further testing in real pilot contexts. The next step is to verify how reliably generated models can be used in the EMOS frontend for the visualisation of measurement points, zones, consumption blocks and operational states.

Fourth, simplified BIM generation from OpenStreetMap and airborne laser scanning data requires further validation. The main open questions are the completeness and accuracy of source data, correct interpretation of storey numbers and heights, and the reliability of generated footprints compared with reference building geometry.

Fifth, the energy BIM model standard must be refined and aligned with the needs of the energy calculation modules and EMOS data layer. The main task is to ensure that the required non-graphical parameters are neither insufficient nor unnecessarily excessive. The standard should remain compact, machine-readable and reusable across different buildings and modelling workflows.

Sixth, integration between BIM outputs, the WP4 data pipeline and the EMOS platform must be further specified and tested. This includes the definition of interfaces, identifier mapping, property-set requirements and procedures for updating or regenerating models without breaking data links.

Seventh, the energy audit module and the module of changes must be transferred from the current Excel-based calculation into the software environment and integrated with the energy BIM model standard (Section 5.5) and the Topic 3 models, so that audits and measure evaluation can be generated automatically from meter data, utility bills and the digital twin. Their outputs also need to be aligned with the what-if and ROI evaluation pursued in T4.4 (chapter 9).

These seven items constitute the Year 3 priorities for Topic 1. They are reflected in the consolidated per-topic priorities of section 3.5 and in the focus for the next period set out in section 10.3.

These activities will ensure that the BIM and visualisation components developed in WP4 can support not only research and testing, but also practical deployment in building energy management workflows.

Functional-strategy validation (TSK).

The communication, visualisation and BIM work of Topic 1 is continuously validated by the Technical Road Administration of Bratislava (TSK) as the prospective end user of the solution. TSK reviews the user roles, the role-based function basket, the audit and module-of-changes outputs and the BIM representations against operational reality, and flags requirements that the research-side modules might otherwise overlook. Its findings feed the requirements iterations of T4.1 and the frontend definition carried into the EMOS platform layer; the validation roles of TSK and DFKI are described together in section 8.7.

6 (Topic 2) — Data preparation pipeline

Lead: Asseco CE

6.1 Introduction and scope

This chapter describes the data preparation pipeline developed in WP4 for converting heterogeneous input data into a controlled and analytically usable form. Its scope covers the full preparation cycle: intake of raw files, technical screening, standardisation of structure and semantics, validation of quality and consistency, targeted modification where required, and handover of prepared outputs as time-series data for downstream analytical and predictive model processes.

Current work is focused on turning this preparation cycle into a dependable operational workflow, combining intake, pre-processing, validation, correction-oriented modification and output preparation into a sequence that remains understandable to both operators and later report users. Topic 2 therefore addresses much more than file ingestion alone: it frames data preparation as a managed end-to-end workflow whose value lies in making incoming data repeatable, auditable, comparable and suitable for later analytical and AI-oriented use. In practical terms the work already supports four core objectives — reliable handling of larger input batches, a traceable pre-processing flow, a stable time-series representation across stages, and quality-controlled artefacts that can support later predictive modelling, reporting and operational decision-making.

Fig.

6.1 Backend components and interactions for Data Import, Validation and Modification modules.

6.2 Data input layer

Owner: Roman Rigo (Asseco CE)

This section describes the core Year 2 methodological work carried out across intake, validation, and modification. Its role is to show how the preparation process was organised in practice and how individual stages were shaped into a coherent workflow.




Fig. 6.2 Backend data preparation flow (Import -> Validation -> Modification).

6.2.1 Data import module

Scope

This workstream focused on the controlled reception of heterogeneous input files and on the separation between lightweight intake and heavier downstream processing.

Current activities

During Year 2, the work established a structured intake layer for heterogeneous inputs and separated lightweight file acceptance from heavier downstream processing. This made it possible to accept files in a responsive manner while preserving traceability of subsequent preparation steps.

The methodological emphasis was on robust technical intake: initial checks, controlled acceptance of supported file types, asynchronous continuation of processing, and explicit recording of progress. This arrangement reduced coupling between user-facing upload actions and computationally heavier transformation stages, which is important for stable operation under variable workload.

Results

The main result of this activity is a more dependable intake stage that supports responsive file acceptance while preserving auditability and operational continuity across the subsequent preparation chain.

Next steps for Year 3

In Year 3, the intake stage should be strengthened through clearer operational metrics, better observability, and tighter linkage between intake decisions and downstream quality-control requirements.

6.2.2 Validation module

Scope

This workstream addresses deterministic quality assessment of prepared series, stored series records, and grouped collections of series. In the current implementation this capability lives in the companion validation service, not in the data-loading service itself. The same workstream also covers an optional human-readable explanation layer and the reuse of the same quality logic after modification of cumulative meter data.

Current activities

The validation module currently exposes series-level checks for required fields, timestamp ordering, timestamp validity, value validity, duplicates, and related quality conditions. The same logic is available for single-series payloads, validated collections, and payloads loaded from storage in the validation service.

Validation is treated as a deterministic quality-control layer that produces a structured result with issues, warnings, and aggregate statistics. It can also run with a selected subset of checks, which makes targeted diagnostics possible. When available, the explanation path passes the deterministic result to a single LLM-backed reporting step and attaches a plain-language summary, likely cause, threshold advice, and model metadata.

Results

The main result of this activity is a reusable quality-assessment layer that can evaluate series data, stored records, and grouped series with the same deterministic rule set while preserving an optional explanation path for human readability. In practice, this means the validation logic can be reused across direct validation, stored-payload checks, and the AI-explained path without changing the deterministic base result.

Next steps for Year 3

In Year 3, validation should be extended with clearer reporting of the selected checks, stronger traceability for stored data assessment, and richer use of the explanation layer only as advice rather than as a decision-maker. Any broader agent-based orchestration should remain a future enhancement, not a current dependency of the validation flow.

6.2.3 Modification module

Scope

This workstream focuses on cumulative-meter modification in the shared transformation pipeline. In the current implementation this capability also lives in the companion validation service, where it supports standalone transformation and the optional pre-assessment modification path.

Current activities

The modification module currently exposes a reusable transformation service and a standalone modification path. Internally, the pipeline can perform spike handling, timestamp deduplication, cumulative to interval conversion, missing-value filling, resampling, and final checking of the modified output.

The current implementation also keeps anomaly flags on the output points so that modification decisions remain traceable.

Results

The main result of this activity is a dedicated transformation pipeline for cumulative-meter series that can be reused both independently and as a pre-assessment step, while keeping the output explainable

through a summary and point-level flags. This is a concrete service capability in bs-data-validation; the loading service only prepares input and context for such downstream use.

Next steps for Year 3

In Year 3, the modification layer should be documented more explicitly around pipeline options, anomaly-flag handling, and the persistence path used when modified output is stored separately. A tighter handoff contract between loading, validation, and modification should also be clarified so the current service boundaries stay explicit.

6.3 Data pre-processing layer

Owner: Petra Hlavinová (Asseco CE)

Introduction and scope

The Data Pre-processing Layer directly follows the outputs of the Data Input Layer ([Section 6.2](#)). It operates on data that has already passed the import, validation, and modification stages. The purpose of this layer is not to revalidate the inputs, but to prepare them for the modelling part of the pipeline: to identify relationships between variables, detect redundancy and multicollinearity, select appropriate data segments for training, and quantify the time lag between inputs and target variables.

Activities and methodology in Year 2

During Year 2, the layer was designed and implemented as three interconnected functional blocks: a correlation analysis block ([Section 6.3.1](#)) an automatic block segmentation block ([Section 6.3.2](#)) and a data shifting block ([Section 6.3.3](#)). Each block addresses a specific aspect of the preparation pipeline: identifying informative predictors and their redundancy, partitioning the dataset into hierarchically coherent blocks, and aligning paired series in time before lag analysis. The detailed methodology of each block is described in the respective subsection.

All three blocks are implemented as asynchronous services, with intermediate results persisted to MongoDB, which allows the reuse of artefacts in both the reporting and modelling steps without recalculation.

Implementation outputs:

The layer is delivered as three standalone analytical services which are currently being prepared for integration into the WP4 application service:

- **Correlation analysis block:** produces correlation matrices, multicollinearity flags, importance rankings, and a tiered feature shortlist (strong / watchlist / drop).
- **Automatic block segmentation block:** produces a reconstructed measurement hierarchy and the resulting block decomposition.
- **Data shifting block:** produces a recommended time shift per driver–target pair, the full lag-response curves, the rolling-window stability profile, and the peak-matching diagnostics. It also delivers the batch-to-timestep redistribution outputs

Each block generates the analysis results in two forms: a machine-readable JSON payload consumed by the modelling pipeline, and a section of the consolidated HTML report rendered for domain experts (heatmaps, scatter grids, hierarchy diagrams, lag curves, rolling-lag heatmaps, peak-matching plots, per-predictor quality assessments).

Validation status:

During Year 2, the layer was tested on three industrial datasets of different scale and signal mix. The methodology and the outputs of all three blocks have been reviewed iteratively under the project lead's methodological and domain oversight. The consolidated HTML report covering all three blocks is the primary medium through which the outputs are reviewed and adjusted iteratively.

The correlation analysis block has been validated functionally through standard sanity checks on its individual components, combined with visual review of its outputs in the HTML report on the real datasets. Integration with the downstream modelling pipeline has not yet been exercised and is part of the Year 3 work.

The automatic block segmentation block has been validated qualitatively on the simpler hierarchies present in the test datasets, where the reconstructed parent-child structure matches the expected topology. For more complex hierarchies the reconstruction does not yet reach the accuracy needed for a meaningful quantitative comparison against a ground-truth topology, and improvement of the scoring schema for these cases is part of the Year 3 work (see (f)).

The data shifting block has been validated qualitatively on the real datasets: the recommended lags, the underlying lag-response curves and the rolling-window have been reviewed visually through the HTML report, and the batch-to-timestep redistribution produces per-timestep production series whose energy alignment can be verified against the underlying energy profile.

Open issues and next steps for Year 3:

1. Completion of in-progress functionality of the three blocks. A subset of the planned functionality of the correlation, segmentation and data shifting blocks is still under development. Year 3 will begin by closing these open items so that each block reaches a stable feature scope before further integration steps.
2. Improvement of the block segmentation on complex hierarchies. The scoring schema used to reconstruct the parent-child measurement hierarchy works well on simpler installations but produces unreliable results on complex ones. The goal is to reach a reliability level that makes quantitative comparison against a ground-truth topology meaningful.
3. Integration of the analytical services into the WP4 service application. The three blocks are currently delivered as standalone analytical services. Year 3 will see them integrated into the main WP4 service application, so the pre-processing layer becomes a callable, versioned part of the production pipeline rather than a standalone analytical workflow.
4. Extension of the validation scope. Validation will be extended to additional industrial datasets covering different installation types, signal mixes and batch reporting conventions. A quantitative validation of the block segmentation and the data shifting blocks is planned once the underlying reliability allows it, and the integration with the modelling stage will provide an additional end-to-end validation channel.

6.3.1 Correlation analysis — preparation for model + reporting

Scope:

The correlation block measures how informative each numeric predictor is about the target — and how much of that information overlaps with the other predictors - so the modelling stage gets a feature selection it can use directly and the report gets an interpretable view of the same picture.

Methodology

Pairwise relationships are evaluated using multiple correlation methods (linear, monotonic and rank-based) to capture different dependency types, and the resulting pairs are visualised in scatter grids. A multicollinearity diagnostic identifies predictors that are strongly collinear and flags them for exclusion from the downstream ranking and from the feature set passed to the model. The remaining predictors are then scored by importance, producing a tiered ranking from highly informative to safely droppable features.

Findings

Combining three correlation methods rather than relying on a single one materially reduces the number of meaningful relationships missed in industrial data, where non-linear and saturated dependencies are common. The multicollinearity diagnostic prevents the importance ranking from being distorted by the same effect being counted several times under different names, which would otherwise inflate the apparent value of redundant predictors. The resulting feature tiering allows the modelling pipeline to consume a stable shortlist directly, instead of redoing pairwise comparisons on the full predictor set.

6.3.2 Correlation — automatic block segmentation

Scope

A full-dataset correlation matrix is often noisy, as high correlations between unrelated signals may simply reflect shared daily cycles. More meaningful relationships usually exist within the measurement hierarchy (e.g., parent meters and their sub-meters or devices in the same functional node). It is appropriate for the analysis to examine the data within this context as well, rather than blindly calculating across the entire dataset.

Methodology

The hierarchy is reconstructed automatically. After cleaning and aligning the data, only valid signals are analysed. Candidate parent–child links are evaluated using physical consistency checks (e.g., child $\leq$ parent, energy balance, spike alignment, rank correlation with small time shifts). A weighted score determines valid edges, which are assembled into a hierarchy and labeled with ok/warning/critical severities. As a result, this service automatically divides the data set into blocks, with each block corresponding to a single branch of the hierarchy (a single device or a group of devices that share a common parent).

Findings

Correlation matrices computed on the full mixed dataset are dominated by coincidental coefficients driven by shared external rhythms; restricting the analysis to reconstructed hierarchical blocks materially improves the signal-to-noise ratio of the resulting feature rankings. The pre-cleaning step also reduces the number of pairwise checks that need to be evaluated, so the reconstruction does not have to test every possible signal pair against every other.

6.3.3 Correlation — data shifting

Scope

Industrial time series are often misaligned in time: batch production is recorded at run completion, sensors sample at different rates, and drivers may lead or lag their effects. The data shifting block aligns paired series and quantifies the time delay between them, so that the downstream analysis can work with comparable, time-consistent inputs rather than being misled by artificial offsets.

Methodology

To uncover the true relationships, the service first distributes the total volume of each batch over the corresponding time window in proportion to the energy consumed at each time point, so that batch-based time series follow the same rhythm per time step as continuously sampled time series. Correlations are then evaluated across a range of positive and negative time offsets by systematically shifting one series against another and recomputing the correlation at each lag. The analysis identifies both the strongest lag and the full lag-response curve. To assess robustness, the same procedure is repeated on overlapping time windows, revealing whether the optimal lag remains stable or changes with the operating regime.

Findings

Batch-reported production data must be redistributed across the underlying timesteps before any lag analysis is meaningful; treating batch totals as point measurements introduces artificial lags into the correlation curves. A non-negligible share of driver–target lag relationships is not stable in time, and the rolling-window view detects regime-dependent drifts that a single-window lag scan would miss. Reporting both the strongest lag and the full lag-response curve provides the modelling stage with the information needed to decide whether a constant shift is sufficient or whether a regime-dependent one is required.

6.4 Pipeline orchestration

This section describes what Topic 2 has delivered as a preparation capability across the full chain — intake, standardisation, validation and modification — rather than implementation detail. In the current codebase the loading service (bs-data-loading) covers intake and standardisation, including deterministic pre-processing and optional schema assistance, while validation and modification are provided by the separate bs-data-validation service (sections 6.2.2 and 6.2.3).

The main implementation output is a multi-stage preparation workflow that connects these stages and the downstream handover within a single traceable process. The workflow preserves a common semantic representation of prepared time-series data — the SeriesSchema data contract specified in section 8.2.3 — so that individual stages can exchange information without losing comparability. It also persists intermediate and final artefacts, so that the preparation history remains auditable, the transformation path can be inspected later, and original and adjusted states can be compared where corrective processing has been applied.

The work also produced operational support outputs: staged processing logic, asynchronous handover between lighter and heavier tasks, progress tracking, structured quality reporting, and diagrams documenting the realised data flow. Together these form a practical preparation baseline that is already suitable for controlled processing and further extension, with clear service boundaries between loading and validation.

For further development, a LangChain/MCP-based orchestration layer is considered as a future extension for coordinating multi-step workflows and standardising inter-service interfaces; the platform-side view of the same direction is given in section 8.2.2. In this roadmap orchestration remains a control-plane enhancement, while deterministic data processing stays the authoritative execution path.

6.5 AI / SLM components

This section summarises the main findings on the role of AI and small-language-model (SLM) components in the preparation workflow, treating them as practical project results rather than architectural intentions. Today that support is deliberately narrow: deterministic loading and standardisation plus optional schema-assist logic in bs-data-loading, and AI-enriched reporting in bs-data-validation.

The Year 2 work produced several practical findings. First, separating initial intake from heavier preparation logic improves responsiveness at the point of upload while allowing the more demanding standardisation, validation and modification steps to proceed in a managed manner across the two services. Second, deterministic preparation logic should remain the core methodological baseline, because rule-based processing provides the repeatability, auditability and operational transparency that the workflow requires. Third, quality assessment is most effective when treated as an iterative process rather than a one-off gate: validation findings can trigger further standardisation or modification, after which the resulting data state is reassessed. Fourth, persisting intermediate artefacts and structured status information strengthens traceability and improves the interpretability of both successful and problematic preparation runs.

Finally, selective intelligent assistance proved useful in carefully bounded scenarios, especially where inputs are ambiguous or difficult to interpret using rules alone — in the present system, optional LLM-based schema help in loading and optional AI explanation in validation, not autonomous AI decision-making. Intelligent support should complement deterministic preparation, not displace it.

From an R&D perspective, the most credible next-year direction is to improve the quality contract between loading and validation before introducing any more advanced AI components, making schema inference, validation explanation and modification handoff more measurable, more explainable and easier to compare against deterministic outcomes. Only once that baseline is stable does it make sense to test limited agent-like orchestration or transformer-based anomaly support, and only as assistive tooling under clear human oversight (section 6.7).

6.6 Performance and scalability findings

This section summarises the current maturity of the validation and modification capabilities described in sections 6.2.2 and 6.2.3, as they are consumed by the broader preparation workflow.

Validation covers direct series checks, collection-level aggregation, stored-data assessment, targeted single-check diagnostics and an optional AI explanation path, returning a structured result with issues, warnings, statistics and, optionally, AI insight. Modification is centred on cumulative-meter series and is exposed both as a standalone transformation path and as a reusable pre-assessment step, handling spike

detection, deduplication, cumulative-to-interval conversion, gap filling, resampling and final value bounds. Both are explicit services within bs-data-validation rather than latent capabilities of bs-data-loading; they share the same canonical data model and are composed together when the combined modification mode is enabled.

6.7 Open issues and next steps

This section consolidates the open issues and forward-looking development directions for Year 3 across the preparation pipeline as a whole; the module-level next steps are given in sections 6.2.1–6.2.3 and the block-level open items of the pre-processing layer in section 6.3.

The main open items concern a deeper consolidation of the deterministic workflow and a clearer transition from a robust deterministic baseline towards a broader hybrid quality framework. The immediate priorities are stronger observability across the full preparation lifecycle, more explicit quality gates and decision checkpoints, a more systematic benchmarking approach for preparation quality and operational stability, continued strengthening of interoperability through a stable data contract, and the formalisation of a controlled propose-and-verify cycle for corrective preparation decisions that makes the decision criteria transparent where several acceptable preparation paths are possible.

AI-assisted capabilities should remain carefully bounded in the near term and continue to serve only as support for ambiguous or difficult cases; at present there is no agentic orchestration layer in the loading or validation services. Agent-based systems are therefore described as a research and development direction for Year 3 and the following period rather than as an already embedded operational capability. If LangChain/MCP orchestration is introduced, its adoption should be gated by explicit criteria: stable benchmark performance, no degradation of data quality indicators, full auditability of decisions and transitions, and clear human override for ambiguous cases.

Transformer-based methods are framed in the same way — as an exploratory track for advanced anomaly detection, whose expected value lies in recognising longer-range temporal irregularities, context-dependent anomalies and multivariate behavioural inconsistencies that local rule-based checks may miss. The most defensible pilot would use historical, labelled quality data to test whether a model can improve anomaly triage or flag suspicious series earlier than the current checks, while leaving the final decision to the deterministic service and the human operator.

Taken together, the Year 3 next steps are to consolidate the existing workflow, formalise the quality-driven propose-and-verify cycle, improve benchmarking and observability, and treat AI agents and transformer-based anomaly detection as structured research directions for the following phase — preserving deterministic and auditable preparation as the authoritative baseline.

Data-strategy validation (DFKI).

The data-preparation pipeline is validated by DFKI in its role as reviewer of the WP4 data strategy. DFKI assesses the methodological choices that govern how data is treated before it reaches the analytical and modelling stages: the standardised time-series representation (SeriesSchema,

section 8.2.3), the validation and modification logic, the correlation, block-segmentation and data-shifting methods, and the planned bounded use of AI/SLM components. Its feedback on data quality and on the propose-and-verify cycle is incorporated into the Year 3 consolidation of the pipeline; the validation roles of DFKI and TSK are described together in section 8.7.

7 (Topic 3) — Predictive and digital twin models

Lead: IT4I (IT4Innovations, VSB – Technical University of Ostrava)

7.1 Introduction and scope

This section summarises the research and development activities related to predictive modelling within WP4. The work focuses on AI-driven energy efficiency control for both industrial and building environments, with the objective of developing modelling approaches that support the identification of inefficient energy use, prediction of future energy demand and automated energy-performance management.

The modelling activities are based on a common concept that can be applied to different types of energy consumption, including electricity, gas, water or other monitored energy and utility flows. This concept is relevant not only for industrial sites, but also for buildings in which a significant part of energy consumption is linked to operational, technological or production-related processes. In this sense, the developed approach is not limited to a single type of facility, but provides a general framework for energy-performance modelling.

Within this framework, two data-driven model types are distinguished. The first model type uses historical data to identify inefficient energy consumption. It establishes expected consumption behaviour under comparable operating conditions and detects deviations that may indicate inefficiencies, anomalies or suboptimal operation. The second model type is focused on future energy consumption prediction. It forecasts expected demand over a defined prediction horizon and provides input for proactive energy management, planning and control.

For building applications, an additional modelling layer is required for thermal management. Thermal energy behaviour in buildings is strongly influenced by building geometry, envelope properties, material parameters, solar gains, shading, weather conditions and thermal storage. These effects are difficult to represent using consumption data alone. Therefore, the building case includes a physics-based thermal model that describes the dynamic thermal behaviour of the building using BIM data, meteorological inputs and analytical equations.

The overall modelling approach therefore combines general data-driven energy-performance models with a specialised physical model for building thermal management. This structure enables WP4 to address both industrial and building environments while maintaining a coherent methodology for AI-driven energy efficiency control.

7.2 General energy-performance model

The general energy-performance model was developed to support the analysis and prediction of energy consumption in environments where sufficient historical operational data are available. The model can be applied to industrial energy consumption as well as to building-related energy or utility consumption, particularly in facilities where operation, production or technological processes significantly influence demand.

The model consists of two main components. The historical data model is used to identify inefficient consumption by comparing measured values with expected behaviour derived from historical data. The predictive model estimates future consumption and supports forward-looking energy management. Together, these components provide both retrospective assessment and predictive capability.

This approach enables energy managers and operators to distinguish between normal consumption variability and deviations that may indicate inefficiency or optimisation potential. It also provides forecasts that can support operational planning, demand management and automated control strategies.

7.2.1 Historical data model

The historical data model evaluates energy consumption inefficiencies using available historical datasets. It analyses the relationship between consumption and relevant explanatory variables in order to determine expected energy use under comparable conditions. When measured consumption deviates significantly from the expected profile, the deviation may indicate inefficient operation, abnormal behaviour or a need for further technical investigation.

Regression-based methods were selected for this task because explainability is a key requirement. In energy management applications, model outputs should be understandable to engineers, facility managers and other technical users. The model should therefore not only indicate that a deviation occurred, but also provide a transparent basis for interpreting the relationship between input variables and resulting consumption.

Although complex artificial intelligence methods, including neural networks and advanced time-series architectures, may achieve high predictive accuracy, they are often less transparent. For historical inefficiency assessment, interpretability was therefore prioritised. Regression-based models provide a suitable balance between analytical capability, robustness and practical applicability.

7.2.1.1 Methodology — algorithm selection, training, validation

Several regression-based algorithms were evaluated for the historical data model. The tested approaches included linear regression, quantile regression, LASSO regression and other regression variants suitable for modelling the relationship between energy consumption and relevant explanatory variables.

The methodology consisted of data preparation, selection of relevant input variables, model training, validation and comparison of model performance. Historical datasets from selected industrial sites were processed and structured into training and validation datasets. The models were then trained to estimate expected energy consumption based on the available input data.

Quantile regression was selected and implemented as the preferred modelling approach. This method was chosen because it allows the model to estimate conditional quantiles of the consumption distribution rather than only the average expected value. As a result, quantile regression can describe a range of expected consumption behaviour and provide a more suitable basis for inefficiency detection.

This is particularly useful in industrial applications, where consumption may vary significantly depending on operating conditions. Quantile regression enables the identification of unusually high consumption values relative to comparable historical situations. Such deviations can then be interpreted as potential indicators of inefficient energy use.

7.2.2 Predictive model

The purpose of the predictive model is to forecast future energy consumption of the industrial site. While the historical data model is focused on evaluating inefficient energy behaviour using past data, the predictive model is focused on estimating future consumption over a defined prediction horizon.

Future consumption prediction supports proactive energy management. By estimating expected demand in advance, the model can support operational planning, peak demand reduction, load management and optimisation of energy-intensive processes. The predictive model is therefore an important component of the AI-driven energy efficiency control framework.

The industrial predictive model treats energy consumption as a time-series forecasting problem. Historical consumption values and other available input variables are used to train forecasting models capable of predicting future demand. The objective is to identify a modelling approach that provides reliable forecasts for industrial operation and can be integrated into decision-support or control workflows.

7.2.2.1 Methodology — algorithm selection, training, validation

Several time-series forecasting methods were tested for future industrial energy consumption prediction. The evaluated methods included


- deep neural network multilayer perceptrons,
- recurrent neural networks,
- long short-term memory networks,
- gated recurrent units,
- transformer-based architectures,
- N-BEATS,
- N-BEATSx,
- N-HITS,
- block variants of N-HITS, LEAR, LightGBM and XGBoost.

The tested transformer approaches included transformer encoder, transformer decoder and transformer encoder-only configurations. The purpose of this broad algorithm comparison was to evaluate both traditional machine learning models and advanced neural architectures for industrial energy consumption forecasting.

The models were tested on datasets from a selected industrial site. The training and validation process focused on the ability of the models to capture temporal dependencies, recurring consumption patterns and operational variability. Model performance was evaluated on validation data representing unseen time periods in order to assess the ability of the models to generalise beyond the training data.

Based on the conducted testing, the best-performing approach was a sequence-to-sequence model using transformer architecture. This model was identified as the preferred solution for future industrial consumption prediction. The sequence-to-sequence transformer model is suitable for multi-step forecasting because it maps a sequence of historical observations to a sequence of future predicted values.

The predictive model therefore provides a basis for forward-looking energy management in industrial environments. In combination with the historical inefficiency detection model, it enables both evaluation of past performance and prediction of future energy demand.



7.3 Building thermal management model

The building thermal management model extends the general energy-performance framework with a physics-based representation of building thermal behaviour. This additional model is needed because building thermal demand is not determined only by historical consumption patterns, but also by the interaction between the building structure, external environment and control actions.

The model uses information from the Building Information Model (BIM) consumption data and analytical equations describing thermal dynamics. The BIM model provides structured information about the building geometry, rooms, walls, windows and relevant material parameters. These data are used to define the physical structure of the thermal model and to calculate the power balance in individual rooms or thermal zones.

The model also incorporates external boundary conditions. Meteorological data for the building location are obtained from Open-Meteo <https://open-meteo.com/>. Shading from surrounding buildings is included using geometry derived from OpenStreetMap <https://www.openstreetmap.org/>, with shading calculations performed in Blender <https://www.blender.org/>. This allows the model to account for the influence of the surrounding built environment on solar gains and indoor thermal conditions.

The resulting model supports the creation of a digital image of the building. This digital representation includes geometry, material properties, external weather conditions, shading effects and thermal-zone behaviour. It can be used to simulate the transient response of the

building and to evaluate how changes in weather, solar gains, internal loads and control settings affect indoor thermal conditions and the resulting heating or cooling demand.

7.3.1 Methodology — algorithm selection, training, validation

The methodology for the thermal management model is based on the integration of building-specific data, external environmental inputs and analytical thermal equations.

First, relevant building information is extracted from the BIM model. This includes the identification of rooms or thermal zones, the geometry of walls and windows, and material parameters such as thermal conductivity and heat capacity. These parameters are required to describe heat transfer through the building envelope, heat storage in building elements and the dynamic thermal response of indoor spaces.

Second, the surrounding environment is represented. Geometry of neighbouring buildings is obtained from OpenStreetMap and used to evaluate shading effects. Blender software is applied to calculate shading over time, enabling the model to estimate how surrounding structures influence solar exposure and solar gains through windows.

Third, meteorological boundary conditions are obtained for the building location from Open-Meteo. These data provide external inputs for the thermal calculations, including outdoor weather conditions relevant to the building energy balance.

Fourth, analytical equations are implemented to calculate the power balance in each room or thermal zone. The balance accounts for heat transfer through walls and windows, thermal storage, solar gains, external boundary conditions and other relevant energy flows. Because the building response is transient, the model solves time-dependent equations rather than a stationary energy balance. Newton's method was implemented to solve the transient thermal problem. This numerical approach enables calculation of the evolution of indoor thermal conditions over time and supports the assessment of how the building responds to changing weather, solar radiation and control actions.

Validation of the thermal management model is based on checking the consistency of BIM-derived inputs, verifying the implementation of the analytical equations and comparing simulated behaviour with available consumption or operational data where applicable. The model provides a physically interpretable basis for automated building thermal management and complements the data-driven prediction models used within the broader WP4 energy-efficiency control framework.

8 (Topic 4) — Research EMOS platform layer

8.1 Introduction and scope

Topic 4 covers the EMOS platform layer — the cloud, multi-tenant environment that acts as the technological and data backbone of WP4. The subject of this chapter is **BackBone**, the research and integration software introduced in section 3.4, in which the process of EMOS analytics is assembled. The following sections describe the target platform architecture (8.2), the BackBone implementation (8.3), the analytics stage (8.4), the end-user frontend (8.5) and the cross-cutting platform concerns (8.6).

As stated in section 3.4, BackBone is not an end-user application: it delivers no final dashboards, reports or wizard services to customers. Its purpose is to provide a shared, executable environment in which the individual steps of the analytical chain — from data ingestion and validation, through correlation and pre-processing, baseline and predictive modelling, up to inefficiency detection — are connected into one continuous pipeline that can be run and debugged end to end. This answers the way WP4 is organised: the modules of the chain are developed in parallel by different consortium members (IT4I, STU, Asseco CE and EMOS), and without a common substrate they would remain isolated prototypes that are difficult to connect, validate end to end or keep in sync.

Technically, the individual functions of the chain are realised as **microservices**; the properties and consequences of this choice are described in section 8.2.1 and the integration mechanism in section 8.3.1. It is this microservice-based composition that lets each team publish its own module independently while consuming the outputs of the others, and that makes BackBone the natural foundation for the later prototyping work in WP8 (D8.4).

8.2 Technological architecture

8.2.1 Microservice architecture

The EMOS platform layer is built as a set of independently deployable microservices, following the principle of one module = one service with a defined interface. Each functional block of the analytical chain — data import, validation, pre-processing, time-series storage and computation, modelling, control and analytics, and the end-user views — is realised as its own service that can be developed, deployed

and upgraded on its own. The services are technologically heterogeneous: the platform and data services are implemented in Node.js/TypeScript (Express, with TSOA generating the OpenAPI interface from the controllers), while the import and modelling services are implemented in Python (FastAPI). They are uniform at the interface level, however — every service exposes an HTTP/JSON API described by an OpenAPI specification, and all data content is exchanged in the one standardised time-series record format (the SeriesSchema data contract, section 8.2.3).

Ownership of the services is distributed across the consortium along the same lines as the research work (section 8.1): each part of the chain is owned by the partner responsible for it — for example the data-preparation services by Asseco CE (chapter 6), the modelling services by IT4I (chapter 7), and the platform, storage, control and frontend services by EMOS. Because the services are decoupled at the network boundary and share only the data contract, a partner can release a new version of its own service — changing its internal implementation, its model or its libraries — without a synchronised release with the others; once published, the new version is consumed by the rest of the chain through the unchanged interface. This independent-versioning property is what keeps the WP4 outputs simultaneously independently developable and fully composable. The concrete integration mechanism (service registration, contracts and the gateway) is described in section 8.3, authentication in 8.6.1, and the build and publication pipeline in 8.6.3.

8.2.2 Orchestration layer (preparation for MCP orchestration)

Above the individual services sits an orchestration layer that drives data through the chain. In its current form this layer is deterministic: stages are connected by message queues and processed by dedicated workers, with the processing state persisted so that a run can be tracked, retried and resumed. The asynchronous import and transformation pipeline (bs-data-loading) uses a Redis-backed work queue and a worker that moves each job through its stages — parsing, transformation and the upsert of the resulting series into time-series storage — applying controlled retries with exponential back-off on failure. The model life-cycle (bs-digital-modeling) is orchestrated with a job system that runs a two-phase training → inference sequence: a training worker submits the fit, polls the model service until the model reaches a terminal state, and on success enqueues a materialisation job that produces the predicted series and flips the model to a ready state; the model status (training → trained → ready / failed) is held durably in the database as an explicit state machine. Analytics that must react to data or configuration changes (the cost calculation in bs-emos-control) are triggered the same way, through stream events consumed by a de-duplicated worker. This is the orchestration seen from the platform side; the data-pipeline view of the same flow is given in section 6.4.

The target orchestration extends this deterministic backbone toward a multi-agent, tool-using workflow, in line with the Year 2 decision to base AI/SLM orchestration on LangGraph and MCP (sections 4.4 and 6.5). The preparatory building blocks already exist in BackBone: platform capabilities are exposed to AI agents as MCP tools — a stdio MCP server (bs-mcp-server) that lets an assistant drive the file-import path

end to end, and an HTTP-transport MCP server (mcp-bs-data-preprocessing) that derives its tool set automatically from a service's OpenAPI specification and keeps per-client session state in Redis — and a first agent prototype wires these tools to a LangGraph reasoning loop. The intended design is hybrid: deterministic rules remain the primary control framework and AI agents are invoked selectively for ambiguous or complex situations, with final approval staying with the responsible user or process role. The full multi-agent orchestration is a Year 3 objective.

8.2.3 Data storage

All data content in the platform is stored in a shared MongoDB instance and is organised around time series. The central store (bs-timeseries) keeps two collections: a values collection holding the individual data points and a metadata collection describing each series. The point schema is deliberately minimal and uniform across every kind of series — a point is a triple of series identifier, timestamp and value, where the timestamp is the single source of truth and the value is either a number (for numeric series) or a text string (for categorical / state series, e.g. a production state or a reserved gap marker). This single semantic model of a data point is the SeriesSchema data contract referred to throughout the report (section 4.4): because every module produces and consumes points in this one shape, the modules are freely combinable. Series metadata carries the descriptive attributes — project scope, unit and source, together with a free-form object holding the input classification (category, sub-category, group) and the import provenance; the closed lists of categories, sub-categories and units that govern validation are encoded as service configuration rather than stored as data.

The platform does not separate raw, processed and result data into physically distinct stores; instead the layering is expressed through series lineage. A raw series records where and how it was imported (the import provenance: file, sheet, column and settings), while a processed or computed series records the series it was derived from and the operations applied to it, so that the chain of raw → validated → computed series is auditable and reproducible. Heavy intermediate artefacts produced during import and validation are held in their own working collections (the transformed payloads) before the resulting series are written back as first-class series, and computed outputs — for example model predictions or derived consumption series — are stored as ordinary series tagged as computed. The same store holds the reference datasets used in Year 2 development and validation (section 8.3.2) and is project-scoped, so reference and production data are distinguished by project scope rather than by separate storage; this is what allows a workflow to be exercised on reference data and then run unchanged on real project data.

8.2.4 Functional inventory

The inventory below maps the functional blocks of the EMOS chain to the services that realise them in BackBone. It is intended as a map of functionality; the substantive description of each block is given in the

chapter responsible for it — the data-preparation blocks in chapter 6 (Topic 2), the predictive and digital-twin models in chapter 7 (Topic 3), the inefficiency analytics in section 8.4, and the visualisation components in section 8.5. Ports are the local development ports; the “Integration” column states whether the service is currently wired into the shared chain through the gateway, reached through the model orchestrator, or at present operated standalone.

#	Functional block	Output / role	Ref.	Service(s)
1	Data input	ingestion layer	6.2	bs-data-loading, bs-data-validation
1a	Raw file digestion	ingests source files of heterogeneous formats and encodings into the pipeline	6.2.1	bs-data-loading
1b	Format transformation	converts heterogeneous inputs into a unified internal representation (SeriesSchema-conformant time series)	6.2.1	bs-data-loading
1c	Validation process	checks consistency and integrity; quality gate before further processing	6.2.2	bs-data-validation
1d	Normalisation	unit/scale alignment, resampling to common time base, gap handling	6.2.3	bs-data-validation
2	Data preprocessing (understanding)	preparation for modelling	6.3	Pre-processing blocks
2a	Correlation analysis	correlation matrices, multicollinearity flags, importance rankings, and a tiered feature shortlist (strong / watchlist / drop)	6.3.1	Pre-processing blocks
2b	Automatic block segmentation	a reconstructed measurement hierarchy and the resulting block decomposition	6.3.2	Pre-processing blocks
2c	Data shifting	time-aligns related signals by detecting and compensating lag/offset between series (e.g. metering vs. driver), producing shift-corrected series for correlation and modelling	6.3.3	Pre-processing blocks
3	Digital model calculation	model layer (see ch. 7)	7	bs-digital-modeling (orchestrator)
3a	Optimal vs. real/planned consumption model	produces the optimal-consumption baseline and contrasts it with the real/planned consumption profile, yielding the deviation used downstream for inefficiency calculation	7.2	bs-model-basic
4	Inefficiency calculation	quantification of inefficiencies by type	8.4	bs-emos-control, bs-timeseries-compute

4a	Procurement	inefficiency arising from energy purchasing / tariff and contract conditions	8.4	— (planned, Year 3)
4b	Technological	inefficiency in the technological / process operation of the site	8.4	bs-emos-control
4c	Performance	inefficiency related to power/demand peaks and load profile	8.4	— (planned, Year 3)
4d	Investment	inefficiency addressable only through capital investment (equipment/retrofit)	8.4	— (planned, Year 3)
5	Visualisation modules (sandbox)	end-user output	8.5	ui-emos-dashboard, bs-bi-view
5a	Sankey	energy-flow visualisation across the reconstructed measurement hierarchy	8.5	ui-emos-dashboard, bs-bi-view
5b	Manager dashboard	structured indicator view (KPI structure with status indicators)	8.5	ui-emos-dashboard, bs-bi-view
5c	Time series	temporal profiles of measured and modelled quantities	8.5	ui-emos-dashboard, bs-bi-view
5d	Overviews	aggregated summary views across sites/periods	8.5	ui-emos-dashboard, bs-bi-view
P	Platform & cross-cutting services	runtime foundation & integration shared by the whole chain	8.2, 8.6	
Pa	API gateway	single authenticated entry point; routing/proxy to all services	8.6.1	bs-gateway (8090)
Pb	Identity & multi-tenancy	accounts, projects, users, API tokens; tenant context; derived-series definitions	8.6	bs-project-registry (8087)
Pc	Time-series store	SeriesSchema storage and retrieval	8.2.3	bs-timeseries (8082)
Pd	Time-series compute	aggregation, derived-series evaluation, cost math	8.4	bs-timeseries-compute (8085)
Pe	File storage	file storage and metadata	8.2	bs-files (8080)
Pf	BI view backend	backend for the BI / dashboard views	8.5	bs-bi-view (8083)
Pg	API spec hub	aggregated Swagger UI across all service specs	8.2	ui-swagger-hub (8091)
Ph	MCP server (stdio)	bridges AI assistants to the platform	8.2.2	bs-mcp-server
Pi	MCP server (HTTP)	introspects a service's OpenAPI to expose tools	8.2.2	mcp-bs-data-preprocessing (3102)
Pj	AI agent prototype	LangGraph agent wiring MCP tools to an LLM	8.2.2	bs-ai
Pk	Earlier preprocessing service	superseded for the upload path	6	bs-data-preprocessing

8.3 Backbone implementation

This section describes how the architecture of section 8.2 is realised in practice: how a partner's module is connected into the shared chain (8.3.1) and how reference data is loaded so that the whole workflow can be run and debugged end to end (8.3.2).

8.3.1 Module integration as microservices

A module is connected into the chain as an HTTP service rather than as a library, in three parts. First, the service is registered with the platform by adding it to the gateway's route configuration; it then becomes reachable under a single, authenticated entry point and inherits the platform's authentication and project-context handling (section 8.6) without implementing them itself, while the gateway proxies the relevant path prefix to it. Second, the service speaks the shared data contract at its boundary: it consumes and produces time series in the SeriesSchema format (section 8.2.3) and describes its HTTP interface with an OpenAPI specification (generated from the controllers for the TypeScript services, provided by the framework for the Python ones), so the contract is explicit and machine-readable. Third, because the services are coupled only through these interfaces, each can be versioned and upgraded independently — a partner rebuilds and republishes its own service image and the chain picks up the new version behind the unchanged interface (the build mechanism is described in 8.6.3). Integration is currently static — services are wired in through configuration rather than through a dynamic service registry — which is appropriate for a research integration environment; several services of the data-preparation layer are still operated standalone and are scheduled for wiring into the chain in Year 3 (section 6.3).

8.3.2 Reference data and end-to-end run

The second function of Backbone is to let the team run and debug the complete workflow on real data. Reference datasets — exports from the pilot and demonstration sites, for example the reference production data and the building datasets — are inserted through the import service (bs-data-loading): a file is uploaded, parsed and transformed into the standardised records, which are then upserted into time-series storage as series scoped to a project. From that point the same chain that would serve production can be executed step by step over the reference data — validation and pre-processing, baseline and predictive modelling, the analytics stage and the end-user views — each step consuming the outputs of the preceding one through the shared store. Running the chain locally is supported by per-service container definitions and a top-level orchestration (the root build script brings up the shared infrastructure — MongoDB and Redis on a common network — and starts each service), so the whole environment can be stood up and a dataset driven through it. This end-to-end capability — develop a module, publish it and immediately see its effect on real reference data flowing through the rest of the

chain — is what made BackBone the central instrument for validating the WP4 outputs in Year 2 (sections 3.4 and 8.1) and the basis for the WP8 prototype (D8.4). End-to-end execution is at present driven manually, through the service APIs and utility scripts; an automated end-to-end test harness is part of the Year 3 work.

8.4 Orchestration and analytics integration

The platform turns the outputs of the data and model layers into the actionable results shown to users. The central analytic of WP4 is the detection and quantification of energy inefficiencies, computed as a stage of the processing chain.

8.4.1 Inefficiency calculation as a pipeline stage

Within BackBone, the calculation of inefficiencies is not a standalone model but a dedicated pipeline stage that consumes the outputs of the model layer (chapter 7) and the prepared, segmented data from the pre-processing layer (chapter 6), and produces the structured inefficiency results rendered by the frontend (8.5). The stage is realised as a microservice like every other block of the chain, so the calculation logic can be developed and versioned independently while remaining composable with the models that feed it. This section describes how inefficiency is computed within BackBone; the underlying models and their validation are described in chapter 7, and the services involved appear in the functional inventory (8.2.4).

The common principle across all inefficiency types is the same: a reference (expected / optimal) value is established for the relevant quantity and period, and the measured / real value is compared against it. The non-negative gap between the two, attributed to a specific cause, is the inefficiency. What differs between the types is the source of the reference, the explanatory variables used and the unit in which the gap is expressed (energy, power, cost or capital). BackBone is designed to orchestrate this as a uniform compute pattern parametrised per type, so that a new inefficiency type can be added by supplying its reference source and attribution rule rather than by building a new module.

The platform currently implements this pattern for two of these comparisons; the remaining ones are designed and scheduled for Year 3. Implemented in Year 2 are (i) an automatic cost computation, which prices each measured interval and aggregates the cost per period, and (ii) an energy-performance (EnPI) over-consumption analysis, which compares measured consumption against the value expected from a trained digital model (chapter 7) together with a reference, and quantifies the excess in both energy (kWh) and cost over the selected period. The four inefficiency types of the target taxonomy relate to these as follows:

- **Procurement** — designed for Year 3, building on the implemented cost computation. The reference is the cost-optimal procurement of the same delivered energy: the measured volume re-priced against the best available tariff / contract structure and time-of-use profile. The gap is the avoidable cost arising from

sub-optimal purchasing conditions rather than from the physical consumption itself, computed at cost level with no change to the consumption profile.

- **Technological** — the implemented over-consumption (EnPI) analysis realises this comparison. The reference is the expected consumption under comparable operating conditions, taken from the historical data model (quantile regression, 7.2.1); measured consumption above the conditional expected band is flagged as technological inefficiency — consumption the process could avoid at unchanged output. Localisation of the gap to a node of the measurement hierarchy via block segmentation (6.3.2) and driver alignment via data shifting (6.3.3) follow once those pre-processing blocks are integrated (section 6.7).

- **Performance** — energy-performance (EnPI) inefficiency — how much energy the site consumes relative to what it produces — evaluated over working hours and separated into a productive and a non-productive case. In the productive case (periods in which output is being produced) the indicator is energy per unit of output, computed with a selectable denominator — product weight, number of products, or workforce headcount from attendance — and compared against the value expected from a trained model (with a selectable model type; chapter 7); a further model automatically yields a further production-performance analysis. In the non-productive case (within working hours but with no output) the target is the idle base load: measured consumption is compared against a simple reference — by default a constant reference of one per hour, or a reference profile supplied by the user — with the acceptable level set by a threshold given as a manual value, a period average (optionally rolling) scaled by a coefficient, or a percentile. In both cases the inefficiency is the energy consumed above the expected or threshold level over the period. The core EnPI comparison this builds on — a trained baseline model and a measured-versus-expected over-consumption test with a tolerance threshold — is already implemented; the productive/non-productive split, the per-unit denominators and the alternative references and thresholds described here are the Year 3 target.

- **Investment** — designed for Year 3. The reference is the consumption of an improved asset state (post-measure), obtained from the module of changes (5.3); the gap is the structural inefficiency removable only through capital measures. Here the reference is counterfactual — a proposed building / technology state rather than a statistical expectation — and the result is paired with the technical-economic evaluation of the measure.

Each analysis produces a consistent result object per node and period — reference value, measured value, attributed gap and type — which is what allows the frontend to present the results together (Sankey, dashboard, time series) and the optimisation layer (chapter 9) to consume them as candidate targets.

Where the measurement hierarchy has been reconstructed, an inefficiency identified at an aggregate node can be decomposed toward the sub-meters that contribute to it.

8.5 Frontend for end users

In this part of WP4 the work is focused on developing individual visualisation modules that could serve to present the outputs of the processing chain — validated data, model results and the calculated inefficiencies — to operators, energy managers and building owners. Rather than building one fixed interface, each candidate visualisation is inserted into BackBone as a **typed sandbox**: a defined frontend component is set up over real pipeline outputs and then evaluated as to whether it is suitable, comprehensible and information-efficient for a potential user. This makes the frontend layer itself a subject of research — the question is not only *what* the chain computes, but *how* the result is best conveyed — while the underlying analytical computation stays in the analytics stage (8.4) and the model layer (chapter 7); the frontend is a presentation and interaction layer over the BackBone pipeline and does not compute its own results.

This layer serves a second, equally important purpose. Because each visualisation sandbox renders the live output of the pipeline, it also lets the members of WP4 see exactly what they need in order to assess the behaviour of their own modules and the effect of successive changes to them. A developer can observe, through the appropriate sandbox, whether a module produces the expected result and how an adjustment propagates through the chain — so the frontend works both as the prospective end-user interface and as a working instrument for module validation within the team. This dual role is consistent with the purpose of BackBone as a shared research environment (8.1).

Unlike the visualisation research carried out under Topic 1 (chapter 5), which investigates how energy and BIM information can be represented in principle, this section concerns the concrete frontend components assembled over the EMOS chain. The catalogue of these components — the visualisation sandboxes — is listed in the functional inventory (8.2.4).

8.5.1 Dashboard and indicator structure

The manager dashboard is the primary entry point. It presents a structured set of indicators (KPIs) with status semantics, organised so that the user moves from an overall assessment of a site toward the specific inefficiencies and their attribution. The indicator structure follows the

inefficiency typology (procurement, technological, performance, investment), so each headline indicator can be expanded toward the underlying nodes and periods. Time-series views show measured versus expected/optimal profiles over a selected period, and overview screens aggregate results across sites, periods or inefficiency types. As a sandbox, the dashboard is evaluated both for end-user clarity and as a means for the team to verify the indicators their modules feed.

8.5.2 Sankey and energy-flow views

The Sankey sandbox visualises energy flows across the reconstructed measurement hierarchy produced by the automatic block segmentation block (6.3.2). It allows the user to follow how energy is distributed between nodes of the hierarchy and to locate where the attributed inefficiency gap (8.4) arises, providing a spatial/structural counterpart to the time-based dashboard views. Because the hierarchy is reconstructed by the pipeline rather than entered manually, the flow view also serves as a direct check on the segmentation module's output.

8.5.3 Roles, multi-tenancy and BIM 3D views

The frontend is role-aware and multi-tenant: different user roles (e.g. operator, energy manager, owner) see scopes and actions appropriate to their responsibility, and tenant isolation ensures each customer sees only its own sites and data. The frontend also integrates the BIM 3D views developed under Topic 1 (chapter 5) as a presentation component, embedding the energy BIM model alongside the analytical results so that inefficiencies can be related to their physical location in the building. The underlying multi-tenancy and access-control mechanisms are described as a cross-cutting platform concern in 8.6; here they are addressed only as they shape what the user sees.

8.6 Cross-cutting platform concerns — security, multi-tenancy, DevOps

Three concerns cut across the whole platform rather than belonging to any single module: securing access to the services, isolating the data of different tenants, and building and deploying the services. Because they apply to the entire BackBone runtime, they are described here once rather than per service. As BackBone is in the first place a shared research and integration environment (section 8.1), these mechanisms are implemented to the extent needed to develop, integrate and validate the chain on real data; their hardening to full production grade is part of the transition to the WP8 prototype.

8.6.1 Security and access control

Access to the platform is centralised at the gateway: every request to a downstream service passes through it and must be authenticated, so the individual services do not implement authentication themselves. Two token-based mechanisms are supported. Interactive users authenticate with a JWT, which the gateway validates locally against a shared secret. Programmatic clients use API tokens (recognisable by an `emos_` prefix), which the gateway validates against the project registry — looking the token up by its hash — and caches in memory for subsequent calls, with an explicit cache-invalidation call when a token is revoked. On success the gateway attaches the resolved user identity to the request (as an `x-user` header) before proxying it, so a downstream service receives an already-authenticated, identified call over the internal network. Identity is managed by the platform itself — the project registry holds users, accounts, projects and API tokens — so no external identity provider is required. Role assignments (admin / write / read, at account and project level) are modelled in the registry; enforcing them uniformly at every service boundary is part of the planned access-control work.

8.6.2 Multi-tenancy and data isolation

The platform is multi-tenant: data belongs to a project (within an account), and a request is always executed in the scope of one project. The frontend identifies the tenant by human-readable account and project keys; the gateway resolves these to internal identifiers — caching the lookups against the project registry — and forwards the resolved account and project as headers. Every data-bearing service applies a project-context step that requires the project header and rejects any request without it, and then constrains all of its database operations to that project: the project identifier is part of every query and is taken from the gateway-provided context rather than from the client payload, so a caller cannot read or write outside its own scope. Isolation is therefore enforced at the query level within a shared database (each record carries its project scope) rather than by giving each tenant a separate database. This complements the frontend-level visibility described in 8.5.3 and rests on the storage model of section 8.2.3.

8.6.3 DevOps and CI/CD

Each service is an independently buildable and deployable unit. For local development and integration, every service ships a container definition and a small set of make targets (start, stop, rebuild, logs, shell), and a top-level orchestration brings the whole chain up against shared infrastructure (MongoDB and Redis on a common Docker network); this is what lets the team run the end-to-end workflow of 8.3.2. For delivery, each service carries its own continuous-integration pipeline that builds a versioned container image — tagged from the git version, with a date-and-commit fallback — publishes it to a shared container registry and deploys it to the integration host. Because each service is built and published on its own, a partner can release a new version of its module independently and the rest of the chain consumes it through the unchanged interface (section 8.3.1), which is the operational counterpart of the independent-versioning principle of 8.2.1. In line with the research-environment scope of BackBone, the current pipeline targets a shared development / integration deployment; production-grade deployment and operations are addressed as part of the WP8 prototyping.

8.7 Open issues and next steps

Several items remain open for the final project period. The most important is the full integration of the chain in BackBone: the data-preparation analytical blocks (correlation, segmentation, data shifting; chapter 6) and the validation step are still partly operated as standalone services and will be wired into the production pipeline as callable, versioned stages, including the write-back of validated and derived series so that the raw → validated → computed lineage is materialised end to end (section 8.2.3). On the orchestration side, the deterministic queue-and-worker layer will be extended toward the multi-agent, MCP-based orchestration prepared in Year 2 (section 8.2.2), keeping the hybrid design in which deterministic rules lead and AI agents assist under human approval.

The inefficiency analytics will be broadened from the currently implemented over-consumption and cost computations to the full inefficiency typology (section 8.4). On the platform-quality side, role-based access control will be enforced uniformly at the service boundary, and both the multi-tenant isolation and the build / deploy pipeline will be hardened from their current research-grade scope toward production, together with an automated end-to-end test harness over reference data (section 8.3.2). This work feeds directly into the WP8 system prototype (D8.4), for which BackBone provides the foundation.

Validation roles of DFKI and TSK.

Two partners contribute to WP4 primarily through validation rather than module development, and the EMOS platform layer is where their two perspectives meet. DFKI validates the data strategy: it reviews the choice of strategies for treating the data — the standardised data contract, the validation and modification logic, and the modelling-oriented preparation that feeds the analytics and digital-twin chain — and confirms that these methodological choices are sound and consistent across the data-preparation pipeline (chapter 6) and the platform that integrates it. TSK validates the functional strategy: as the prospective end user, it assesses whether what WP4 builds makes sense and corresponds to real operational needs, reviewing the dashboards, indicators, energy-flow views and role-based access of the EMOS frontend (sections 8.5–8.6) and the communication and BIM concepts of Topic 1 (chapter 5) against the requirements of owners, operators and facility managers. Because the EMOS platform layer brings together the prepared data, the analytics and the end-user frontend, both validation strands are exercised here on the integrated chain running on real pilot data in BackBone: DFKI confirms that the data and modelling strategy is methodologically sound, and TSK confirms that the resulting solution is meaningful and usable in practice. Their findings are fed back into the respective modules (chapters 5, 6 and 7) and into the WP8 prototyping that builds on BackBone.

9 Optimisation and Dynamic Control

This chapter reports the early-stage activities of T4.4 (operation optimisation) and the preparatory work on the Action Layer. As R&D here is only beginning, the substantive results are expected in Year 3 (D4.3).

9.1 Optimisation module — early progress

T4.4 (Asseco-led) follows the three-layer optimisation framework from D4.1 — static/investment, day-ahead and what-if. In Year 2 (~30 %) the work focused on the what-if layer, with reinforcement learning explored for automating recommendations.

9.2 Dynamic microgrid control — preparatory work

Preparatory design of an MPC core and of the interfaces for simulating the optimal operating profile of BESS, controllable consumption and sources. At this stage the layer is planned as a simulation that computes the optimal operating profile rather than as direct control.

9.3 Interfaces between Optimisation, Models and EMOS

Optimisation and control consume the Topic 3 models and Topic 4 analytics, exchange data via SeriesSchema, and run as microservices in BackBone. This keeps the modules independently developable yet composable, with integration consolidated in T4.5 toward M24.

10 Pilot deployment, verification and outlook

10.1 Progress with pilot data sites

During Year 2 the team established cooperation with five pilot sites and obtained the operational data needed to verify the WP4 blocks. The main acquisition results were:

- cooperation was established with five sites — Hobra, the Lasselsberger production site, the Mondi production site, the Aupark office and shopping centre, and the .53 office centre — with a non-disclosure agreement signed for each of them;
- data were obtained for the reference apartment building Střešovičky and for RD Zobor; BIM models were created for both locations and a reference comparison of the model against the recorded data is in progress;
- the Lasselsberger and Hobra datasets were used both for the general energy-performance model, and the basic types of inefficiency were processed on them; and for the verification of the pre-processing pipeline (correlation analysis, hierarchy reconstruction, lag analysis)
- metering devices were procured through public tender VO1 and are used to extend the set of measurement points and improve measurability in line with the research needs; the data from these devices is used for the preparation of the models and their validation;
- the datasets were extended several times, adding data according to the needs of the individual teams (for example higher granularity down to detailed consumption lines, and additional reference filtering by operating state and by product).

These pilot cases cover the full range of BIM representation required by Chapter 5, from lightweight box models for visualisation to detailed energy BIM models containing the non-graphical parameters needed for energy assessment. RD Zobor was used for detailed BIM modelling from a georeferenced terrestrial laser-scanning point cloud, Střešovičky for the application of the 3D energy BIM model standard, and further building examples for the simplified BIM generation approaches (IFC4 box-model generator, OpenStreetMap footprint retrieval and airborne laser scanning).

10.2 Deployment of first functional blocks

The first functional blocks of WP4 were deployed on the pilot data in two complementary ways:

- the data-preparation, processing, modelling and inefficiency-detection blocks are deployed as microservices in the BackBone research software (sections 3.4 and 8.3), where team members work with the shared datasets and build on the work of the other teams; each team upgrades its own modules, so the latest version is continuously available to the whole consortium;
- outside the BackBone environment the work focuses on the BIM modelling modules, which have been successfully deployed across a wide range of situations (see further in this chapter).

The Year 2 work confirmed that no single BIM creation approach is suitable for all practical situations. The verification therefore focused on the portfolio of approaches described in Section 5.4:

- custom-built IFC4 box model for rapid EMOS visualisation;
- simplified BIM generation from OpenStreetMap or airborne laser-scanning data where no documentation is available;
- BIM reconstruction from minimal 2D documentation;
- semi-automatic creation of space-based box models from indoor point clouds;
- detailed BIM modelling from terrestrial laser-scanning point clouds.

10.3 Functional feedback and focus for the next period

Following the topic structure defined in Section 4.1, this section combines the functional feedback from Year 2 — what confirmed or refined the requirements set in D4.1 and which gaps were identified — with the focus proposed for the remaining period of the project (towards D4.3 and D4.4 at M24, and the consolidation phase up to M30).

Topic 1 — Communication, visualisation interface and BIM

The energy BIM model standard (Section 5.5) was tested on the RD Zobor and Střešovičky cases. The verification confirmed that dedicated property sets are more suitable than broad, software-

generated material sets, as they keep the IFC model compact, and that parameters must be assigned to the correct object groups (walls, floors, slabs and partitions; windows and doors; external wall surfaces; room equipment; spaces and functional zones). The identified gap is that the required parameter set is not yet finalised.

Focus for the next period: finalise the required parameter set so that it is sufficient for energy modelling without unnecessary complexity, and complete the validation of the BIM-to-EMOS visualisation logic (colouring, filtering, spatial mapping) on real pilot data.

Topic 2 — Data preparation pipeline

The data input and pre-processing layers (Sections 6.2–6.3) were verified on the Lasselsberger and Hobra datasets, which were repeatedly extended on request. The pre-processing layer already delivers pairwise correlation analysis (Pearson / Spearman / Kendall) with VIF multicollinearity diagnostics, seasonality and periodicity profiling (STL decomposition, annual K-factor), automatic reconstruction of the measurement hierarchy, symmetric lag analysis with rolling-window stability, batch-to-timestep redistribution of production data, and a consolidated HTML report. The identified gap is the segmentation accuracy for complex measurement hierarchies.

Focus for the next period: improve block-segmentation accuracy for complex hierarchies, integrate the three analytical services into the WP4 service application, and keep the SeriesSchema data contract stable as the number and heterogeneity of pilot datasets grows.

What was confirmed

- The preparation process works as an end-to-end chain from intake through standardisation, validation, and modification to downstream handover.
- The deterministic baseline is in place and remains the authoritative decision path for quality assessment and corrective processing.
- Asynchronous intake and staged processing improved responsiveness and made handling larger batches more practical.
- A shared semantic time-series representation works across stages and supports subsequent analytical and predictive model use.
- Intelligent Inputs (Data Ingestion): heterogeneous and partially unstructured inputs are handled in the implemented intake flow, where JSON preparation is deterministic and the language model is used only as bounded assistance in ambiguous cases.
- Interactive Validation: automatic deterministic checks are implemented and operational across validation workflows.

- Automatic Transformation: batch-oriented processing into a shared time-series representation is implemented. Interpolation is applied conditionally, only when required by the quality/modification path.

What was not confirmed or had to be revised

- The earlier expectation that a language model would prepare production JSON as a primary path was revised; the implemented path is deterministic, while the language model remains advisory only.
- Interactive querying during entry is only partially confirmed; automatic checks are present, but an explicit operator dialogue loop for ambiguity resolution is not yet complete.
- End-to-end benchmark maturity is not fully confirmed; quality evidence exists, but still needs stronger formalisation.

Identified gaps

- Maintain deterministic JSON preparation as the authoritative path and keep language-model outputs advisory, with human review for ambiguous cases.
- Complete interactive ambiguity handling at intake (guided prompts and operator confirmation for unclear mappings, units, and timestamps).
- Improve lifecycle observability and benchmark discipline so quality improvements and regressions are demonstrable.
- Clarify rule-selection, reporting, and modification writeback conditions to reduce ambiguity in operational interpretation.
- Maintain robust preprocessing and, where available, redundancy in data sources to mitigate insufficient source-data quality.

Topic 3 — Predictive and digital twin models

The general energy-performance model and the thermal-management model (Sections 7.2–7.3) were verified against the Lasselsberger and Hobra data and the BIM-derived geometry of the reference cases. A regression model for estimating inefficiencies, a seq2seq model for future energy consumption and a basic physical model for building energy consumption are in place; validation of the solar-load module and a wind-speed-dependent refinement of the heat-transfer coefficient are in progress. This confirmed that the generic consumption-versus-reference formulation is necessary at industrial sites and that the physics-based thermal model can be driven from BIM geometry and open meteorological data.

Focus for the next period: complete the digital-twin thermal models, improve forecast accuracy, extend the recognised inefficiency types, and explore advanced AI methods (PINNs, neural

operators) for a detailed building digital twin; move from validation on isolated datasets to continuous comparison against recorded pilot data.

Topic 4 — EMOS platform layer

Deploying the first functional blocks confirmed that integrating data-preparation, modelling and inefficiency-detection modules as microservices into the BackBone environment is a workable model for cross-team collaboration. The identified requirement is that the platform layer and the Pset_EMOS structure must keep every measurement point, consumption value and operational state addressable through a stable spatial hierarchy.

Focus for the next period: consolidate the data-ingestion and EMOS platform blocks, and ensure that the results from Topics 1–3 are surfaced in the correct spatial context in the EMOS frontend, ready for pilot validation in T4.6.

Cross-cutting — Optimisation and dynamic control

Optimisation and dynamic control (Chapter 9) are still in an early phase, so the Year 2 feedback is mainly architectural: the interfaces between the optimisation layer, the models and EMOS must be defined now, while the pipeline and models are being verified.

Focus for the next period: further develop the optimisation algorithms and the what-if simulation capability, explore reinforcement-learning-based automation, and validate the Action Layer and management strategies on real pilot data, including simplified strategies for operation-critical components able to run when the central software is offline.

Overall, the Year 2 verification confirmed the assumptions made in D4.1 across all topics: BIM models and complete documentation cannot be assumed for all buildings; different use cases require different levels of detail; and stable identifiers together with clearly defined property sets and data contracts are essential for keeping models, datasets and the EMOS visualisation consistent. All WP4 tasks conclude at M24; the period up to M30 will focus on consolidating results, quantifying benefits and finalising deliverables D4.3 and D4.4.

10.4 Practical lessons learned

The key practical lessons learned in Year 2 are:

- securing data access and agreeing on data granularity with the pilot sites under NDA proved more time-consuming than the technical processing itself;
- datasets must be treated as living inputs — they were extended several times to add consumption-line granularity and operating-state filtering, so the pipeline must absorb such changes without rework;
- industrial datasets vary not only in scale but also in the cleanliness of their measurement hierarchy and in their batch-reporting conventions, so the pre-processing pipeline must absorb all kinds of variability without manual configuration per site
- correlation analysis on the full mixed dataset is dominated by coincidental coefficients driven by shared external rhythms (daily cycles, working shifts), so restricting the analysis to automatically reconstructed installation hierarchies is meaningful;
- automated BIM reconstruction methods (OpenStreetMap, airborne and terrestrial laser scanning) reduce modelling time but always require manual verification and enrichment before use in energy assessment;
- the BIM-to-EMOS integration (Section 5.6) was verified conceptually, confirming the importance of stable IFC GUIDs, the Pset_EMOS structure, a consistent spatial hierarchy from building to storey, space, zone and equipment, clear mapping between measurement points and BIM objects, and the ability to colour, highlight, filter or hide elements in EMOS;
- the BIM model should therefore be treated not as geometry alone but as a spatial and semantic framework linking energy data, sensor measurements and user-facing visualisation;
- the energy BIM model standard, the data contracts and the optimisation interfaces must all be validated on real pilot data and user workflows, not only conceptually.

11 Risks and mitigation

Rank	Risk	Impact	Mitigation	Status
1	Data sparsity and missing reference (ground-truth) vectors	Constrains training, calibration and predictive accuracy of monitoring and digital-twin models	Progressive acquisition of operational data under NDA; preprocessing in T4.2 (interpolation, outlier detection, removal of invalid segments)	Managed
2	Unavailability of BIM models / documentation for pilot buildings	Blocks BIM-based visualisation, energy modelling and spatial linking of sensor data	Portfolio of parallel BIM creation approaches (box model, OSM/ALS simplified model, 2D-to-3D reconstruction, point-cloud modelling) for use with and without a pre-existing BIM	Managed
3	BIM-to-EMOS integration not yet fully validated in the frontend	Models may be technically valid but not yet usable in EMOS workflows	Pilot testing of box models and energy BIM models in EMOS; validation of colouring, filtering and spatial mapping on real pilot data	Managed
4	Insufficient correlation between related signals	Distorts the evaluation of relevant reference quantities	Data adjustment, piecewise interpolation, modification of logging logic	Mitigated

12 Publications, dissemination and community building

12.1 Publications and conference presentations

P. Hlavinová. *Automated Data Preparation and Sensitivity Analysis of Energy Consumption. Conference presentation at AI Synergy: Connecting Research, Industry, and Regional Innovation (2nd InnovAlte Conference), Žilina, Slovakia, 29 April 2026.*

The presentation focused on automated data preparation and sensitivity analysis of energy consumption data. It addressed automatic hierarchical ordering, data structuring, and the handling of misalignment in industrial energy datasets. For the InnovAlte project, the contribution is relevant as a basis for improving the quality, consistency and usability of energy data before their further application in analytics, modelling and AI-supported energy management.

R. Fiala. *Next-Generation Energy Analytics: Not a Single kWh Wasted! Conference presentation at AI Synergy: Connecting Research, Industry, and Regional Innovation (2nd InnovAlte Conference), Žilina, Slovakia, 29 April 2026.*

The presentation addressed advanced approaches to energy analytics with a focus on reducing inefficiencies and energy losses. It introduced a typology of energy losses related to purchasing, technology, consumption and investment, as well as maturity levels of energy management approaches ranging from passive invoice-based control to AI-driven dynamic microgrid operation. For the InnovAlte project, the contribution provides a conceptual basis for developing analytics workflows, reference models, inefficiency metrics, what-if simulations and AI orchestration mechanisms for intelligent energy management.

T. Funtík, R. Varga, L. Michalák, M. Šveda. *Vývoj systému digitálneho manažmentu budov s podporou AI a BIM modelov. 12. vedecko-odborná konferencia Energetický manažment 2026: Rozumne a efektívne v energeticky náročných krízových časoch, Horný Smokovec, 18–19 June 2026. In: Zborník prednášok z konferencie Energetický manažment 2026. ISBN 978-80-8284-058-5.*

The article describes the development of a digital building management system based on the integration of static BIM data, dynamic sensor time series and artificial intelligence methods. It presents the methodology for creating BIM models from limited input data and partial outputs in the areas of visualisation, data infrastructure and energy management. For the InnovAlte project, the contribution provides a direct basis for the development of an intelligent building management ecosystem using digital twins, BIM models and AI-supported analytical tools.

R. Varga, L. Michalák, D. Olexík. Využitie dát z BIM modelu pri energetickom manažmente budov. 12. vedecko-odborná konferencia Energetický manažment 2026: Rozumne a efektívne v energeticky náročných krízových časoch, Horný Smokovec, 18–19 June 2026. In: Zborník prednášok z konferencie Energetický manažment 2026. ISBN 978-80-8284-058-5.

The paper deals with the use of data from Building Information Models in the open IFC format for building energy management during the operational phase of the building life cycle. It analyses the connection between static BIM models and real buildings through the concept of digital twins and the Internet of Things, as well as barriers to interoperability between BIM, energy management and simulation tools. For the InnovAlte project, the article is relevant as a basis for using structured semantic data, machine learning and deep learning methods for energy consumption prediction and real-time operation optimisation.

M. Krajčík, A. Václavová, R. Ingeli. Schopnosti a požiadavky systémov na manažment budov na Slovensku – výsledky dotazníkového prieskumu. 12. vedecko-odborná konferencia Energetický manažment 2026: Rozumne a efektívne v energeticky náročných krízových časoch, Horný Smokovec, 18–19 June 2026. In: Zborník prednášok z konferencie Energetický manažment 2026. ISBN 978-80-8284-058-5.

The article presents preliminary results of a survey conducted among building managers in Slovakia. The aim was to map the current level of use of building management systems and to identify practical requirements for improving existing platforms. The survey focused on technical systems, operation and maintenance, assessment of the current state and preferences for modernisation. For the InnovAlte project, the findings provide an initial overview of user needs and practical requirements relevant for the development of advanced digital building management solutions.

M. Krajčík, A. Václavová, R. Ingeli. Current Building Management System Capabilities and Requirements in Slovakia – Preliminary Results from a Questionnaire Survey. Proceedings of the 13th International Scientific Conference Environmental Engineering, Vilnius, Lithuania, 23–24 April 2026. eISSN 2029-7092, eISBN 978-609-476-444-8. DOI: 10.3846/enviro.2026.1885.

The paper builds on preliminary findings from a questionnaire survey of building managers in Slovakia and provides a more detailed analysis of the current use of building management systems and requirements for their improvement. The survey included 75 respondents representing residential, office and public buildings and focused on technical systems, monitoring, operation and maintenance, user assessment and modernisation preferences. For the InnovAlte project, the findings are relevant for defining functional requirements, identifying barriers to wider deployment and supporting the development of a comprehensive digital solution for building management.

R. Honti, J. Erdélyi, T. Funtík. Evaluation of Selected 3D Data Acquisition Systems for Capturing the Geometry of Building Structures. In: Acta Polytechnica Hungarica, Vol. 22, Iss. 8, 2025, pp. 203–223. ISSN 1785-8860. DOI: 10.12700/APH.22.8.2025.8.12.

The article complements the research on digital building management by addressing the acquisition of accurate spatial data for existing buildings. It evaluates selected 3D scanning systems, including Matterport-based solutions and SLAM technology, in comparison with reference terrestrial laser scanning data. For the InnovAite project, the results are relevant as a basis for selecting suitable technologies for digital building documentation, creation of spatial data layers and integration of geometric information into advanced data-driven building management and modernisation processes.

J. Erdélyi, T. Funtík, R. Honti, M. Marčíš. Building Information Modeling: A Case Study of the Natural History Museum in Bratislava. In: Contributions to International Conferences on Engineering Surveying, INGENO 2025. Springer Proceedings in Earth and Environmental Sciences. Springer, Cham, 2026, pp. 125–140. Online ISBN 978-3-032-12070-0, Print ISBN 978-3-032-12069-4. DOI: 10.1007/978-3-032-12070-0_11.

The article presents a case study focused on the creation of a Building Information Model of the Natural History Museum in Bratislava. It describes the process of defining information requirements, acquiring spatial data using terrestrial laser scanning and UAV photogrammetry, and developing an information model of an existing building. For the InnovAite project, the article provides a methodological basis for digital documentation, spatial data acquisition and the integration of building geometry into advanced data-driven building management solutions.

R. Honti, J. Erdélyi, T. Funtík. Evaluation of Point Cloud Generation from 360° Videos of Indoor Spaces and Heritage Sites. In: 11th International Workshop 3D-ARCH: 3D Virtual Reconstruction and Visualization of Complex Architectures. 1st ed. Hannover: International Society for Photogrammetry and Remote Sensing, 2026, online, pp. 199–206. ISSN 1682-1750. DOI: 10.5194/isprs-archives-XLVIII-2-W12-2026-199-2026.

The article complements the InnovAite project by examining affordable methods for acquiring spatial data from existing indoor environments. It evaluates point clouds generated from 360° video and photogrammetric processing in comparison with reference terrestrial laser scanning data. The results indicate that this approach is suitable for rapid and accessible documentation of building interiors, while its limitations in terms of accuracy, noise and edge definition must be considered. For the InnovAite project, the findings provide a basis for deciding when low-cost image-based capture can be used for digital building documentation, spatial data layers and further integration into data-driven building management solutions.

Annex

A.1 About the project

The InnovAlte Slovakia project was launched in response to the call "Transformation and Innovation Consortia" announced by the Government Office of the Slovak Republic under the Recovery and Resilience Plan of the Slovak Republic (Component 9 – Investment 2: Supporting cooperation between companies, academia, and research and development organisations). The call aims to drive systemic transformation and increase the added value of key sectors of the Slovak economy through intensive collaboration among research institutions, innovative enterprises, the public sector, and internationally renowned partners.

InnovAlte Slovakia seeks to build a dynamic and sustainable innovation ecosystem in the field of artificial intelligence (AI), effectively linking cutting-edge research with practical applications. The project emphasizes the development of AI solutions that are not only technologically advanced but also ethical, environmentally sustainable, and socially beneficial. The consortium brings together leading research centres, universities, and businesses from Slovakia, Germany, and the Czech Republic.

Key focus areas include the development of AI algorithms for improving building energy efficiency, enhancing traffic safety through video analytics, automating software development, driving digital transformation in the insurance sector, and validating functional prototypes. Special attention is also given to AI education, talent development, and the incubation of startups.

The project's outcomes will significantly strengthen Slovakia's position in research and innovation, reduce environmental burdens, increase the competitiveness of the national industry, and modernize public services. By bridging the public and private sectors with international research excellence, InnovAlte Slovakia stands as a key instrument in addressing the social and economic challenges of today.

"Funded by the EU NextGenerationEU through the Recovery and Resilience Plan for Slovakia under the project No. 09I02-03-V01-00029"

www.innovaite.sk

Innovaite
SLOVAKIA



www.innovaite.sk



Financované
Európskou úniou
NextGenerationEU

PLÁN [OBNOVY]



ÚRAD PODPRESEDU VLÁDY
SLOVENSKEJ REPUBLIKY
PRE PLÁN OBNOVY
A ZNALOSTNÚ EKONOMIKU



VÝSKUMNÁ
A INOVAČNÁ
AUTORITA