



transpAIrent.energy

Transparent AI Forecasts for Green
Energy in Austria

Project number: 910239

D4.1 Application requirements and use case specification

WP4 – Platform development and implementation

28.05.2025

Authors

Stefan AISTLEITNER
Thomas BLEIER
Dominik PUECKER
Dennis SCHMIDT
Hannes Koller
Stefan STRÖMER

Organisation

B-SEC
B-SEC
B-SEC
B-SEC
AIT
AIT



Funding Disclaimer

This project is being carried out as part of the 2023 “AI for Green” call for proposals from the Federal Ministry for Climate Protection, Environment, Energy, Mobility, Innovation and Technology (BMK). The processing is carried out on behalf of the BMK by the Austrian Research Promotion Agency (FFG). The project is funded as part of the topic of digital technologies, an initiative of the BMK, under the grant agreement number FO999910239. The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the texts, or in the decision to publish the results.

Table of Contents

Funding Disclaimer.....	1
Table of Contents	2
Executive Summary	3
1. Introduction.....	4
2. Use Cases	4
2.1. Use Case 1: Accessing charts displaying high resolution data	4
2.2. Use Case 2: Knowledge Compendium: The Austrian electricity system	5
2.3. Use Case 3: Interfaces, documentation, and digital communication.....	5
2.4. Use Case 4: Operational optimization configurator	5
3. Non-function application requirements.....	6
3.1. Purpose of the Platform	6
3.2. User Interface and Usability	6
3.3. Scalability and Reliability.....	7
3.4. Transparency and Data Integrity	7
3.5. Monitoring and Analytics	7
4. Architecture Design	7
4.1. Automation and Reproducibility, Versioning	7
4.2. Resilience	8
4.3. Traceability of Events and Monitoring.....	8
4.4. Security Aspects	8
4.5. Scalability.....	9
4.6. Architecture.....	9
4.6.1. Functional Components	9
4.6.2. Operational Components	10
5. Chart Specifications.....	0
5.1. General.....	0
5.2. Quality Tables	0
5.3. Bar Charts.....	0
5.4. Line Charts	1
5.5. Distribution Chart	1
5.6. Heatmap	1

Executive Summary

This document describes the functional requirements for the implementation of the transpAlrent.energy platform. The goal of this deliverable is to collect and document functional requirements for the platform during the initial phase of the project based on the abstract motivation and use case specification.

Chapter 2 describes the use cases that will be implemented within the transpAlrent.energy project. In Chapter 3, the non-functional Application Requirements are described. In Chapter 4 the architecture of the transpAlrent.energy platform describes the infrastructure that is used to run the services that make up the platform. Finally, in chapter 5 the charts that are used for displaying dynamic data are specified.

1. Introduction

This document outlines the application specifications and use cases for the "transpAIrent.energy" platform. The platform is designed to provide a transparent and user-friendly interface that delivers general information about:

- Relevant energy/electricity system related information – general explanations and an overview of various markets and products, available data sources, or similar topics that may be of interest for the target user group, linking to important (information) providers such as APG, E-Control, or others.
- High-level insights into publicly available data sources – informative and consistent visualizations of data points of common interest (e.g., day-ahead electricity market prices), including a data quality indication (e.g., frequency of missing data points) and links to data sources and including data protocols, API documentations, or similar information necessary for accessing these.
- Showcases of probabilistic live forecasts for energy system-related parameters such as electricity prices and CO₂ intensities, that are generated as part of the project (in work package 3).

The following sections detail the requirements, features, and use cases that need to be addressed within this milestone.

2. Use Cases

To clarify the platform's requirements and specifications, several key use cases are outlined below.

2.1. Use Case 1: Accessing charts displaying high resolution data

Description: A user wants to view a chart that displays data. This data consists partly of the live forecasts created in the project, e.g., for CO₂ intensity, electricity prices, or other forecasted parameters for the next 24 hours. In addition, supplementary (relevant) publicly available data can also be displayed. Lastly, the quality (which is data in itself) of the displayed forecasts and other various data, as well as linked existing metadata, can be viewed.

Target group: Public users who require transparent and validated information. Especially local/decentral producers and consumers who want to integrate this data into their strategies and real-time decision process.

Functional Requirements:

- The platform provides a list and/or a map of available regions where the user can select the desired region.
- Forecasts, quality information, and other data are displayed in charts that offer (at least) hourly resolution.
- The user can choose from various time horizons (e.g., 24 hours, 7 days) to access the relevant data.
- The chart will be displayed using an internal chart API, which is specified in chapter 5

Details: Information and data displayed are not limited, but highly focused on, to those created as part of this project. Existing public information and data will only be included or re-published if embedding the original chart, etc., or linking to the publication is not considered feasible or does not satisfy expected user requirements. "High resolution" data hereby refers to allowing users, as far as possible, to view full resolution data that is not aggregated or summarized. Where applicable options to aggregate/summarize will be made available.

Relevant examples or

related information:

- Chart types are outlined in detail in Chapter 5
- [Electricity Maps](#)

2.2. Use Case 2: Knowledge Compendium: The Austrian electricity system

Description: “High level explanation” of the foundations and various topics around the Austrian electricity system, grid, and market, as well as related and important general energy topics, including links to relevant publicly accessible sources.

Target group:

- “Industry newcomers”, e.g., asset owners of smaller generation/consumption plants who are interested in the topic and/or more details about certain aspects.
- Current or future asset operators/owners without specific domain knowledge.

Functional Requirements:

- CMS based website hosted by AIT

Details: The scope is further detailed to cover “*information on relevant topics related to one of: Markets & Products, Mechanisms, Market Premiums & Renewable Energy Support, Stakeholder Descriptions*”. Stakeholders could be flexibility service providers (FSPs), the transmission system operator (TSO), and many more.

Relevant examples or related information:

- [Marktprämien-Verordnung](#)
- [Österreichisches Strommarktmodell](#)

2.3. Use Case 3: Interfaces, documentation, and digital communication

Description: Knowledge collection about technical interfaces, protocols, available data and data formats, including information about where data exchange takes place, etc. related to the Austrian electricity system, grid, and market operation.

Target group:

- Software developers who are not familiar with the operation of the electricity system, grid, or market but want to create applications or other software related to these.

Functional Requirements:

- CMS based website hosted by AIT

Details:

- High level description of the relevant APIs, protocols and data formats
- Links to detailed API documentation
- Usage examples (request / response examples, code)

Relevant examples or related information:

- [Spezifikation API "Elektronische Gebotsabgabe"](#)
- [Power Monitor REST API](#)

2.4. Use Case 4: Operational optimization configurator

Description: The transpAIrent.energy web application will provide an operational optimization configuration feature. It gives users a basic evaluation of their own assets’ optimized operation

based on the generative AI forecasts created in the project. Parameters such as the types of assets and system sizes can be easily specified by the user based on simplified reconfigured models of energy system components.

Target group:

- “Low volume” asset operators who want to estimate the benefit of applying operational optimization to their own assets

Functional Requirements:

- Simplified models of energy system components
- GUI embedded in the transpAIrent.energy platform, that allows users the configuration of a/their asset or (virtual) plant
- Quick assessment functionality, showing estimated performance results (e.g., annual revenue streams) for different use of an asset (e.g., day-ahead market participation)

Details: No details are available yet because they will be specified in Task 4.3, which will only begin in May 2025. Task 4.3 will focus on the implementation of the operational optimization configurator, including the requirements specification, software implementation of both the frontend UI as well as backend APIs to the databases and the model, and user acceptance testing with a selected focus group of relevant stakeholders and interested users.

Relevant examples or related information:

- [Watts.happening](#)

3. Non-function application requirements

This section outlines the non-functional platform requirements, focussing on aspects such as its purpose, user interface and usability, scalability and reliability, transparency and data integrity, and monitoring and analytics.

3.1. Purpose of the Platform

The transpAIrent.energy platform is built with the motivation to provide valuable insights about energy related metrics and AI generated forecasts to the broader public, and support users unfamiliar with details by supplying general background information. Another key aspect of the platform is to make the processes how energy metrics are calculated, and forecasts will be generated, as transparent as possible to the end user. To archive this goal the project is hosting a content-management-system (CMS) driven website which acts as the main entry point to the platform and additionally provides a compendium with information about energy related terms, processes, and technologies. Additionally, a web application consisting of a webserver and a database backend is used to visualize the dynamic elements of the energy information, which are then embedded into the user-facing platform.

The platform must meet several key requirements to ensure effective use by end-users and external systems. These specifications focus on usability, security, and scalability while maintaining transparency and reliability:

3.2. User Interface and Usability

- **Responsive Design:** The platform must provide a fully responsive design, ensuring usability across various devices, including desktops, tablets, and smartphones. It should support multiple data visualization formats, such as charts, tables, and maps, allowing users to access forecast data quickly and easily.

- **Intuitive Navigation and Layout:** The platform must feature an intuitive and clean design, ensuring that users can easily find and interact with the information and tools they need. This includes clear menus, consistent design elements, and logical organization of content. Users should be able to complete tasks with minimal clicks and without the need for extensive training or prior knowledge. Providing internal links to supporting material, further detailed information, or term definitions helps users quickly navigate through unfamiliar topics.

3.3. Scalability and Reliability

- **Load Handling:** The platform must be able to handle the envisaged number of users and data volumes without significant performance issues.
- **Redundancy and Backup:** Regular backups and data redundancy measures must be implemented to prevent data loss in the event of system failures or cyberattacks.

3.4. Transparency and Data Integrity

- **Method and Data Transparency:** To build trust, the platform must clearly document the methods, data sources, and algorithms used. This includes the ability for users to compare historical data with forecasts, assessing accuracy and reliability. Where applicable existing and commonly used services (e.g., GitHub, Zenodo) will be utilized and linked to.
- **Data Validation:** All incoming data must undergo validation checks to ensure accuracy and consistency.

3.5. Monitoring and Analytics

- **System Monitoring:** Continuous monitoring of system performance, user activity, and data flows should be implemented to detect and address issues proactively.
- **Usage Analytics:** The platform should track and analyse user behaviour to provide insights into platform usage, identifying popular features and areas for improvement. This should follow a non-intrusive approach, collecting as little data as necessary, and ensure conformity with European and Austrian data privacy regulations.
- **Feedback Mechanisms:** Users should have the ability to provide feedback, helping identify bugs, suggest new features, or report security concerns.

4. Architecture Design

The architecture design will represent the fundamental structure on which the whole platform will be operated. Therefore, the architecture design will play an important part in terms of security, overall user experience, operational comfort and further development possibilities. Therefore, the specifications in this chapter will provide an overview, how the development and structure of the architecture will be structured and handled. It is important to note that all the mentioned aspects are equally important and rely, respectively influence each other.

4.1. Automation and Reproducibility, Versioning

- **Reproducibility:** All aspects regarding the platforms infrastructure and internal software configurations should and must be defined and executed on basis of code wherever possible. Besides other aspects, this ensures reproducibility and reliable configurations.
- **Versioning:** All code that defines the infrastructure and its configuration must be uploaded in a central repository for versioning, peer review opportunities, and transparency.

- **Automation:** The first two requirements complement the usage of automation mechanisms regarding the rollout of the entire infrastructure, where it can be used.

With this approach it is also possible to roll out the entire infrastructure within a short time, promoting automation. This ensures reliable configuration results with no room for unknown or in-transparent settings which could possibly broaden the attack surface of the entire platform.

- **Rules and Structure for Configuration Changes:** Changes on the infrastructure must only be deployed from the used automation tools to ensure that the code basis can be seen as the “single source of truth” which reflects the whole existing infrastructure and its configuration. This especially applies to the live and production system.

To enable such a work methodology, open source automation tools, like Ansible, will be used for the platform implementation, further enabling the developer teams to effectively collaborate, review each other’s code, and iteratively improve upon.

4.2. Resilience

- **Resilience through Redundancy:** For selected time-critical components, in case of a component failure there should be another node or component who can still handle the request and data which is sourced or destined to the platform, which promotes resilience.

4.3. Traceability of Events and Monitoring

The tracing of events within the backend platform is a key aspect to sustainably strengthen and improve the operational excellence of the platform.

- **Traceability:** All log messages generated by the backend platform should be sent and stored in a central place to be able to analyse them, to detect anomalies or errors, and to be able to explore and solve unknown problems or even performance bottlenecks.
- **Monitoring and Alerting:** All infrastructure components should be monitored on their state and health in a central place with the possibility of IT personal being alerted to take further action.

A central key aspect to achieve the centralisation of all logging activity will be the software Elasticsearch which is able to cover all the specified criteria. On the other hand, Zabbix will play an integral role to monitor the state of every involved IT-infrastructure component involved in the platform’s operations.

4.4. Security Aspects

Security will also be one of the focus points in regards of the design process. We strive to design the platform right from the beginning around the concept “Security by Design”, reducing the risk of finding major security flaws in advanced phases of the development, which are significantly harder to fix and would require exponentially more time. Therefore, the following key requirements will be adhered to:

- **Secure Communication and Protocols:** Every communication between infrastructure nodes must be established based in trust via TLS certificates supported by a certificate authority which stems from a Trusted Certificate Root Authority (e.g., Let’s Encrypt).
- **Input Validation:** Any user input must be validated to prevent the possibility of code manipulation or sending/injecting malicious code which could harm the integrity of data or causes downtime of the entire platform.
- **Network Segmentation:** communication between different components in the platform must be carried out in a way that

they are only exposed where needed and isolated wherever possible (e.g private Networks, Firewall Settings., ACLs).

- **Encryption:** Data must be encrypted at transit and at rest wherever reasonable.
- **Traceability and Accountability:** Actions carried out on the platform should be able to be assigned to a user to address the security goal of accountability.

4.5. Scalability

As our platform will gain more attention, it has to manage more and more concurrent user requests. Therefore, it must be ensured that the platform’s underlying infrastructure is able to grow with the demand and on the other side provide a resource-efficient implementation if demand is low. This must be ensured to be on one side economically efficient and on the other side to also consider being sustainable with the needed IT infrastructure and its effects on the environment.

- **Flexible and Dynamic Infrastructure:** The IT infrastructure should be able to scale up and down with the experienced demand of computing power.

4.6. Architecture

The diagram (see Figure 1) shows the basic architecture of the infrastructure used for the transpAlrent.energy platform, based on the foundational requirements described in Chapters 3.1 to 3.5.

In the following, detailed explanations of all components will be given, highlighting how they contribute to the whole platform, differentiating between functional and operational components and addressing their purpose.

4.6.1. Functional Components

The described components in the following provide the functionality of the whole transpAlrent.energy platform and contribute to the realization of the initial motivation.

transpAlrent.energy Website: The CMS based site is the first point of entry for the platform and provides documentation how the platform works and how the different energy metrics are gathered, processed, visualized and how forecasts are generated. It also includes a compendium with descriptions of energy related terms, processes, and relevant technologies to provide users – who are new or not very experienced in the energy sector – an introduction and efficient way to gather all the required knowledge to understand the data provided on the platform.

Webservers: On the webservers the dynamic content of the platform will be generated. For the beginning we will start with one webserver, and if demand grows over time a load balancer will be added and load can be distributed between multiple webservers. This step ensures that even with higher user rates every visitor has a smooth experience while browsing through the platform.

Relational Database: The relational database will handle the main data storage and processing for the web application. For this project it is planned to use MariaDB as is it a reliable and open-source database. As demand grows a second database node would be added to implement an active-passive database cluster to increase availability and resilience.

Time Series Database: The time series database InfluxDB is used to store and process time series data like energy metrics and the forecasts. Data sourced to this timeseries database will be used to fill the plots with actual data and to display it to the end user. Similarly, to the relational database the project will start with one InfluxDB node. In the future with growing visitor numbers there will be also redundancy be added. To achieve that a second node will be planned which replicates produced time series data in real time and enables disaster recovery.

4.6.2. Operational Components

In order to be able to operate the platform in a reliable way, monitoring and logging components need to be added to provide performance analysis, debugging and anomaly detection. Therefore, the following components will be utilized to enable efficient and secure operation.

Logstash: Logstash acts as a central point where all logs of the whole infrastructure will be accumulated. This component is the single point of entry for logs and acts as a gateway for further processing of the logs in the Elastic Cloud. In addition, it only allows trusted and verified clients to send their data to prevent malicious activities.

Elastic Cloud: The Elastic Cloud only allows the described Logstash instance to send the accumulated logs to it. The Elastic Cloud is responsible for processing the logs in a way so that they can be analysed, and potential undetected misconfigurations or unexpected behaviour can be detected and be solved. It also allows the visualization of the logs for simple monitoring.

Zabbix: Zabbix allows detailed monitoring and alerting for the whole infrastructure. Through Zabbix the operations team can closely monitor the behaviour and health of each node and is immediately notified in case of changes which need attention or need to be resolved in order to have a satisfying experience for the end users.

The combination of the described components allows the transpAIrent.energy project to implement the decided functionality in terms of the technical aspects and provides a solid plan for further scaling as the platform is expected to grow in the future. In addition, also important measures have been set to ensure the reliant operation of the platform by providing technical solutions which allow easy observation and quick debugging capabilities.

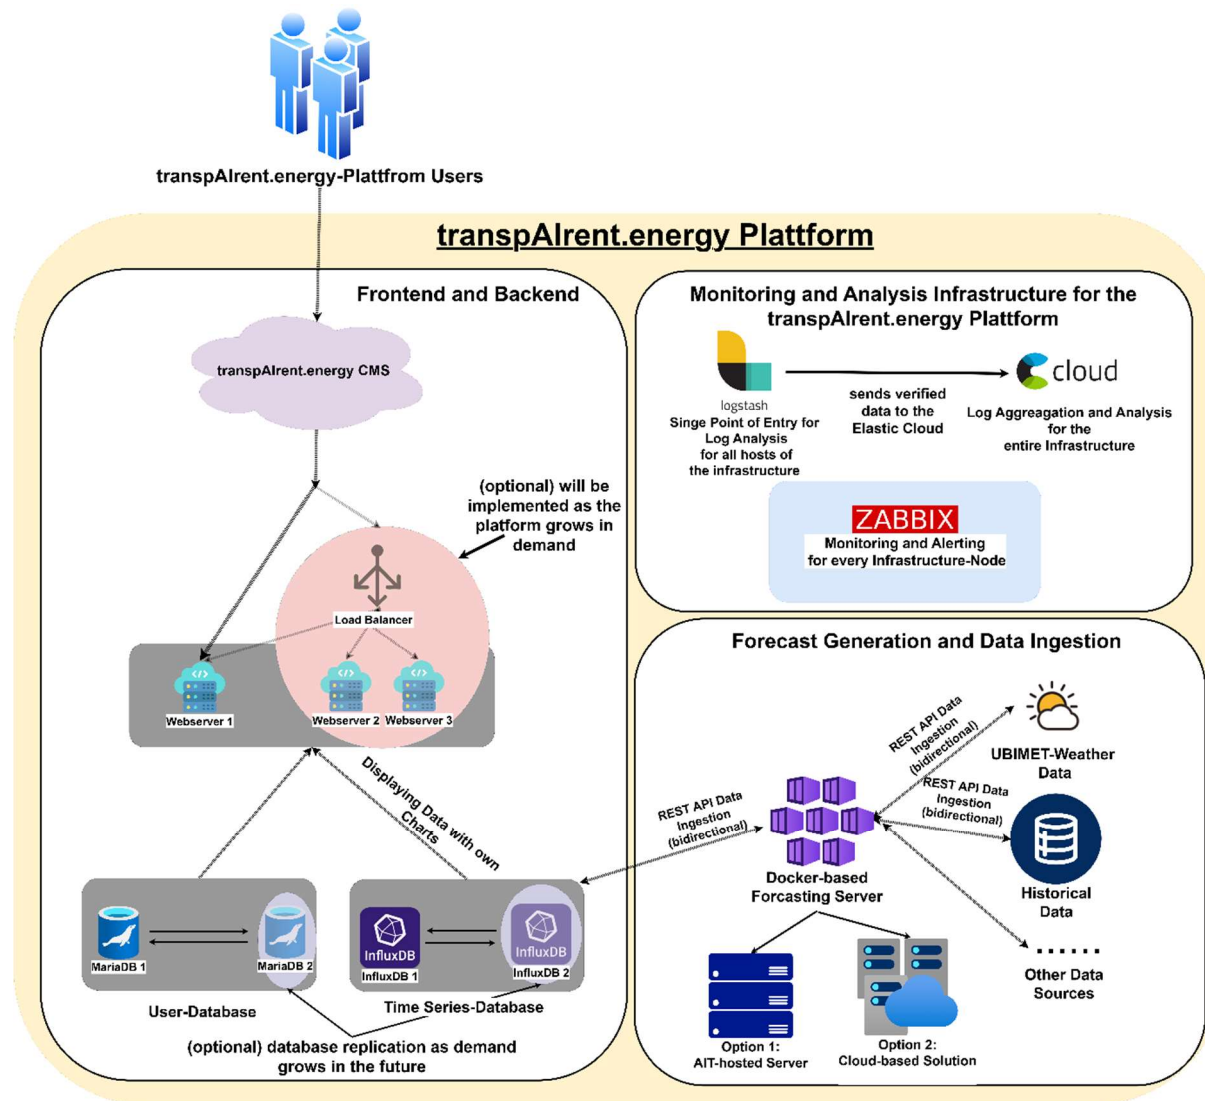


Figure 1: Planned architecture for the transpAlrent.energy platform

5. Chart Specifications

This chapter specifies the internal chart functionality that is used to display dynamic data.

5.1. General

Chart Features:

- **Hover Tooltip:** Displays values, unit, timestamp (or aggregated range), and publication time.
- **Legend:** Includes title.
- **Axis Labels:** Show unit.
- **Embedded Information:** Includes last update, sources, disclaimer, and a "read more details" link to the compendium page.
- **Fixed Palette:** Especially for energy carriers.
- **Download Option:** Allows download as PNG (and possibly other formats).
- **Horizontal Rulers:** Option to add (e.g., for average long-term values).

Parameters:

- **General Type/Subtype** (Or a specific set with IDs).
- **Trace Features:** Customizable name and colour for each trace.
- **Aspect Ratio and Legend Position**
- **User Controls Above Each Chart:**
 - **Time-Range Selection:** Calendar picker or predefined ranges.
 - **Time-Scale Selection/Aggregation:** Options include original, hourly, daily, weekly, grouped means, with error bars if not stacked & aggregated.

5.2. Quality Tables

The "Quality Tables" feature provides a compact, mini-dashboard view of both data and forecast quality for selected features. These tables, designed for easy embedding, display key metrics like data availability, publication time, and forecast accuracy, including relative MAE, update frequency, look-ahead periods, and the 90% quantile range of errors. Each table includes tooltips for deeper insights, linking to detailed explanations in the compendium.

Used for: TBD

5.3. Bar Charts

For bar charts, the design varies based on the data structure, supporting configurations with or without spacing, and allowing grouping by timestamp for multiple sources or stacking for one source and multiple sub-features or multiple sources and multiple sub-features. These charts can properly handle both positive and negative values, as well as the overlay of a single line chart with a secondary y-axis for added context, e.g. overlaying demand data on generation mix charts.

Used for:

- Day-ahead price
- Demand
- CO2 intensity
- Balancing (FCR, aFRR, mFRR, CAI, prices), "activation probabilities"
- Generation mix (APG-3type-forecast; stacked+grouped)

5.4. Line Charts

Line charts offer interpolation methods, including linear or stepped approaches, and supports multiple stroke widths. They are used to visualize historic and live forecasts, with actual values as baselines and options for displaying multiple forecast versions. For live forecasts, special indicators mark the last known entry and the current time, with confidence intervals or "spaghetti" plots illustrating different realizations. Features can be displayed as single lines or with min/max bounds, depending on the context.

Used for:

- Demand

5.5. Distribution Chart

A histogram with an approximated density curve, is intended to support forecast line charts.

Used for:

- For all forecasted features

5.6. Heatmap

Heatmaps provide a 2D visualization with dual time axes (e.g. x=day, y=hour), useful for analysing patterns like day-ahead prices or CO2 intensity over different time frames.

Used for:

- Day-ahead prices
- CO2 intensity
- Balancing "activation probabilities"

Project coordinator:

Klara Maggauer, M. Sc., B. Sc.

AIT Austrian Institute of Technology GmbH

Center for Energy

Giefinggasse 4

A-1210 Vienna

E-mail: klara.maggauer@ait.ac.at