



Learning to be Green

An AI-enhanced multilingual online course for developing a green entrepreneurial mindset

PROJECT DELIVERABLE

Conversational Agent Implementation



**Co-funded by
the European Union**

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.



Alliances for Innovation – Lot 1: Alliances for Education and Enterprises in the European Union’s Erasmus+ Programme

Call: ERASMUS-EDU-2023-PI-ALL-INNO-EDU-ENTERP

Funding period: 1 April 2024 – 31 March 2027 (36 months)

Grant Agreement no. 101139849

Deliverable D4.2

Conversational Agent Implementation

Leading Editor Apostolos Mavridis, Stergios Tegos

Reviewer(s) Pantelis M. Papadopoulos, Ahsen Çini, Alina
Kopkow, Viktoria Pammer-Schindler

Date January 2026

Dissemination Level Public

Status Final version (1.0)

The L2BGreen Consortium

No.	Organisation full name	Short name	Country
1	Universiteit Twente (coordinator)	UT	Netherlands
2	Berner Fachhochschule	BFH	Switzerland
3	Universität des Saarlandes	SU	Germany
4	Technische Universität Graz	TUG	Austria
5	Enchatted P.C.	ENC	Greece
6	Združenje Delodajalcev Obrti In Podjetnikov Slovenije Giz	ZDO	Slovenia
7	Eesti Rohetehnoloogia Liit MTÜ	ERL	Estonia
8	Verein zur Entwicklung der Digitalen Zukunftsbildung	VDZ	Austria
9	Scholes Delta Voreiou Ellados Anonymi Etaireia	DEL	Greece
10	Sihtasutus Estonian Business School	EBS	Estonia

Contributors

Name	Organisation
Stergios Tegos	Enchatted P.C.
Apostolos Mavridis	Enchatted P.C.
Konstantinos Manousaridis	Enchatted P.C.

Legal Notice

The information in this document is subject to change without notice. The Organisations of the L2BGreen consortium make no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

The Organisations of the L2BGreen consortium shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

The information and views set out in this deliverable are those of the contributing author(s) and editor(s) and do not necessarily reflect the official opinion of the European Union. Neither the European Union institutions and bodies, nor any person acting on their behalf, may be held responsible for the use which may be made of the information contained herein.

Executive Summary

This document, Deliverable 4.2, provides a comprehensive account of the technical implementation of the L2BGreen learning ecosystem, translating the architectural blueprints established in Deliverable 4.1 into a fully functional, integrated, and deployed platform. As an output of Work Package 4 (WP4), led by Enchanted P.C., this deliverable marks the transition from design specification to working software, realizing the project's vision of an AI-enhanced, multilingual learning experience for aspiring green entrepreneurs across Europe.

The implemented ecosystem is built on a three-part distributed architecture designed for flexibility, security, and long-term sustainability. The Learning Hub, a customized Moodle LMS instance, serves as the primary user-facing portal. The Interaction Engine, a custom-built Laravel application, houses the project's core technological innovations. These two environments are connected through a Secure Bridge implemented using the LTI 1.3 standard, ensuring seamless and secure interoperability with any compliant Learning Management System and avoiding vendor lock-in.

At the pedagogical heart of the system is Gaia, the L2BGreen conversational agent, implemented as a sophisticated hierarchical multi-agent framework. An Orchestrator Agent coordinates the learning flow, delegating tasks to four specialized worker agents — the Scenario Presenter, Knowledge Scaffolder, Application Coach, and Argumentation Feedback Agent — each precisely designed to execute a distinct phase of the instructional design defined in D2.1. To ensure the quality and pedagogical soundness of AI-generated feedback, the implementation incorporates advanced techniques including Chain of Thought (CoT) prompting and a self-correcting Evaluator-Optimizer workflow. A dedicated moderation layer further safeguards the learning environment by screening all learner input before it reaches the core agents.

The Educator and Administration Platform provides educators with an intuitive control panel for the entire ecosystem. A step-by-step scenario builder enables the creation and deployment of AI-powered learning activities, while a robust Role-Based Access Control (RBAC) system supports collaborative, institutional use. A key innovation of this platform is its advanced learning analytics dashboard, which moves beyond descriptive reporting to deliver prescriptive, actionable insights through metrics such as Engagement Score, Argumentation Depth, Misconception Flags, Dropout Risk, and Recommended Interventions.

This deliverable also documents the complete technical implementation of the Moodle LMS integration via LTI 1.3, including secure user authentication, context exchange, and the automatic return of grades via LTI Assignment and Grade Services (AGS). Underpinning the entire system is a robust approach to data privacy and security, implemented in full compliance with GDPR and the project's Data Management Plan, incorporating pseudonymization, AES-256 encryption at rest, and TLS/SSL encryption in transit.

The functional platform documented herein provides the operational foundation for stakeholder engagement activities in WP5 and will be the subject of D4.4 "System evaluation and public deployment," the next critical milestone in the project's trajectory.

Table of Contents

Executive Summary	4
1. Introduction.....	7
1.1. Project Vision and Context	7
1.2. Purpose and Scope of this Deliverable	7
1.3. Relation to Other Work Packages and Deliverables	8
1.4. Guiding Principles for Implementation	8
2. System Architecture Implementation	10
2.1. Overview of the Implemented Architecture	10
2.2. Technical Architecture Components	11
2.2.1 Moodle Server (Learning Management System).....	11
2.2.2 L2BGreen Application Server (Interaction Engine).....	14
2.2.3 Moderation Layer	14
2.2.4 LLM API Endpoints (External AI Service).....	15
2.2.5 User's Browser (User Interaction)	15
2.3. Technology Stack Implementation	15
2.4. Data Flow Implementation	16
2.4.1 Primary Interaction Flow	17
2.4.2 Analytics Data Pipeline	17
2.4.3 Data Privacy and Security Implementation	18
3. Implementation of the Gaia Conversational Agent.....	19
3.1. Realization of the Multi-Agent Architecture	19
3.1.1 The Orchestrator Agent.....	19
3.1.2 Specialized Worker Agents	20
3.2. Implementation of Advanced Prompting Techniques.....	21
3.2.1 Chain of Thought (CoT) Prompting for Feedback Generation	21
3.2.2 Few-Shot Prompting for Application Questions.....	22
3.2.3 Socratic Tutoring Implementation.....	22
3.3. Implementation of the Evaluator-Optimizer Workflow	23
3.3.1 Workflow Architecture	23
3.3.2 Iteration Limits and Latency Management.....	24
3.4 Implementation of the Moderation Layer.....	24
3.4.1 Content Classification and Analysis	24
3.4.2 Calibrated Response Mechanisms.....	25
3.5 Learner Interface Implementation	26
3.5.1 Visual Design and Layout.....	26

3.5.2 Structured Interaction Flow	27
3.5.3 Accessibility and Responsive Design	31
4. Implementation of the Educator and Administration Platform	31
4.1. Scenario Orchestration and Content Management	31
4.1.2 Scenario Builder Workflow	32
4.1.3 Activity Details and Conversation Monitoring.....	33
4.2. User and Team Management Implementation	34
4.2.1 User Profile Management	35
4.2.2 Hierarchical Team Structure	36
4.2.3 Role-Based Access Control (RBAC) Implementation	37
4.3. Implementation of the Learning Analytics Dashboard	38
4.3.1 Data Aggregation and Sources	38
4.3.2 Key Metrics and Visualization.....	38
4.3.3 Analytics Metrics Implementation	39
5. Moodle LMS Integration and Configuration.....	40
5.1. LTI 1.3 Protocol Implementation.....	40
5.1.1 Security Features	40
5.1.2 User Authentication and Context Exchange.....	40
5.1.3 Automatic Grade Return	41
5.2. Moodle Course and Activity Design.....	41
5.2.1 Course Structure Template	41
5.2.2 LTI Tool Configuration	42
5.3. Multilingual User Interface Implementation.....	44
5.3.1 Static Content Translation	44
5.3.2 Dynamic Interaction Translation	44
6. Conclusion	47
6.1. Summary of Implemented Components	47
6.2. Alignment with Project Objectives	48
6.3. Next Steps and Future Work	48

1. Introduction

1.1. Project Vision and Context

The L2BGreen project is founded on the ambitious vision of cultivating a green entrepreneurial mindset across a broad and diverse European audience. As articulated in the project proposal, this initiative aims to "leverage scalable, digital ways of learning as the means to move forward the green transition through increased awareness of and knowledge on green entrepreneurship." The target audience is explicitly pan-European, encompassing higher education students, vocational trainees, and working professionals from a wide array of cultural and linguistic backgrounds.

Central to achieving this vision is the development of an innovative, AI-enhanced, multilingual online course. The project's core output is designed to be both highly interactive and broadly accessible, combining cutting-edge conversational AI technology with pedagogically grounded learning experiences. This document, Deliverable 4.2 "Conversational agent implementation," marks a critical milestone in the project's trajectory: the transition from comprehensive design specifications to a fully functional, deployed system.

1.2. Purpose and Scope of this Deliverable

This deliverable, documents the complete technical implementation of the L2BGreen learning ecosystem as specified in Deliverable 4.1 "Conversational Agent Design." As an output of Work Package 4 (WP4 – Technology Development: Tutorial Agent and Multilingualism), led by Enchanted P.C. (ENC), this document provides a detailed account of how the architectural blueprints established in D4.1 have been realized as working software components.

The scope of this deliverable encompasses the implementation of three interconnected technological components:

The Gaia conversational agent, the AI-powered pedagogical heart of the system, has been implemented as a sophisticated, hierarchical multi-agent framework. This implementation realizes the design of the Orchestrator Agent and the specialized "worker" agents (Scenario Presenter, Knowledge Scaffolder, Application Coach, and Argumentation Feedback Agent) as detailed in D4.1. The implementation includes the advanced prompting techniques, specifically Chain of Thought (CoT) prompting and the Evaluator-Optimizer workflow for quality assurance, that enable Gaia to provide transparent, pedagogically sound, and self-corrected feedback to learners.

The Educator and Administration Platform has been developed as the operational control panel for the L2BGreen ecosystem. This implementation delivers the scenario orchestration and content management capabilities, the user and team management system with Role-Based Access Control (RBAC), and the learning analytics dashboard designed to provide educators with actionable insights through metrics such as Engagement Score, Argumentation Depth, and Misconception Flags.

The Moodle LMS integration has been fully implemented using the Learning Tools Interoperability (LTI) 1.3 standard. This secure connection enables the L2BGreen platform to be embedded as a native activity within existing Moodle course environments, supporting secure user authentication via OAuth2 and JSON Web Tokens (JWT), as well as the automatic return of grades via LTI Assignment and Grade Services (AGS).

1.3. Relation to Other Work Packages and Deliverables

The implementation documented herein is the direct result of a carefully coordinated sequence of work packages and deliverables. This deliverable translates design into reality, drawing upon and integrating the outputs of preceding work packages.

Deliverable D2.1 "Interactive online course instructional design" established the foundational pedagogical framework, including the case-based learning methodology and the structured learning phases (Introduction, Sub-topic Introduction, Application, and Argumentation) that the Gaia agent is designed to execute. The implementation of Gaia's multi-agent architecture directly realizes this instructional design, with each specialized worker agent corresponding to a specific phase of the learning script.

Deliverable D3.1 "Learning materials" provided the rich, case-based content that powers the learning experience. The implementation of the content management system and database architecture has been specifically designed to store and deliver this material, including the business case studies such as "ReNuTex – Rethinking Textiles Sustainability," sub-topic knowledge inputs, and application questions.

Deliverable D4.1 "Conversational Agent Design" served as the primary technical specification for this implementation. The architectural diagrams, component descriptions, data flow specifications, and user interface wireframes presented in D4.1 have guided every technical decision documented in this deliverable.

Deliverable D4.3 "Multilingualism Support" details the two-pillar strategy for providing comprehensive multilingual functionality. The implementation described herein incorporates the Translation Service Layer and the 'Translate-Process-Translate' pipeline that enables Gaia to interact with learners across the six official project languages: English, German, Dutch, Greek, Slovenian, and Estonian.

This deliverable, in turn, provides the functional platform that will be the subject of D4.4 "System evaluation and public deployment," enabling the consortium to proceed with stakeholder engagement activities (WP5) on a fully operational system.

1.4. Guiding Principles for Implementation

The implementation of the L2BGreen platform has been guided by five core principles that have shaped every architectural and technical decision. Together, these principles ensure that the delivered system is not only technically robust but also pedagogically grounded, privacy-aware, and built for long-term sustainability beyond the initial funding period.

Pedagogy-First Development: Throughout the implementation, technology has been conceived as a servant to pedagogy. Every technical feature has been developed with direct reference to its support for the instructional design goals outlined in D2.1. The implementation of the multi-agent architecture, for example, was driven by the need for specialized, nuanced feedback at each stage of the learning journey, rather than by purely technical considerations.

Fidelity to Design Specifications: The implementation has adhered closely to the comprehensive design specifications established in D4.1. This principle ensures consistency and traceability between

the design phase and the implemented system, allowing for clear verification that the project's stated objectives have been met.

Scalability and Sustainability: In alignment with the critical requirement for long-term viability articulated in D4.1, the implementation has prioritized architectural choices that support broad adoption and future maintainability. The use of the LTI 1.3 standard, for instance, ensures the platform can be integrated with any compliant Learning Management System, avoiding vendor lock-in and maximizing the project's potential for widespread exploitation and sustainability beyond the initial funding period.

Data Privacy and Security by Design: The implementation incorporates robust data privacy and security measures as a core component, not an afterthought. In full compliance with the project's Data Management Plan (DMP) as outlined in D1.1, the system applies a two-tier data protection model. For analytics, research, and any data shared beyond the immediate educational context, all personally identifiable information is pseudonymized at the point of collection. Within the operational educator view, identifiable learner data (such as names and email addresses) remains accessible to authorized educators under the lawful basis of the educational relationship (GDPR Article 6(1)(e)), governed by the platform's Role-Based Access Control (RBAC) system. Across both tiers, data is encrypted in transit and at rest using industry standards (TLS/SSL and AES-256), and all data processing adheres strictly to GDPR requirements.

Modularity and Maintainability: The system has been implemented using a modular architecture that decouples the core conversational AI logic from the Learning Management System interface. This design, which separates the "Interaction Engine" (the L2BGreen Platform) from the "Learning Hub" (Moodle LMS), enables specialized development, independent updates, and the flexibility to adapt to evolving AI technologies without requiring a complete system overhaul.

2. System Architecture Implementation

2.1. Overview of the Implemented Architecture

The L2BGreen learning ecosystem has been implemented as a cohesive, three-part architecture designed for flexibility, security, and a seamless user experience. As specified in D4.1, this is not a monolithic application but rather a carefully orchestrated ecosystem where each component has a distinct role, and they are interconnected through industry-standard protocols to ensure robustness and future-proofing. This architectural separation allows for focused development on each part of the system while ensuring they work together in harmony to deliver a unified learning journey.

The decision to implement a distributed architecture rather than a single integrated platform was driven by several strategic considerations derived from the project's core principles. First, the need for specialized development meant that the sophisticated AI capabilities of Gaia required different technologies and expertise than standard LMS functionality. Second, the requirement for independent updates allowed the AI components to evolve rapidly without disrupting the stable course management infrastructure. Third, the goal of long-term sustainability mandated that the platform avoid vendor lock-in by using open standards for integration.

The implemented architecture comprises three primary layers, each fulfilling a critical function within the ecosystem:

The Learning Hub (Moodle LMS): The L2BGreen-branded Moodle instance has been deployed and configured to serve as the primary user-facing portal. It provides the foundational layer of the user experience, managing user authentication through a dedicated login page, presenting the course catalog, and structuring the overall learning content. This approach leverages Moodle's robust capabilities in user and course management, allowing learners to begin and end their journey within a familiar and well-established environment. The Moodle instance is administered directly by the consortium, enabling full control over configuration, security settings, and the installation of required plugins for LTI integration.

The Interaction Engine (L2BGreen Platform): A custom-built, standalone web application has been developed to house the project's core technological innovations: the Gaia conversational agent and the Educator and Administration Platform. This platform functions as the "Interaction Engine," driving the interactive, AI-powered learning scenarios. By decoupling this engine from the LMS, the project gains significant advantages, including the ability to pursue specialized development, push independent updates, and utilize cutting-edge technologies that are best suited for the demanding requirements of a sophisticated conversational AI system. The platform handles all real-time communication with the User's Browser via REST APIs to power the dynamic chat interface for Gaia. Crucially, it also makes secure API calls to external Large Language Model (LLM) providers to generate Gaia's responses. Beyond direct interaction, this server processes and stores all learner interaction data. For analytics and research purposes, personally identifiable information is pseudonymized at the point of collection. Within the operational context, authorized educators retain access to learner identities for their own courses through the platform's RBAC system, in accordance with GDPR provisions for educational data processing. It also hosts the analytics data pipeline responsible for aggregating raw data into meaningful metrics for educators.


The Secure Bridge (LTI 1.3): The secure and seamless connection between the Moodle LMS and the L2BGreen Platform has been achieved through the implementation of the LTI 1.3 standard. This critical industry standard, known as Learning Tools Interoperability, enables the Moodle hub to securely launch and communicate with the L2BGreen interaction engine. The external L2BGreen platform is embedded directly within a Moodle course, appearing to the user as a native activity named "Gaia Agent Activity." This strategic choice avoids vendor lock-in and ensures that the L2BGreen platform can, in the future, be integrated with other LTI-compliant Learning Management Systems such as Canvas, Blackboard, or Brightspace. This inherent interoperability maximizes the project's potential for widespread exploitation and long-term sustainability beyond the initial funding period.

2.2. Technical Architecture Components

The technical architecture of the L2BGreen learning ecosystem, as illustrated in D4.1 Figure 1, is a meticulously designed structure, composed of several key components that communicate over secure protocols to deliver a complete and cohesive learning experience. This architecture, functioning as a synergistic relationship between a "Learning Hub," an "Interaction Engine," and a "Secure Bridge," ensures flexibility, security, and a seamless user experience. The following subsections detail the implementation of each primary component.


2.2.1 Moodle Server (Learning Management System)

This component serves as the "Learning Hub," providing the foundational layer of the user experience. As the primary user-facing portal, the L2BGreen-branded Moodle instance manages user authentication through a dedicated login page, presents the course catalog, and structures the overall learning content. The implementation includes a customized Moodle theme featuring L2BGreen branding elements, ensuring visual consistency across the learning experience.



The screenshot shows the L2BGreen Moodle LMS login page. At the top is the L2BGreen logo, which consists of a green leaf icon with a white checkmark inside, followed by the text 'L2BGreen' in a bold, sans-serif font. Below the logo are two input fields: 'Username' and 'Password'. Under the 'Password' field is a blue 'Log in' button. Below the button is a link that says 'Lost password?'. A horizontal line separates the login section from the guest access section. Below the line is the text 'Some courses may allow guest access'. Under this text is a grey button that says 'Access as a guest'. Another horizontal line is below this button. At the bottom is a grey button that says 'Cookies notice'.

Figure 1: The L2BGreen Moodle LMS login page, providing secure access to the platform



L2BGreen

Home Participants Question bank Content bank

About L2BGreen Project

L2BGreen offers online green entrepreneurship training to European citizens.

L2BGreen will develop:

- Introductory learning materials that teach about what green entrepreneurship means, and how this relates to a wide variety of sectors and job types (Interactive online course including text, videos, quizzes)
- A chatbot with whom European citizens can discuss case studies (real-world examples), and plans for green entrepreneurship action within their own sphere of influence.
- A community by connecting matching users to each other to continue discussions; as well as face-2-face events in which the learning materials, the chatbot, and interested other citizens can be interacted with.

As Europe is a multilingual continent, learning materials and the chatbot will be available in multiple languages (Dutch, English, Estonian, German, Greek, and Slovenian are planned).

L2BGreen project partners include training institutions for apprentices, for green entrepreneurs, and Start-Ups, as well as universities.

Available courses

ReNuTex: Navigating Sustainability, Strategy, and Greenwashing in the Textile Industry

This course delves into the critical challenges and opportunities of sustainability in modern business, using the "ReNuTex: Rethinking Textiles Sustainably" teaching case as its core. Participants will explore how companies, particularly in the textile industry, can embed sustainability into their strategic vision, foster a truly green corporate culture, and navigate the pitfalls of greenwashing. Through a blend of case analysis, theoretical insights, self-assessment quizzes, and collaborative discussions, you will develop a comprehensive understanding of what it takes to build a truly sustainable and reputable business in the 21st century. Prepare to analyze real-world dilemmas, critically evaluate corporate communications, and propose actionable strategies for ethical and impactful sustainability leadership.

Figure 2: The L2BGreen Moodle home page displaying the project description and available courses, including the ReNuTex case study

The Moodle server has been configured with the following key features:

- **User Management:** Full integration with institutional authentication systems where applicable, with support for self-registration for open courses. User profiles are linked to pseudonymized identifiers for data protection.
- **Course Organization:** Courses are structured using Moodle's topic format, with collapsible sections that guide learners through the material sequentially. Each section can contain native Moodle resources (PDFs, pages, forums) alongside LTI external tool links to Gaia activities.
- **Gradebook Integration:** The Moodle gradebook has been configured to receive scores from the L2BGreen platform via LTI Assignment and Grade Services (AGS), enabling seamless tracking of learner progress across both native and external activities.
- **Language Support:** Moodle's built-in multilingual capabilities have been leveraged to provide interface translations in all six project languages, with language packs installed and configured for English, German, Dutch, Greek, Slovenian, and Estonian.

2.2.2 L2BGreen Application Server (Interaction Engine)

This custom-built, standalone web application functions as the "Interaction Engine" and runs the backend logic for the project's core technological innovations: the Gaia conversational agent and the Educator and Administration Platform. By decoupling this engine from the LMS, the project gains significant advantages, including specialized development, independent updates, and the ability to utilize cutting-edge technologies required for a sophisticated conversational AI system.

The L2BGreen Application Server has been implemented with the following responsibilities:

- **LTI Launch Handling:** The server implements the LTI 1.3 Tool Provider specification, receiving and validating cryptographically signed launch requests from Moodle. Upon successful validation, it authenticates users based on their pseudonymized LTI user ID and initializes appropriate learning sessions.
- **Real-time Chat Management:** All real-time communication with the User's Browser is handled via REST APIs and WebSocket connections, powering the dynamic chat interface for Gaia. The server manages conversation state, message queuing, and response streaming.
- **AI Orchestration:** The server coordinates the multi-agent Gaia system, determining which specialized agent should handle each interaction based on the current learning phase. It constructs appropriate prompts, manages conversation context, and routes requests to external LLM providers.
- **Data Management:** All learner interaction data is processed and stored securely, with personally identifiable information pseudonymized at the point of collection. The server maintains separate databases for operational data and analytics.
- **Analytics Pipeline:** A dedicated background process aggregates raw interaction data into meaningful metrics, transforming low-level logs into the high-level insights displayed on the Educator Analytics Dashboard.

2.2.3 Moderation Layer

Positioned to process user input, this layer is responsible for moderating content provided by the user before it is sent to the core Gaia agents, ensuring it adheres to safety and appropriateness guidelines. The moderation layer operates as a critical safeguard within the system's processing pipeline.

The implementation leverages an LLM-based classification model that systematically evaluates all incoming participant content against predefined guidelines for appropriate digital communication and academic conduct. The system is designed to identify and mitigate material that could be deemed harmful, offensive, irrelevant, or otherwise unsuitable within an educational framework, including content exhibiting aggressive intent, discriminatory language, sexually suggestive or explicit material, incitement to violence or self-harm, and factual inaccuracies introduced with malicious intent.

Should the analysis indicate that a participant's contribution falls outside acceptable parameters, the system initiates a calibrated response. This may involve preventing the questionable content from being passed to the primary AI agents for further interaction, issuing automated warnings or guidance to the participant, or in more severe cases, flagging the interaction for human review by platform administrators. The dynamic nature of online communication necessitates an ongoing refinement of these moderation capabilities to address emerging challenges and ensure sustained effectiveness.

2.2.4 LLM API Endpoints (External AI Service)

These endpoints are provided by Anthropic. They supply the core generative AI capabilities essential for Gaia's functionality. The L2BGreen backend makes secure API calls to Anthropic to generate the complex, nuanced responses that Gaia delivers to learners.

The implementation supports multiple LLM providers to ensure resilience and flexibility:

- **Provider Abstraction:** The system implements an abstraction layer that allows switching between LLM providers without changing the core application logic. This enables the project to leverage improvements from different providers and ensures continuity if any single provider experiences issues.
- **Secure Communication:** All API calls to external LLM services are made over HTTPS with proper authentication using API keys stored securely in environment variables. Rate limiting and retry logic ensure robust operation under varying load conditions.
- **Response Streaming:** The implementation supports streaming responses from LLM providers, enabling Gaia's messages to appear progressively in the chat interface rather than requiring users to wait for complete response generation.

2.2.5 User's Browser (User Interaction)

This represents the client-side interface that renders the user interface and facilitates the learner's interaction with the L2BGreen ecosystem. It sends learner input to the L2BGreen backend and receives real-time communication for the dynamic chat interface. The browser-based implementation ensures broad accessibility without requiring software installation.

The client-side implementation includes:

- **Responsive Design:** The interface adapts to different screen sizes, ensuring usability on desktop computers, tablets, and mobile devices. This is particularly important for the diverse European audience targeted by L2BGreen.
- **Real-time Updates:** WebSocket connections enable immediate display of Gaia's responses as they are generated, creating a natural conversational flow.
- **Accessibility Features:** The interface includes accessibility features such as keyboard navigation, screen reader compatibility, and sufficient color contrast to ensure usability for learners with disabilities.

2.3. Technology Stack Implementation

The implementation of the L2BGreen platform is built upon a modern, robust technology stack selected for its performance, scalability, and alignment with the project's long-term sustainability goals. The following technologies form the foundation of the implemented system:

Backend Framework (Laravel): The L2BGreen Application Server is built using Laravel, a PHP framework known for its elegant syntax, strong security features, and comprehensive ecosystem. Laravel provides built-in support for database management through Eloquent ORM, authentication and authorization, job queuing for background processing, and real-time event broadcasting. The framework's extensive package ecosystem enabled rapid development while maintaining code quality

and security standards. Laravel's built-in support for encryption, CSRF protection, and input validation ensures that security is integrated throughout the application.

Frontend Technologies: The L2BGreen platform employs a dual frontend architecture tailored to different user needs. The learner-facing Gaia conversational agent interface is built using React, providing a highly responsive, component-based single-page application (SPA) experience optimized for real-time chat interactions. The Educator Platform and administrative interfaces are built using HTML, Bootstrap, Livewire, and JavaScript, leveraging Laravel's server-side rendering capabilities with Livewire for dynamic interactivity without the complexity of a separate SPA framework. This combination ensures a consistent and modern visual design across all interfaces while enabling rapid development.

Database (MySQL): A MySQL relational database is used to persistently store all application data, including user profiles, team configurations, scenario content, learner interaction logs, and analytics data. MySQL was chosen for its reliability, widespread adoption, excellent Laravel integration through Eloquent ORM, robust transaction support, and strong performance at scale. The database schema has been designed to support efficient querying for both real-time operations and analytics reporting.

AI/LLM Integration: The core generative AI capabilities for Gaia are provided by Anthropic's Claude Sonnet 4.5 model through secure API endpoints. The L2BGreen backend makes secure API calls to Anthropic to generate Gaia's complex, nuanced responses. This approach allows the system to leverage state-of-the-art AI capabilities and benefit from Anthropic's continuous model improvements without requiring internal model training, hosting, or the significant infrastructure investment that would entail. The system architecture maintains a provider abstraction layer, allowing future flexibility in LLM selection if needed.

Real-time Communication (Laravel Echo and Pusher): Real-time chat functionality is implemented using Laravel Echo and Pusher (or a compatible WebSocket server such as Laravel Reverb). This enables the dynamic, conversational experience with Gaia, where responses are streamed to the user's browser as they are generated. The WebSocket implementation ensures low-latency communication essential for natural conversational interactions.

Hosting and Infrastructure: The L2BGreen Platform is hosted on a cloud infrastructure that provides high availability, automatic scaling, and robust security. The infrastructure includes load balancing for handling concurrent users, automated backups for data protection, and monitoring systems for performance and availability. The Moodle LMS instance is hosted separately, managed by the consortium, with the LTI 1.3 connection enabling secure communication between the two environments over encrypted HTTPS connections.

2.4. Data Flow Implementation

The flow of data through the implemented ecosystem has been designed to support both the pedagogical functions of the platform and the research-oriented analytics, all in full compliance with the project's Data Management Plan (DMP). The implementation ensures that data flows securely between components while maintaining user privacy and enabling the analytics necessary for continuous improvement.

2.4.1 Primary Interaction Flow

The primary interaction flow, from user action to AI response, proceeds through the following implemented steps:

Step 1 - Initial Access and Launch: A user, authenticated within the Moodle LMS through the dedicated login page, navigates to their course and clicks on a "Gaia Agent Activity" link. This action triggers Moodle to initiate a cryptographically signed LTI 1.3 Launch Request to the L2BGreen Application Server. This secure request authenticates the user and passes essential contextual information about the course and the specific activity being launched, including the user's pseudonymized ID, their role in the course (e.g., learner, instructor), and identifiers for the course and activity context.

Step 2 - Session Initialization: The L2BGreen Application Server validates the signed request using public key cryptography, authenticates the user based on their pseudonymized ID, and initializes a new learning session. The server retrieves the scenario configuration, including the language setting, from the Content Database and loads the appropriate learning materials and system prompts for the Gaia agents. The interface is rendered in the user's browser, embedded within the Moodle course page via an iframe.

Step 3 - Conversational Turn (User Input): When a learner types a response into the Gaia chat interface, the text is sent from the User's Browser to the L2BGreen backend via a REST API call. The input passes first through the Moderation Layer, which evaluates the content against predefined guidelines for appropriate digital communication and academic conduct before it is sent to the core Gaia agents for processing.

Step 4 - AI Processing: The moderated input, along with the conversation history and the relevant system prompt for the current learning phase, is sent via a secure API call to the external LLM provider. The multi-agent system determines which specialized worker agent (Scenario Presenter, Knowledge Scaffolder, Application Coach, or Argumentation Feedback Agent) should handle the response based on the current phase. For complex feedback in the Argumentation phase, the Evaluator-Optimizer workflow may invoke multiple LLM calls to ensure quality. The LLM generates Gaia's response, which is then streamed back to the L2BGreen server.

Step 5 - Response Delivery and Logging: The L2BGreen server relays Gaia's response to the User's Browser via the real-time WebSocket connection, enabling the response to appear progressively as it is generated. Simultaneously, the full interaction—including the learner's input, Gaia's response, timestamps, the current learning phase, and any agent-specific metadata—is stored in the project's dedicated database. This data is categorized as "Participant action data" and "Participant state data" as defined in the DMP, and all personally identifiable information is pseudonymized at the point of collection.

Step 6 - Grade Return (Upon Completion): When a learner completes a Gaia-powered activity, the L2BGreen platform utilizes the LTI Assignment and Grade Services (AGS) component to securely and programmatically send a completion status or a calculated score back to the appropriate column in the Moodle gradebook. This eliminates the need for manual grade entry and ensures a fully integrated workflow for educators.

2.4.2 Analytics Data Pipeline

A dedicated analytics data pipeline, running as a background process on the L2BGreen server, periodically aggregates the raw, pseudonymized interaction data. This pipeline transforms low-level logs of clicks and chat messages into meaningful, high-level metrics. These metrics include

engagement scores based on time-on-task and interaction patterns, competency mastery levels based on performance, and argumentation quality scores derived from the Evaluator agent's assessments.

Once calculated, these processed metrics are stored in a separate analytics database, optimized for fast querying and ready to be visualized on the Educator Analytics Dashboard. It is important to note that the analytics pipeline operates exclusively on pseudonymized data; no identifiable learner information is present in the aggregated metrics or dashboard visualizations. The pipeline is designed for efficiency, processing data incrementally to minimize system load while ensuring that educators have access to timely insights about their learners' progress.

2.4.3 Data Privacy and Security Implementation

Underpinning the entire system is a robust approach to data privacy and security, implemented as a core architectural requirement rather than an afterthought. In full compliance with the project's Data Management Plan (DMP) as outlined in D1.1, the implementation incorporates the following measures:

- **Pseudonymization:** The system implements a two-tier approach to personally identifiable information. For analytics, aggregated reporting, and any data shared beyond the immediate educational context, all PII is pseudonymized at the point of collection. Users are identified within the analytics pipeline by pseudonymous identifiers that cannot be traced back to individuals without access to the separate key held by consortium data controllers. Within the operational educator interface, authorized educators can view learner identities (name, email) for their own course context, governed by the platform's RBAC system and justified under the lawful basis of the educational relationship (GDPR Article 6(1)(e)). The use of the LTI 1.3 standard ensures that only the minimum necessary user information is passed from Moodle to the L2BGreen platform, including the user's role in the course and the course context.
- **Encryption in Transit:** All data is encrypted in transit using industry-standard TLS/SSL protocols. Communication between the user's browser and Moodle, between Moodle and the L2BGreen platform, and between the L2BGreen platform and external LLM providers all occurs over HTTPS connections with modern cipher suites.
- **Encryption at Rest:** Sensitive data stored in the database is encrypted at rest using AES-256 encryption. Database backups are similarly encrypted to protect against unauthorized access.
- **GDPR Compliance:** The system is designed to adhere strictly to GDPR requirements. All data processing is based on either the explicit consent of the user or a clearly defined legitimate interest, as documented in the project's DMP. Users have the right to access, rectify, and delete their data through documented procedures.
- **Access Control:** Role-Based Access Control (RBAC) ensures that users can only access data appropriate to their role. Educators can view aggregated and anonymized analytics for their courses but cannot access raw data from other courses or institutions

3. Implementation of the Gaia Conversational Agent

The Gaia conversational agent stands as the pedagogical heart of the L2BGreen project, embodying its core educational philosophy and delivery mechanism. Far from being a singular, undifferentiated artificial intelligence, Gaia has been meticulously engineered as a sophisticated, hierarchical system comprising multiple collaborating agents. Each of these specialized agents is precisely designed to execute a distinct and integral part of the overarching educational script, as meticulously defined and detailed in Deliverable D2.1 Interactive online course instructional design.

This innovative agentic architecture is not merely an arbitrary design choice; rather, it provides a multitude of critical benefits, including unparalleled robustness in handling diverse educational scenarios, a high degree of specialization that allows for nuanced and targeted interactions, and, crucially, the inherent capacity for advanced reasoning capabilities necessary for truly effective and adaptive learning experiences. This modular and interconnected design ensures that Gaia can seamlessly guide learners through complex topics, adapt to individual learning styles, and provide targeted feedback, thereby maximizing the pedagogical impact of the L2BGreen initiative.

3.1. Realization of the Multi-Agent Architecture

The Gaia conversational agent has been implemented as a sophisticated, hierarchical system of collaborating AI agents, precisely as designed in D4.1. This agentic architecture is central to the pedagogical effectiveness of the L2BGreen system, providing robustness, specialization, and the capacity for advanced reasoning that a single, monolithic AI could not achieve. Each specialized agent has been implemented as a distinct module with a highly tuned system prompt and a well-defined function.

3.1.1 The Orchestrator Agent

The Orchestrator Agent serves as the "brain" of the conversational flow, functioning as the master agent that coordinates all interactions within a learning session. This agent does not directly generate pedagogical content but instead manages the state of the learner's interaction and ensures the educational journey follows the intended path.

Its primary responsibilities, realized through its system prompt and programmatic logic, are:

- **Phase Tracking:** The Orchestrator tracks the learner's progress through the predefined learning phases (Introduction, Knowledge Acquisition, Application, Argumentation) as defined in D2.1. It maintains awareness of which phase the learner is currently in and which content has been delivered.
- **Conversation History Management:** The agent maintains the full conversation history, ensuring that context is preserved across interactions and that responses are coherent with previous exchanges.
- **Agent Delegation:** Based on the current phase and the learner's input, the Orchestrator delegates tasks by invoking the appropriate specialized worker agent at the correct time. This delegation is rule-based for phase transitions but can be adaptive within phases based on learner responses.

- **Flow Control:** The Orchestrator ensures the conversation follows the intended pedagogical path, as evidenced by the clear progression through "PART 1," "PART 2," etc., in the user interface. It manages transitions between parts and ensures learners have engaged with required content before advancing.

The implementation of the Orchestrator follows a state machine pattern, where the current state determines which actions are available and which worker agent should be invoked. State transitions are triggered by specific events such as learner confirmation buttons, completion of question sequences, or reaching the end of content delivery for a particular phase.

3.1.2 Specialized Worker Agents

The four worker agents specified in D4.1 have been fully implemented, each as a distinct LLM-powered module with a unique system prompt tailored to its function. This specialization allows for more precise and reliable performance compared to a single, general-purpose prompt, as each agent can be optimized for its specific pedagogical task.

The following sections describe the key worker agents within the Gaia framework, detailing their triggering phases, primary functions, inputs, outputs, and key prompting techniques employed:

Scenario Presenter Agent: This agent is triggered during Phase 1 (General Case Introduction) of the learning script. Its primary function is to present the business case from D3.1 Learning materials and ask initial comprehension questions to ensure learners have engaged with the foundational content. The agent receives the Case ID as input and produces formatted case text along with initial comprehension questions as output. It uses zero-shot prompting to deliver the case content clearly and engagingly, introducing itself as Gaia and welcoming the learner to the activity. The ReNuTex case, for example, is presented with rich contextual information about the company's sustainability journey and the tensions that arose from marketing claims.

Knowledge Scaffold Agent: This agent is active during Phase 2 (Sub-topic Introduction), delivering the two sub-topic knowledge inputs from D3.1 in structured parts and asking knowledge questions to check understanding. The agent receives the Case ID and Sub-topic ID as input and produces formatted knowledge text and knowledge questions as output. It employs a combination of zero-shot and few-shot prompting techniques, and crucially implements Socratic tutoring capabilities. When a learner provides an incorrect answer, the Knowledge Scaffolder does not simply state the answer is wrong; instead, it provides a contextual hint that guides the learner to self-correct. For example, when a learner incorrectly answers a question about ReNuTex's marketing mistake, the agent might respond: "That's an interesting thought, but let's look closer at the case. It mentions that sustainability experts warned the claims were 'exaggerated and potentially misleading.' Does this suggest they hyped it too much or not enough? Please try again." This approach promotes deeper learning through guided discovery.

Application Coach Agent: Triggered during Phase 3 (Application), this agent poses application questions that require learners to connect the sub-topics to the case, demonstrating their ability to apply theoretical knowledge to practical situations. The agent receives the learner's prior answers, case context, and application questions with initial feedback templates as input. It uses few-shot prompting, providing the LLM with examples of good analysis to guide its evaluation of learner responses. When a learner provides a strong answer connecting sustainability strategy to the ReNuTex case, the Application Coach affirms the reasoning: "That's an excellent approach. Setting clear 'green

goals' for each department is a powerful way to ensure cross-functional integration, making sustainability everyone's responsibility, not just a single team's."

Argumentation Feedback Agent: Active during Phase 4 (Argumentation), this composite agent provides detailed, structured feedback on the learner's arguments. It receives the learner's argument, case context, and argumentation rubric as input and produces structured feedback as output. This agent is implemented using the Chain of Thought (CoT) prompting and Evaluator-Optimizer workflow described in detail in subsequent sections, ensuring that feedback is transparent, pedagogically sound, and self-corrected. The agent evaluates arguments against established frameworks such as the Toulmin model, providing learners with insight into the strengths and weaknesses of their reasoning.

3.2. Implementation of Advanced Prompting Techniques

To move beyond simple question-answering and provide genuinely educational interactions, the implementation incorporates the advanced prompting techniques specified in D4.1. These techniques are particularly crucial for the high-stakes task of providing feedback on a learner's reasoning, where generic or inaccurate feedback could undermine the learning experience.

3.2.1 Chain of Thought (CoT) Prompting for Feedback Generation

The generation of feedback, especially on complex argumentation, must be transparent and pedagogically sound. Simply providing a score or a generic comment is insufficient for learning. Chain of Thought (CoT) prompting is a technique that forces the LLM to "show its work" by generating a step-by-step reasoning process before delivering a final answer. This approach makes the AI's "thinking" visible, allowing the learner to understand why their argument is strong or weak, which is fundamental to metacognitive development.

The system prompt for the Argumentation Feedback Agent has been implemented to command a specific, structured reasoning process. Rather than simply asking for feedback, the prompt instructs the LLM to follow a multi-step process:

- 1. Identify the student's primary claim:** The agent first extracts and articulates the main assertion the learner is making in their argument.
- 2. List the evidence used from the case study:** The agent identifies specific facts, examples, or data from the ReNuTex case (or other case material) that the learner cited to support their claim.
- 3. Evaluate the logical connection between claim and evidence:** The agent assesses whether the evidence actually supports the claim, referencing the Toulmin model of argumentation. This includes considering the claim (the assertion being made), the data (the evidence supporting it), and the warrant (the underlying reasoning or principle connecting the data to the claim).
- 4. Consider any overlooked counterarguments or alternative perspectives:** The agent identifies potential rebuttals or alternative interpretations that the learner might have addressed to strengthen their argument.
- 5. Synthesize these points into a final piece of constructive feedback:** The agent combines its analysis into encouraging, educational feedback formatted in Markdown for clear presentation.

To enhance academic rigor and ensure explicit integration of argumentation theory, the system prompt provides the LLM with a concise, operational definition of the Toulmin model's key components. For instance, the prompt specifies: "When evaluating the logical connection between the claim and the evidence, explicitly consider the student's claim (the assertion being made), data (the evidence supporting the claim, drawn from the case study), and warrant (the underlying reasoning or principle connecting the data to the claim). Assess if the warrant is clear, justifiable, and relevant within the context of green entrepreneurship. Also, consider any implicit backing for the warrant, and whether qualifiers (e.g., 'usually', 'sometimes') or rebuttals (potential counterarguments) are acknowledged where appropriate, indicating a nuanced understanding."

This refinement guides the LLM to systematically apply the Toulmin model's framework during its reasoning process, leading to more structured and academically robust feedback that directly addresses the components of a well-formed argument. The feedback to the learner then not only states that the Toulmin model was referenced but implicitly demonstrates its application by critiquing or reinforcing the student's claim-data-warrant connections.

3.2.2 Few-Shot Prompting for Application Questions

The Application Coach agent employs few-shot prompting to guide its evaluation of learner responses to application questions. Unlike zero-shot prompting where the model receives only instructions, few-shot prompting includes examples of high-quality responses and feedback within the prompt, providing the LLM with concrete patterns to emulate.

The implementation includes carefully crafted examples that demonstrate:

- **Good Analysis Structure:** Examples show how strong responses connect theoretical concepts from the sub-topics to specific elements of the business case.
- **Appropriate Feedback Tone:** Examples illustrate how to affirm correct reasoning while gently redirecting incomplete or incorrect analysis.
- **Reference to Case Material:** Examples demonstrate how to ground feedback in specific details from the learning materials, reinforcing the case-based learning methodology.

3.2.3 Socratic Tutoring Implementation

The Knowledge Scaffold agent implements Socratic tutoring principles, particularly for handling incorrect answers to comprehension questions. Rather than simply indicating that an answer is wrong or immediately providing the correct answer, the agent guides learners through a process of self-discovery.

The implementation follows a structured approach:

1. **Acknowledge the Response:** The agent first acknowledges the learner's attempt, avoiding dismissive language that might discourage engagement.
2. **Provide a Contextual Hint:** The agent references specific information from the case material that points toward the correct answer without stating it directly.

3. Prompt Reconsideration: The agent explicitly invites the learner to reconsider their answer in light of the hint.

4. Affirm Correct Responses: When the learner arrives at the correct answer (whether initially or after hints), the agent provides positive reinforcement that explains why the answer is correct, deepening understanding.

This approach is evidenced in the user interface, where a learner who incorrectly identifies ReNuTex's marketing mistake might receive feedback such as: "That's an interesting thought, but let's look closer at the case. It mentions that sustainability experts warned the claims were 'exaggerated and potentially misleading.' Does this suggest they hyped it too much or not enough? Please try again." When the learner correctly identifies that the company overclaimed their environmental achievements, the agent responds: "That's exactly right. They exaggerated their achievements, which could be seen as greenwashing."

3.3. Implementation of the Evaluator-Optimizer Workflow

While Chain of Thought prompting improves the quality of a single AI-generated response, it does not guarantee its correctness or pedagogical appropriateness. To address this, the Argumentation Feedback Agent has been implemented with an internal, self-correcting loop as specified in D4.1. This Evaluator-Optimizer workflow creates a layer of automated quality assurance, ensuring the feedback delivered to learners is not only accurate but also constructive, appropriately referenced, and aligned with learning objectives.

3.3.1 Workflow Architecture

The implemented workflow proceeds through the following stages:

Optimizer Agent (First Pass): The first AI call receives the learner's argument and the CoT prompt described above, generating an initial draft of the feedback. This draft incorporates the step-by-step reasoning process and produces structured feedback based on the argumentation analysis.

Evaluator Agent (Quality Check): A second, distinct AI call receives the learner's original argument, the Optimizer's feedback draft, and a specialized "rubric prompt." This rubric prompt instructs the Evaluator to critically assess the feedback against a quality checklist. The Evaluator considers the following criteria:

- Is the feedback constructive and encouraging, avoiding overly negative language?
- Does it accurately reference the case material from D3.1 Learning materials (such as the ReNuTex case)?
- Does it align with the learning objectives defined for the activity?
- Is it specific enough to be actionable, avoiding generic comments?
- Does it correctly interpret subtle nuances in the learner's argument?

The Evaluator responds with a single word, 'ACCEPT' or 'REJECT', followed by a one-sentence justification explaining its decision.

Control Logic: A programmatic loop implemented in the L2BGreen Application Server manages this interaction. If the Evaluator returns 'ACCEPT', the feedback is deemed high-quality and is sent to the learner. If it returns 'REJECT' (e.g., "The feedback is too generic and fails to mention the PESTEL analysis" or "The feedback incorrectly attributes the learner's point to a different concept"), the justification is appended to the original prompt, and the Optimizer is called again to generate a revised, improved response that addresses the identified deficiency.

3.3.2 Iteration Limits and Latency Management

To manage latency and computational cost, the refinement loop is capped at a maximum of two iterations. The decision to cap at two iterations is a pragmatic balance between achieving high-quality feedback and managing computational latency and cost. In practice, the first iteration of the Optimizer Agent, guided by CoT prompting, typically produces a solid initial draft. The Evaluator Agent often identifies common failure modes that necessitate a second pass, such as feedback that is too generic, avoids negative language at the expense of accuracy, fails to accurately reference the case material, or misinterprets subtle nuances in the learner's argument.

These quality criteria ensure the feedback is not only constructive but also specific, accurate, and aligned with learning objectives. The second iteration allows the Optimizer to integrate this targeted critique, aiming to resolve any specific deficiencies and achieve the desired quality standards without incurring excessive processing time or expense for marginal further gains.

This layered architecture, combining a multi-agent framework with CoT and an Evaluator-Optimizer loop, represents a systematic approach to de-risking the use of generative AI in a sensitive educational context. The multi-agent design breaks down complexity, CoT ensures transparent reasoning, and the Evaluator-Optimizer workflow provides a final quality gate, all working in concert to produce reliable and pedagogically valuable feedback.

3.4 Implementation of the Moderation Layer

To cultivate a secure, respectful, and pedagogically sound learning environment, the L2BGreen ecosystem has implemented a robust and adaptive content moderation capability as specified in D4.1. This essential layer forms a critical part of the system's advanced mechanisms for maintaining quality and integrity during interactive sessions. Its primary function is the proactive analysis and filtering of participant contributions, encompassing both textual inputs and any associated media, before these reach the core generative AI components for processing.

3.4.1 Content Classification and Analysis

Operating as a sophisticated safeguard, this moderation capability systematically evaluates all incoming participant content against a comprehensive set of predefined institutional guidelines for appropriate digital communication and academic conduct. The system leverages an LLM-based classification model that performs a multi-faceted assessment of the submitted content, classifying the nature of the material and providing a degree of confidence regarding its potential misalignment with established policies.

The aim is to promptly identify and mitigate material that could be deemed harmful, offensive, irrelevant, or otherwise unsuitable within an educational framework. This encompasses a broad spectrum of potentially problematic expressions, including but not limited to:

- Content exhibiting aggressive intent or harassment
- Discriminatory language targeting protected characteristics
- Incitement to violence or self-harm
- Sexually suggestive or explicit material
- Factual inaccuracies introduced with malicious intent
- Any content that aims to disrupt the learning process or undermine a constructive dialogue

3.4.2 Calibrated Response Mechanisms

Should the analysis indicate that a participant's contribution falls outside acceptable parameters, the L2BGreen system is engineered to initiate a calibrated response. The response is proportionate to the severity of the violation:

- **Content Blocking:** For clearly inappropriate content, the questionable material is prevented from being passed to the primary AI agents for further interaction. The learner receives a message indicating that their input could not be processed and is invited to rephrase.
- **Automated Warnings:** For borderline content or first-time minor violations, the system issues automated warnings or guidance to the participant, explaining what type of content is expected in the educational context.
- **Human Review Flagging:** In more severe cases or for repeated violations, the interaction is flagged for human review by platform administrators, who can take appropriate action including communication with the learner or, in extreme cases, account restrictions.

The dynamic nature of online communication necessitates an ongoing refinement of these moderation capabilities to address emerging challenges and ensure sustained effectiveness. This diligent and adaptive screening process is paramount to preserving a positive and productive learning atmosphere, reinforcing the project's unwavering commitment to ethical AI deployment. By proactively analyzing and filtering participant contributions, this layer ensures a safe, respectful, and supportive environment, which fosters sustained engagement and enhances overall learning motivation throughout the interactive learning journey.

3.5 Learner Interface Implementation

The design of the Gaia learner interface has been implemented to prioritize clarity, focus, and a guided conversational experience, directly following the specifications in D4.1. The user interface translates the complex backend operations of the multi-agent system into an intuitive and interactive learning journey, enabling learners to focus effectively on the content of the interaction and maximizing pedagogical impact and user engagement.

3.5.1 Visual Design and Layout

The user interface is clean and minimalist, using a standard chat format with distinct bubbles for Gaia's messages and the learner's responses. This familiar paradigm minimizes cognitive load, allowing the learner to focus on the content of the interaction rather than navigating an unfamiliar interface. Key visual elements include:

- **Dark Header Theme:** A dark-themed header provides visual contrast and displays the current activity context.
- **Agent Avatar:** Gaia's messages are accompanied by a friendly avatar, personalizing the interaction and distinguishing agent messages from learner input.
- **Checkmarks for Completion:** Green checkmarks indicate completed sections or correctly answered questions, providing visual feedback on progress.
- **Action Buttons:** Clearly labeled buttons (such as "Let's begin," "Got it," "I have read the above, let's continue") guide learners through structured interactions, ensuring they consciously engage with content before proceeding.
- **Text Input Area:** A clean text input area at the bottom of the interface allows learners to compose their responses, with a send button for submission.




Figure 3: The Gaia learner interface showing the initial greeting and Part 1 (Introducing the Business) of the ReNuTex case study activity

3.5.2 Structured Interaction Flow

The interaction flow is structured into discrete, numbered "Parts" that directly correspond to the learning phases defined in D2.1 Interactive online course instructional design and are managed by the Orchestrator Agent. A typical activity, such as the "ReNuTex" case, unfolds through a sequence of clearly labeled parts:

Introduction (Part 1 - Introducing the Business): The activity begins with Gaia introducing herself and the case study. The Scenario Presenter agent delivers a concise block of text summarizing the

business case, including an illustrative image. For the ReNuTex case, learners are introduced to a European textile company based in Portugal, producing clothing using recycled fibers and low-impact dyeing technologies. The narrative explains the company's sustainability journey, the creation of a Sustainability Task Force, and the tension that arose when marketing claims were deemed exaggerated. The learner signals readiness to proceed by clicking a "Let's begin" or "Got it" button, ensuring they have actively engaged with the initial context.

Knowledge Comprehension: Following the case introduction, the agent asks a series of closed-ended questions to check for basic comprehension of the case. Questions address key facts such as what type of products ReNuTex produces, what the main goal of their Sustainable Transformation Strategy 2030 was, and what mistake the marketing team made. When a learner provides an incorrect answer, the Knowledge Scaffold agent demonstrates its Socratic tutoring capability by providing a contextual hint that guides the learner to self-correct. Correct answers are met with positive reinforcement ("Excellent, that's the right idea") before moving to the next question.





Figure 4: The Gaia conversational interface demonstrating Socratic tutoring: guiding a learner through knowledge comprehension questions on the ReNuTex case

Knowledge Delivery (Parts 2 & 3 - Sub-topic Content): The Knowledge Scaffold agent presents the core learning material from D3.1. Each sub-topic is delivered as a separate, clearly labeled "Part"

(e.g., "PART 2: Strategic Thinking and Fostering Sustainability in Corporate Culture" and "PART 3: Greenwashing and its Risks in the Context of Strategic Leadership"). The content is presented in structured text blocks with numbered key points and practical examples tied to the ReNuTex case. To ensure the learner has processed the information, they must click a confirmation button like "I have read the above, let's continue" before the activity can advance. This deliberate pacing prevents learners from passively scrolling through content without engagement.



PART 2  Strategic Thinking and Fostering Sustainability in Corporate Culture

Building a sustainable corporate culture requires more than mission statements and PR campaigns. It demands strategic thinking, integrated decision-making, and leadership commitment. Sustainability must become a lens through which the organization sees everything—from product design to logistics, employee engagement, and investor relations. Sustainability is not a side strategy—it is a core business strategy that enhances innovation, resilience, and stakeholder trust.

Key Aspects of Strategic Sustainability Leadership

- 1 Purpose-driven vision: A strong sustainability culture starts with clear strategic intent. In ReNuTex's case, the "Sustainable Transformation Strategy 2030" gave a long-term vision, aligning the company's values with environmental and social goals.
- 2 Cross-functional integration: Sustainability is not only the job of the sustainability officer. Strategic thinking means embedding it across departments—from HR to Finance, from Marketing to Logistics.
- 3 Empowering employees: ReNuTex offered sustainability workshops and learning labs for staff. This built ownership and helped employees understand how their everyday actions contribute to the broader mission.
- 4 Transparent systems: KPIs, sustainability dashboards, and open internal communication help track progress and build a culture of accountability. ReNuTex's tracking system ensured employees at all levels had visibility on impact metrics.
- 5 Adapting business models: Strategic sustainability can lead to new models (e.g. circular fashion, repair services, or product-as-a-service models). Strategic thinking explores how to pivot operations toward these new models.
- 6 Courage to rethink communication: When ReNuTex realized their marketing campaign was not aligned with actual performance, leadership paused the campaign and brought together all departments to realign the message with the strategy. This strategic pause helped prevent reputational harm.

Results at ReNuTex

ReNuTex didn't just reduce emissions or increase recycling—it shifted how its people thought and acted. Team meetings began including sustainability updates. Employee rewards included "impact champions" awards. Strategic sustainability became part of everyday decisions, not a separate track. This shift increased employee retention, built brand trust, and helped ReNuTex attract sustainable investment funds.

Figure 5: PART 2 of the Gaia activity: 'Strategic Thinking and Fostering Sustainability in Corporate Culture' — knowledge content delivered with a structured read-and-confirm interaction

Application and Argumentation (Part 4 - Application Questions): The final phase shifts to open-ended questions posed by the Application Coach. Questions challenge learners to apply their knowledge, such as: "If you were the CEO of ReNuTex, how would you structure internal processes to ensure all departments stay aligned on sustainability goals?" and "Can you describe a real case of greenwashing? What could the company have done differently?" The learner types their analysis and arguments into a text box and submits their response. The Argumentation Feedback Agent then

provides evaluative feedback that affirms correct reasoning ("That's an excellent approach"), connects the learner's ideas back to the case material and underlying concepts, and may reference real-world examples (such as the Volkswagen 'Dieselgate' scandal as a textbook example of greenwashing). The conversation proceeds sequentially through the application questions, guided by Gaia's prompts.

PART 4
Application and Argumentation Questions

If you were the CEO of ReNuTex, how would you structure internal processes to ensure all departments stay aligned on sustainability goals?

If I was CEO, I'd make sure every department had clear "green goals" linked to our main plan. Also I would implement regular check-ins and shared dashboards so everyone can see our progress.

That's an excellent approach. Setting clear "green goals" for each department is a powerful way to ensure cross-functional integration, making sustainability everyone's responsibility, not just a single team's.

And you're right, shared dashboards and regular check-ins are crucial for creating the kind of transparency and accountability mentioned in the case. This ensures everyone is aligned and moving in the same direction.

Let's move to the next question.

Can you describe a real case of greenwashing? What could the company have done differently?

A real example of greenwashing was when Volkswagen got caught cheating on emission tests for their "clean diesel" cars. Instead of cheating the numbers they should have been honest about them and focus on truly developing cleaner technology.

That's a classic and powerful example. The Volkswagen "Dieselgate" scandal is a textbook case of greenwashing, where the company deliberately deceived the public and regulators by marketing cars as "clean diesel" while they were, in fact, emitting pollutants far above the legal limits.

Your suggestion for what they could have done differently is spot on. Instead of cheating, a strategic and honest approach would have involved being transparent about their technological limitations and communicating a clear, credible roadmap for developing genuinely cleaner vehicles. This builds long-term trust, which is far more valuable than any short-term marketing win.

Well done. Let's move on.

Imagine you're a marketing manager—what tools or checks would you put in place before making environmental claims about a product?

Type your answer...

Figure 6: PART 4 of the Gaia activity: Application and Argumentation Questions, showing learner responses and AI-generated feedback, including the Volkswagen 'Dieselgate' greenwashing example

3.5.3 Accessibility and Responsive Design

The interface has been implemented with accessibility in mind, incorporating features such as keyboard navigation support, sufficient color contrast ratios for readability, screen reader compatibility, and responsive design that adapts to different screen sizes. These features ensure that the L2BGreen learning experience is accessible to the diverse European audience targeted by the project, including learners with disabilities and those accessing the platform from various devices.

4. Implementation of the Educator and Administration Platform

While Gaia provides the direct learning experience, the Educator and Administration Platform provides the tools for educators to manage, orchestrate, and analyze the learning process. This platform has been implemented as the "control panel" of the L2BGreen ecosystem, designed for usability by non-technical users and scalability for institutional deployment. It empowers educators to create and deploy sophisticated, AI-powered learning scenarios using the content developed in WP3 Learning materials on green entrepreneurship, while providing actionable insights through advanced learning analytics.

4.1. Scenario Orchestration and Content Management

The platform's primary function is to empower educators to create and deploy sophisticated, AI-powered learning scenarios using the content developed in WP3 Learning materials on green entrepreneurship. The implementation delivers a complete workflow for scenario management, from initial creation through deployment and ongoing maintenance.

4.1.1 Dashboard User Interface Design

The user interface has been designed for simplicity and clarity, following modern UX principles to ensure accessibility for educators regardless of their technical background. Upon logging in, the educator is presented with a main dashboard providing an at-a-glance view of all available "Case Activities," presented as visual cards. Each card displays:

- The case title (e.g., "ReNuTex – Rethinking Textiles Sustainability")
- Its primary topic displayed as a tag (e.g., "Sustainability")
- A brief description of the activity's focus
- Key metrics like the number of interactions and indicative duration
- A "Details" link for accessing the full activity configuration and analytics

A prominent "Add a new case activity" button guides the user toward content creation, while a toggle allows filtering to show or hide inactive activities. The dashboard also features a header with the L2BGreen branding, navigation to the Home view and Contact pages, and user account access in the upper right corner.




Figure 7: The Educator and Administration Platform dashboard, showing deployed Case Activity cards including the ReNuTex scenario

4.1.2 Scenario Builder Workflow

The implementation guides educators through a simple, step-by-step process to configure a new learning activity. This workflow has been designed to minimize complexity while providing sufficient flexibility for customization:

1. Create Scenario: The educator initiates the process by clicking "Add a new case activity" from the main dashboard. This opens a dedicated form for configuring the new activity.

2. Configure General Information: A dedicated form prompts the educator for the core details of the activity. This includes a unique Title for the activity, selecting an associated Course Name from available options, providing an optional Activity Description using a rich text editor with formatting options (bold, italic, underline, lists, links, and headings), and setting an Indicative Duration in minutes to help learners plan their time.

3. Set Language: A mandatory language selector allows the educator to define the language for the scenario from the six supported project languages (English, German, Dutch, Greek, Slovenian, Estonian). This selection is saved as an attribute of the activity and dictates the language of interaction with the Gaia agent, as detailed in D4.3 Multilingualism Support.

4. Add Visuals: To enhance the learner experience, the educator can upload a Case activity image via a simple drag-and-drop interface or by browsing their local files. The interface also supports adding media from URL for externally hosted images. The image appears on the activity card in the dashboard and within the Gaia interface during the case introduction.

5. Deploy: Once saved, the new case activity becomes available on the main dashboard. A processing period is required before the activity becomes available for use, ensuring all content is properly indexed and the AI agents are configured. The activity can then be linked from the Moodle LMS via the LTI external tool configuration.





Figure 8: The 'Add a New Case Activity' form in the Educator Platform, enabling step-by-step scenario configuration including title, course, description, and case image

4.1.3 Activity Details and Conversation Monitoring

Clicking on "Details" for any case activity opens a detailed view that provides educators with comprehensive information about the activity and its usage. This view includes:

- **Activity Preview:** A live preview window showing the Gaia interface as learners would see it, allowing educators to test the activity flow.
- **About This Activity:** A text section describing the activity's educational purpose and approach.
- **Activity Details:** A structured list of the specific Learning Objectives for the case, allowing educators to directly connect learner performance to the pedagogical goals defined in D3.1 Learning materials.
- **Conversations Table:** A log of all learner interactions with Gaia for this activity. As an operational tool designed to support direct pedagogical follow-up, the table displays the learner's ID, name, email, and a preview of the conversation messages. Access to this identifiable data is restricted by the platform's RBAC system to authorized educators within the specific course context, under the lawful basis of the educational relationship (GDPR

Article 6(1)(e)). This operational view is distinct from the analytics dashboard, which operates exclusively on pseudonymized data. Educators can click on any conversation to view the full transcript, enabling detailed review of individual learner progress and identification of common difficulties.



ReNuTex
Script Template: A1

Strategic thinking

Hi, I am Gaia
I am an AI assistant specializing at green entrepreneurship training.

When you are ready to start the case-based learning activity (**ReNuTex**), which consists of 4 parts, click on the below button.

Let's begin

About this activity

This activity uses the "ReNuTex - Rethinking Textiles Sustainably" case study to explore the implementation of sustainability within a business context. Participants analyze how ReNuTex, a European textile company, attempts to embed sustainability into its core strategy and operations through its "Sustainable Transformation Strategy 2030". The case highlights key aspects of strategic sustainability leadership, fostering a sustainable corporate culture, and the critical link between internal strategy and external communication. It specifically addresses the risks of greenwashing when marketing claims potentially outpace actual performance, prompting reflection on transparency and strategic alignment. The activity involves reading the case, answering comprehension and knowledge questions, and engaging with application questions to deepen understanding.

Activity Details

Learning Objectives:

- Understand the role of strategic thinking in embedding sustainability into corporate culture (Main Subtopic 8.2).
- Identify key actions for fostering a sustainable corporate culture (e.g., cross-functional integration, employee empowerment, KPIs).
- Define greenwashing and recognize its common forms and risks (Further Subtopic 6.2).
- Analyze the relationship between strategic alignment and truthful sustainability communication to avoid greenwashing.
- Apply concepts of sustainable strategy and communication ethics to the ReNuTex case and other scenarios.

Indicative Duration: 30 minutes

Conversations

ID	NAME	EMAIL	MESSAGES
1	Apostolos Mavridis	apostolos@enchatted.com	Hello! I'm Gaia, your virtual assistant...
2			Hello! I'm Gaia, your virtual assistant...
3			Hello! I'm Gaia, your virtual assistant...
4			Hello! I'm Gaia, your virtual assistant...
5			Hello! I'm Gaia, your virtual assistant...
6			Hello! I'm Gaia, your virtual assistant...

Report

© 2025, made with ❤ by L2BGreen

License

Figure 9: The Activity Details view showing the ReNuTex scenario with Learning Objectives, an embedded Gaia preview, and the Conversations log table for educator monitoring


4.2. User and Team Management Implementation

To support use within institutions, the platform includes a robust system for managing users and organizing them into collaborative teams, implementing the Role-Based Access Control (RBAC) system specified in D4.1. This system ensures appropriate access levels for different users while facilitating collaboration within educational institutions.

4.2.1 User Profile Management

Each user has a comprehensive profile page where they can manage their personal information and account security. The profile management interface includes:

- **Profile Information:** Users can update their account's profile information and email address. A profile photo can be selected to personalize their account.
- **Update Password:** A dedicated section allows users to ensure their account security by using a long, random password. The interface requires the current password for verification and confirmation of the new password.
- **Two-Factor Authentication (2FA):** The platform provides an option to add additional security to accounts using two-factor authentication. When enabled, users are prompted for a secure, random token during authentication, which can be retrieved from a phone's Google Authenticator application or similar 2FA app.
- **Browser Sessions:** Users can manage and log out their active sessions on other browsers and devices. The interface shows a list of all active browser sessions, including the browser type, operating system, and last active time. If users feel their account has been compromised, they can log out all other browser sessions and update their password.
- **Delete Account:** For GDPR compliance, users have the ability to permanently delete their account and all associated data.



Profile Information
Update your account's profile information and email address.

AM
Select A New Photo

Name
Apostolos Mavridis

Email
apostolos@enchatted.com

Save

Update Password
Ensure your account is using a long, random password to stay secure.

Current Password

New Password

Confirm Password

Save

Two Factor Authentication
Add additional security to your account using two factor authentication.

You have not enabled two factor authentication.

When two factor authentication is enabled, you will be prompted for a secure, random token during authentication. You may retrieve this token from your phone's Google Authenticator application.

Enable

Browser Sessions
Manage and log out your active sessions on other browsers and devices.

If necessary, you may log out of all of your other browser sessions across all of your devices. Some of your recent sessions are listed below; however, this list may not be exhaustive. If you feel your account has been compromised, you should also update your password.

OS X - Chrome
212.205.251.184, **This device**

OS X - Chrome
87.202.101.83, Last active 1 minute ago

Log Out Other Browser Sessions

Delete Account
Permanently delete your account.

Once your account is deleted, all of its resources and data will be permanently deleted. Before deleting your account, please download any data or information that you wish to retain.

Delete Account

© 2025, made with ❤️ by L2BGreen

[License](#)

Figure 10: The User Profile page, providing account management including profile information, password update, Two-Factor Authentication, browser session management, and GDPR-compliant account deletion

4.2.2 Hierarchical Team Structure

The platform has been built for collaborative work, recognizing that educational content is often developed and managed by teams rather than individuals. An administrator can create a Team, giving it a name and assigning a Team Owner. This allows an institution, department, or project group to manage its users and content centrally.

The team creation interface allows administrators to:

- Assign a Team Owner who has overall responsibility for the team

- Set the Team Name for easy identification
- Add a description of the team's purpose or scope


4.2.3 Role-Based Access Control (RBAC) Implementation

Within a team, administrators can add new members by email and assign them specific roles. The platform defines two primary roles with clearly differentiated permissions:

Administrator: Users with the Administrator role can perform any action within the team context. This includes managing team settings and information, adding and removing team members, assigning roles to team members, creating, editing, and deleting all case activities, viewing full analytics for all activities, and configuring team-level preferences.

Editor: Users with the Editor role have the ability to read, create, and update case activities but cannot manage team settings or members. This role is appropriate for content creators and educators who need to develop and maintain learning activities without administrative responsibilities. Editors can view analytics for activities they have access to but cannot modify team membership or settings.

This RBAC system ensures that different users have appropriate levels of access, which is crucial for security and organized content management within an educational institution. The interface for adding team members displays clear descriptions of each role's capabilities, helping administrators make informed decisions about role assignments.



The screenshot displays the L2BGreen Team Settings interface. On the left, a sidebar contains the L2BGreen logo and navigation links for 'Home' and 'Contact'. The main content area is divided into two primary sections. The top section, titled 'Team Name', includes a subtitle 'The team's name and owner information.' and displays the 'Team Owner' as 'Apostolos Mavridis' with the email 'apostolos@enchanted.com'. Below this, there is a text input field for the 'Team Name' which currently contains the word 'Enchanted', and a 'Save' button. The bottom section, titled 'Add Team Member', has a subtitle 'Add a new team member to your team, allowing them to collaborate with you.' and a note stating 'Please provide the email address of the person you would like to add to this team. The email address must be associated with an existing account.' This section contains an 'Email' input field, a 'Role' dropdown menu, and an 'Add' button. The 'Role' dropdown is open, showing two options: 'Administrator' (selected) and 'Editor'. Below each role, a brief description of its capabilities is provided: 'Administrator users can perform any action.' and 'Editor users have the ability to read, create, and update.'

Figure 11: The Team Settings interface, demonstrating the Role-Based Access Control (RBAC) system with Administrator and Editor roles and team member management for the Enchanted team

4.3. Implementation of the Learning Analytics Dashboard

A key innovation of the platform is its advanced analytics dashboard, designed to provide educators with actionable insights that go far beyond the descriptive reports available in a standard Moodle installation. Unlike standard descriptive reports, this dashboard delivers prescriptive and actionable insights through aggregated, pseudonymized data from both Moodle and the L2BGreen database. The implementation delivers the data aggregation, metric calculation, and visualization capabilities specified in D4.1.

4.3.1 Data Aggregation and Sources

The dashboard's backend securely pulls and aggregates pseudonymized data from two primary sources:

- **Moodle Gradebook:** Scores on native Moodle Quizzes and activities are retrieved via the LTI connection, providing data on learner performance in traditional assessment formats.
- **L2BGreen Database:** Detailed interaction data with Gaia, including conversation transcripts, response times, agent assessments, and completion status, is stored in the L2BGreen database and aggregated for analytics.

The detailed view for each case activity reveals the raw data source: a "Conversations" table logging every interaction between a learner and Gaia. This table includes learner identifiers, email addresses, and message previews, allowing educators to drill down into individual interactions when needed.

4.3.2 Key Metrics and Visualization

The dashboard offers several interactive views to help educators understand learner progress, organized in a hierarchy from cohort-level overview to individual learner detail:

Cohort Overview: A high-level view displaying key metrics for the entire class or cohort, such as overall completion rates, average time spent on activities, distribution of engagement scores, and identification of activities with high dropout rates. This view enables educators to quickly assess the health of their course delivery.

Case-Specific Analysis: Educators can drill down into a specific case study to see detailed performance metrics. This view includes an "About this activity" section summarizing the activity's purpose and a list of the specific "Learning Objectives" for the case. This allows educators to directly connect learner performance to the pedagogical goals defined in D3.1 Learning materials. The view also features analytics derived from the raw conversation logs, showing patterns in learner responses and common areas of difficulty.

Learner-Level Deep Dive: For targeted intervention, educators can view a detailed report for any individual learner. This includes their full conversation transcript with Gaia, time spent on each phase, performance on comprehension questions, and quality assessments of their argumentation responses. This level of detail supports personalized feedback and identification of learners who may need additional support.

4.3.3 Analytics Metrics Implementation

The implementation includes the calculation and display of the metrics outlined in D4.1 Table 2 (Educator Analytics Dashboard Metrics), moving from simple description to prediction and prescription. These metrics empower educators to make data-informed pedagogical decisions:

Engagement Score (Descriptive): A composite score based on time-on-task, number of interactions, and response length. This metric provides a quantitative measure of how actively learners are engaging with the material, with data sourced from Participant Action Data. Higher engagement scores correlate with deeper interaction with the learning content.

Argumentation Depth (Diagnostic): An internal classification performed based on the input of the Evaluator agent during the Argumentation phase. This metric assesses the quality of learner arguments, considering use of evidence from the case study, logical structure of arguments, acknowledgment of counterarguments, and sophistication of reasoning. The data is derived from Participant State Data captured during the Evaluator-Optimizer workflow.

Misconception Flags (Diagnostic): This metric identifies recurring incorrect answers to specific knowledge questions across the cohort. By aggregating patterns of incorrect responses from Participant State Data, educators can identify concepts that are commonly misunderstood and may require additional explanation or alternative teaching approaches. For example, if multiple learners incorrectly identify the marketing mistake in the ReNuTex case, this flags a potential gap in understanding of greenwashing concepts.

Dropout Risk (Predictive): A predictive score based on low engagement, poor performance, and long periods of inactivity. This metric combines data from all sources to identify learners who may be at risk of not completing the course. Early identification enables proactive intervention before learners disengage completely.

Recommended Intervention (Prescriptive): AI-generated suggestions for the educator based on analysis of learner data patterns. Examples include "Post a forum clarification on 'Triple Bottom Line' concept" when misconception flags indicate widespread confusion, or "Reach out to [learner] who shows declining engagement" when dropout risk increases. These prescriptive recommendations transform data into actionable next steps for educators.

5. Moodle LMS Integration and Configuration

The final critical piece of the implementation is the seamless and secure integration with the Moodle LMS. This integration ensures that the powerful L2BGreen tools are easily accessible within the environments where educators and learners already work, eliminating the need for users to navigate to separate systems and maintaining a cohesive learning experience.

5.1. LTI 1.3 Protocol Implementation

The integration is based on the Learning Tools Interoperability (LTI) 1.3 standard, as specified in D4.1. In this relationship, the L2BGreen platform functions as the LTI Tool Provider, and the Moodle instance acts as the LTI Tool Consumer. LTI 1.3 is the latest version of the standard, providing a significant security upgrade over previous versions by using the industry-standard OAuth2 and JSON Web Tokens (JWT) for authentication and secure message exchange.

5.1.1 Security Features

The LTI 1.3 implementation incorporates several security features that protect user data and ensure secure communication:

- **Public Key Cryptography:** Launch requests are signed using RSA keys, allowing the L2BGreen platform to verify that requests genuinely originate from the configured Moodle instance. Key pairs are exchanged during the initial tool registration process.
- **OAuth2 Authentication:** The authentication flow follows the OAuth2 standard, ensuring that user credentials are never exposed to the external tool. Instead, short-lived tokens are used to authorize specific actions.
- **JSON Web Tokens (JWT):** Launch data is encoded in JWT format, providing a secure and standardized method for transmitting user context, course information, and activity parameters.

5.1.2 User Authentication and Context Exchange

When a learner clicks a Gaia activity link within their Moodle course, Moodle constructs and sends a cryptographically signed LTI launch request to the L2BGreen platform. This request securely transmits:

- The user's pseudonymized ID (not their actual username or personal information)
- Their role in the course (e.g., learner, instructor, teaching assistant)
- The context of the course and specific activity being launched
- Return URLs for grade passback and navigation

The L2BGreen platform validates the signed request, authenticates the user based on their pseudonymized ID, and launches the appropriate Gaia chat interface for that specific learner and

scenario. The external platform is embedded directly within the Moodle course page via an iframe, appearing to the user as a native Moodle activity.

5.1.3 Automatic Grade Return

A critical feature for educators is the automatic return of grades via LTI Assignment and Grade Services (AGS). Upon a learner's completion of a Gaia-powered activity, the L2BGreen platform utilizes the AGS component to securely and programmatically send a completion status or a calculated score back to the appropriate column in the Moodle gradebook. This eliminates the need for manual grade entry and ensures a fully integrated workflow for educators. The grade return can include:

- A simple completion status (complete/incomplete)
- A calculated score based on performance in comprehension questions
- A composite score incorporating engagement and argumentation quality metrics.

5.2. Moodle Course and Activity Design

The Moodle course environment has been structured to provide a clear and cohesive learning journey that blends native Moodle resources with the innovative Gaia agent activities. This blended approach leverages the strengths of both platforms, using Moodle for content organization, forums, and traditional assessments while using Gaia for interactive, AI-powered learning experiences.

5.2.1 Course Structure Template

The "ReNuTex: Navigating Sustainability, Strategy, and Greenwashing in the Textile Industry" course serves as the design template for L2BGreen courses. The course is organized into logical, collapsible sections that guide the learner sequentially through the material:

1. General Section: Contains course announcements and a General Information PDF providing an overview of the course structure and expectations.

2. Lesson Introduction & Objectives: Introduces learners to the case-based learning approach with a welcome page ("Welcome to ReNuTex: Navigating Sustainable Business") that sets expectations and outlines learning objectives.

3. Part 1: Introducing the Business: A dedicated section introduces the business case using native Moodle resources. This includes a PDF or page resource ("Part 1: ReNuTex – The Business of Sustainable Textiles") providing background reading before the interactive activity.

4. Gaia Agent Activity: The core interactive component is presented as a single, clearly labeled "Gaia Agent Activity" link. This is the LTI external tool link that launches the conversational experience within the L2BGreen platform, embedded within the Moodle page.

5. Extra Courses Material: Following the agent activity, the course provides supplementary materials in the form of PDFs ("Extra material 1," "Extra material 2"), allowing learners to deepen their understanding of the concepts covered in the case.

6. ReNuTex - Strategic Application & Ethical Debate: The course concludes with further activities designed to foster peer discussion and reflection. This can be a Moodle Forum where learners discuss application of concepts, debate ethical dimensions of sustainability marketing, or share additional examples they have encountered.




Figure 12: The ReNuTex course structure in Moodle LMS, showing the full learning journey from case introduction through Gaia Agent Activity to supplementary materials and discussion forum

5.2.2 LTI Tool Configuration

The L2BGreen external tool is configured within Moodle by a site administrator, utilizing the site administration panel. This one-time setup involves:

- 1. Tool Registration:** Registering the L2BGreen platform as an external LTI tool, providing the tool URL, client ID, and deployment ID.
- 2. Key Exchange:** Configuring the public key URL for JWT signature verification or manually entering the public key.

3. Service Configuration: Enabling the LTI Assignment and Grade Services for grade passback.

4. Privacy Settings: Configuring privacy settings to ensure only the minimum necessary user data is shared during the LTI launch, in accordance with GDPR principles.

For educators, adding a Gaia activity to their course is as simple as adding any other Moodle activity and selecting "Gaia Agent Activity" from the list of available external tools. The familiar Moodle workflow makes adoption straightforward for educators already comfortable with the LMS.




Figure 13: The Moodle Site Administration panel, used for configuring language packs, LTI external tool registration, and AI provider settings for the L2BGreen platform

5.3. Multilingual User Interface Implementation

As detailed in D4.3 Multilingualism Support, the entire L2BGreen ecosystem, including the Moodle LMS interface, has been configured to support the six official project languages: English, German, Dutch, Greek, Slovenian, and Estonian. The implementation follows the two-pillar approach outlined in D4.3 to ensure comprehensive multilingual functionality.

5.3.1 Static Content Translation

All static UI elements have been professionally translated and are managed using a standard internationalization (i18n) framework:

- **Moodle Interface:** The Moodle installation includes language packs for all six project languages. The user-facing language selection menu provides a straightforward way for learners to choose their preferred language for the platform's user interface. Menu items, navigation elements, and system messages all appear in the selected language.
- **L2BGreen Platform:** The Educator and Administration Platform uses JSON-based locale files for internationalization. All UI strings—button labels ("Add a new case activity," "Save," "Delete"), menu items, form labels, and system notifications—have been professionally translated.
- **Gaia Interface:** The learner-facing Gaia chat interface includes translated UI elements, including phase labels ("PART 1," "PART 2"), button texts, and system messages.

5.3.2 Dynamic Interaction Translation

For real-time conversational interactions with Gaia, the 'Translate-Process-Translate' pipeline has been implemented as described in D4.3. This pipeline operates as follows:

- 1. Input Translation:** When a learner submits input in their native language, the text is first translated to English using the Translation Service Layer.
- 2. Processing:** Gaia processes the translated input and generates a response in English. The use of English as the processing language ensures pedagogical consistency, as the AI agents and their prompts have been optimized for English.
- 3. Output Translation:** The English response is translated back to the learner's selected language before being displayed in the chat interface.

This approach localizes the complexity of translation while maintaining the quality of AI-generated pedagogical content. The Translation Service Layer integrates with machine translation APIs while implementing quality checks to ensure translations are appropriate for the educational context.

The Moodle LMS dashboard and the L2BGreen Administration Platform have both been successfully rendered in all six project languages, demonstrating the successful implementation of static content translation and a responsive design that accommodates diverse text lengths. The interface adapts to different text lengths common across languages, ensuring that German compound words and Greek diacritics display correctly.




Figure 14: The Moodle LMS language selector displaying all six supported project languages: English, Deutsch, eesti, Nederlands, Slovenščina, and Ελληνικά




Figure 15: The Moodle LMS dashboard rendered in Estonian (eesti), demonstrating complete static UI translation across the platform

6. Conclusion

This document, Deliverable 4.2 "Conversational agent implementation," has provided a comprehensive account of the technical realization of the L2BGreen learning ecosystem. The implementation documented herein represents a significant milestone for the project, successfully translating the detailed design specifications of D4.1 into a fully functional, integrated, and deployed platform ready to serve the project's target audience of higher education students, vocational trainees, and working professionals across Europe.

6.1. Summary of Implemented Components

The core innovations of the L2BGreen project have been brought to life through this implementation. Each component contributes to the overall vision of fostering a green entrepreneurial mindset across Europe through scalable, digital learning.

The Gaia Conversational Agent: This component has been implemented as a sophisticated, hierarchical multi-agent system that stands as the pedagogical heart of the L2BGreen system. The Orchestrator Agent coordinates the flow of learning while the four specialized worker agents—Scenario Presenter, Knowledge Scaffolder, Application Coach, and Argumentation Feedback Agent—work in concert to deliver a structured, pedagogically sound learning experience aligned with the phases defined in D2.1. The implementation of advanced techniques, including Chain of Thought (CoT) prompting and the Evaluator-Optimizer workflow, ensures that the feedback provided to learners is not only accurate but also transparent, self-corrected, and aligned with the learning objectives. The Socratic tutoring capabilities of the Knowledge Scaffolder and the structured argumentation analysis using the Toulmin model demonstrate the sophisticated pedagogical approach embedded in the AI system.

The Educator and Administration Platform: This component has been developed as the essential control panel for the L2BGreen ecosystem, designed for usability by non-technical users and scalability for institutional deployment. The implementation delivers intuitive tools for scenario orchestration and content management, enabling educators to create and deploy AI-powered learning activities through a simple workflow. The robust user and team management system with Role-Based Access Control (RBAC) ensures appropriate access levels within educational institutions. The advanced learning analytics dashboard represents a key innovation, providing metrics such as Engagement Score, Argumentation Depth, Misconception Flags, Dropout Risk, and Recommended Interventions. These metrics empower educators to make data-informed pedagogical decisions, moving from passive data collection to active intervention and continuous improvement of teaching and learning.

The Moodle LMS Integration: This component has been seamlessly achieved through the implementation of the LTI 1.3 standard, ensuring easy accessibility within existing educational environments. The L2BGreen platform acts as an LTI Tool Provider, allowing Moodle to securely launch Gaia activities and exchange essential user context. The integration includes automatic return of grades via LTI Assignment and Grade Services (AGS), streamlining the workflow for educators. The Moodle course environment has been structured to blend native Moodle resources with Gaia agent activities, providing a cohesive learning journey. This strategic choice of LTI 1.3 integration maximizes the project's potential for future scalability and adoption across diverse educational landscapes, avoiding vendor lock-in and promoting long-term sustainability.

Comprehensive Multilingual Support: This component has been integrated across the entire system, as detailed in D4.3. The implementation of the Translation Service Layer and the 'Translate-Process-Translate' pipeline enables learners to interact with Gaia in any of the six official project languages: English, German, Dutch, Greek, Slovenian, and Estonian. Both static UI translation through i18n frameworks and dynamic conversational translation ensure a fully localized experience while maintaining pedagogical quality through English-based AI processing.

6.2. Alignment with Project Objectives

The implementation directly contributes to the overarching objectives of the L2BGreen project as established in the project proposal and implementation plan (D1.1). By creating a functional, AI-enhanced, multilingual platform, this deliverable provides the technological foundation necessary to achieve the project's central vision: "to leverage scalable, digital ways of learning as the means to move forward the green transition through increased awareness of and knowledge on green entrepreneurship."

The implemented system addresses the project's key objectives:

- **Scalability:** The architecture separating the Learning Hub (Moodle), Interaction Engine (L2BGreen Platform), and Secure Bridge (LTI 1.3) enables the platform to scale to serve large numbers of learners while maintaining performance and quality.
- **Accessibility:** The use of open standards, responsive design, and integration with existing LMS infrastructure ensures that the L2BGreen learning experience is accessible to the diverse European audience targeted by the project.
- **Pedagogical Effectiveness:** The sophisticated multi-agent architecture, advanced prompting techniques, and structured learning flow directly implement the instructional design from D2.1, ensuring that technology serves pedagogical goals rather than dictating them.
- **Data-Driven Improvement:** The implementation of comprehensive analytics enables continuous improvement of both the learning content and the AI agents based on actual learner performance and engagement data.

Privacy and Security: The implementation of pseudonymization, encryption, and GDPR-compliant data handling ensures that learner privacy is protected throughout the learning experience.

6.3. Next Steps and Future Work

With the completion of this implementation, the project is now positioned to proceed with the subsequent phases of Work Package 4 and the broader project plan. The functional platform documented herein will be the subject of D4.4 "System evaluation and public deployment," which represents the next critical milestone in the project's trajectory.

The next phase will involve:

- **User Testing:** Rigorous testing with target users including higher education students, vocational trainees, and working professionals during stakeholder engagement activities (WP5).

- **Feedback Integration:** Gathering systematic feedback from learners, educators, and institutional stakeholders to inform iterative improvements to both the AI agents and the user interfaces.
- **Performance Optimization:** Monitoring system performance under realistic usage conditions and optimizing for latency, reliability, and scalability.
- **Content Expansion:** Working with WP3 to expand the learning materials and case studies available through the platform.
- **Public Deployment:** Preparing the system for its full public launch, including documentation, training materials for educators, and support infrastructure.

In conclusion, the comprehensive implementation detailed in this deliverable lays a solid foundation for a scalable, engaging, and impactful learning experience. By strategically combining cutting-edge AI technology, intuitive educator tools, and seamless LMS integration, the L2BGreen ecosystem is poised to effectively cultivate the next generation of green entrepreneurs across Europe, fulfilling the project's central vision of advancing the green transition through innovative digital learning.