

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Issued by: Credential Engine

Release Date: July 9th, 2026

Contacts:

Jeanne Kitchens, CTSO, Credential Engine jkitchens@credentialengine.org

Rohit Joy, Director of Engineering, Credential Engine rjoy@credentialengine.org

Overview	1
Background	2
CTDL xTRA	2
CTDL is SOLID Data	5
Credential Registry	6
RFP: Scaling CTDL xTRA	7
Collaborative Development Model	8
CTDL xTRA: Priorities and Areas of Interest	9
High Level Multi Year Direction	14
Timeline and How to Respond	15
Appendix: CTDL xTRA Screenshots	19

Overview

This RFP invites responses from consultants with expertise in AI/ML-assisted extraction systems, scalable platform architecture, semantic technologies, workflow orchestration, linked open data, and sustainable AI/ML implementation approaches to help scale the CTDL xTRA platform.

The purpose of this RFP is to identify and evaluate organizations with the technical expertise, implementation experience, and collaborative approach needed to help mature CTDL xTRA into a sustainable, scalable, configurable, and reliable platform capable of significantly increasing the availability of high quality Credential Transparency Description Language (CTDL) linked open data nationally and, over time, internationally.

This RFP does not constitute a procurement commitment or contractual offer. However, it is Credential Engine's intent to retain consulting services to support CTDL xTRA scaling, architecture, infrastructure, orchestration, and implementation activities associated with the goals described in this RFP.

The objective is to identify a collaborative implementation partner, not simply a provider of architectural recommendations. Respondents should be prepared to work directly with Credential Engine's Director of Engineering and assigned engineering staff to iteratively design,

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

implement, test, validate, and deploy production-ready capabilities. Recommended approaches should emphasize rapid implementation, measurable quality improvements, integration with existing Credential Registry publishing workflows, and long-term sustainability within Credential Engine's open source development model.

Background

Credential Engine (<https://credentialengine.org>) is a nonprofit organization dedicated to improving the transparency, comparability, and usefulness of credential and skills information through open standards, linked open data, and interoperable infrastructure.

Our organization maintains the CTDL, the Credential Registry, and a growing ecosystem of publishing tools and services that support the structured description and exchange of credential related information across education and workforce systems. CTDL enables the structured description of credentials, learning programs, courses, competencies, occupations, jobs, pathways, transfer information, quality assurance information, and related entities as linked open data. The Credential Registry provides the infrastructure that supports this information to be connected, discovered, compared, and reused across systems and applications.

Our mission includes reducing barriers to publishing high quality CTDL linked open data and supporting the large-scale availability of transparent, interoperable credential and skills information across education and workforce ecosystems.

CTDL xTRA

To scale the availability of CTDL data, Credential Engine partnered with Learning Tapestry to develop CTDL xTRA (eXtensible Extract and Transformation Assistant), an open source, LLM-assisted structured extraction and transformation platform designed to convert publicly available web content into CTDL-aligned structured data suitable for publishing to the Credential Registry.

Development of CTDL xTRA is ongoing as Credential Engine expands extraction capabilities, improves quality assurance and reproducibility, strengthens workflow orchestration, and advances the platform toward scalable national and international deployment.

CTDL xTRA is intended only for processing publicly available information and does not ingest or process learner records, personal information, or private institutional data.

Pilot and testing activities have been completed with selected partners, and the platform has demonstrated the ability to extract and transform publicly available information into publish-ready CTDL-aligned formats compatible with Credential Registry publishing workflows. Existing CTDL xTRA extraction workflows are functioning to varying degrees of maturity,

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

automation, and accuracy. The objective of this engagement is to improve and expand these existing capabilities rather than replace the current platform.

Additional information about the current CTDL xTRA platform, architecture, and ongoing development work is available through the public GitHub repository.

- CTDL xTRA GitHub Repository:
<https://github.com/CredentialEngine/ctdl-xtra>

The repository contains the current open source codebase, which is under active development. CTDL xTRA is an operational platform that has been successfully demonstrated through pilot implementations and supports working extraction workflows. However, it should not be considered production-grade software or a finalized production architecture. One objective of this engagement is to help mature the platform by improving the quality, accuracy, scalability, reproducibility, and maintainability of existing CTDL xTRA extraction workflows while expanding support to additional CTDL classes and semantic relationships identified in this RFP. All new capabilities should integrate cleanly with Credential Engine's production architecture, engineering practices, and publishing workflows.

LLM Integrations: CTDL xTRA currently uses an OpenAI API for LLM-assisted extraction and transformation workflows. Development is underway to support agentic workflows using Anthropic, configurable model provider orchestration, including the ability to switch between AI/LLM providers.

Current CTDL xTRA Coverage: CTDL xTRA currently supports extraction and transformation workflows for multiple CTDL related classes, including the following classes. These links reference the CTDL schema definitions rather than CTDL xTRA output examples.

(1) CTDL Class	(2) CTDL Handbook Section	(3) CTDL Schema
Course	CTDL Handbook Learning Opportunity Classes	ceterms:Course
Learning Program	CTDL Handbook Learning Opportunity Classes	ceterms:LearningProgram
Competency Framework and Competency	CTDL Handbook Competency Framework Classes CTDL Handbook Competency Classes	ceasn:Competency Framework ceasn:Competency
Credential	CTDL Handbook Credential Classes	ceterms:Credential

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

The CTDL xTRA platform currently supports:

- Public URL ingestion
- Extraction from public web pages
- IP rotation
- AI/LLM-assisted structured extraction, transformation, and linked data generation into CTDL-aligned structures
- Validation workflows
- CSV generation to a pre-defined publish-ready template for Credential Registry publishing
- Human review and publishing workflows

Additional current platform capabilities include:

- Storage using PostgreSQL and management of extraction data and workflow outputs
- User interfaces for reviewing extraction activity and status
- Visibility into extraction successes and failures
- Estimated ML and operational cost tracking
- Workflow monitoring and extraction reporting
- Auditability and review support for extraction workflows

Human review and validation workflows remain important components of the platform's quality assurance and linked data integrity strategy.

Transformed Output: CTDL xTRA currently extracts and generates CSV files based on predefined CTDL-aligned CSV templates supported by Credential Engine's Credential Registry publishing system. This workflow ensures that generated files are structured for immediate upload and publication readiness.

The Registry publishing infrastructure then transforms uploaded CSV data into CTDL JSON-LD for validation, linking, and publication into the Credential Registry.

This architecture allows CTDL xTRA to integrate with existing Credential Engine publishing workflows while supporting scalable transformation of publicly available information into structured linked open data.

CSV output examples:

(1) CTDL Class	(2) CTDL xTRA Output: Bulk Upload Template	(3) Publishing System: Bulk Upload Template (All Properties)
Course and Learning Program	 CTDL xTRA Course Extract...	 Course and Learning Progr...

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

	 CTDL xTRA Learning Progr...	
Competency Framework and Competency	 CTDL xTRA Comptency Fr...	 CompetencyFramework Bu...
Credential	 CTDL xTRA Credential Bul...	 Credential Bulk Upload Sa...

References:

- CTDL xTRA: <https://credentialengine.org/toolkit/credential-engine-ctdl-xtra-tool/>
- CTDL xTRA Sequence Diagram: [SequenceDiagram.org](https://sequence-diagram.org/)

CTDL is SOLID Data

Credential Engine advances the use of the Credential Transparency Description Language (CTDL) as the foundation for publishing structured, linked open data across education and workforce ecosystems.

CTDL is the only comprehensive open standard for describing and linking credentials, learning, competencies, pathways, occupations, quality assurance information, transfer information, and related learn and work ecosystem data.

Credential Engine defines this approach as SOLID data. To meet growing national and global demand for trustworthy education and workforce information, public data must be:

- Structured data uses standardized, machine-readable formats that support reliable interpretation, comparison, and reuse.
- Open data supports transparency, innovation, and broad reuse across commercial and non-commercial applications.
- Linked data connects information through persistent identifiers and semantic relationships so data can be discovered, combined, and understood across credentials, skills, jobs, pathways, and related entities.
- Interoperable data uses shared technical and semantic standards that allow reliable exchange and use across platforms, systems, and jurisdictions.
- Durable data uses persistent identifiers, governance, and long-term infrastructure to support stability, trust, and referenceability over time.

By using CTDL, information published through Credential Engine infrastructure becomes SOLID data that can be connected, reused, and scaled across education, workforce, policy, research, and technology systems.

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

CTDL continues to evolve as new data types, use cases, and interoperability needs emerge. Current priorities include scaling the infrastructure, tooling, workflows, and services needed to support large-scale publishing, linking, validation, discovery, interoperability, and AI/ML-assisted use of CTDL linked open data.

References:

- Credential Engine Technical Site: <https://credreg.net>
- CTDL Handbook: <https://credreg.net/ctdl/handbook>
- SOLID Data Overview: <https://credentialengine.org/resources/solid-data-what-it-means/>

Credential Registry

The Credential Registry is Credential Engine's linked open data registry for publishing, managing, connecting, and accessing CTDL data across education and workforce ecosystems.

Organizations use the Registry to publish structured CTDL linked open data describing credentials, learning programs, courses, competencies, transfer information, quality assurance information, pathways, occupations, and related entities.

The platform supports persistent identifiers such as CTIDs, semantic relationships, linked data interoperability, and large-scale discovery and reuse of CTDL data across systems and applications.

Credential Engine maintains and operates the Registry infrastructure, including publishing workflows, APIs, linked data services, and data consuming pipelines. The Registry also supports multiple data consuming options and services that allow organizations, applications, researchers, policymakers, and technology providers to access and reuse CTDL linked open data.

CTDL xTRA is intended to support and scale the publication of high quality CTDL linked open data by reducing the technical and operational barriers associated with transforming publicly available information into structured and linked CTDL data.

CTDL xTRA is expected to become increasingly integrated into Credential Engine's publishing infrastructure and data pipelines to support scalable extraction, transformation, validation, and publication workflows.

References:

- Credential Registry Overview: <https://credentialengine.org/credential-registry/>
- About the CTID: <https://credreg.net/ctdl/ctid>
- Credential Finder: <https://credentialfinder.org>
- Credential Engine Technical Site: <https://credreg.net>

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

RFP: Scaling CTDL xTRA

Credential Engine is seeking consultant support to accelerate the continued evolution of CTDL xTRA into a scalable, production-ready platform that supports Credential Engine's long-term architecture and large-scale CTDL linked open data publishing.

Responses should address:

- Architectural recommendations
- Implementation strategies
- Operational approaches
- Workflow designs
- Sustainability models
- Phased scaling recommendations
- Infrastructure and orchestration approaches
- AI/LLM implementation recommendations
- Testing, benchmarking, validation, and quality measurement approaches
- Semantic interoperability and linked data management strategies
- Strategy for achieving production-ready capabilities within the first six months

Multiple technical and implementation approaches may satisfy the objectives of this RFP. Respondents are encouraged to recommend the approach they believe best balances quality, scalability, maintainability, and cost.

Credential Engine expects all future platform enhancements, workflows, infrastructure improvements, and related technical implementations developed through this effort to remain aligned with Credential Engine's open source development approach and GitHub-based collaboration workflows. The IP will be owned by Credential Engine, code will be made available via our GitHub repository under an Apache License for commercial and non-commercial use.

Foundational platform objectives include:

- Reusable and maintainable open source architecture
- Transparent and well documented development practices
- Modular and configurable workflows
- Sustainable long-term maintainability
- Community-oriented contribution model
- GitHub based collaboration and version management
- The engineering environment utilizes technologies including:
 - Visual Studio
 - .NET
 - Python
 - JavaScript

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

- Microsoft Azure
- Azure Blob Storage
- PostgreSQL
- Kubernetes
- Argo Workflows
- Keycloak
- Reduction of unnecessary vendor lock-in and proprietary dependency constraints

Collaborative Development Model

Credential Engine intends this engagement to augment, not replace, its internal engineering capabilities. Credential Engine's Director of Engineering and engineering team will retain responsibility for platform architecture, application infrastructure, authentication, publishing workflow integration, deployment, and ongoing platform development. Credential Engine will assign a full-time software engineer to work directly with the selected consultant throughout the engagement and to lead implementation of platform components, workflow integration, and publishing system integration.

The selected consultant is expected to provide specialized expertise in AI/LLM technologies, including AI/ML-assisted extraction, LLM integration, model evaluation, workflow optimization, testing methodologies, and implementation of agreed-upon AI/LLM enhancements. The consultant's primary responsibility is the design and implementation of AI/LLM enhancements that improve extraction quality, accuracy, scalability, reproducibility, and operational efficiency. Credential Engine engineering staff will implement or co-implement the platform integrations required to incorporate these AI/LLM capabilities into the production environment. New capabilities should be designed for iterative implementation, testing, and deployment through Credential Engine's existing development, test, and production workflows.

All AI/LLM capabilities developed through this engagement must be implemented as integrated components of the existing CTDL xTRA platform and must work seamlessly with Credential Engine's platform architecture, authentication infrastructure, publishing workflows, operational environment, and open source development practices. Standalone or isolated AI/LLM solutions are not within the scope of this engagement.

Implementation responsibilities will be shared throughout the engagement. The consultant is expected to work closely with Credential Engine's engineering team through collaborative design, implementation, code reviews, testing, validation, documentation, and knowledge transfer to ensure that ongoing ownership, maintenance, and future platform evolution remain with Credential Engine.

This engagement requires specialized AI/ML expertise that complements Credential Engine's existing engineering capabilities rather than a fully outsourced software development team. Respondents should propose lean, collaborative teams that maximize knowledge transfer,

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

leverage Credential Engine's engineering resources, and deliver practical, production-ready improvements through an incremental implementation approach.

This RFP is intended to identify qualified implementation partners and understand their technical approaches. Credential Engine anticipates an initial implementation engagement focused on the highest-priority AI/LLM enhancements, with subsequent phases to be prioritized collaboratively based on results, organizational priorities, and available resources. Selected respondents will be invited to participate in follow-up technical discussions to refine scope, implementation priorities, division of responsibilities, and estimated level of effort prior to development of a detailed implementation proposal.

Successful outcomes from this engagement include measurable improvements in extraction quality, scalability, operational efficiency, and integration with Credential Engine's publishing workflows while maintaining a sustainable, open source platform that Credential Engine can continue to evolve independently.

CTDL xTRA: Priorities and Areas of Interest

The following technical areas represent the highest implementation priorities for this engagement. Respondents should prioritize recommendations that can be implemented rapidly, integrate with Credential Engine's existing open source platform and Credential Registry publishing workflows, and support iterative development alongside Credential Engine's Director of Engineering and engineering team.

Sustainable AI/LLM Operations

- Sustainable and cost-conscious AI/LLM model usage
- Reduced token usage and extraction waste
- Cost optimization strategies for high-volume extraction workflows
- Reproducible and auditable extraction workflows
- Deterministic validation and quality assurance controls
- Approaches that maximize extraction consistency and minimize hallucinations
- Hybrid extraction architectures that combine deterministic methods with AI/LLM-assisted extraction
- Structured extraction benchmarking and evaluation methodologies
- Precision, recall, extraction quality, and semantic accuracy measurement approaches
- Operational transparency, monitoring, and extraction auditability
- Production-ready implementations that support rapid deployment, iterative testing, and continuous quality improvement

Extraction and Workflow Improvements

- Workflow orchestration and modular extraction pipelines

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

- Public web page extraction and transformation
- Public PDF ingestion, OCR, and static document extraction
- Improved handling of dynamic content, lazy loading, and pagination
- Scalable extraction management, monitoring, and scheduling
- Human-in-the-loop review and validation workflows
- Automated re-extraction and data currency workflows
- Reusable and configurable extraction workflows
- Validation and quality assurance throughout the extraction lifecycle
- Improved extraction reporting, export capabilities, and operational visibility
- Integration with Credential Registry publishing workflows to support efficient production publishing

Platform Architecture and Infrastructure

Recommendations should assume that Credential Engine engineering staff will lead platform infrastructure implementation.

- Reusable and maintainable open source architecture
- Modular and configurable platform design
- Architecture that supports future multi-tenant deployment
- Deployment portability across major cloud platforms
- Integration with Credential Engine authentication, publishing workflows, and Credential Registry infrastructure
- Logging, monitoring, auditability, and operational reporting
- Role-based access control and organizational management
- Sustainable long-term maintainability
- Compatibility with Credential Engine's open source GitHub development workflows and Apache-licensed codebase
- Architectural approaches that support rapid implementation, iterative development, automated testing, and continuous deployment

Extraction Configuration and User Experience

- Intuitive extraction configuration and workflow design
- Visual tools for configuring, testing, and managing extraction workflows
- Approaches that reduce or eliminate manual extraction configuration where appropriate
- Automated identification of extraction patterns, document structures, and semantic relationships
- Comprehensive testing, benchmarking, and validation workflows
- Regression testing approaches that ensure extraction quality over time
- Reusable extraction configurations, templates, or equivalent approaches
- Improved documentation, onboarding, and developer experience

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

- Better visibility into extraction logic, workflow behavior, AI decision-making, and operational performance
- Improved information architecture and overall user experience

AI/ML and Model Flexibility

CTDL xTRA currently uses OpenAI APIs for AI/ML-assisted extraction and transformation workflows. Credential Engine is expanding the platform to support configurable model orchestration, including commercial APIs, self-hosted and open-weight models, and future AI/ML providers.

Priority architectural objectives for AI/ML and model flexibility include:

- Flexible model orchestration and provider abstraction
- Evaluation of commercial, open-weight, open-source, and custom AI/ML approaches
- Self-hosted AI/ML deployment where appropriate
- Support for specialized AI/ML and NLP models
- Appropriate use of smaller language models to improve cost, performance, reproducibility, and sustainability
- Provider independence and reduced vendor lock-in
- Comparative evaluation of model quality, performance, cost, and operational characteristics
- Repeatable testing and evaluation methodologies for comparing AI/ML approaches against representative extraction workloads
- Architecture that enables new model providers to be incorporated with minimal changes to extraction workflows
- Long-term operational sustainability and technology evolution

Respondents are encouraged to recommend implementation approaches that balance AI/ML-assisted extraction with deterministic workflows, validation controls, semantic interoperability, linked data integrity, and long-term maintainability. Phased implementation recommendations that allow Credential Engine to evaluate and adopt evolving AI technologies while preserving platform stability are encouraged.

Internationalization and Future Expansion

Priority areas for internationalization and future expansion include:

- Internationalization and multilingual extraction support
- Localized extraction workflows
- Expansion beyond English language environments
- Future support for international qualification and credential systems
- Long-term global interoperability and publishing support

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Credential Engine recognizes that respondents may recommend different implementation sequences, architectural approaches, workflow strategies, and prioritization models to support these objectives.

Linked Data and CTDL Expansion

- Expansion to additional CTDL classes and linked data relationships
- Semantic interoperability across CTDL classes
- Linking across credentials, courses, competencies, learning programs, assessments, support services, outcomes, pathways, occupations, and jobs
- Relationship validation and integrity checking
- Persistent identifiers and durable references
- Expansion of CTDL property and term coverage
- Scalable linked data generation and maintenance

The existing CTDL xTRA workflows described earlier in this document are expected to continue evolving throughout this engagement. The additional CTDL classes listed below represent priority areas for functional expansion rather than replacements for the current extraction capabilities.

Priority CTDL classes identified for expansion include the following. These represent the initial priority areas for functional expansion based on current organizational priorities and the availability of publicly available information suitable for extraction. Final implementation sequencing, prioritization, and selection of additional CTDL classes will be determined collaboratively by the CTDL xTRA Product Owner, the project team, and the selected consultant based on technical dependencies, implementation complexity, public data availability, organizational priorities, project outcomes, and emerging opportunities.

Because public sources are unlikely to include all or most CTDL properties, recommended approaches should emphasize the ability to identify and extract relevant CTDL properties when they are present in unstructured public content, rather than assuming complete property coverage for every resource.

CTDL Class	CTDL Handbook	Example Bulk Upload Templates
Support Service	CTDL Handbook: Support Service	 SupportService Bulk U...

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

CTDL Class	CTDL Handbook	Example Bulk Upload Templates
Occupation	CTDL Handbook: Occupation	 Occupation Bulk Uplo...
Job	CTDL Handbook: Job	 Job Bulk Upload Sample
Metric and Data Set Profile	CTDL Handbook: Quantitative Data	 Metric Bulk Upload Sa...
Action	CTDL Handbook: QA Actions	 Action Bulk Upload Sa...
Assessment Profile	CTDL Handbook: Assessment Profile	 Assessment Bulk Uplo...

- Extraction of multiple CTDL entity classes from shared public sources
- Identification and preservation of semantic relationships between entities
- Relationship validation and integrity checking
- Reusable linking and entity resolution workflows
- Persistent identifiers and durable references
- Cross-entity orchestration and workflow management
- Scalable linked data generation and maintenance
- Relationship-aware update and re-extraction workflows
- Semantic interoperability across CTDL classes
- High-quality linked open data publishing rather than disconnected record generation

Examples of important linked relationships may include:

- Courses linked to competencies
- Learning programs linked to courses
- Credentials linked to competencies
- Credentials linked to learning programs
- Assessments linked to competencies
- Transfer information linked to courses and credentials
- Outcomes linked to credentials or programs
- Pathways linked across credentials, occupations, jobs, and learning opportunities

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Linked data relationships should be treated as foundational infrastructure requirements rather than secondary enrichment processes.

High Level Multi Year Direction

Credential Engine welcomes phased implementation recommendations and iterative scaling approaches aligned with the following high level direction. Credential Engine maintains governance authority over CTDL, the Credential Registry, and related publishing infrastructure.

Year 1 - Primary focus areas include:

- Expanded production-ready extraction capabilities within the first six months
- Significant CTDL coverage expansion within the United States
- Continued focus on extraction from publicly available web URLs
- Expansion from public web URL extraction into public PDF ingestion, OCR, and static document extraction
- Continued improvement of existing CTDL extraction workflows and expansion of CTDL class coverage across credentials, learning programs, courses, competencies, occupations, jobs, assessments, support services, and related linked entities.
- Expansion of semantic linking across CTDL entities and relationships
- Workflow orchestration and scalable extraction infrastructure
- Data currency, automated re-extraction, and update cycle workflows
- Quality assurance, benchmarking, and reproducibility workflows
- Sustainable and cost-conscious AI/LLM model usage
- Integration into Credential Engine publishing infrastructure and publishing pipelines
- Platform integration of AI/LLM capabilities with Credential Engine infrastructure
- Foundational self-service and partner enablement capabilities
- Logging, auditability, operational reporting, and workflow monitoring
- Integration with Credential Engine authentication and publishing infrastructure
- Support for deployment portability across major cloud platforms
- Initial multilingual and international extraction experimentation
- Open source maintainability, modularity, and workflow configurability

Year 2 - Primary focus areas include:

- Expanded production capabilities
- Additional CTDL class coverage and linked data expansion
- Multilingual extraction support
- Self publishing and partner-oriented workflows
- Broader deployment across U.S. states and statewide transparency initiatives
- Increased automation and operational scalability
- Expanded international extraction pilots and testing

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Year 3 - Primary focus areas include:

- Large-scale national deployment and publishing support
- Expanded international and multilingual deployment capabilities
- Sustained high volume linked open data publishing
- Matured semantic linking and interoperability workflows
- Long-term operational sustainability and cost optimization
- Expanded global interoperability and publishing support
- Mature monitoring, governance, and update cycle workflows

Timeline and How to Respond

Credential Engine requests responses to this RFP by August 14th, 2026.

Responses and questions should be submitted electronically in Microsoft Word, Google Doc, or PDF format to: Jeanne Kitchens, CTSO, Credential Engine jkitchens@credentialengine.org

To support transparency and consistency, Credential Engine will respond to questions via email and post questions and responses on the CTDL xTRA RFP.

Response Template

Responses should be concise, focused, and easy to read. Recommended word counts are provided as general guidance rather than minimum or maximum requirements. Clear, well-organized responses are preferred over lengthy narratives.

Respondents are encouraged to provide practical recommendations, implementation approaches, examples of relevant work, and realistic phased strategies aligned with the goals described in this RFP.

Responses do not need to be highly formal or procurement-oriented. Credential Engine is primarily interested in identifying organizations with relevant expertise, scalable approaches, and strong alignment with the long-term goals of the CTDL xTRA platform.

Because this engagement is expected to involve ongoing collaboration with Credential Engine leadership and engineering staff, respondents should be able to support regular collaboration during U.S. business hours.

Responses from organizations located both within and outside the United States are welcome. However, preference may be given to respondents who can demonstrate strong operational alignment with U.S.-based collaboration, communication, and implementation workflows.

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Following review of submitted responses, Credential Engine expects to invite selected respondents to participate in one or more technical discussions focused on architecture, implementation approaches, collaboration, and solution fit. These discussions will help refine project scope, implementation priorities, and the division of responsibilities before requesting a detailed implementation proposal or statement of work.

1. Organization Overview

Include the following:

- Organization legal name:
- Organization website:
- Organization Type (Tax Status):
- Primary contact information (Full name, position, email, phone number):
- Organization's relevant areas of expertise:
- Team size and structure proposed for this engagement:
- Relevant technical capabilities:
- Geographic location(s) (Full Address):
- Relevant GitHub repository links if applicable:
- 1 year estimated cost:

2. Relevant Experience

Recommended length: 500 words

Please describe relevant experience related to:

- Experience transitioning AI/LLM solutions from pilot or proof-of-concept implementations into production software platforms.
- AI/ML-assisted extraction systems
- Linked open data and semantic technologies
- Experience integrating AI/LLM capabilities into existing production software platforms and engineering environments.
- Workflow orchestration
- AI/ML evaluation, benchmarking, or optimization
- Open source software development
- Education, workforce, credentialing, or interoperability ecosystems
- Large-scale public data infrastructure
- AI/ML-assisted document extraction, PDF extraction, or OCR workflows

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

3. Recommended Technical Approach

Recommended length: 500 words

Identify which recommendations you believe can be implemented immediately and which would benefit from additional collaborative discovery with Credential Engine.

- Scalable architecture and orchestration
- AI/LLM implementation strategies
- Approaches for improving existing CTDL xTRA extraction workflows while expanding support for the additional CTDL classes identified in this RFP.
- Sustainable and cost-conscious AI/LLM model usage
- Deterministic extraction and validation approaches
- Semantic interoperability and linked data management
- Integration of AI/LLM capabilities into existing production platforms and engineering workflows
- PDF and OCR extraction strategies
- Internationalization and multilingual support
- Long-term maintainability and operational sustainability
- Integration with publishing pipelines and authentication infrastructure

Credential Engine welcomes phased implementation recommendations and iterative scaling approaches.

4. Linked Data and Semantic Relationship Strategy

Recommended length: 500 words

Please describe approaches for:

- Extracting multiple CTDL entity classes from shared public sources
- Preserving semantic distinctions between entities
- Managing and validating linked relationships
- Entity resolution and persistent identifiers
- Relationship-aware update and re-extraction workflows
- Scalable linked open data publishing

5. Collaboration and Engagement Approach

Recommended length: 500 words

Please describe:

- Experience working collaboratively with internal engineering teams

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

- Open source collaboration practices
- Documentation and knowledge transfer approaches
- Code review, testing, and collaborative development practices
- Communication cadence and preferred engagement model
- Availability for collaboration during U.S. business hours

6. Implementation Approach

Recommended length: 500 words

Please include:

- Preliminary phased implementation approach
- Proposed technical leadership and key roles
- Recommended priorities for the initial implementation phase
- The first technical capabilities you would recommend implementing during the first 90 days to enable expanded production use of CTDL xTRA within the first six months, including the rationale for those priorities
- How you would establish confidence that these capabilities are ready for expanded production use, including your approach to testing, benchmarking, validation, and measuring success
- Key assumptions or dependencies
- Information needed before preparing a detailed statement of work or implementation proposal

7. References and Relevant Projects

Please include:

- At least two professional references for successful projects
- Links to relevant projects or implementations
- Relevant GitHub repositories where applicable
- Examples of comparable technical implementations
- Open source contributions if applicable

8. Optional Supporting Materials

Respondents may optionally include:

- Architecture diagrams
- White papers
- Technical documentation
- Relevant case studies

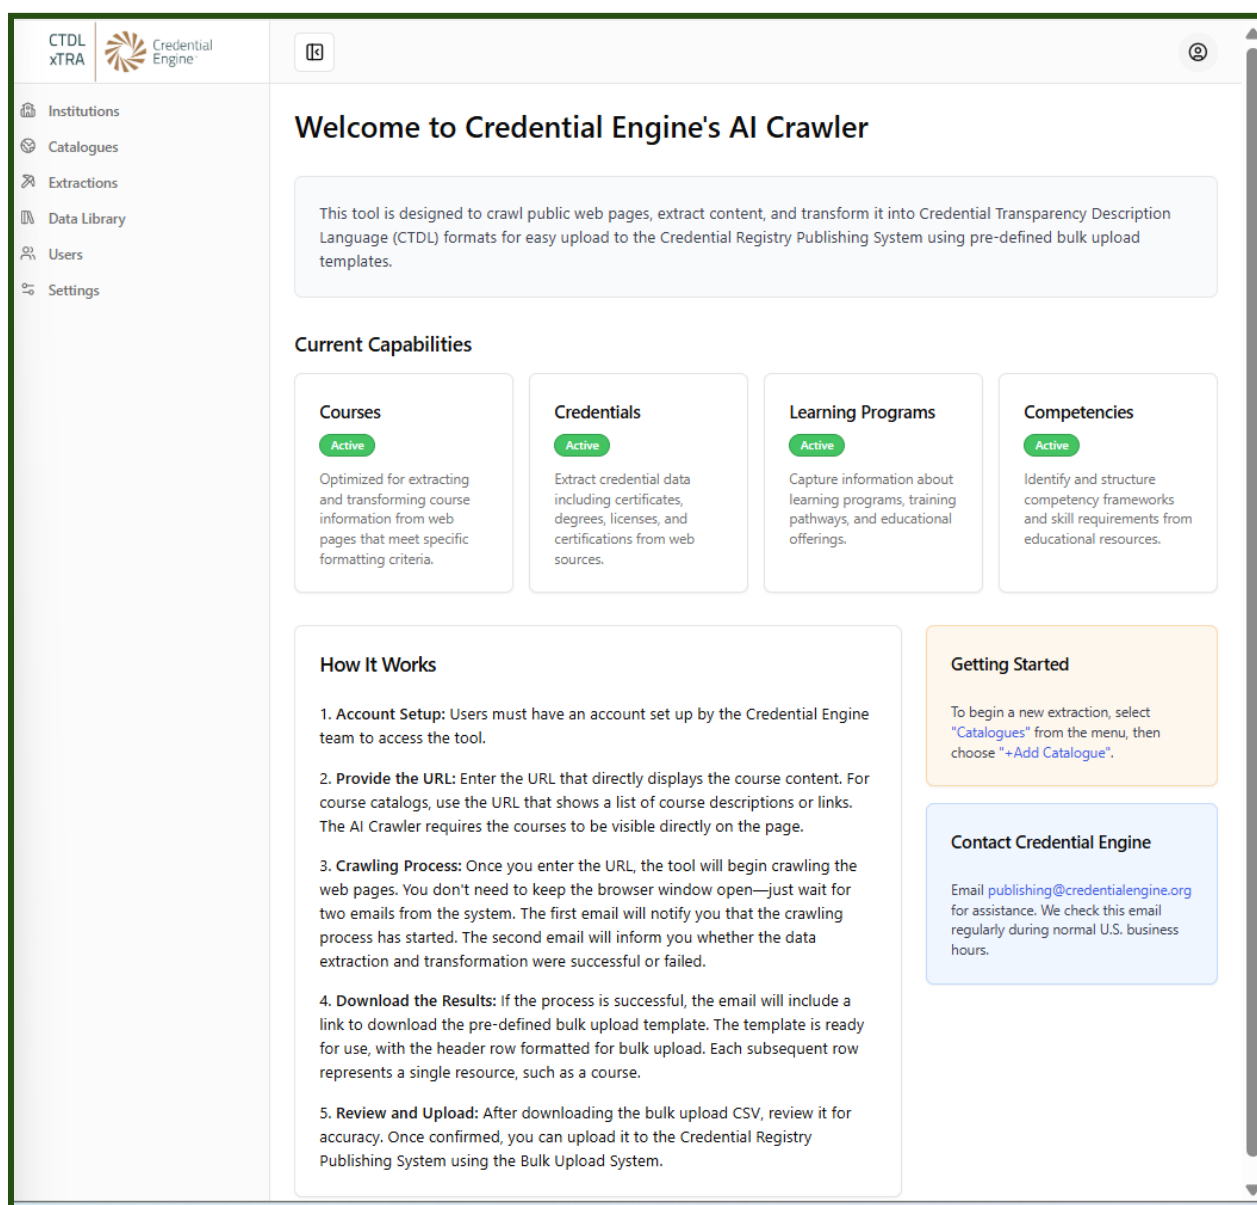
Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

- Open source project references
- Sample implementation approaches
- Pricing models or engagement structures

Credential Engine may conduct follow up discussions with selected respondents after reviewing submitted materials.

Appendix: CTDL xTRA Screenshots

Landing page



The screenshot displays the landing page of the CTDL xTRA platform. The page features a dark green header with the 'CTDL xTRA' and 'Credential Engine' logos on the left, a user profile icon on the right, and a central navigation menu. The main content area is white and includes a welcome message, a description of the tool's purpose, and a section titled 'Current Capabilities' with four cards: 'Courses', 'Credentials', 'Learning Programs', and 'Competencies'. Below this is a 'How It Works' section with five numbered steps, and a 'Getting Started' section with instructions on how to begin a new extraction. A 'Contact Credential Engine' section at the bottom right provides an email address for assistance.

CTDL xTRA Credential Engine

Welcome to Credential Engine's AI Crawler

This tool is designed to crawl public web pages, extract content, and transform it into Credential Transparency Description Language (CTDL) formats for easy upload to the Credential Registry Publishing System using pre-defined bulk upload templates.

Current Capabilities

- Courses** (Active)
Optimized for extracting and transforming course information from web pages that meet specific formatting criteria.
- Credentials** (Active)
Extract credential data including certificates, degrees, licenses, and certifications from web sources.
- Learning Programs** (Active)
Capture information about learning programs, training pathways, and educational offerings.
- Competencies** (Active)
Identify and structure competency frameworks and skill requirements from educational resources.

How It Works

- Account Setup:** Users must have an account set up by the Credential Engine team to access the tool.
- Provide the URL:** Enter the URL that directly displays the course content. For course catalogs, use the URL that shows a list of course descriptions or links. The AI Crawler requires the courses to be visible directly on the page.
- Crawling Process:** Once you enter the URL, the tool will begin crawling the web pages. You don't need to keep the browser window open—just wait for two emails from the system. The first email will notify you that the crawling process has started. The second email will inform you whether the data extraction and transformation were successful or failed.
- Download the Results:** If the process is successful, the email will include a link to download the pre-defined bulk upload template. The template is ready for use, with the header row formatted for bulk upload. Each subsequent row represents a single resource, such as a course.
- Review and Upload:** After downloading the bulk upload CSV, review it for accuracy. Once confirmed, you can upload it to the Credential Registry Publishing System using the Bulk Upload System.

Getting Started

To begin a new extraction, select "Catalogues" from the menu, then choose "+Add Catalogue".

Contact Credential Engine

Email publishing@credentialengine.org for assistance. We check this email regularly during normal U.S. business hours.

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Institutions Page

CTDL
xTRA



Credential
Engine™

Institutions

Catalogues

Extractions

Data Library

Users

Settings

Institutions

+

Add institution

Q

Search institutions by name or domain...

Institution	Domains	Catalogues
Union College of Union County (UCNJ)	ucc.edu	1 catalogues
Moreno Valley College	mvc.edu	3 catalogues
Hudson County Community College	hccc.edu	3 catalogues
Purdue University Global	catalog.purdueglobal.edu	1 catalogues
Bristol Community College	catalog.bristolcc.edu	3 catalogues
St. Cloud State University	catalog.stcloudstate.edu	1 catalogues
Colorado State University	catalog.colostate.edu	1 catalogues
Baruch College (CUNY)	baruch-undergraduate.catalog.cuny.edu	1 catalogues
Asheville-Buncombe Technical Community College	abtech.edu	2 catalogues
NorthWest Arkansas Community College	coursecatalog.nwacc.edu	1 catalogues
Lake Region State College	lrsc.edu	1 catalogues
Accelerate Education	accelerate.education	1 catalogues
United States University	usuniversity.smartcatalogiq.com	1 catalogues
University of Montana	programfinder.umd.edu	1 catalogues
Sinclair Community College	sinclair.edu	2 catalogues
Weber State University	continue.weber.edu catalog.weber.edu	2 catalogues

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Catalog Page

CTDL xTRA

Credential Engine™

Institutions

Catalogues

Extractions

Data Library

Users

Settings

Catalogues

+

Add catalogue

Q

Search catalogues by name or URL...

All catalogue types

▼

Catalogue	Type	Institution	
 Course Descriptions - Weber State University - Modern Campus Catalog™	COURSES	Weber State University	 Extract
 Brookdale Community College Catalog (Credit)	COURSES	Brookdale Community College	 Extract
 Taft College Credential Extract	CREDENTIALS	Taft College	 Extract
 Antelope Valley College - Credentials	CREDENTIALS	Antelope Valley College	 Extract
 Bergen's CE Catalog	COURSES	Bergen Community College	 Extract
 UCNJ NJ Continuing Education and Workforce Development Courses Catalog	COURSES	Union College of Union County (UCNJ)	 Extract
 El Camino - Credentials	CREDENTIALS	El Camino College	 Extract
 Norco - Credentials	CREDENTIALS	Norco College	 Extract
 Taft - Program Competencies	COMPETENCIES	Taft College	 Extract
 catalog.gfcmsu.edu	COURSES	Great Falls College MSU	 Extract
 eConnect - Fall 2024 - Credit Class Schedule	COURSES	Dallas College	 Extract
 Camden non-credit	COURSES	Camden County College	 Extract

Page 21

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Extractions Page

CTDL
xTRA

 Credential
Engine™

 Institutions

 Catalogues

 Extractions

 Data Library

 Users

 Settings

Extractions

Export

From

mm/dd/yyyy

To

mm/dd/yyyy

Catalogue	Recipe	Type	Status	Stats	Cost	Extracted Items	Date
Credentials	Recipe #399	CREDENTIALS	WAITING	Pending	-	0	06/11/2026, 02:35 AM
 City College of San Francisco - Competencies	Competencies (Programs)	COMPETENCIES	STALE	100% 100%	\$1.08	1,370	05/10/2026, 11:47 PM
Taft - Program Competencies	Recipe #398	COMPETENCIES	STALE	100% 84%	\$0.46	184	05/08/2026, 04:57 AM
Taft - Program Competencies	Recipe #398	COMPETENCIES	STALE	100% 76%	\$0.42	162	05/07/2026, 06:15 AM
Taft - Program Competencies	Recipe #398	COMPETENCIES	COMPLETE	0% 0%	\$0.00	0	05/06/2026, 02:10 AM
Norco - Credentials	Recipe #397	CREDENTIALS	COMPLETE	0% 0%	\$0.00	0	05/06/2026, 02:08 AM
 Norco College - Courses - 2025-26	Norco	COURSES	COMPLETE	0% 0%	\$0.00	0	05/04/2026, 11:49 AM
Moreno Valley Course Catalog	Moreno Valley College	COURSES	COMPLETE	0% 0%	\$0.00	0	05/04/2026, 11:23 AM
Norco - Credentials	Recipe #397	CREDENTIALS	COMPLETE	100% 100%	\$1.24	357	05/04/2026, 04:17 AM
Norco - Credentials	Recipe #397	CREDENTIALS	CANCELLED	0% 0%	\$0.00	0	05/04/2026, 04:12 AM
Moreno Valley College - Program Competencies	Recipe #396	COMPETENCIES	COMPLETE	0% 0%	\$0.00	0	04/22/2026, 12:47 AM
Moreno Valley College - Program Competencies	Recipe #396	COMPETENCIES	COMPLETE	0% 0%	\$0.00	0	04/22/2026, 12:25 AM

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Data Library Page

CTDL
xTRA

 Credential
Engine™

 Institutions

 Catalogues

 Extractions

 Data Library

 Users

 Settings

13



Data Library

Catalogue

Courses A-Z < University of Pennsylvania



Professional Development Courses for Educators - Michigan Virtual

Forsyth Tech Catalog



Career and Technical Education - Arkansas Department of CorrectionsAR DEPT CORRECTIONS_WHITEInfo IconWeb IconThumbs Up Icon



Course Descriptions | Texas A&M International University Academic Catalog



Courses < University of Wisconsin-Madison



HCC HOUSTON COMMUNITY COLLEGE Course Descriptions - Houston Community College - Modern Campus Catalog™



IVY TECH COMMUNITY COLLEGE Course Descriptions - Ivy Tech Community College - Acalog ACMS™



IVY TECH COMMUNITY COLLEGE Course Descriptions - Ivy Tech Community College - Modern Campus Catalog™2024-2025



COMMUNITY COLLEGE OF DENVER Courses < Community College of Denver

Course Catalog Grades 6-8 | Accelerate Education | Blended Learning

Course Descriptions - Vincennes University - Modern Campus Catalog™



CENTURY COLLEGE Course Descriptions and Outlines - Century College - Modern Campus Catalog™



Thomas Edison State University | Undergrad Course Offerings



Course Catalog - College of Lake County



B Baruch College Catalog

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Settings

CTDL
xTRA

 Credential
Engine[™]

Institutions

Catalogues

Extractions

Data Library

Users

Settings

Settings

OpenAI API Key

The API key that will be used for OpenAI requests.
The key is currently set to sk-...

New key

Update

Max Extraction Budget

Maximum USD budget allowed per extraction. (Current: \$150)

New budget (USD)

Update

Proxy Settings

Activating a proxy is helpful in overcoming rate limits or blocking by anti-scraping measures.
The proxy is currently **enabled**.

Disable Proxy

Proxy URLs

Add new proxy URLs or remove existing ones. Credentials are stored encrypted and never exposed to the UI.

Configured proxies:






Add proxy URL

Page 24

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Course Catalog Extraction Interface



The screenshot shows the 'Add new catalogue' interface within the CTDL xTRA platform. The interface is divided into a left sidebar and a main content area. The sidebar contains a list of navigation items: Institutions, Catalogues (highlighted), Extractions, Data Library, Users, and Settings. The main content area is titled 'Add new catalogue' and contains several form fields: a URL field with the placeholder 'http://www.example.com/course_catalogue', a Name field with the placeholder 'Catalogue Name', a Type dropdown menu set to 'Courses', and an Institution dropdown menu with the placeholder 'Select an institution'. Below the Institution dropdown, there is a note: 'Institutions that match the catalogue URL appear first. [Create a new institution](#) if you don't see yours.' A 'Submit' button is located at the bottom of the form.

CTDL xTRA | Credential Engine[™]

Institutions

Catalogues

Extractions

Data Library

Users

Settings

Add new catalogue

URL

URL for the catalogue

Name

Display name for the catalogue

Type

Type of catalogue to create

Institution

Institutions that match the catalogue URL appear first. [Create a new institution](#) if you don't see yours.

Submit

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Learning Program Extraction Interface

CTDL
xTRA

 Credential
Engine[™]

Institutions

Catalogues

Extractions

Data Library

Users

Settings

13

Add new catalogue

URL

URL for the catalogue

Name

Display name for the catalogue

Type

Learning Programs

Type of catalogue to create

Institution

Select an institution

Institutions that match the catalogue URL appear first. [Create a new institution](#) if you don't see yours.

Submit

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Competencies Extraction Interface

CTDL
xTRA

 Credential
Engine[™]

Institutions

Catalogues

Extractions

Data Library

Users

Settings

13

?

Add new catalogue

URL

URL for the catalogue

Name

Display name for the catalogue

Type

Competencies

Type of catalogue to create

Institution


Select an institution

Institutions that match the catalogue URL appear first. [Create a new institution](#) if you don't see yours.

Submit

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Credentials Extraction Interface




The screenshot displays the 'Add new catalogue' form within the Credential Engine xTRA interface. The left sidebar contains navigation links: Institutions, Catalogues (highlighted), Extractions, Data Library, Users, and Settings. The main content area is titled 'Add new catalogue' and includes the following fields:

- URL**: A text input field containing the example URL 'http://www.example.com/course_catalogue'.
- URL for the catalogue**: A label for the URL field.
- Name**: A text input field containing the placeholder 'Catalogue Name'.
- Display name for the catalogue**: A label for the name field.
- Type**: A dropdown menu currently set to 'Credentials'.
- Type of catalogue to create**: A label for the type field.
- Institution**: A dropdown menu with the placeholder 'Select an institution'.
- Institutions that match the catalogue URL appear first. [Create a new institution](#) if you don't see yours.**: A note below the institution dropdown.
- Submit**: A dark button at the bottom of the form.

Credential Engine Request for Proposals: Scaling CTDL xTRA Platform

Extraction Details Interface

The extraction details page uses the same template for each CTDL class.



The screenshot displays the 'Extraction Details' page for a CTDL xTRA class. The interface is divided into several sections:

- Left Sidebar:** Contains navigation links for Institutions, Catalogues, Extractions (highlighted), Data Library, Users, and Settings.
- Extraction Details (Main Content):**
 - Status:** Started at 05/04/2026, 04:17 AM. **COMPLETE** (in a black pill). Completed at 05/04/2026, 09:20 AM. Took 5hr2mins.
 - Last Action:** Alex - RESUME (05/04/2026, 09:11 AM) with a menu icon.
 - View Data:** Latest dataset from 05/04/2026, 04:17 AM.
 - Rerun data extraction:** A black button.
 - Models used:** gpt-5.4-mini (in a black pill).
 - Recipe #397:** CATEGORY_LINKS (no pagination) > DETAIL (no pagination). URL: https://rccd.curriqnet.com/catalog/alias/nc-catalog/iq/5367/5467...
 - Delete extraction:** A red button.
- Extraction Stats (Right Panel):**
 - 357** Items extracted
 - 141/141** Downloads (with a black progress bar)
 - 141/141** Page Extractions (with a black progress bar)
 - Cost Breakdown:** Estimated Total Cost: \$1.24
- Extraction Steps (Bottom Section):**
 - Buttons: Table (selected), Hierarchy, Sample.
 - | Step ID | Type | Parent | Configuration | Started At | Items |
|---------|-------------|------------|---------------|------------------|--------------------------------|
| 3776 | FETCH_ROOT | | JSON | 2026-05-04 04:17 | View 1 items |
| 3777 | FETCH_LINKS | FETCH_ROOT | JSON | 2026-05-04 04:26 | View 141 items |