




Digitale Verwaltung Schweiz
Administration numérique suisse
Amministrazione digitale Svizzera

Pre-Study: Swiss Ecosystem Discovery Catalog

Pre-study on the development of a
Use Case and Credential Catalog for
the Swiss verifiable data ecosystem
(formerly “VC Catalog”)



Pre-Study: Swiss Ecosystem Discovery Catalog

Date June 30 2026


Authors Accelerate.swiss, mandated by the working group iSVC (inhaltliche Standardisierung Verifiable Credentials), Digital Public Services Switzerland (DPSS)

Table of contents	Executive Summary	3
	1 Context and Motivation	6
	2 Methodology of the Pre-Study	7
	2.1 Scoping and ecosystem consultations	7
	2.2 Review of Swiss artefacts and international references	7
	2.3 Applied expertise and interpretive lens	8
	2.4 Flow-based design synthesis	8
	3 Vision and Conceptual Framework	8
	3.1 High-Level Vision for the Catalog	8
	3.2 The Catalog in the Broader Ecosystem	9
	3.3 Conceptual Model	10
	3.3.1 The Parts of the Conceptual Model	10
	3.3.2 Catalog User Archetypes	12
	4 Functional Requirements, Illustrative Views and Flow-Based Findings	14
	4.1 Functional Requirement Baseline	14
	4.2 Illustrative Use-Case View	15
	4.3 Conceptual Visualisations and User Journeys	15
	4.4 Key Findings and Considerations from the Flow Work	18
	5 Initial Ecosystem Landscape and Ecosystem Fit	19
	5.1 Federation with Ecosystems, Registries and Public Catalogs	19
	5.2 Verifier Governance	19
	5.3 DVS4U, Credential Hub and the Proposed Catalog	20
	6 Technical Approach	20
	6.1 High Level Architectural Blueprint	21
	6.1.1 Supporting components	22
	6.2 Technical Standards	23
	7 Governance Model & Processes	23
	7.1 Catalog Governance Model Structure Design	24

7.2	Key roles	24
7.3	Governance Scope and Institutional Fit	25
7.4	Processes Surrounding the Ecosystem Discovery Catalog	26
7.5	Adjacent Ecosystem Governance-Orchestration Function	26
8	Implementation & Operating Model Options	26
8.1	Core operating activities	26
8.2	Initial operating capacity	27
8.3	Institutional setup and neutrality	28
8.4	Role of ecosystem orchestrators	29
8.5	Technical operation	29
8.6	Alignment and federated operating logic	29
8.7	Funding model	29
8.8	Scalability and onboarding	30
8.9	Process design	30
8.10	Progression	31
9	Recommendation and Phased Approach	32
9.1	Overall recommendation	32
9.2	Strategic execution principles	32
9.3	Functional priorities for the MVP	32
9.4	Phased Approach	33
9.4.1	Phase 1: Decision and Procurement Readiness for the Ecosystem Discovery Catalog MVP	34
9.4.2	Phase 1.5: Procurement	34
9.4.3	Phase 2: MVP Implementation and Pilot Activation	35
9.4.4	Phase 3: Version 1 Development, Deployment and Early Onboarding	36
9.5	Future Considerations: Interoperability and Later-Stage Capabilities	37
9.5.1	Conformance testing	37
9.5.2	Dynamic, context-driven data exchange conditions	37
9.5.3	Runtime Binding (Making Governance User-Comprehendable)	37
9.6	Closing recommendation	37

Pre-study on the development of a Use Case and Credential Catalog for the Swiss verifiable data ecosystem (formerly “VC Catalog”)

Executive Summary



The “Ergebnisbericht der DVS Arbeitsgruppe inhaltliche Standardisierung Verifiable Credentials (iSVC)” (published in July 2024 in German and French) of the DPSS iSVC Working Group recommends specifying the publication platform (e.g. VC Catalogue, I14Y) and assigns the working group responsibility for implementing and maintaining a VC Catalogue (credential catalogue). As the scope, effort, and framework conditions were still undefined, the working group commissioned a pre-study for the VC Catalogue.

The pre-study discussions showed that the term “credential catalog” is still interpreted inconsistently across stakeholders. We recommend developing the Ecosystem Discovery Catalog as Switzerland’s coordination, discovery and transparency layer for the verifiable data ecosystem. The name reflects its actual role: making use cases, credential supply and demand, ecosystem actors, governance conditions and reuse potential legible and actionable across all three ambition levels (AL).

With swiyu planned to go live in the first half of 2027, the foundations will be in place. What will be missing at the launch and during operation of the ecosystem, is transparency about which credentials are already issued and which additional credentials are needed, who issues or verifies them, which governance conditions apply, where reuse is possible and where coordinated action is required. Without this layer, ecosystem development will remain fragmented, difficult to discover and hard to scale.

The examined vehicle leasing and motor vehicle claims processing scenarios made these gaps concrete. They are organisational and informational in nature: demand in credential ecosystems is typically not visible. Without the Ecosystem Discovery Catalog, metadata and trust references remain fragmented and must be manually assembled for the relevant use context. In addition, verifier needs are not yet connected to issuer action - and vice versa.

The first build phase should deliver a focused MVP with five core capabilities:

- (i) structured use-case entries
- (ii) visibility of existing and missing credentials
- (iii) issuer and verifier discovery
- (iv) demand-and-supply signalling
- (v) lifecycle and maturity visibility for use cases and Catalog entries, with federated references to I14Y, trust registries and governance information.

Schema creation and schema lifecycle management should remain outside the Catalog and with I14Y, eCH-linked processes or recognised sectoral repositories. Runtime policy execution and conformance testing may become later-stage capabilities.

Product, governance and operating model should be developed together. A credible model combines public anchoring, stable technical operation, curation capacity, ecosystem activation and clear interfaces to formal standardisation, publication and trust-governance bodies. DPSS is well placed to anchor public-interest funding and AL2 activation. Technical operations should be assessed for placement with a stable operator, such as the Federal Office of Information, Technology, Systems and Telecommunications (FOITT) or another suitable public operator. Private-sector and PPP participation should support AL3 activation and cross-sector use-case development.

The immediate next step is a focused three-month Decision and Procurement Readiness Phase. It should produce a defined MVP scope, draft technical architecture, governance and curation model, operating model, priority use cases and a decision-ready procurement and funding basis. This reduces delivery risk while preserving momentum and gives Switzerland a clear path toward a publicly anchored MVP that is small enough to build, clear enough to procure and useful enough to demonstrate coordination value across AL2 and selected AL3 contexts.

1 Context and Motivation

The use of VCs marks a significant shift in how trusted, reusable data can be exchanged between individuals, organisations and public authorities. In Switzerland, the Federal Office of Justice (FOJ) is leading the implementation of the national public trust infrastructure, swiyu, to support the issuance, holding, presentation and verification of such credentials across three ambition levels: Ambition Level 1 (AL1) for the e-ID, AL2 for public-sector credentials and AL3 for credentials issued by private and other non-state actors.

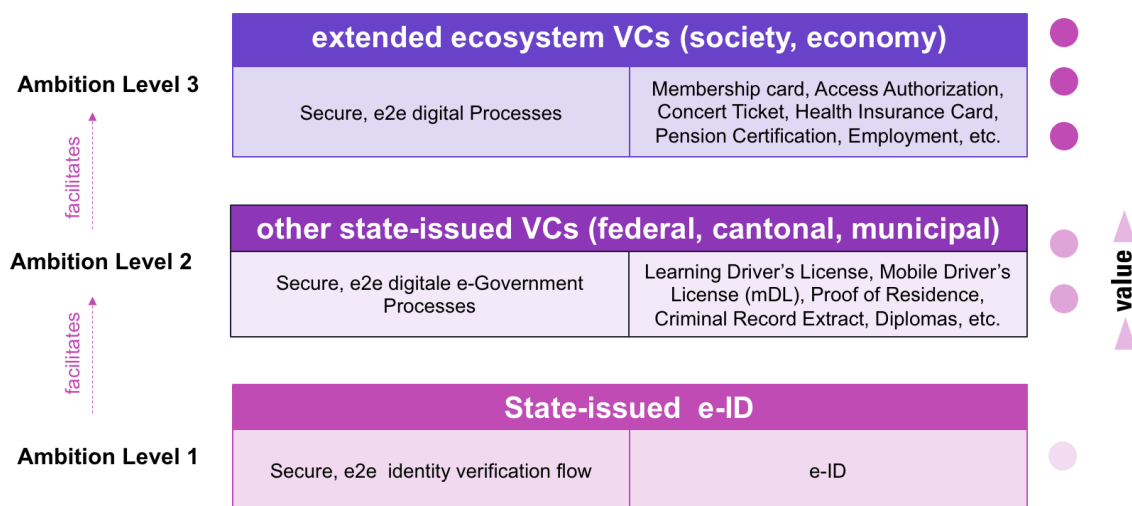


Fig. value contribution of "ambition" levels

For simplicity, this study continues to use the AL1–AL3 terminology for ecosystem segmentation. Further, "verifiable data ecosystem" is used as the broader term, while "verifiable credential ecosystem" refers to the credential-based implementation pattern currently enabled through swiyu.

Verifiable credentials can operationalise the *Once-Only* Principle across domains by allowing individuals and organisations to reuse original, trusted data from authoritative sources instead of submitting similar but self declared information to different verifiers and making individual data interfaces from verifiers to authoritative sources obsolete. They also enable data minimisation by design, allowing verifiers to request only what is needed for a given use case while still relying on the authenticity of the presented attributes.

This also strengthens digital sovereignty: individuals, organisations and public authorities can exchange and verify data under their own control, based on shared rules and trusted sources. The infrastructure therefore supports not only the digitalisation of services, but also autonomy, resilience and long-term competitiveness by enabling applications and services to build on shared trust infrastructure

AL1 provides the initial trust anchor; AL2 and AL3 will develop across federal, cantonal, municipal, sectoral and private contexts. In a wallet-based, privacy-preserving and federated

ecosystem, credentials, issuer activity, verifier demand and reuse potential are not easily visible by default. This creates a clear risk of fragmented development, redundant schemas, uneven governance paths and limited reuse. Realising the broader value of swiyu therefore depends on making reuse and coordination visible across actors and domains.

2 Methodology of the Pre-Study

The findings and recommendations draw on expert interviews, DIDAS community input, a World Café-style consultation, desktop research, Swiss artefact review, international reference analysis and flow-based design work. Together, these inputs shaped the proposed role, scope, MVP priorities, governance implications and operating model of the Ecosystem Discovery Catalog.

2.1 Scoping and ecosystem consultations

At the outset of the pre-study, an initial scoping workshop was held with the commissioning authority, namely the head of the DPSS Working Group “inhaltliche Standardisierung Verifiable Credentials” (iSVC), to align on scope, objectives and expected outcomes. This was followed by five expert interviews with relevant ecosystem actors, a group consultation in the context of the federal ecosystem exchange meeting of 27 February 2026 and exchanges within the DIDAS community. These consultations provided qualitative insights into current adoption barriers, operational requirements, interoperability considerations and governance challenges across public- and private-sector settings. At the same time, the work has shown that, at national level, the conceptual understanding of ecosystem enablement remains uneven. This also affects how the role of an Ecosystem Discovery Catalog is understood.

2.2 Review of Swiss artefacts and international references

The review covered Swiss artefacts such as the DPSS iSVC final report, DIDAS VC Catalog v1.4 (2022), VC blueprint schema templates, iSVC standardisation guidelines and earlier national credential-catalog architecture concepts. These materials remain relevant as early design references for Swiss credential standardisation and ecosystem coordination.

Additional references included I14Y as the Swiss interoperability platform and publicly available swiyu trust infrastructure artefacts, which form part of the emerging operational foundation of the Swiss verifiable credential ecosystem.

Relevant international references included The European Digital Identity (EUDI) Architecture and Reference Framework¹, the EUDI Technical Specification TS 11 v1.0.² the EUDI Attestation Rulebooks Catalog³, Community governance and ecosystem tooling for decentralized trust

¹ Available at: <https://eudi.dev/latest/architecture-and-reference-framework-main/>

² Available at: <https://github.com/eu-digital-identity-wallet/eudi-doc-standards-and-technical-specifications/blob/main/docs/technical-specifications/ts11-interfaces-and-formats-for-catalogue-of-attributes-and-catalogue-of-schemes.md>

³ Available at: <https://github.com/eu-digital-identity-wallet/eudi-doc-attestation-rulebooks-catalog/blob/main/rulebooks/pid/pid-rulebook.md>

ecosystems⁴, European procurement initiatives related to ecosystem management platforms for digital wallet infrastructures⁵ and the European Interoperability Framework (EIF)⁶.

2.3 Applied expertise and interpretive lens

The pre-study also draws on the direct experience of the Accelerate.swiss team, as authors of this study, in the design, governance and ecosystem enablement of verifiable credential infrastructures and applications in Switzerland and internationally. It helped connect analytical observations with operational feasibility, likely coordination bottlenecks and concrete implications for the Catalog's design, governance and operating model.

2.4 Flow-based design synthesis

Concrete flow design was used to test the Catalog logic under realistic ecosystem conditions. Verifier and issuer journeys in vehicle leasing and motor vehicle claims processing surfaced practical questions around reuse, demand visibility, issuer incentives, lifecycle progression and federation across systems.

Stakeholder input confirmed that "credential catalog" is interpreted inconsistently, ranging from schema registry to governance tool to ecosystem coordination layer. Concrete use-case flows, mockups and visualisations therefore became essential to make the Catalog tangible and reveal ecosystem questions that remained hidden in abstract discussion.

3 Vision and Conceptual Framework

3.1 High-Level Vision for the Catalog

The Ecosystem Discovery Catalog should become a web-based coordination and discovery platform for credential use across all three ambition levels.

"The Catalog does not create or manage authoritative data, credential schemas, or trust decisions. It federates and contextualises authoritative sources and creates visibility of relationships between them in the context of use cases."

Its architectural role is to make the conditions for trustworthy reuse and coordinated adoption visible.

It addresses the ecosystem's market-coordination challenge by clarifying where publication is needed, where verification demand exists, which dependencies matter and which incentives can support participation. In doing so, the Catalog can reduce redundancy, improve data quality,

⁴ Available at: <https://fides.community/community-tools/>

⁵ Available at: <https://oeffentlichevergabe.de/ui/de/search/details?noticeId=b2690f7d-b65c-4628-bdf6-2cb75bb4061d&lotId=LOT-0001>

⁶ Available at: <https://interoperable-europe.ec.europa.eu/collection/iopeu-monitoring/european-interoperability-framework>

lower transaction costs and enable greater automation across public services and the wider digital economy.

At its core, it enables participants to understand:

- Use-cases and their potential cross-sectoral context
- which credentials exist or are emerging
- which organizations issue or verify them
- which governance frameworks apply
- where reuse, extension, or collaboration may be worthwhile

This broader framing reflects stakeholder feedback that the need extends beyond a technical registry toward adoption, coordination, reuse and ecosystem value.

3.2 The Catalog in the Broader Ecosystem

The figure and table below position the Catalog relative to these adjacent functions and infrastructures.




Fig. VC Catalog and the broader ecosystem

Component	Core role
I14Y	Swiss interoperability platform for publishing and reusing semantic artefacts, attributes, schemas and cross-domain concepts.
swiyu Trust Registry	Machine-readable trust layer for issuer recognition, declared data requests and, where applicable, declared data requests and purposes.
swiyu Swiss Profiles	Technical profiles defining how credentials are issued, held, presented, verified and how actors are anchored in the Swiss ecosystem.
Catalog Curation / Operations Team	Operational function responsible for maintaining the catalog, enforcing quality and coherence and running day-to-day catalog operations.

3.3 Conceptual Model

The Catalog is structured around the conceptual parts and the fundamental roles model (issuer-holder-verifier) of the Swiss verifiable data ecosystem. The model connects concrete use cases with the trust, semantic and operational information needed to make credentials reusable across contexts.

The model defines what the Catalog must make visible: what needs to be proven, by whom, under which governance conditions, using which credentials or schemas and with which actors involved.

Across all parts, the Catalog acts as a cross-cutting coordination and discovery function. Its role is to make the relationships between use cases, credentials, trust references, schemas and visible and reusable.




Fig. Ecosystem Discovery Catalog as cross-cutting coordination layer across the parts of the verifiable credential ecosystem relevant for the catalog, making use cases, trust conditions, schemas, operational patterns and roles visible and actionable for different user archetypes.

3.3.1 The Parts of the Conceptual Model

Part 1 - Decision Context: Trust Domains and Use Cases

Trust Domains provide the legal, governance and assurance context in which credentials are issued, presented and relied upon. They include the relevant legal basis, policy rules,

accreditation mechanisms, liability considerations, sectoral governance frameworks and applicable trust anchors.

Use Cases describe the concrete interaction scenarios in which verifiable data is exchanged to achieve a specific outcome, such as onboarding, eligibility verification, compliance, service access or claims processing. They define which claims must be proven, which credentials may be required, which actors are involved and which governance context applies.

Catalog function: In this part, the Catalog makes use cases and their trust context - within and potentially across domains - discoverable, along with the coordination needs they entail (as it does in every following part).

Part 2 - Trust Resolution

Trust Resolution describes how a relying party determines whether a credential can be accepted for a specific decision, based on trust ecosystem frameworks, that include issuer recognition, assurance context, registry references, status mechanisms and policy constraints.

In the Swiss ecosystem, this part relates to the swiyu Base Registry, the swiyu Trust Registry and, where relevant, sectoral or international trust frameworks. Over time, cross-ecosystem trust resolution may also require references to interoperable mechanisms such as trusted lists, trust registries or federation protocols.

Catalog function: In this part, the Catalog makes trust frameworks, registry references, assurance context and issuer-recognition logic visible.

Part 3 - Semantic Layer: Claims and Credential Schemas

The semantic layer describes what a credential means. Claims represent the individual assertions made about a subject, while credential schemas define the structure and expected content of a credential.

In the Swiss context, semantic alignment may involve eCH standards, I14Y publications, OCA overlays and sector-specific data models. Reuse depends not only on technical compatibility, but also on a shared understanding of attributes and claims.

Catalog function: In this part, the Catalog makes verifiable discoverable, comparable, reusable and gap-visible.

Part 4 - Operational Interaction Patterns

Operational interaction patterns describe how credentials are used in practice and what life cycle they have: issuance, holding, verification and revocation. These patterns connect the use case, trust context and semantic structure to an actual ecosystem interaction.

For each use case, actors need to understand which credentials are issued, requested or presented, which parties perform which role and which external infrastructure components are involved.

Catalog function: In this part, the Catalog makes credential flows, operational dependencies, reuse opportunities and implementation dependencies visible.

Roles

The roles box identifies the actors involved in a credential ecosystem. Issuers provide credentials, holders receive and present them, verifiers rely on them and governing authorities define or maintain the conditions under which credentials may be issued and accepted.

In practice, one organisation may play several roles within one single and/or across different use cases. A public authority may be an issuer in one context and a verifier in another. A private-sector actor may create demand for a credential while also becoming an issuer of another.

Catalog function: In this part, the Catalog makes actor roles, dependencies and coordination needs visible.

Interaction Flow

The model becomes operational when a use case is connected across the four conceptual parts and the relevant roles. A governing authority or domain body defines the relevant trust context. An issuer provides a credential under that context. A holder receives and presents the credential. A verifier assesses whether the credential can be accepted for a specific decision. Trust, semantic and operational dependencies are resolved through the relevant registries, schemas, governance references and status mechanisms.

The Catalog supports this flow by making the required knowledge discoverable before implementation and reusable across use cases. It helps actors identify which credentials exist, which are missing, who may issue or verify them, which governance conditions apply and where existing artefacts can be reused instead of recreated.

3.3.2 Catalog User Archetypes

The Catalog addresses four ecosystem roles - issuers, holders, verifiers and governing authorities - while supporting distinct user archetypes within and across these roles. The term "Holder" is used both as an ecosystem role and, in the archetype model, as the perspective of a person seeking to understand credential governance, transparency, privacy context and reuse possibilities.

Six user archetypes are relevant for product design and access control. They cut across ecosystem roles: for example, a compliance officer at an insurance company and a legal advisor at a cantonal authority may share the same need to understand legal and trust context.

Credential User / Holder: Uses the Catalog to understand which credentials exist, what they contain, which organisations issue and accept them and in which situations they can be used.

Policy Navigator: Uses the Catalog to understand legal basis, governance context, trust frameworks and where policy clarification or formal change may be required.

Schema Builder: Uses the Catalog to discover, compare, request and reference credential schemas and related semantic artefacts. Schema definition, publication and lifecycle management primarily take place in I14Y or other recognised semantic repositories.

Demand / Supply Signaller: Makes visible both; what the ecosystem needs and what it may be able to inform before use case implementation or extension starts. This includes signalling missing credentials, schemas and highlighting adoption or implementation intent as well as indicating where organisations with relevant data or mandate may be willing and able to issue and/or verify under defined conditions.

Implementer: Uses the Catalog to access the technical and operational detail needed to build and operate real integrations.

Curator: Manages the Catalog itself by validating entries, moderating processes, advancing lifecycle phases where authorised and maintaining coherence in day-to-day operation.

Archetype	Primary catalog interface / focus	Main conceptual model parts	Typical governance role(s)
Credential User / Holder	Credential transparency view · data disclosure overview · accepted-use view · request context and terms consent information	P1 Decision Context · P3 Semantic Layer · P4 Operational Interaction Patterns	Community Participant
Policy Navigator	Use case index · trust framework references · legal basis view · version history, decisions and amendments	P1 Decision Context · P2 Trust Resolution	Community Participant · Amendment Commenter
Demand / Supply Signaller	Gap board · issuer / verifier signalling · readiness declaration · implementation intent view	P1 Decision Context · Roles	Community Participant · Ecosystem Contributor
Schema Builder	Lifecycle tracker · semantic artefact view · schema request view · federation / publication reference panel	P3 Semantic Layer · P4 Operational Interaction Patterns	Ecosystem Contributor · Domain (data) Steward
Implementer	Schema detail view · technical specification view · conformance status · onboarding path	P4 Operational Interaction Patterns · Roles	Ecosystem Contributor
Curator	Governance dashboard · role registry · moderation and amendment tools · lifecycle control	All parts	Catalog Curator · Standards Authority

4 Functional Requirements, Illustrative Views and Flow-Based Findings

4.1 Functional Requirement Baseline

The research and flow work translate into the following functional requirements.

The figure below shows the early functional view used to test the Catalog's discovery and ecosystem-signalling logic with participants.

SWIYU VC Community "Marketplace" – Web Portal, Catalog Search & API access to:

- | | | |
|----------------------------|---------------------------------|-------------------------|
| • Which VCs exists | • VC Issuers | • VC Attributes |
| • Which VCs are needed | • VC Verifiers | • Semantic Schemas |
| • Thematic Events Calendar | • VC LoA & Governance | • Editing and Discovery |
| • Use-Case Showcase | • PII & Compliance Declarations | |

Adoption Awareness Use-Cases Acceptance Transparency Best Practices Standards

Fig. Our early overview of key functionality and ecosystem's success factors we evolved from

These enabling conditions imply the following functional requirements for the Catalog:

Ecosystem Legibility: Provide structured visibility into credentials, issuers, verifiers, governance frameworks and use-case contexts.

Reuse and Convergence: The Catalog should enable actors to discover existing use cases, schemas, attributes and governance approaches before creating new ones. It should support comparison, reference and reuse of existing artefacts in order to reduce duplication and promote convergence where appropriate.

Coordination and Ecosystem Formation: The Catalog should make visible where demand exists, where supply may emerge and where coordination between actors is needed. It should support signalling of needs, opportunities, dependencies and areas where shared solutions may be viable.

Trust, Transparency and Accountability: The Catalog should expose the conditions under which credentials are issued, accepted and relied upon. This includes visibility into issuers, verifiers, applicable governance frameworks, assurance context and relevant declarations concerning sensitive attributes and compliance conditions.

Continuous Ecosystem Learning: Support structured contribution, curation, validation and maintenance of use cases, credential needs, governance references, schema references and emerging requirements.

Standards and Ecosystem Maturity: Make technical and semantic building blocks discoverable, comparable and traceable across maturity and standardisation pathways.

Together, these requirements define the functional baseline: the Catalog must make the Swiss verifiable data ecosystem comprehensible - and capable of maturing over time.

4.2 Illustrative Use-Case View

The illustrative view below shows how these requirements can appear in a concrete Catalog interface, combining process flow, available and missing credentials, actors and governance context.

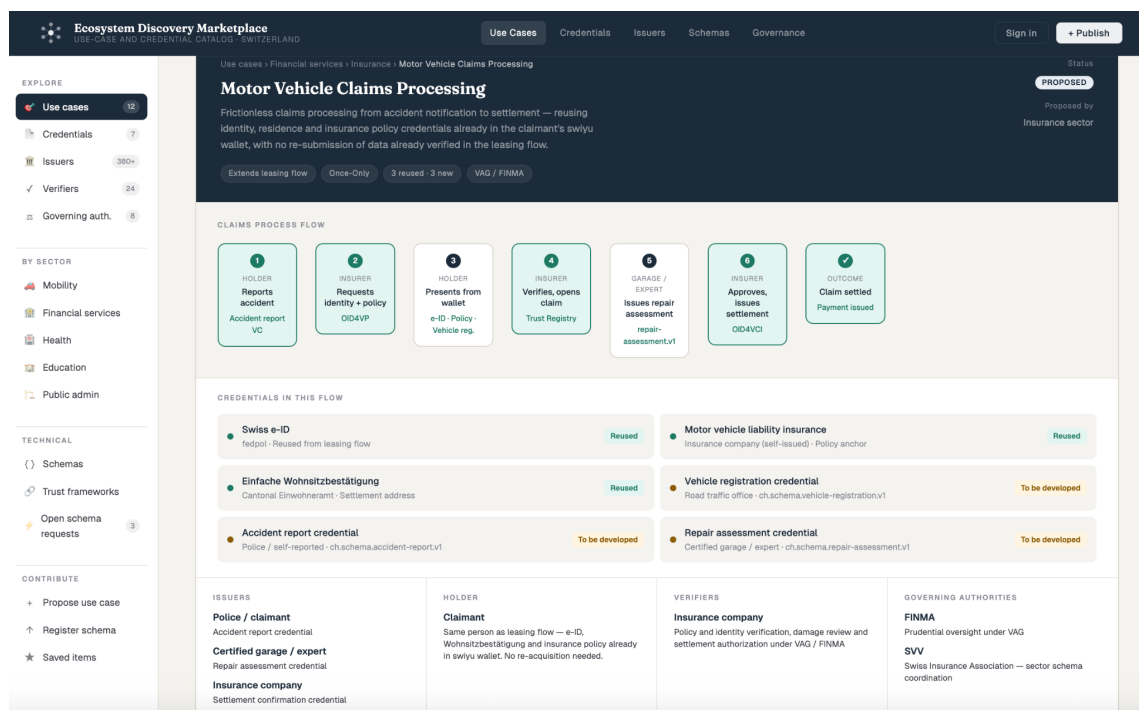


Fig. Illustrative example of a use-case view in the Ecosystem Discovery Catalog, showing process flow, credential reuse, missing credentials and actor context

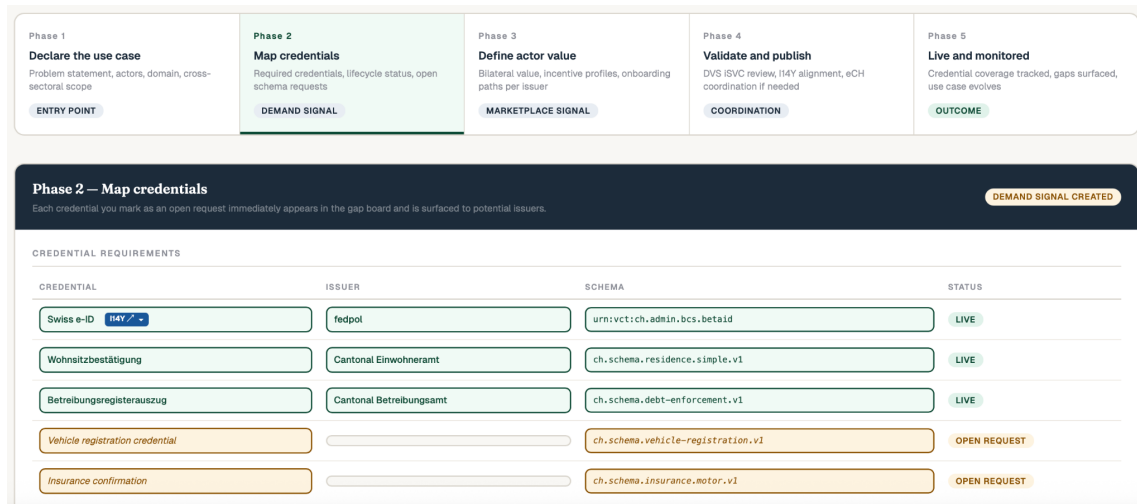
4.3 Conceptual Visualisations and User Journeys

The mockups tested whether use cases, credential needs, reuse potential, gaps and issuer opportunities can be made understandable in a shared interface.

The Vehicle Claims Processing example demonstrates cross-sectoral credential reuse, schema reuse and trust resolution across federal, cantonal and private-sector actors.

Use-case credential mapping mockups

The first mockups show how a use-case owner or Domain Steward can map existing and missing credentials against a concrete process. This makes available credentials visible, identifies missing credentials and turns unmet needs into explicit demand signals.



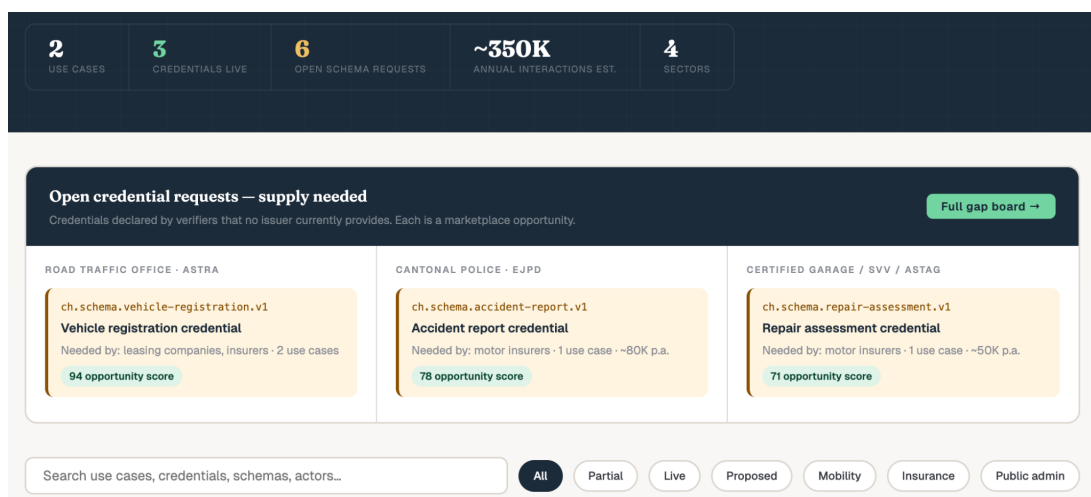
The mockup shows a five-phase process for credential mapping. Phase 2, 'Map credentials', is the active stage, indicated by a green bar and a 'DEMAND SIGNAL' label. It displays a table of 'CREDENTIAL REQUIREMENTS' with columns for CREDENTIAL, ISSUER, SCHEMA, and STATUS. Existing credentials are marked 'LIVE', while missing ones are marked 'OPEN REQUEST'.

CREDENTIAL	ISSUER	SCHEMA	STATUS
Swiss e-ID (14Y / -)	fedpol	urn:vct:ch.admin.bcs.betaid	LIVE
Wohnsitzbestätigung	Cantonal Einwohneramt	ch.schema.residence.simple.v1	LIVE
Betreibungsregistrauszug	Cantonal Betreibungsamt	ch.schema.debt-enforcement.v1	LIVE
Vehicle registration credential		ch.schema.vehicle-registration.v1	OPEN REQUEST
Insurance confirmation		ch.schema.insurance.motor.v1	OPEN REQUEST

Fig. Illustrative use-case credential mapping view, showing how existing and missing credentials can be mapped against a concrete process.

Gap-board and supply-needed mockups

The gap-board mockups show how these demand signals can be aggregated across use cases and surfaced as issuer opportunities. Missing credentials can be grouped by use case, sector or likely issuer type and enriched with signals such as estimated interaction volume, blocked use cases, lifecycle state and opportunity score.



The mockup displays a dashboard with key metrics: 2 USE CASES, 3 CREDENTIALS LIVE, 6 OPEN SCHEMA REQUESTS, ~350K ANNUAL INTERACTIONS EST., and 4 SECTORS. Below, it shows 'Open credential requests — supply needed' for three issuers: ROAD TRAFFIC OFFICE - ASTRA, CANTONAL POLICE - EJPD, and CERTIFIED GARAGE / SVV / ASTAG. Each issuer's section lists specific credentials, their need, and an opportunity score.

ISSUER	CREDENTIAL	NEED	OPPORTUNITY SCORE
ROAD TRAFFIC OFFICE - ASTRA	ch.schema.vehicle-registration.v1 Vehicle registration credential	Needed by: leasing companies, insurers - 2 use cases	94 opportunity score
CANTONAL POLICE - EJPD	ch.schema.accident-report.v1 Accident report credential	Needed by: motor insurers - 1 use case - ~80K p.a.	78 opportunity score
CERTIFIED GARAGE / SVV / ASTAG	ch.schema.repair-assessment.v1 Repair assessment credential	Needed by: motor insurers - 1 use case - ~50K p.a.	71 opportunity score

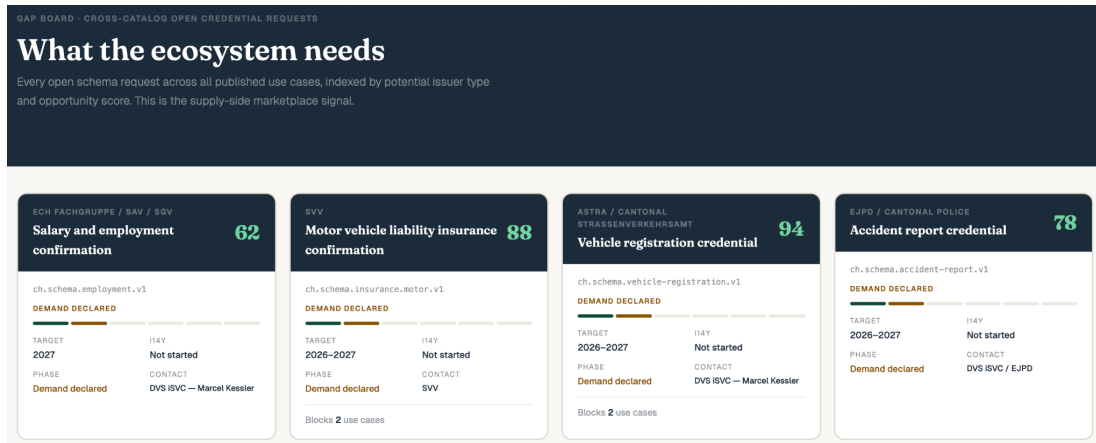


Fig. Illustrative gap-board view, aggregating missing credentials and surfacing them as potential issuer opportunities.

Issuer-facing discovery mockups

The issuer-facing mockups show how a potential issuer can enter through its mandate or data authority rather than through abstract technical schema categories. Based on the issuer profile, the Catalog can surface relevant open requests, affected use cases, anticipated relying parties and governance context.

01

ENTRY POINT — ISSUER PERSPECTIVE

What kind of issuer are you?

The marketplace surfaces opportunities relevant to your data assets and organizational mandate.

Public authority Federal, cantonal or municipal body holding authoritative data — civil registries, vehicle registration, criminal records, public registers	Certified professional body Licensed or certified organizations that can attest professional qualifications, assessments, or compliance status — garages, experts, notaries	Private sector issuer Companies issuing credentials within their domain — insurance policies, employment records, qualifications, financial status
---	--	---

02

DEMAND DISCOVERY

Open credential requests — what the ecosystem needs from you

These requests were declared by verifiers who have published use cases. Each one represents committed demand, not speculation.

3 open schema requests matching your profile

Requested by the insurance sector as part of the motor vehicle claims processing use case, which extends the vehicle leasing and immatriculation flow. All three schemas are needed to complete an end-to-end claims flow that is currently blocked on missing credentials.

4

Requesting verifiers

2

Blocked use cases

~180K

Est. annual interactions

94 OPP.

ch.schema.vehicle-registration.v1

Vehicle registration credential

Machine-readable attestation of vehicle ownership and registration issued by the cantonal road traffic office at the point of immatriculation — naturally downstream of an existing statutory process.

Insurance sector Leasing companies Mobility SVG / VZV - ASTRA Statutory process

~50K ANNUAL INTERACTIONS EST.

OPEN REQUEST

Collapse

Fig. Illustrative issuer-facing discovery flow, matching open credential requests to an organization's likely mandate, data assets and institutional role.

4.4 Key Findings and Considerations from the Flow Work

The verifier and issuer flow work produced evidence-informed design implications for the MVP. Five findings are particularly relevant.

1. The primary bottleneck is coordination, not core technology

The VC technology stack is sufficiently mature for the tested leasing and motor vehicle claims journeys. The bottleneck lies in making demand visible, connecting it to potential issuers and coordinating next steps. The MVP should therefore prioritise coordination and discovery value over deeper technical functionality.

2. Issuers lack decision-ready signals

Potential issuers lack decision-ready signals on expected demand, relying-party uptake, implementation effort and near-term value. Open credential needs and semantic-gap requests should therefore be presented as structured opportunities, with enough context to assess relevance, mandate fit and feasibility.

3. Reuse only works when it is made explicit

Several credentials needed in the motor vehicle claims scenario were already available from an underlying registration and leasing flow. Reuse is therefore a major source of ecosystem value, but only when visible at use-case level. The Catalog should show reused, newly required and deliberately excluded credentials so that reuse and data minimisation become operational.

4. Relevant information remains fragmented across systems

Actors still need to assemble information across schema publication points, trust references, overlays and documentation. This fragmentation obscures maturity, provenance and governance. The Catalog should federate these references into one coherent view.

5. Opportunity signals help non-technical actors act

Open semantic or credential gaps become more actionable when they highlight likely demand, affected use cases and practical relevance. The Catalog should therefore translate ecosystem needs into structured opportunity signals and support matching between demand, potential issuers and relevant actors.

5 Initial Ecosystem Landscape and Ecosystem Fit

This chapter positions the Catalog within adjacent Swiss ecosystem components and clarifies where it should federate, reference, or complement existing infrastructure.

5.1 Federation with Ecosystems, Registries and Public Catalogs

The Catalog should support federation across sectoral ecosystems, public-sector domains and private-sector governance contexts. Different domains will develop their own schemas, governance models, operational processes and trust references. The Catalog should therefore organise use cases and credentials by ecosystem while making cross-domain reuse and interoperability visible where appropriate.

The Catalog references authoritative sources such as I14Y, swiyu registries, sectoral repositories and recognised trust frameworks and makes their relevance understandable in use-case context. I14Y may also be relevant for private-sector contexts. Where I14Y is not available or not suitable for private or sectoral semantic artefacts, complementary repositories may need to be established, for example as I14Y-compatible or I14Y-derived instances, to support structured publication, reuse and federation beyond the public sector.

5.2 Verifier Governance

The Catalog should surface verifier governance as a visibility and transparency function. Following the Federal Council's additional measures of 25 February 2026⁷, the relying-party side of the ecosystem becomes more explicit: providers are expected to register intended data queries and purposes in a public federal register; unauthorized or excessive requests trigger wallet warnings and serious non-compliance can lead to exclusion from the e-ID system and trust infrastructure by the FOJ. The Catalog should therefore make verifier status, declared purpose, applicable

⁷ <https://www.eid.admin.ch/en/akzeptanz-der-e-id-mit-zusaetzlichen-massnahmen-staerken-e>

governance basis and relevant attribute-request context visible, so that credential use is easier to assess from both a regulatory and trust perspective.

5.3 DVS4U, Credential Hub and the Proposed Catalog

DVS4U is part of the DPSS 2026 execution plan⁸. It is the proposed authority-side capability for issuing, managing, revoking and verifying verifiable credentials in an interoperable way. The Credential Hub can be understood as a future access and retrieval capability: a user-facing layer through which holders can identify and obtain available credentials through simple, reusable flows, potentially based on credentials they already hold.

The Catalog adds the coordination and discovery function: it makes visible which credential types exist or are missing, which use cases drive demand, who may issue or verify, which governance and trust references apply and where reuse across public and private contexts is possible.

This ecosystem discovery function is relevant for public, private, self-governing and mixed public-private environments, including banking, insurance, healthcare, education, retail and e-commerce and sport.

This differentiation can be stated simply:

- **DVS4U** = operational authority-side capability
- **Credential Hub** = future open-source access / self-service capability
- **Catalog** = coordination and discovery layer to drive value and adoption

6 Technical Approach

The Catalog is envisioned as a modular web-based application that can progressively provide additional functionality for AL2 and AL3 communities.

The application should serve technical, business, legal, academic and policy users.

Additionally, the catalog should provide a set of APIs to enable automated discovery of the defined credentials along with their governance, issuers and other characteristics.

Finally, the catalog should support integration with the swiyu registries, I14Y schema repository and store and manage Catalog-specific metadata.

It is important to note that given the early stage of the swiyu ecosystem and the evolving nature of VC technology, we expect user requirements to go through a rapid evolution. It is therefore

⁸https://www.digitale-verwaltung-schweiz.ch/application/files/8317/6163/7496/Umsetzung-splan_2026_27_Oktober_2025_Webversion.pdf

important to design the system's architecture in a modular and expandable fashion, enabling iterative approach and rapid deployment of the new and updated features.

6.1 High Level Architectural Blueprint

To meet these requirements the system will be composed of a number of loosely coupled subsystems as follows:

- **Front-End** - a modular, extensible Web-UI subsystem that supports access to various functionalities of the catalog, tailored for various user roles.
- **Content Management Subsystem (CMS)** - manages the storage, delivery and editing of a variety of content types that will be presented to the users of the catalog. E.g. articles, blog posts, news, events calendar, etc. However the support of schema creation and /or credential definition and their lifecycle management will be out of scope for this subsystem.
- **Search Subsystem** - supports search and discovery functionality for credentials, schemas, issuers, verifiers, use-cases, etc. Abstracts internal and evolving storage mechanism of the catalog for the Front End. Provides core enabling functionality for the API subsystem.
- **Integration Subsystem** - extensible subsystem that supports the rest of the catalog and enables data retrieval and/or ETL processes from external sources, such as swiyu Trust and Base Registry, credential schema catalogs/directories (e.g. I14Y) and similar third-party data sources.
- **Metadata Storage Subsystem** - extensible storage subsystem that allows the catalog to store its specific metadata. E.g. Level of Assurance definitions, structured governance information, etc. Might be used to serve as a cache to store schema definitions, etc. from I14Y and third-party sources, as needed. Or provide the primary storage facility for the schema definitions for parties that don't have their own facility and there's no other Catalog that is suitable for such data.
- **API Subsystem** - subsystem targeted at serving external applications that need to have a way to programmatically discover information about credentials, their schemas, issuers, etc.




Fig. High Level Architectural Blueprint

The Fig. High-Level Architecture Blueprint illustrates the catalog's high level composition and highlights the need for integration with a potentially expanding set of external data sources. A critical design boundary that must be resolved in Phase 1, concerns what the Catalog stores natively versus what it federates from adjacent systems. As a working assumption for the next phase, the following distinction is proposed:

Artefact type	Catalog role
Use-case entries, demand signals, opportunity scores	Stored natively in Catalog
Schema metadata (for formally published AL2 schemas)	Federated from I14Y
Schema metadata (for AL3 / pre-publication schemas)	Preferably entire lifecycle in and federated from I14Y, if otherwise, stored natively until I14Y publication
Trust registry references (issuers, verifiers)	Linked to swiyu Trust Registry; not duplicated
Governance framework references	Linked to authoritative source; summary metadata stored natively
Use Case Lifecycle and maturity state	Stored natively in Catalog
Schema Lifecycle and maturity state	Federated from I14Y
OCA overlays	Linked to source repository; not stored natively

This boundary is indicative and must be confirmed with I14Y and swiyu operators in Phase 1. The guiding principle is: the Catalog federates authoritative data from its canonical source rather than duplicating it and stores only Catalog-native metadata that has no equivalent home elsewhere.

The MVP does not need to implement a schema authoring workspace, runtime policy resolution, or a conformance test harness.

6.1.1 Supporting components

In addition to the mentioned high-level components of the Catalog, the following subsystems and technical considerations will be addressed:

- User Management and IAM - support roles, access rights and permissions for different user roles and the Catalog's curation team.
- Operational Tooling - a set of tools and systems used to monitor, troubleshoot, maintain and secure the Catalog. Deployment tooling, log monitoring, back-up, etc.

The exact functionality and technical design of these components should be explored and defined in the next phases of the Catalog's development.

6.2 Technical Standards

This standards overview is intentionally preliminary. A detailed standards profile should be developed during Phase 1. The Catalog should support established technical standards where possible, while remaining flexible enough to accommodate emerging standards over time. This is particularly important in areas related to credential definitions, schema representation, trusted lists and interoperability across ecosystems.

Established standards likely relevant in initial phases include:

- Credential formats: SD-JWT, mDL
- Schema and semantic definition standards: OWL, Turtle, SHACL
- Trusted list / trust framework standards: swiyu Trust Protocol, ETSI TS 119 612, ETSI TS 119 602

Emerging standards that may become relevant over time include:

- Credential formats: W3C VC Data Model 2.0 / JSON-LD, ACDC
- Semantic / presentation standards: OCA 2.0
- Trust and federation standards: TRQP, OpenID Federation, EBSI issuer trust model

The Catalog will also rely on established implementation technologies such as OpenID Connect for IAM and standard API technologies. The detailed standards profile should be refined during the detailed architectural design phase.

7 Governance Model & Processes

The Catalog operates in a multi-stakeholder environment with different mandates, incentives and levels of authority. A federal authority, a cantonal police office, an insurance company and a garage network may all have a legitimate interest in the same credential type, but they will not assess relevance, urgency, quality, or acceptable change in the same way. Governance is therefore needed because publication, prioritization, lifecycle progression and formal approval cannot be left implicit.

Any viable model must handle four tensions:

- Legitimacy vs. speed: formal standardization creates legitimacy, but it is slow; community-driven publication is faster, but risks fragmentation.
- Openness vs. quality: broad participation increases coverage, but weak entries reduce trust in the Catalog as a whole.
- Neutrality vs. coordination: no single actor should control the Catalog, but a purely passive model will not produce convergence where convergence is needed.
- Transparency vs. process overhead: visible review and decision paths build trust, but excessive process slows contribution and adoption.

The governance model should therefore keep entry simple where possible and make authority explicit where necessary.

7.1 Catalog Governance Model Structure Design

The governance model can be structured across four governance dimensions. The table below makes this operational by showing, for each major artefact type, who may submit it, who curates it, who approves lifecycle transitions and who holds formal authority:

Artefact type	May submit	Curates	Approves lifecycle transition	Formal authority
Use-case entry	Ecosystem Contributor, Domain Steward	Catalog Curator	Domain Steward (within domain); Catalog Curator (cross-domain)	Standards Authority where normative
Demand / supply signal	Community Participant, Ecosystem Contributor	Catalog Curator (quality check)	Automatic on submission	n/a
Schema request	Ecosystem Contributor, Domain Steward, Data Steward	Data Steward + Catalog Curator	Domain Steward or Data Steward	Standards Authority for formal publication
Schema (AL2)	Data Steward, Domain Steward	Data Steward	Catalog Curator → I14Y / eCH handover	eCH / I14Y-linked body
Schema (AL3)	Data Steward, Domain Steward	Data Steward	Domain Steward (sector-governed); Catalog Curator for handover where relevant	Sector governance body or none
Trust framework reference	Domain Steward, Standards Authority	Catalog Curator	Standards Authority	Standards Authority
Governance reference	Domain Steward	Catalog Curator	Catalog Curator	Standards Authority where applicable

7.2 Key roles

The governance model relies on defined participation roles that determine contribution rights, decision rights and authority boundaries.

- **Community Participant:** may upvote demand, comment and propose use cases.
- **Ecosystem Contributor:** belongs to a verified organisation and may formally declare schema requests, submit use-case entries and vote on endorsements.
- **Domain Steward:** holds recognised responsibility for a business, policy, or sector domain and curates related use cases, credential needs, actors, demand signals, lifecycle status and governance references.
- **Data Steward:** holds recognised responsibility for semantic quality, reuse and alignment of claims, attributes, data models, schemas and metadata. In public-sector contexts this role is usually linked to I14Y, eCH, FOS, or another formal data-governance structure; in private-sector contexts it may sit with recognised sectoral bodies.

- **Catalog Curator:** validates entries, resolves editorial disputes, ensures consistency across domains and maintains the coherence of the Catalog as a shared discovery and coordination layer.
- **Standards Authority:** provides formal approval where required for schema publication, trust framework references, interoperability specifications, or other normative artefacts. Depending on the subject matter, this role may be exercised by eCH, I14Y-related bodies, DPSS, a sectoral governance body, or another recognised authority.

The user archetypes described earlier do not replace this governance model. They complement it by translating formal authority into concrete interaction patterns and interface needs.

7.3 Governance Scope and Institutional Fit

The governance model must provide legitimacy, usability and scalability across AL1, AL2 and AL3 without introducing unnecessary centralisation. The Catalog's own roles, processes and interfaces must fit credibly within a landscape where governance conditions differ across ambition levels.

Three aspects require particular attention:

- **Governance anchoring for AL3.** In AL2, Domain Steward authority is anchored in an institutional mandate - a federal office, cantonal authority, or recognised public-sector body. Lifecycle transitions for AL2 artefacts require Catalog Curator sign-off and the formal publication path runs via I14Y or eCH. For AL3, one proposed approach is that Domain Steward authority be held by a recognised ecosystem orchestrator for that sector, confirmed through a light-weight participation agreement with the Catalog. Lifecycle transitions are lighter and sector-governed; formal publication is optional. Where no recognised orchestrator exists for a private-sector domain, the Catalog Curator holds provisional Domain Steward authority until one is confirmed. This distinction must be reflected in both the governance model and the operating model.
- **Linkage to existing structures.** Clear interfaces are needed to eCH, I14Y and the swiyu trust infrastructure, including handover of maturing artefacts into formal maintenance where appropriate. The Catalog should not duplicate these processes but connect to them at defined handover points.
- **Dispute resolution.** Where two Domain Stewards contest ownership of a cross-sector use case, or where a Standards Authority and an Ecosystem Contributor disagree on a schema lifecycle transition, a recommended escalation path would be: Catalog Curator mediation first; if unresolved within a defined period, the matter is referred to the ecosystem orchestrator secretariat responsible for the relevant ambition level; final authority rests with the Standards Authority for normative artefacts and with the Catalog Curator for non-normative artefacts. This escalation path must be documented in the governance model published alongside the Catalog.

The Catalog must be governed in a way that remains institutionally compatible with the broader Swiss ecosystem and extendable beyond it.

7.4 Processes Surrounding the Ecosystem Discovery Catalog

The Catalog creates value only when embedded in clear and repeatable surrounding processes: intake, qualification, demand signalling, curation, lifecycle progression, governance visibility and handover into formal pathways. Some of these processes already exist in more formalized form, particularly in the public-sector context. The Catalog should therefore connect to existing institutional workflows where they exist, while also supporting lighter and more plural coordination paths where they do not. The detailed process design is described in the Operating Model section later in this document.

7.5 Adjacent Ecosystem Governance-Orchestration Function

The Catalog exposes use cases, credentials, schemas, actors, lifecycle states and governance references while formal authority remains with the relevant governance bodies. Especially in AL3 and cross-sector contexts, an adjacent governance-orchestration function may be needed to convene actors and structure governance questions that the Catalog itself should not decide. The Catalog can expose such artefacts, their status and responsible roles without determining their legal or normative validity.

8 Implementation & Operating Model Options

The Ecosystem Discovery Catalog needs an operating model that assigns functions, responsibilities, interfaces and maturation paths.

The operating model should be neutral, credible to public and private actors, sustainable and lean enough to support iteration across AL1, AL2 and AL3. The basic premise is broad participation and lightweight signalling by default, with stronger curation and governance as artefacts mature or become institutionally relevant.

The Catalog combines a digital product, a curation function, a coordination mechanism and interfaces into formal standardisation, publication and trust-governance processes.

8.1 Core operating activities

A viable operating model should cover the following core functions:

Use-case intake and structuring: Capturing use cases in a consistent format, including involved actors, required credentials, relevant governance context, expected value and reuse potential.

Credential and schema request management: Recording existing, emerging and missing credentials and schemas and turning unmet needs into visible demand signals for potential issuers, data stewards and standardisation bodies.

Curation and quality assurance: Reviewing entries for completeness, consistency, clarity, duplication and semantic coherence, without turning early contribution into a heavy approval process.

Supply-demand coordination: Connecting verifier demand, issuer capability and domain stewardship so that relevant actors can recognise opportunities, assess mandate fit and coordinate next steps.

Governance and trust-reference maintenance: Maintaining references to applicable trust frameworks, assurance context, verifier governance, issuer recognition, legal basis and other relevant governance conditions.

Federation with external systems: Linking to authoritative or adjacent systems such as I14Y, swiyu registries, OCA-related resources, eCH standards, sectoral repositories and other relevant catalogs.

Lifecycle and maturity visibility: Making visible whether an artefact is an early signal, curated candidate, qualified reference, implemented credential, deprecated object, or formally maintained standard.

Preparation for formal maintenance or standardisation: Preparing mature artefacts for hand-over into formal processes where required, for example through I14Y, eCH, DPSS structures, sectoral governance bodies or other recognised authorities.

Implementation feedback and continuous improvement: Capturing lessons from pilots, integrations and operational use and using them to improve metadata, processes, governance references and future Catalog functionality.

Communication and ecosystem enablement

Translate use cases, opportunities, governance context and progress signals into actionable information for technical and non-technical stakeholders. Ecosystem orchestrators should play an active role in onboarding, alignment and sector engagement.

For segregation of duties, ecosystem curation, technical platform operation and formal approval should remain institutionally distinct.

8.2 Initial operating capacity

Initial operation requires a small but clearly defined delivery team. For MVP planning, this implies: a product / operating lead to steer scope, priorities and institutional alignment; a Catalog curator / ecosystem coordinator to manage intake, curation, coordination and lifecycle progression; a semantic / federation lead to support schema-related structuring, reuse and qualification; technical operations and platform support for system reliability, releases, access management and integrations; and part-time governance / legal support to help maintain governance references and prepare handover into formal processes where required. A realistic initial staffing model would therefore be in the order of 2.5 to 4 FTE for a bare-minimum MVP-stage setup and

closer to 4 to 5 FTE if active AL3 onboarding, sector coordination and technical operations are not largely covered by existing institutions.

One possible MVP-stage baseline, used here as an indicative planning assumption, could be:

Role	Indicative FTE	Core responsibility
Product / Operating Lead	0.6–0.8	Overall operational ownership, scope and priority steering, institutional alignment and coordination of delivery across product, curation and governance
Catalog Curator / Ecosystem Coordinator	0.8–1.0	Use-case intake, curation, coordination of supply and demand, lifecycle progression and day-to-day ecosystem-facing operation
Semantic / Federation Lead	0.4–0.8	Schema-related structuring, liaison to data stewards and other collaborative groups, reuse analysis, semantic alignment, qualification support and preparation of artefacts for formal maintenance or standardisation where needed.
Technical Operations / Integration Support	0.4–0.8	Platform reliability, release support, access management, technical maintenance and integration with external systems such as I14Y and trust infrastructure components
Governance / Legal / Standards Liaison	0.2–0.4	Maintenance of governance references, support for governance fit and compliance questions and liaison with standards, legal and formal approval processes
Role	Indicative FTE	Core responsibility
Product / Operating Lead	0.6–0.8	Overall operational ownership, scope and priority steering, institutional alignment and coordination of delivery across product, curation and governance

The exact allocation depends in particular on the scope of the MVP itself; how far technical operation is centralised, the initial governance established, how much semantic work can build on existing I14Y or eCH processes and to what extent ecosystem coordination is shared across public and private orchestrator functions, etc.

8.3 Institutional setup and neutrality

The most credible model is layered: community contribution, domain and data stewardship, ambition-level orchestration, whole-catalog coordination and stable public-interest technical operation.

Content is contributed by issuers, verifiers, domain actors and other ecosystem participants across AL2 and AL3. Curation authority sits with the relevant domain and data stewards, who hold the closest mandate and responsibility for quality within their scope. DVS coordinates this for AL2. For AL3, recognised sector bodies hold curation authority where they exist; otherwise, the designated neutral cross-sector AL3 orchestrator acts as supporting curator until sector stewardship is established.

An independent whole-catalog Curator function ensures coherence across ambition levels, handles cross-AL use cases, mediates boundary disputes and maintains the integrity of the Catalog as a shared discovery layer. It does not override domain curation decisions within their own

scope. This function may sit with one of the orchestrators or be operated as a shared function between them.

Two structural rules follow: (1) The AL3 orchestrator should be a neutral cross-sector body; if it is affiliated with a single sector, its role as gap-filler across AL3 communities would be difficult to sustain credibly; (2) technical platform operation sits with a stable public operator, independent of the orchestrators and the Curator.

8.4 Role of ecosystem orchestrators

Ecosystem orchestrators translate the Catalog's coordination function into practical ecosystem activation. For AL2, this means connecting the Catalog to public-sector priorities, I14Y/eCH processes and DVS structures. For AL3, it means surfacing market demand, identifying cross-sector use cases, convening sectoral actors and translating business needs into structured credential and schema requests. The designation and precise mandate of the AL3 orchestration function should be confirmed during Phase 1.

8.5 Technical operation

Technical operation should be assessed for placement with a stable public-interest operator such as FOITT. This includes hosting, security, access control, reliability, release management and operation of the core system.

8.6 Alignment and federated operating logic

The Catalog should build on existing Swiss infrastructure and avoid creating a parallel institutional stack. Authoritative artefacts should remain in their canonical systems where such systems exist, especially I14Y, swiyu registries, eCH processes and sectoral repositories. The Catalog should store Catalog-native metadata, early-stage use cases, demand signals and exploratory artefacts only where no authoritative external home exists. This allows the Catalog to remain a discovery and coordination layer rather than becoming a monolithic repository.

8.7 Funding model

The study suggests that baseline development and operation of the Catalog should be funded primarily by the public sector, since it performs a foundational coordination function for the trust infrastructure and serves an important public-interest role for ordinary swiyu users, particularly around orientation, value discovery and transparency in relation to personal data use.

Long-term funding should reflect the public-private nature of the ecosystem. The Catalog should therefore be treated as part of the enabling layer of Switzerland's emerging public-private digital public infrastructure.

A plausible funding model can be structured across three tiers:

Funding tier	Primary funder	Covers	Governance condition
Baseline	Federal / cantonal public sector	Platform development, technical operation, baseline curation capacity,	DVS mandate; standard public procurement rules

		AL2 integration with public infrastructure	
Activation	PPP (DVS + sector bodies / orchestrators)	Sector-specific curation, AL3 onboarding, cross-sector use-case development, ecosystem activation	Participating body must have a transparent, purpose-bound governance instrument (statutes or equivalent)
Premium track	Private sector direct	Higher-effort AL3 implementations, sector acceleration, advanced integration	Contract-based, ring-fenced; no editorial influence on Catalog governance

This tiering preserves the public-interest foundation while enabling credible private-sector co-investment. The premium track in particular must be designed so that financial contribution cannot be converted into governance influence over curation, lifecycle decisions, or trust references. In that sense, the funding model is itself part of the infrastructure: without a credible public-private funding logic, the broader public-private DPI ambition of swiyu will remain difficult to operationalise.

8.8 Scalability and onboarding

The operating model must scale beyond a small group of early participants. It should therefore support low-threshold participation for signaling, contribution and early progression, while reserving stronger controls for lifecycle advancement, formal declarations and authoritative maintenance. This is especially relevant for AL3, where distributed issuer landscapes require onboarding and coordination models that work at network level rather than institution by institution. AL3 sectoral ecosystems onboard via the designated AL3 orchestrator, which holds umbrella Domain Steward responsibility at the AL3 level and is complemented by sector-specific stewardship to recognised sectoral bodies where they exist. The AL3 orchestrator signs a participation agreement with the Catalog operator and commits to a minimum curation cadence. Any issuer, verifier, or other ecosystem participant can contribute directly to the Catalog. Sector bodies and the AL3 orchestrator provide coordination and curation support, but do not gate participation. This keeps contribution friction low while ensuring that each domain has a named responsible party for content quality and lifecycle progression.

8.9 Process design

The surrounding processes must be clear, lightweight and repeatable: use-case intake, supply-demand signalling, governance visibility, lifecycle progression and formal handover where required.


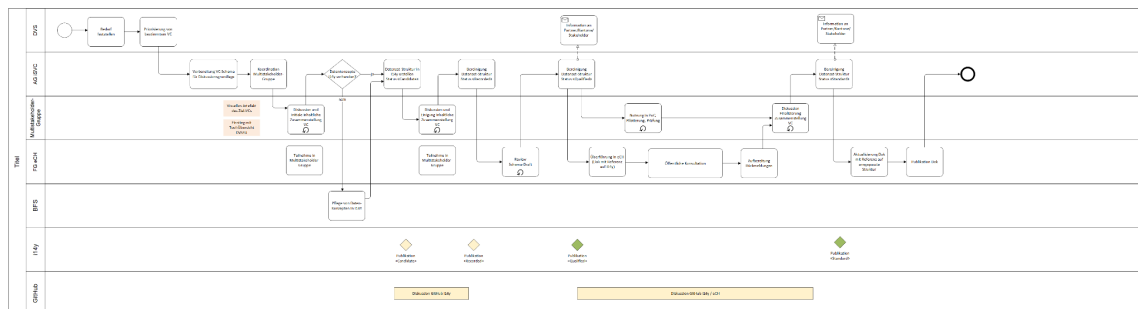


Fig. Reference process for I14Y- and eCH-linked schema maturation and publication

The diagram in the figure above illustrates an existing reference process for schema maturation in the Swiss interoperability context. It shows how an identified need can move through multi-stakeholder discussion, I14Y publication states, review, consultation and eventual publication or standardization. The figure is especially relevant for more formalized public-sector pathways. For schema maturation and publication, the Catalog should rely and reuse established I14Y processes and tooling wherever possible and not recreate the process. Where this proves insufficient for early-stage ecosystem needs, limited Catalog-level curation functionality may still be required through integration with existing I14Y capabilities. This should be clarified in the next phase.

The process design must connect to formal pathways where they exist and support lighter coordination paths where they do not. This is especially relevant for AL3-related content and for private-sector schemas that may not follow the same publication and maintenance route as public-sector artifacts. Several surrounding processes therefore require explicit definition.

8.10 Progression

Progression within the Ecosystem Discovery Catalog should follow explicit but lightweight maturity stages. Use cases, credential needs, demand signals and referenced artefacts should be able to evolve from early signals into curated, qualified and, where appropriate, formally maintained artefacts.

A practical progression path includes:

- **Initial signal:** a proposed use case, credential need or schema request is submitted.
- **Curated candidate:** the entry is reviewed for clarity, completeness and duplication.
- **Qualified reference:** demand, governance context, semantic alignment and implementation relevance are sufficiently clear.
- **Implemented / active:** the artefact is used in one or more operational or pilot contexts.
- **Formally maintained:** the artefact is maintained in a recognised external pathway, such as I14Y, eCH or a sectoral repository and referenced by the Catalog.

- **Deprecated / superseded:** the artefact remains visible for traceability but is no longer recommended for new implementations.

Transitions between stages should be guided by transparent criteria such as demonstrated demand, mandate fit, implementation readiness, semantic consistency, governance clarity and re-use potential. This allows artefacts to mature without premature formalisation or unnecessary friction.

9 Recommendation and Phased Approach

9.1 Overall recommendation

We recommend proceeding with the next phase of the Ecosystem Discovery Catalog.

The next step should be a focused Decision and Procurement Readiness Phase. Its purpose is to translate the findings of this pre-study into an implementation-ready basis: a defined MVP scope, draft technical architecture, curation and governance model, operating model, priority use cases and decision-ready procurement and funding basis.

The Catalog should advance in synergy with DVS4U and a possible future Credential Hub. These capabilities are complementary and should be developed accordingly.

The Catalog should therefore be treated as a shared ecosystem capability. Its first implementation should focus on practical coordination value: making relevant use cases, credentials, actors, dependencies, maturity states and governance references easier to discover and act upon.

The first build phase should stay focused on practical coordination value. Full schema governance, runtime policy execution and comprehensive conformance testing belong on the roadmap, not in the MVP.

9.2 Strategic execution principles

Four principles should guide the next phase: develop product and operating model together; treat federation as both a technical and institutional design principle; validate through concrete multi-actor journeys; and enable broad participation at the input layer while applying stronger controls to maturity progression, formal publication and authoritative maintenance.

9.3 Functional priorities for the MVP

The MVP should focus on discovery, coordination, federation and lifecycle visibility. Priority should be given to:

- structured use-case entries
- visibility of existing, emerging and missing credentials
- issuer and verifier discovery
- demand and supply signalling
- lifecycle and maturity visibility

- federation of governance, trust and schema references
- structured opportunity signals for potential issuers
- explicit visibility of credential reuse and deliberate exclusion at use-case level

Open credential needs or semantic-gap requests should contain enough context for potential issuers to judge whether action is justified. This includes expected relying parties, indicative scale, affected use cases, mandate fit, likely value, governance context and a plausible onboarding path.

The Catalog should also distinguish clearly between two different lifecycles. Use cases mature through demand validation, stakeholder alignment, pilot activity and operational adoption. Schemas mature through semantic review, publication, reuse, versioning and possible formal standardisation. The Catalog should make both lifecycles visible and connect them where dependencies exist.

A further distinction is required between domain stewardship and semantic stewardship. Domain stewardship concerns the business, policy, sectoral and adoption context of use cases and credentials. Semantic stewardship concerns the quality, reuse and alignment of data elements, claims, schemas and metadata. The Catalog should connect both responsibilities without collapsing them into one role.

9.4 Phased Approach

Execution should proceed through three consecutive phases, each producing a usable outcome and a decision basis for the next step. The work must cover technical development, governance, curation, operating model design and ecosystem activation.

The MVP represents a deliberately limited, pilot-oriented implementation focused on proving coordination and discovery value. Version 1 represents the first stable, onboarding-ready release with expanded coverage, refined governance and operational scalability.

The initial roadmap:



9.4.1 Phase 1: Decision and Procurement Readiness for the Ecosystem Discovery Catalog MVP

The goal of this phase is to define the decision-ready scope of the first running Minimum Viable Product (MVP) for both the technical system and the curation and governance mechanism and to prepare the basis for a subsequent implementation decision and procurement.

Duration: 3 Months

Deliverables:

- Draft requirement specifications for the MVP Web Portal
- Draft technical architecture and deployment options, including key technical components, standards, interfaces and data assumptions
- Draft curation design for the MVP
- Draft governance design for the MVP
- Initial operational model, including staffing assumptions, responsibilities and operating options
- Assessment of operating model options
- Refined MVP scope
- Indicative forward view for Version 1 of the Catalog
- Refined implementation effort estimate for the MVP
- Identification of priority AL2 credentials and use cases
- Identification of 1–2 priority AL3 sectoral ecosystems
- Decision-ready procurement basis for Phase 1.5, including draft tender documentation, evaluation criteria and contracting assumptions

9.4.2 Phase 1.5: Procurement

The goal of this phase is to select the implementation partner and establish the contractual and operational basis for MVP development.

Duration: 3–4 months (may overlap partially with the tail of Phase 1)

Activities:

- Publish tender documentation based on Phase 1 deliverables, including MVP scope, draft architecture, governance and curation model and operating model
- Run procurement in accordance with Swiss public procurement law (BöB / IVöB as applicable)
- Evaluate bids against technical, operational and governance criteria defined in Phase 1
- Award contract and establish onboarding basis for the implementation partner
- Confirm operating model assignments: technical operator, AL2 orchestrator, AL3 orchestrator and Catalog Curator function
- Finalise integration assumptions with I14Y and swiyu operators ahead of Phase 2 kick-off

Outcome: signed implementation contract, confirmed operating model and a ready-to-start Phase 2.

9.4.3 Phase 2: MVP Implementation and Pilot Activation

The goal of this phase is to deliver a working, populated and governable MVP that demonstrates coordination value across at least three AL2 use cases and two AL3 demand signals. This includes the finalisation of the MVP specification with the selected provider or implementation team before full implementation. Success is defined by: a deployed web portal with core discovery and signalling functions; a live federation link to I14Y and the swiyu Trust Registry; a functioning governance and curation model with named role holders; and a pilot activation package validated with at least one Domain Steward and one Implementer. Activities outside these outcomes should be deferred.

Duration: 9 Months

Deliverables:

Design and specification

- Final requirement specifications for the MVP Web Portal
- Final mockups of the use-case UI
- Refined forward scope for Version 1 of the Catalog

Technical implementation

- Updated technical architecture, including deployment setup and key technical components
- Implemented data structures, metadata model and relevant standards / interface mappings
- Deployed MVP Web Portal

Operations, governance and curation

- Implementation of the Phase 1 operating model assumptions for MVP pilot operation
- Named role holders for MVP curation, governance coordination, technical operation and pilot support
- Tested curation and governance workflows under MVP pilot conditions
- Lessons learned and refinement inputs for the Version 1 operating model

Ecosystem rollout and pilot readiness

- Extend priority AL2 credentials and use cases to 5+
- Identify and structure priority use cases for selected engaged players in AL3
- Identify the next 2 priority sectoral ecosystems for AL3, bringing the total to 4
- Initial Catalog population for the agreed MVP baseline
- Pilot activation package, including acceptance criteria, handover and launch readiness. MVP acceptance is defined by the following minimum criteria:

- a. ≥ 10 structured use-case entries, each with full actor, credential and governance context;
 - b. ≥ 3 AL2 credentials with complete lifecycle metadata and a live federation link to I14Y;
 - c. ≥ 2 AL3 demand signals with at least one matched issuer response;
 - d. federation link to swiyu Trust Registry operational for issuer status lookup;
 - e. governance dashboard showing current lifecycle state and role assignments for all catalog entries;
 - f. at least one end-to-end user journey validated with a Domain Steward and an Implementer archetype representative.
- Initial monitoring, feedback and iteration setup for pilot operation
 - Refined effort estimation for subsequent Version 1 expansion and next releases

9.4.4 Phase 3: Version 1 Development, Deployment and Early Onboarding

This phase moves from a validated MVP to a stable Version 1 ready for broader onboarding. Success means: AL2 use case coverage expanded to 10+ entries; at least four AL3 sectoral ecosystems active with named Domain Stewards; governance and curation processes refined based on MVP learnings; and an onboarding model in place that allows new AL3 participants to join without requiring case-by-case coordination. Deliverables not contributing to these outcomes should be deferred to subsequent releases.

Duration: 7-9 Months

Deliverables:

- Deployed Version 1 release of the Catalog
- Updated technical architecture
- Refined operational model, version 1.x
- Refined governance model, version 1.x
- Refined curation design and curation workflow, version 1.x
- Onboarding approach, onboarding workflow and onboarding materials for first AL3 users
- Initial support and operating procedures for early Version 1 operation
- Extend priority AL2 credentials and use cases to 10+
- Identify the next set of priority use cases for active AL3 players, bringing the total to 5+
- Identify the next 2 priority sectoral ecosystems for AL3, bringing the total to 6
- Expanded Catalog population for the agreed Version 1 baseline
- Initial monitoring, feedback and iteration setup for Version 1 operations
- Updated roadmap for the next phases of Catalog development

Once Version 1 is delivered, the project can re-evaluate priorities of the system, reflect on the lessons learned and revisit the updated Roadmap.

Detailed estimations of the required FTE and other resources can be provided on request.

9.5 Future Considerations: Interoperability and Later-Stage Capabilities

This paper does not attempt a full interoperability blueprint. However, several interoperability-related capabilities should already be recognized as roadmap items because they will become important once the Catalog moves beyond discovery and into ecosystem coordination at scale.

9.5.1 Conformance testing

After the MVP has proven its basic discovery and coordination value, the Catalog should progressively address interoperability and conformance aspects. Schemas and implementations should be testable against relevant standards and governance requirements so that issuers and verifiers can assess whether their implementations are compatible with the broader ecosystem. This should be treated as a later-phase capability, not as an MVP requirement.

9.5.2 Dynamic, context-driven data exchange conditions

Another roadmap topic is the structured expression of context-dependent data exchange conditions. Over time, the Catalog should help surface not only static schema definitions, but also the conditions under which data may be requested, disclosed, relied upon, or restricted in specific contexts. This is particularly relevant where exchange conditions depend on role, purpose, regulatory basis, or sectoral governance logic.

9.5.3 Runtime Binding (Making Governance User-Comprehendable)

Runtime binding should become a later-stage capability, linking catalogued governance conditions more directly to live credential interactions. This would make elements such as trust framework references, issuer requirements, status constraints, disclosure limitations and applicable policy rules available in forms that systems can process and users can understand at the point of exchange.

This would reduce ambiguity for implementers and verifiers, improve transparency for holders and help prevent wallet interactions from becoming formal but poorly understood consent flows. Runtime binding should therefore not be treated as an MVP feature, but as a later capability building on the Catalog's initial discovery and coordination role. It would leverage standardised protocol-level mechanisms that make governance information verifiable and comprehensible at the point of interaction, including across interoperable third-party wallets.

9.6 Closing recommendation

We recommend moving forward with the development of the Ecosystem Discovery Catalog. The Catalog can become a foundational resource for verifiable-credentials-based digitalisation and service design in Switzerland.

The next step is to launch the Decision and Procurement Readiness Phase as the bridge from pre-study to implementation.

This phase should produce the concrete MVP decision basis: scope, architecture, operating model, governance and curation setup, integration assumptions, priority use cases, procurement logic and funding pathway. It should be treated as implementation preparation, not as another conceptual study.

A readiness phase reduces delivery risk while preserving momentum. It creates a path toward a publicly anchored MVP that is focused, procurable and able to demonstrate coordination value across AL2 and selected AL3 contexts.

We therefore recommend funding and executing the Decision and Procurement Readiness Phase as soon as possible, in view of the planned swiyu go-live in H1 2027.