Building a Trusted Environment

Event-based Attribute Assurance

Alpha Project Report

Written by Ben Helps, Factern Ltd

October 2020

FACTERN

OPEN IDENTITY EXCHANGE
White Paper

PROJECT PARTICIPANTS

The Open Identity Exchange (OIX) is a non-profit, technology agnostic, collaborative cross sector membership organisation with the purpose of accelerating the adoption of digital identity services based on open standards. OIX's broad membership and independent nature have seen it develop a significant body of digital identity research, and it is a significant influencer working towards the development of a digital identity market.

The following individuals and organisations participated in the OIX Alpha Project that formed the basis of this Report:

- Factern: Ben Helps, Robbie Fraser
- GDS: Chris Allgrove
- Food Standards Agency: Sid Kalita
- Future Borders: Ross McDonald, Stephen Cowx, Tom Robinson, Paul Evans, Jake Luscombe
- HMRC: Nick Davies, Wayne Robinson
- Idemia: David Rennie, Richard Thompson
- Inrupt: Paul Worrall
- Oxford Semantic Technologies: Peter Crocker
- University of Lincoln: Steve Brewer

OIX and the authors would like to thank all the Alpha Project team members for their considerable contribution to the work.

CONTENTS

1. BACKGROUND

During the summer of 2019, the Open Identity Exchange (OIX) ran a Discovery Project that sought to explore event-based data assurance.

During the course of day-to-day business, entities interact with other entities (such as their customers, staff, suppliers, etc.) in ways that can be used to assure the quality of the data associated with the latter. The Discovery Project's incoming hypothesis was that managed access to a digital record of such interactions would constitute a valuable service to others.

The findings, conclusions and recommendations from the Discovery Project were published as an OIX Whitepaper "Building a Trusted Environment: Event-based Attribute Assurance" which can be found on the OIX website.

The Whitepaper recommended the establishment of an OIX Alpha Project with the objective of agreeing a set of common requirements to support the provision, exchange and consumption of Events. The common requirements would be adopted by collaborating entities willing to share the value of the processes, activities and claims that the Events represent, using them to meet their own – potentially different – data assurance needs.

Following the completion of the Discovery Project, several organisations put forward resources to deliver an OIX Alpha Project to meet this brief, with the intention of taking the work forward into a subsequent Beta Project, where the generalised standards and rules of engagement developed in the Alpha could be deployed in the context of a specific use case that would be of value to those organisations, to act as a reference ecosystem deployment.

The OIX Alpha Project was launched towards the end of January 2020 and ran until the end of March. The target content was discussed over the course of nine workshops, the last of which took place on 25th March. A summary of each workshop can be found in the Appendix to this report.

Over the course of the summer, a smaller group of project participants then set out to detail the "Ecosystem Toolkit", which comprises both a Technical Specification and – importantly – Rules of Engagement to provide a configurable framework for collaboration.

2. AN INTRODUCTION TO THE ECOSYSTEM TOOLKIT

PURPOSE OF THE ECOSYSTEM TOOLKIT

There are many situations and contexts in which entities that share a common domain of interest could create value by collaborating with one another. Identity and eligibility assurance is one such use case, in which the benefits of collaboration are clear cut[1].

The primary objective of the Alpha Project was to develop a generalised and repeatable toolkit that can be easily deployed by *any* given entity to facilitate collaboration with other *any* other entity over *any* given domain of interest.

This is referred to as the Ecosystem Toolkit.

The ideas contained within the Ecosystem Toolkit aim to connect expertise across collaborating entities without unnecessary friction. The Ecosystem Toolkit aims to remove barriers to collective action that prevent the emergence of collaborative ecosystems, reducing the incremental effort required either to set up a new ecosystem or to extend an existing one. It also allows for a greater level of interoperability *between* ecosystems.

We envisage a future state in which the Ecosystem Toolkit helps significantly more collaborative ecosystems to emerge – and thereby create value.

[1] The Discovery Project sets out how Events can be used to benefit data assurance, and why different entities might participate in an ecosystem that facilitates the exchange of Events that assure data quality. We do not seek to repeat these arguments here. The Alpha Project assumes that – in many different situations and contexts – good reasons exist for otherwise separate entities to collaborate either bilaterally or multilaterally and to exchange data about the domain of interest that they have in common, including a rich layer of metadata, such that the latter accretes over time into an increasingly strong basis for trust.

ECOSYSTEM TOOLKIT RELATIVE TO THE OIX GUIDE TO TRUST FRAMEWORKS

The Alpha Project which developed the Ecosystem Toolkit ran in parallel to the development of the OIX Guide to Trust Frameworks and Interoperability.

The OIX Guide articulates a super-set of the areas of concern to be addressed when implementing an individual Trusted Digital Identity ecosystem (see Figure 1 below), while acknowledging that each framework may implement a sub-set to meet its specific needs.

Figure 1: OIX Trust Framework

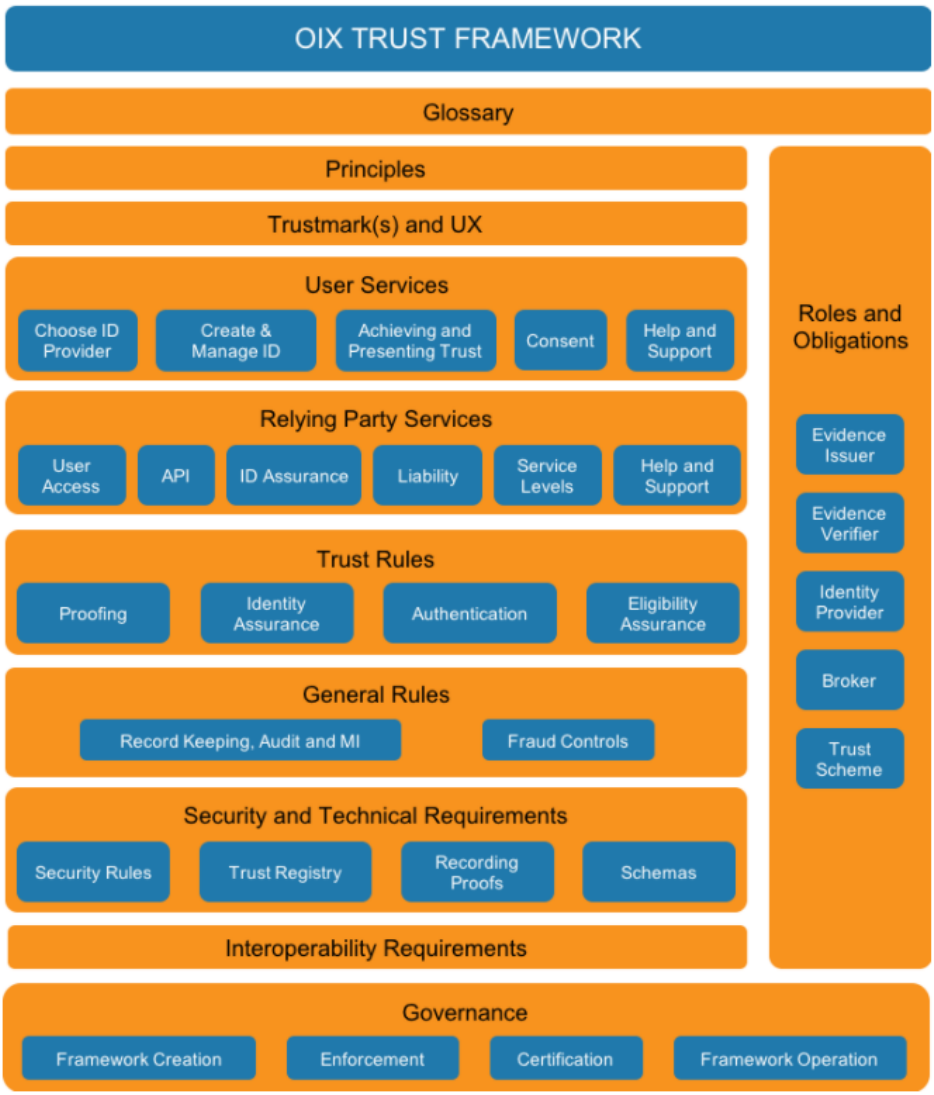

The Ecosystem Toolkit specifies an approach that can be used to implement a sub-set of these areas of concern. While the Ecosystem Toolkit introduces a number of requirements and constraints, it does not mandate particular implementation choices.

The OIX Guide might therefore be considered as a set of *business requirements*: it articulates the different areas of concern to be addressed as part of a Trust Framework.

In the same way, the Ecosystem Toolkit can be seen as setting out an *architectural style* that can be used to implement a Trust Framework, constraining some design decisions in order to elicit a number of beneficial qualities, as described in Chapter 4 “Design Goals”.

Neither the OIX Guide nor the Ecosystem Toolkit should be confused with an *instance* of a trust framework or scheme that has been deployed to deliver a particular use case, and which has made its own specific implementation choices.

COMPONENTS OF THE ECOSYSTEM TOOLKIT

Many initiatives set out to establish technical standards[2]. However, a key feature of the Ecosystem Toolkit is that it seeks to specify an approach for tackling those elements of collaboration and trust that go *beyond* the deployment of technical standards.

The heterogeneity of the different situations and contexts which might benefit from collaboration inevitably means that the legal, operational and behavioural aspects require as much – if not more – effort to specify and configure as any technical standard.

The Ecosystem Toolkit therefore comprises two major artefacts:

– **RULES OF ENGAGEMENT** to facilitate the *agreement* required for scalable, interoperable and extensible collaboration between participants within a given ecosystem.

[2] Examples include OpenID Connect for Identity Assurance and the W3C Verifiable Credentials Data Model

- A **TECHNICAL SPECIFICATION**, offering a generic approach based on open standards, to implement the *execution* of scalable, interoperable and extensible collaboration over the Events that – for example – might be used to record proof of identity or eligibility.

Broadly speaking, the Rules of Engagement intersect with the top half of the OIX Trust Framework as described in Figure 1 (i.e. Trust Rules, Relying Party Services, User Services) while the Technical Specification intersects with the bottom half (i.e. Interoperability Requirements, Security and Technical Requirements, General Rules).

Like the OIX Trust Framework, the Rules of Engagement comprises of different areas of concern, referred to in the Rules of Engagement itself as components: legal basis (e.g. consent), liability framework, assurance levels, service levels and issue resolution.

Each of these components is configurable (e.g. different levels of liability can be set), and the Rules of Engagement includes a default configuration for a minimum set of components which provides a basic framework agreement for collaboration between entities.

Of course, this basic framework will not meet the needs of entities looking to collaborate in situations that require relatively high levels of trust. The Rules of Engagement allow for new components to be added and for different configurations of each component.

Importantly, specific instances (e.g. a defined level of liability, or a given issue resolution mechanism) of any component can be easily referenced by the machines (known as **EXPERT SYSTEMS**) that implement the Technical Specification.

The intention is for easily repeatable patterns to emerge, as ecosystems compose and configure components to meet their needs for specific levels of assurance and trust. These patterns can be exposed in a machine-readable way to other entities looking to participate in the ecosystem, or by other ecosystems operating in a different situation or context.

EVENT SCHEMA

One of the most important constraints imposed by the Technical Specification is the requirement to implement a universal and standardised Event schema.

The Event schema is divided into Header and Body, where the Header captures the provenance of the Event and the Body expresses its content in the form of an a RDF[3] triple.

All Events therefore take the form of an attestation ("X says that Y is true"), whose provenance provides a basis for trust: at its simplest, an Event Consumer can decide how much confidence to place in the content of the Event based on the entity attesting to it[4].

This standardised Event schema also means that:

- The information exchanged between Expert Systems can be about anything – as a result, ecosystems are not restricted to any particular domain of interest
- A universal API can be used to move information between Expert Systems – there is no need to re-work or extend the client which consumes the API as circumstances change

These properties hugely increase interoperability (both internally within an ecosystem and externally between ecosystems) and extensibility of solutions that are deployed in a way that conforms to the Ecosystem Toolkit. See Chapter 4 "Design Goals" for more detail.

ROLES WITHIN AN ECOSYSTEM

The standardised Event schema makes three roles explicit:

- Event Producer: the entity that generated the assertion being made
- External Feed: the external source used to input the assertion into the Expert System
- Rights Owner: any legal entity or entities (other than the Event Producer) that exercise rights over the assertion

[3] Resource Description Framework. An RDF triple takes the form Subject-Predicate-Object

[4] Information in this form is often referred to as a credential or a verifiable credential. Credentials have become integral to the self-sovereign identity movement, as the holder of a credential is able to present information about themselves not just as attributes that they *claim* to be true but as attributes that an identifiable third party has *attested* to. Although digital signing does not form a part of the Technical Specification within the Ecosystem Toolkit, it is an obvious potential future extension.

Of these, the Technical Specification that the role of Event Producer is systematically captured as an integral part of the Event Header. The Technical Specification optionally allows for the Header to identify entities that play the roles of External Feed or Rights Owner, as the inclusion of these is context specific – see the example scenarios cited below.

Three other roles are implicit in the architectural style being advocated:

- Owning Entity: the legal entity on whose behalf an Expert System acts
- Event Provider: the Expert System managing access to an Event
- Event Consumer: the Expert System securing access to an Event

The interaction between these roles is summarised in Figure 2 below.

Figure 2: Roles within an Ecosystem

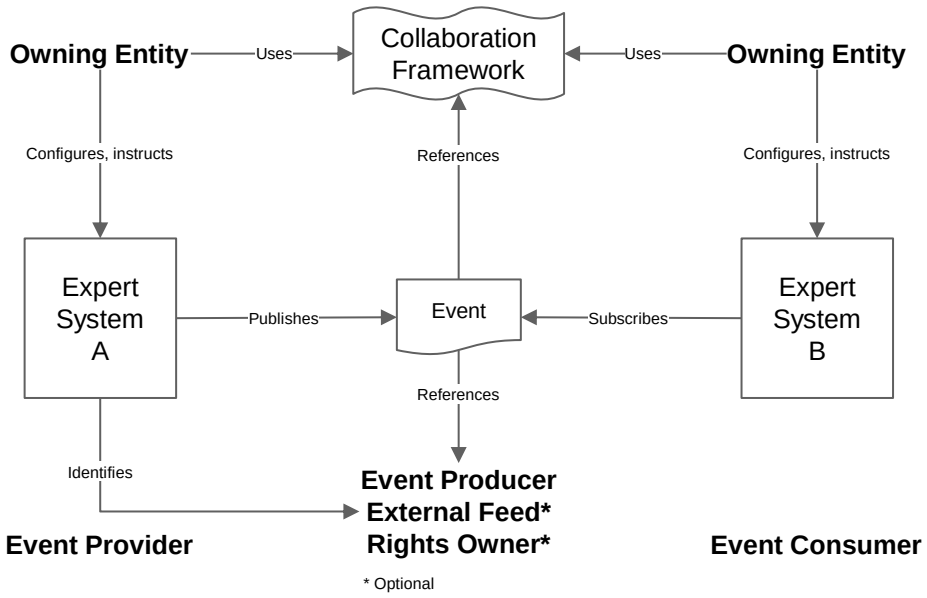

Example Scenario 1

A private citizen uses an Expert System under their control both to assert a piece of information to configure entitlements such that an institution (such as a Government department) can access the information.

In this scenario, the private citizen plays the role of Event Producer, and the two Expert Systems play the roles of Event Provider and Consumer respectively. The exchange of information is governed by a collaboration framework that must be identified and agreed to by both the private citizen and the Government Department as part of the interaction.

Whether the information exchanged is about the private citizen themselves (i.e. a claim) or about someone/something else is a function of the Event content (e.g. who is identified as the Subject?). It is up to the Event Consumer to interpret the content given its provenance – see Appendix A (Roles and Obligations vs. Provenance) for more detail.

Example Scenario 2

A commercial data provider adds value by aggregating data about a UK legal entity from multiple public data sources, including Companies House. The packaged information is re-sold to clients of the commercial data provider.

In this scenario, the commercial data provider is the Owning Entity for the Expert System which plays the role of Event Provider, and identifies Companies House as Event Producer. If needed, they can also identify the service used to source the data from Companies House as the External Feed. The UK legal entity is the Subject of the Event content.

As in Scenario 1, it is up to the client of the commercial data provider (as the Owning Entity of the Event Consumers) to agree to the collaboration framework (such as a commercial agreement) which governs the exchange.

Example Scenario 3

A commercial data provider adds value by undertaking an investigation to assure the quality of data about a UK legal entity sourced from Companies House. The assurance is sold alongside the data to clients of the commercial data provider.

In this scenario, the commercial data provider plays the role of Event Producer. However, instead of the UK legal entity being the Subject, other Events are (i.e. those containing the data from Companies House).

The Events containing the assurance information can be passed to the client alongside the Events containing the data from Companies House. However, the collaboration framework which they reference may be different, to reflect a different service, with a different level of liability, value exchange, etc.

Example Scenario 4

A commercial service provider adds value by authenticating that a private citizen is the holder and subject of a physical Passport document. The service uses this link to assure the quality of information about the citizen's name, date of birth and nationality. The assurance is sold to a Government department with which the private citizen is interacting, and which is seeking to establish the identity of the private citizen.

In this scenario, the commercial service provider will pass information in the form of Events constructed in a similar way to those illustrated in either Scenario 2 and 3 (or both).

However, to ensure that they with their obligations under data privacy laws, the commercial servicer provider also identifies the private citizen as a Rights Holder to the Events, and configures their Expert System so that the Events cannot be accessed until suitably the Event Consumer has been authorised by the private citizen.

Alternatively, the commercial service provider may share the Events with an Expert System controlled by the private citizen under a different collaboration agreement, leaving the private citizen free to make them available to the Government department.

MOTIVATIONS AND RESPONSIBILITIES

The Ecosystem Toolkit aims to motivate Event Producers to participate in collaborative ecosystems by establishing a highly scalable, fully configurable, easily extensible mechanism through which to 'write' information which can then be made available to Event Consumers.

By doing so, it is hoped that Event Consumers will be in a position to 'read' more information, from different Event Producers, relevant to their domain of interest. However, it becomes their responsibility to interpret that information, and aggregate it in such a way that meets their own assurance standards or requirements.

The ecosystem itself provides the basis for collaboration, however. This may take the form of a scheme operating around a specific domain of interest, or of a market (in which loosely coupled 'read' and 'write' may need to be actively matched for the market to function).

3. AN OVERVIEW OF THE ECOSYSTEM TOOLKIT

This chapter provides a short overview of the components that comprise each of the two major artefacts of the Ecosystem Toolkit.

For more detail, please refer to the artefacts themselves, which are stored as separate documents in the OIX library.

RULES OF ENGAGEMENT

As indicated above, the Rules of Engagement set out a composable and configurable collaboration framework. This comprises of ten separate components, together with a set of default parameters that provide a *minimum* basis for collaboration between entities.

Some form of legal instrumentation is required to govern a given instance of collaboration. This may take the form of a bilateral contract between two entities, a multi-lateral collaboration agreement between several entities, a series of contracts with a legal entity tasked with operating a scheme, etc.

The ten components (which represents a non-exhaustive list) are:

- **LEGAL BASIS:** the legal basis on which entities are collaborating over information. The default setting is consent.
- **LIABILITY FRAMEWORK:** the level of liability linked to the quality of information being shared. The default setting is zero.
- **ASSURANCE FRAMEWORK:** the level of assurance over the quality of information being shared. The default setting is zero.
- **ISSUE RESOLUTION:** the mechanisms used to resolve issues and disputes, as well as sanctions for ecosystem members that fail to observe them. The default is the general body of national and international law and regulation relevant to each participant.
- **TERMS OF SERVICE:** these might include service levels, feedback loops, revocation rights, etc. There are no default terms of service.

- **Scope of Service:** the default scope of service is the provision of access to write information in the form of Events.
- **Ecosystem Intent:** intent goes beyond the immediate transaction and may be used to capture what is expected of participants. There is no default expression of intent.
- **Membership Administration:** the people, entities, policies, mechanisms and tools used to administer membership of an ecosystem. There are no default settings.
- **Value Exchange Framework:** the motivation for collaboration, expressed as a value on the information being shared. The default motivation is altruism, and the default value is therefore zero.
- **Settlement Framework:** the mechanism through which value is exchanged to settle for the sharing of information. The default setting is psychic reward, meaning that unless configured otherwise no settlement framework is in place.

The defaults contained within the Rules of Engagement provide a basic pattern for such legal instrumentation. The intention is to test a more evolved collaboration agreement – which can in turn be published as a pattern of collaboration – in a subsequent Beta Project.

TECHNICAL SPECIFICATION

The Technical Specification sets out the technical considerations that must be deployed in order to implement an Expert System. An Expert System can be implemented by any entity that sees value in collaborating with other entities also deploying Expert Systems.

The Technical Specification contains the following responsibilities:

- **Exchange:** the gateway that connects the Expert System to the outside world (including other Expert Systems), and through which the Owning Entity (which operates an Expert System) configures and instructs that Expert System
- **Administration:** the instruction set used to provision the Expert System, configure the Owning Entity's entitlements, instantiate an Ecosystem and specify a knowledge base
- **Refinery:** the transformation of information (including Administration instructions) into a machine-readable representation in the form of an Event

- **EVENT STORE:** the persistent storage of Events written to the Expert System
- **REGISTRY:** a registry of all the 'things' that have been declared (i.e. nodes), including the Participants that form part of an Ecosystem and their entitlements
- **ONTOLOGY:** a conceptual model of the domain of interest shared by the Participants of an Ecosystem, expressed in a machine-readable logic-based language

These responsibilities may be implemented to varying levels depending on the requirements of any given ecosystem. With the exception of a number of clearly articulated **ARCHITECTURE CONSTRAINTS**, implementation details are not part of the Technical Specification.

4. DESIGN GOALS OF THE ECOSYSTEM TOOLKIT

The Ecosystem Toolkit aims to elicit five beneficial qualities:

- Scalability: after the initial investment to configure it, an Expert System must have minimal marginal costs for executing the tasks of collaboration within an ecosystem
- Configurability: Expert Systems must be configurable across different contexts, fostering reuse and discouraging repeated development of the same core functionalities
- Extensibility: Expert Systems must be able to easily accommodate an extension – or indeed change – of the domain of interest over which entities choose to collaborate
- Interoperability: Expert Systems that have been deployed using different underlying technologies must be able to communicate with each other
- Traceability: every actor and action within an ecosystem must be traceable through the Event history, so that there is a basis for trust that is accessible to other participants

SCALABILITY

Machine-to-machine collaboration is inherently scalable. The Ecosystem Toolkit therefore sets out to machines to execution collaboration between entities, by acting on behalf of the entities that are responsible for them.

However, as noted above, a technical specification alone is not enough to define the legal, procedural and behavioural aspects that govern the collective action within an ecosystem.

The Alpha Project sought to address this by requiring that the Rules of Engagement be formulated as a set of configurable components, instances of which can be easily referenced by the Expert Systems which deploy the Technical Specification.

This is intended to encourage the emergence of easily repeatable patterns – starting with the default collaboration agreement contained within the Rules of Engagement – that can be scaled as part of machine-to-machine interactions.

CONFIGURABILITY

Just as the Rules of Engagement must be configurable, so too the entitlements that govern access to information must be configurable. Configuration of entitlements is therefore at the epicentre of the Administration of an Expert System.

The Owning Entity (i.e. the entity operating an Expert System) needs to be able to specify the entitlements of the other entities that it is collaborating with. The scope of action that is afforded to each ecosystem participant can therefore be tailored to each specific need.

Entitlements can be codified as a distinct data shape. Entitlements can therefore not only be deployed by an Expert System but they can also easily be shared with other Expert Systems, reducing the need to rebuild the same set of entitlements for different entities or contexts.

Just as the Rules of Engagement encourage patterns of collaboration agreement, so too the Technical Specification encourages the configuration of entitlements into repeatable data shapes that can be easily shared and implemented as a matter of best practice.

EXTENSIBILITY

The real world involves an endless flow of information across and between different contexts, and along many different extended value chains and user journeys. A “user” wants to be able constantly to assert and re-assert core attributes or properties about themselves as they move from one context to the next, *without* having to re-establish trust in them.

As importantly, a “user” will acquire new attributes and properties as they interact with entities in one context, and the accretion of these new attributes and properties can be used as ‘signals’ that streamline their interaction with entities in a subsequent context.

As discussed above, the Ecosystem Toolkit requires all information to be expressed using an Event, conforming to standardised schema. This allows information of any kind to flow to wherever it is needed, with both the content of the information and the *basis of its assurance* able to permeate different contexts.

Practically speaking, this implies that the Ecosystem Toolkit must have the capability to create and edit a machine-readable Ontology: i.e. the ability to develop, evolve and communicate from one machine to another the *meaning* of the information being shared.

INTEROPERABILITY

The Technical Specification imposes "API First" as an Architectural Constraint on the deployment of Expert Systems. All interaction with the Expert System takes place through a gateway referred to as the Exchange, which must be deployed as an API that it is self-describing, negotiable and evolvable (borrowing elements from REST, and minimising the need for out of band information). All interactions must use the standardised Event schema.

This means that Owning Entities can deploy Expert Systems using different technical topologies, and that they will still be able to communicate with other Expert Systems. The "API First" approach ensures that Expert Systems can cope with today's rapidly evolving technical landscape, and the standardised Event schema protects against the systemic fragility inherent in the context-specific data schema often used as API standards.

To achieve these benefits, the Expert System must provide the ability to convert information that is pushed into it (from outside) into the standardised Event schema. This responsibility is referred to as the Refinery.

TRACEABILITY

As explored in detail in the Whitepaper, one of the biggest barriers to collective action is the ability to agree *how* to trust the quality of the information over which entities collaborate.

Most efforts to address this facet of collective action have focused on the definition of assurance standards, whereby entities either converge on a common due diligence process or rely on an "authoritative source" that undertakes such processing on their behalf.

In this approach, the frameworks that govern the relationships between collaborating entities are paramount. The information over which the entities collaborate (e.g. the “user” attributes and properties) often does not need to be supported by any further metadata (i.e. information about the information) – it is usually sufficient to know that it has come from an authoritative source that delivers against a known assurance level.

This approach places a high burden on those entities acting as authoritative sources: their operations must align with a given standard, and they must be sufficiently financially robust to warrant the high levels of trust placed in them to meet assurance standards.

The reliance on such authoritative sources can act as a barrier to widespread collective action. Where there is a strong business case, entities (usually large institutions) are willing to act as authoritative sources. However, it is very difficult to extend or replicate the same level of reliance into other contexts where the business case may not be so clear cut.

Although the Ecosystem Toolkit does not preclude the role of authoritative sources, it encourages a fundamentally different model of assurance based on traceability.

The Technical Specification requires that each and every step in the Administration of every Expert System (such as its provisioning and configuration, the declaration of other entities participating in an ecosystem, the editing of an ontology, the creation of a knowledge base, etc.) is recorded as an Event, and persisted within an Event Store.

This means that every actor and action within an ecosystem is traceable through the history of Events, and can always be linked to the Events containing the “user” attributes and properties information over which collaboration primarily occurs.

Since entities deploying the Ecosystem Toolkit are able to collaborate over *any* Event, they can collaborate over the history of Events that led to the assertion of a particular “user” attribute or property.

Entities that are in a position to contribute signals into an ecosystem (e.g. because they have interacted with the “user”) are able to do so, with their assertions supported by a fully traceable audit trail of Events.

The metadata-centric nature of the Ecosystem Toolkit *amplifies the quality* of signals from a (much wider) network of (potentially less trusted) sources, without introducing the burdens imposed by the need to become an authoritative source.

In contrast, the burden shifts to the Event Consumer, who are made responsible for interpreting the information they receive, including the audit trail of Events. To do so, they must deploy reasoning models to determine if that audit trail is sufficient to satisfy the assurance levels that they require.

This approach – in common with other event-driven architectures – allows Event Producers and Event Consumers to scale independently of each other, and so reduces barriers to collective action[5].

In addition, the Ecosystem Toolkit actively encourages (without imposing) the emergence of assurance models that draw on the fully traceable audit trail of events to derive a basis for trust. Such assurance models leverage network effects and pattern recognition, providing powerful and robust alternatives to those that are anchored on authoritative sources.

[5] See Appendix A for a fuller explanation

APPENDIX A – KEY DIFFERENCES

The OIX Guide to Trust Frameworks and Interoperability and the Ecosystem Toolkit described in this Alpha Project Report clearly intersect in their respective areas of concern.

However, as noted in Chapter 2, the objectives of each piece of work are different, with the OIX Guide presenting a comprehensive set of the "business requirements" required to establish a Trust Framework, while the Alpha Project prescribes an "architectural style" to be used when establishing a collaborative ecosystem.

Below, we set out some of the key differences in terminology as a way of emphasising and explaining the differences in objectives and perspectives.

TRUST FRAMEWORK VS. ECOSYSTEM TOOLKIT

The requirement to structure Events as assertions means that the Ecosystem Toolkit *explicitly enforces* an important separation of concerns: every ecosystem has a *write* side (i.e. entities involved in production and publication of Events) and a *read* side (i.e. entities that are entitled to consume and reason an Event, e.g. for use in a Reasoning Engine).

The Ecosystem Toolkit treats the *write/produce/publish* and *read/consume/reason* sides of the ecosystem independently of each other – they are only "loosely coupled" by the nature of any given collaboration. This architectural paradigm[6] is integral to event-driven architectures[7] and crucially important for the removal of barriers to collaboration.

As a result, each side of the ecosystem is infinitely scalable: any entity can (cheaply) provide assertions that are refined into Events, ready to be consumed within any ecosystem. At the same time, any entity can (cheaply) deploy Reasoning Engines that evaluate whether the

[6] Sometimes referred to as CQRS (Command-Query Responsibility Separation) or asynchronous message patterns, this paradigm is central to event-driven architecture.

[7] Event-driven architectures often refer to the *publish* and *subscribe* (aka pub/sub) sides of an ecosystem.

Events that they are entitled to access meet the assurance levels (whether for identity or eligibility) that their circumstances require.

In contrast, the *implicit assumption* of the OIX Guide to Trust Frameworks is that both sides of the ecosystem are "tightly coupled" – there is no suggestion that the *write* side (broadly captured in the General Rules, Security and Technical Requirements and Interoperability Requirements) can evolve independently of the *read* side (broadly captured in the Trust Rules, Relying Party Services, User Services and Trustmark).

ROLES AND OBLIGATIONS VS. PROVENANCE

The OIX Guide to Trust Frameworks specifies a number of distinct roles on the *write* side of the ecosystem (Evidence Issuer, Evidence Verifier, Identity Provider, Broker). As indicated by Figure 1, the implication is that the exact responsibilities of these roles will always be defined in the context of each Trust Framework, based on the requirements of the *read* side of the ecosystem (i.e. "tightly coupled").

In contrast, by specifying a requirement to capture the provenance of an Event, the Ecosystem Toolkit makes the distinct roles on the *write* side of the ecosystem explicit in every single piece of information. Consider the following:

- Ben says that Ben's UK Passport is valid
- Company A says that Ben's UK Passport is valid
- UK Passport Office says that Ben's UK Passport is valid

In the Ecosystem Toolkit, the role being played by Ben (Claimant), Company A (Evidence Verifier) and the UK Passport Office (Evidence Issuer) can be inferred from their relationship with Ben's UK Passport. The responsibility for making this inference – indeed, any reasoning – is transferred from the *write* to the *read* side of the ecosystem (i.e. "loosely coupled").

This means that two different entities on the *read* side of the ecosystem (i.e. playing the role of Relying Party, in the terminology of the OIX Guide) can come to different conclusions, based on the same piece of information. For example, the first entity may have configured its Reasoning Engine to trust assertions made by Company A, but not the second.

The uncoupling of the *write* and *read* sides of the ecosystem is – in the opinion of the Alpha Project – vital to addressing the design goals of scalability, interoperability and configurability.

TRUST FRAMEWORK VS. COLLABORATION FRAMEWORK

The OIX Guide takes Trusted Digital Identity as its primary use case. The implicit assumption is that the *read* side of the ecosystem demands high assurance levels (e.g. as set out in the Trust Rules that support the Relying Party Services). Hence the central concept of a Trust Framework.

In contrast, the Ecosystem Toolkit deliberately uses the term Collaboration Framework, as there are many use cases and contexts that may be predicated on an absence of trust. The default settings of the Rules of Engagement are an example of such collaboration: i.e. collaboration has been configured on a "take it or leave it" basis.

The components of contained in the Rules of Engagement can be added to and the default settings can be configured "upwards", such that it becomes a Trust Framework.

APPENDIX B – WORKSHOP SCHEDULE

The following workshops were held to discuss various aspects of the Alpha Project, and formed the basis of subsequent work to develop the Ecoystem Toolkit, in the form of the Rules of Engagement and Technical Specification.

1. KICK OFF

Held at 12-3pm on 27/01/2020 in AgFe Offices.

Agenda

- Introductions
- Recap: why are we here?
    - Problem statements
    - Incoming hypotheses
    - Common language / mental models
- Perspectives: what matters to you?
    - Questions
    - Expressions of interest
- Planning: what do we need to do next?
- Resourcing for Workshops #1 and 2

2. STANDARDS

Held at 1-4pm on 05/02/2020 in AgFe Offices.

Agenda

- (Re)introductions

- Recap actions and notes from Kick Off
- Semantic Web standards and tools
- A worked example
- Open knowledge base (GitHub)
- Look forward to Workshop #2

3. METADATA MODELS (1)

Held at 12-3pm on 12/02/2020 in AgFe Offices.

Agenda

- Recap actions and notes from Workshop 1
- Proposal: reschedule Security Workshop
- Ontology – Structure
- Ontology – Initial Draft
- Look forward to Workshop #3

4. METADATA MODELS (2)

Held at 12-3pm on 18/02/2020 in AgFe offices.

Agenda

- Recap from Workshop #2
- Domain knowledge we need to represent
- Components of the toolkit we need to specify
- Look forward to Workshop #4

5. METADATA MODELS (3)

Held at 2-5pm on 26/02/2020 in AgFe offices.

Agenda

- Defining what the simplest possible systems context diagram would look like.
- Stating what the minimum requirements are for 'refined' events in that system.

6. SEMANTIC AND WEB STANDARDS EXPERT INPUT

Held at 1-4pm on 04/03/2020 in AgFe offices.

Agenda

- Expert feedback on conclusions so far, specifically the structure of an event and the specification of a system context approach

Expert feedback was provided by Peter Crocker of Oxford Semantic Technologies and Paul Worrall of Inrupt, both of whom joined the session in person.

7. GOVERNANCE AND RULES OF ENGAGEMENT (1)

Held at 2-5pm on 11/03/2020 in AgFe offices.

Agenda

- Recap from previous session
- Discuss OEF rules of engagement
    - Boundary of open alliance
    - Scalable / mutual contractual agreements for value exchange
    - Value exchange mechanisms
    - Publication of / subscription to events

  - Permissioning (including consent)
  - Use of standards and metadata model
  - Governance
- Target use case - importing wine into the UK

Further expert input was also sought from Paul Evans, Security Architect at Future Borders.

8. GOVERNANCE AND RULES OF ENGAGEMENT (2)

Held at 1-3pm on 18/03/2020 via Microsoft Teams.

Agenda

- Different forms of collaboration and legal instrumentation to support collaboration
- How to define the “minimum set”?

This working session was attended by Ross McDonald, Sid Kalita and Ben Helps only.

9. GOVERNANCE AND RULES OF ENGAGEMENT (3)

Held 11-12.30pm on 24/03/2020 via Microsoft Teams.

Agenda

- Navigation through the (work in progress) material on GitHub
- Take immediate feedback (and invite detailed comments offline)
- Agree outstanding actions needed to complete the Alpha
- Any other business

More detail, including meeting attendees and presentation material, can be found here.

Subsequent working sessions involving Ross McDonald, Sid Kalita and Ben Helps were held remotely over the course of April and May, to iterate and finalise the Ecosystem Toolkit.