PNNL-32403

Universal Utility Data Exchange (UUDEX) – Functional Design Requirements – Rev 1

Cybersecurity of Energy Delivery Systems Research and Development

December 2021

SR Mix
SA Neumann
S Sridhar
C Gonzalez-Perez
C Peloquin
MA Rice
CM Schmidt
SV Singh
ML Cohen
T Bobka

Prepared for the U.S. Department of Energy
under Contract DE-OE0000190

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights**. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <https://www.ntis.gov/about>
Online ordering: http://www.ntis.gov

Universal Utility Data Exchange (UUDEX) – Functional Design Requirements – Rev 1

Cybersecurity of Energy Delivery Systems Research and Development

December 2021

SR Mix
SA Neumann
S Sridhar
C Gonzalez-Perez
C Peloquin
MA Rice
CM Schmidt
SV Singh
ML Cohen
T Bobka

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99354

Revision History

Revision	Date	Deliverable (Reason for Change)	Release #
0	11/2018	Initial Release	PNNL-28207
0.1		Updated NERC Functional Model diagram, Misc. non-technical edits	internal
1	102021	Updates based on Implementation	PNNL-32403

Summary

This document contains the set of functional and architectural requirements for building the Universal Utility Data Exchange (UUDEX) Framework.

Section 1.0 contains an introduction to UUDEX, including a description of the project scope, a discussion of how UUDEX will support existing and emerging utility communications and infrastructures, an overview of the proposed UUDEX architecture, and a high-level overview of the lifecycle of UUDEX data.

Section 2.0 contains 12 use cases, based primarily in electric entity interactions, that document how UUDEX could be used to provide a communications structure for both operational data (such as ICCP), large data files (like power system model updates), incident and event reporting (such as OE-417 reports or information sharing with the E-ISAC), mass alert notifications (such as NERC alerts), and temporary *ad hoc* connections with first responders (such as FEMA during a hurricane response).

Section 3.0 contains functional descriptions of operational and cyber security data showing the breadth of data UUDEX is capable of communicating. Operational data types include ICCP, RCIS, power system model updates, synchrophasor, disturbance files, operations planning, and asset management. Cybersecurity data includes incident reporting, indicator of compromise sharing, guidance (e.g., firewall rule sharing), conformance reports (e.g., verification that patches have been installed), patch availability notification, vulnerability disclosure reporting, and threat notification.

Section 4.0 introduces the specific functional requirements, including a taxonomy of roles and terms used to provide functional descriptions of the interactions between various users of UUDEX, as well as a discussion of data exchange architectures and requirements, including information flow, identity, data storage, testing, and data lifecycle.

Section 5.0 introduces the UUDEX message and data exchange formats and describes them at a functional level.

Section 6.0 describes how UUDEX is proposed to be a hybrid publish-subscribe and query-response architecture supporting real-time and near-real-time notification of data elements available to data subscribers (following the current ICCP model), as well as the capability for querying a database of stored information (such as available software patch updates).

Section 7.0 contains an overview of security threats and mitigations planned for inclusion in the final product, including information disclosure, information corruption, denial of service, identity spoofing, and trust relationships, all with functional descriptions of how UUDEX will mitigate the threats or address the issues.

Section 8.0 contains a list of reference documents relevant to UUDEX.

Appendix A contains a set of notional data element characteristics considered during the development of the functional specification.

Appendix B contains a summary of the changes between this document and the original version of this document published in November 2018.

The functional specification addresses a number of areas that were raised in discussions with industry, both during information gathering phases early in the project, as well as later conversations with the project's Industrial Advisory Board.

These functional design requirements are intended to be used in future more detailed design documents related to this project.

Acronyms and Abbreviations

ACL	access control list
API	application programming interface
BA	Balancing Authority (NERC term)
BES	Bulk Electric System (NERC term)
CIM	Common Information Model, as defined by IEC 61970 and IEC 61968 series of standards
DHS	U. S. Department of Homeland Security
DOE	U. S. Department of Energy
E-ISAC	Electricity Information Sharing and Analysis Center
EMS	energy management system
ESB	enterprise service bus
FERC	Federal Energy Regulatory Commission
GPS	Global Positioning System
ICCP	Inter-Control Center Communications Protocol, as defined by IEC 60870-6
ICS-CERT	Industrial Control Systems Cyber Emergency Response Team
IDS	intrusion detection system
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IoC	indicator of compromise
IPS	intrusion prevention system
IROL	Interconnected Reliability Operating Limit
ISAC	Information Sharing and Analysis Center
IT	information technology
NERC	North American Electric Reliability Corporation
NIST	U. S. National institute of Standards and Technology
OMS	outage management system
OT	operations technology
PCAP	packet capture
PMU	phasor measurement unit (also known as a synchrophasor)
QoS	quality of service
RC	Reliability Coordinator (NERC term)
RCIS	Reliability Coordinator Information System
SCADA	supervisory control and data acquisition
SOL	system operating limit
STTP	streaming telemetry transport protocol
TOP	Transmission Operator (NERC term)

U-Administrator	UUDEX Administrator
U-API	UUDEX Application Programming Interface
U-Client	UUDEX Client
U-Component	UUDEX Component
U-Connection	UUDEX Connection
U-Consumer	UUDEX Consumer
U-Data Element	UUDEX Data Element
U-Data Manifest	UUDEX Data Manifest
U-Data Type	UUDEX Data Type
U-Endpoints	UUDEX Endpoints
U-Exchange	UUDEX Exchange
U-Framework	UUDEX Framework
U-Header	UUDEX Header
U-Identity Authority	UUDEX Identify Authority
U-Identity Object	UUDEX Identity Object
U-Infrastructure	UUDEX Infrastructure
U-Instance	UUDEX instance
U-Message	UUDEX Message
U-Message Envelope	UUDEX Message Envelope
U-Notification	UUDEX Notification
U-Participants	UUDEX Participants
U-Participant Administrator	UUDEX Participant Administrator
U-Payload	UUDEX Payload
U-Producers	UUDEX Producers
U-Protocol	UUDEX Protocol
U-Server	UUDEX Server
U-Subject	UUDEX Subject
U-Subscription	UUDEX Subscription
UUDEX	Universal Utility Data Exchange
UUID	universally unique identifier, as defined by IETF RFC 4122
XML	eXtensible Markup Language, as defined by the World Wide Web Consortium (W3C)

Contents

Figures

1.0 Introduction

The purpose of this document is to provide a functional specification for the Universal Utility Data Exchange (UUDEX).

1.1 Project Scope

UUDEX describes a communications architecture and protocol suite that allows energy sector control centers and related organizations, referred to as UUDEX Participants (U-Participants[1]), to exchange data and information. It does this by defining relations between a set of client nodes, referred to as UUDEX Endpoints (U-Endpoints), that share data via a set of server nodes, collectively called the UUDEX Infrastructure (U-Infrastructure).

The primary focus of UUDEX is to facilitate information sharing between control centers, operations centers, and other trusted organizations. This involves the communication mechanisms necessary for reliable and secure operations of an energy delivery system. The most common usage would be for the conveyance of measurements, calculations, and schedules between entities and applications that are responsible for managing the electrical grid at both the transmission and distribution levels. UUDEX can be used for communications between utility organizations and non-utility organizations, including government organizations and commercial enterprises, for example between a utility organization and its Information Sharing and Analysis Center (ISAC), the E-ISAC for the electricity sub-sector. UUDEX could also be used to send event and outage information to the U.S. Department of Energy (DOE) following requirements of the OE-417 reporting criteria. It could be used to coordinate information dissemination about line outage and restoration information between a utility and local first responders or the Federal Emergency Management Administration. Or, UUDEX could be used to provide near-real-time alerts from security service providers pertaining to vulnerability disclosure or patch availability.

UUDEX could also be used by an organization to coordinate internal data communications and some aspects of application integration. For example, it could be used as an intermediary for exchange of network models between systems such as a network model manager, geographic information system, energy management system (EMS), distribution management system, or outage management system (OMS); to transmit current operations information from a control system to a market system; or to pass data from a protected enclave at a control center to a server on a business network for non-real-time use.

UUDEX could also be leveraged to facilitate communications between a market operator and market participants. Although each market often establishes its own data communication protocols, market information could be exchanged using the same interfaces as for other UUDEX data exchanges.

The focus of UUDEX is an information exchange mechanism, not a mechanism for issuing control commands. Therefore, UUDEX is not designed to communicate control instructions to field devices (e.g., distributed energy resources) or other locations (e.g., peer control centers).

[1] Note – Throughout the document, the notation "U-" will be used to denote UUDEX specific terms for readability purposes. Thus, UUDEX Participants becomes U-Participants, etc.

In general, the use of UUDEX for direct communications to end devices in the field is beyond the scope of this document.

For organizations in the North American electricity sector, Figure 1-1 shows a high-level overview of the potential uses of UUDEX between electric sector organizations described by the North American Electrical Reliability Corporation (NERC) functional model.[2] In the figure, the solid lines show logical communications that take place between the control centers or centralized control systems of various functional organizations that are within scope of UUDEX, while the dashed lines show communications between some functional organization's control centers and field devices that are out of scope for UUDEX. Not all communications interactions are shown in the figure, but it is clear that a significant number of existing communications interactions could make use of UUDEX. Ultimately, UUDEX could be applied to communications within other energy sectors such as oil and gas delivery systems in a similar manner, but initially UUDEX is focused on the electricity sector.

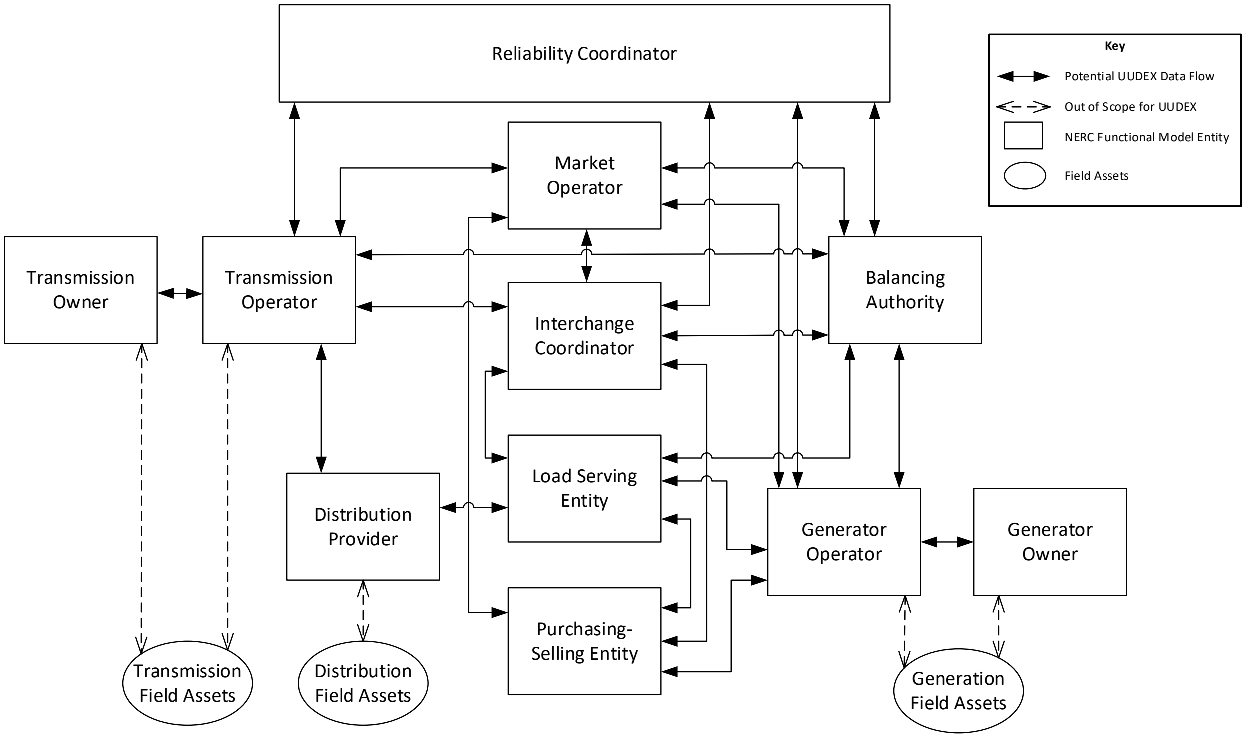

Figure 1-1: NERC Functional Model Notional Relationships

UUDEX is designed to support the transfer of most any type of data, but especially data that use formats commonly exchanged by control centers. There are some types of information exchanged by control centers that do not have a standardized data format or use a format that is inherently tied to the network protocol used to deliver it. In these cases, UUDEX may define a data model that can convey this information. However, UUDEX itself is not tied to the use of any particular data models and can facilitate the exchange of data that are structured in any way. This is necessary to ensure use of UUDEX will be supportive of future needs to exchange new types of data or by the evolution of how existing data are expressed.

[2] See https://www.nerc.com/pa/Stand/Pages/FunctionalModel.aspx

Another key aspect of UUDEX is that communications will be “secure by design.” A compliant UUDEX implementation will be deployed with security enabled by default, while still allowing the communications and data exchanges to be diagnosed in the event of errors or data mismatches. The security will provide for integrity and confidentiality of data transmissions. All this will be done with minimal disruption to the availability of the links and general flow of information.

In addition to the security of the data in transit, UUDEX will provide access security of the data within the system, allowing UUDEX Producers (U-Producers) to specify or agree to how their data can be accessed. UUDEX will support a robust access control language with an extensible set of primitives to allow simple expression of common access statements.

UUDEX is designed to be “transport agnostic,” in that the underlying physical communications infrastructure is largely irrelevant to how UUDEX works. UUDEX can be implemented using utility-owned infrastructure, communications leased from a common carrier, the public internet, cloud infrastructure, or any combination. UUDEX’s secure-by-design philosophy allows a common implementation to simultaneously use any combination of physical infrastructure, subject to the risk appetite and security requirements of the utility organization.

UUDEX is also designed to ensure that necessary communications remain possible even in conditions of network congestion or connectivity loss. UUDEX supports message prioritization schemes that help ensure high-priority data are given preference over lower priority data in the event that congestion precludes transmission of all data. In addition, UUDEX includes support for server and communications redundancy, which can make implementations more resilient. Operators are not required to deploy such redundancy (since redundancy can be costly and not all UUDEX Instances (U-Instances) will necessarily have the same criticality), but components are required to be able to support redundancy should operators wish to make use of it.

A single organization might participate in multiple U-Instances. By design, every U-Instance is segregated from every other U-Instance (and from the broader Internet) regardless of whether it shares a network with those other entities. Nonetheless, because UUDEX standardizes the interfaces and operations used to interact, a common suite of tools could support an organization's interactions across all of the U-Instances that it utilizes. This standardization can reduce costs and improve efficiencies significantly compared to needing to host and maintain unique solutions for different relationships and exchanges.

1.2 Coordination with Existing Communications and Initiatives

The UUDEX project will support a number of existing and emerging utility communications infrastructures and technologies, including but not limited to the following:

- Inter-Control Center Communications Protocol (ICCP) – the ICCP is well established, tracing its roots to the desire for standardized control center communications in the mid-1990s. ICCP uses a subset of the Manufacturing Message Specification protocol to provide messaging and allow control commands to be passed from one organization to another. ICCP has been augmented to include a secure communications option, but has seen limited use of the secure options, at least in the U.S.
- DOE Electric Emergency Incident and Disturbance Report Form OE-417 – the DOE OE-417 report form (referred to as either “OE-417 Report” or “DOE-417 Report”) is required to be submitted to DOE following the occurrence of a specific disturbance or incident as outlined on the form. There are currently 24 categories of incidents that trigger submission on the web

version of the form and 12 categories on the PDF version. Currently, the form can be submitted either by filling out a web form (preferred) or completing a PDF form and emailing it to DOE.

- NERC Standard CIP-008 revisions – in 2018, the Federal Energy Regulatory Commission (FERC) released a directive[3] to revise the reporting requirements in NERC Standard CIP-008 to include specific required fields when reporting cybersecurity incidents. These additional fields are included in CIP-008-6.
- Cybersecurity notifications – many formats already exist for the exchange of cybersecurity-related information, whether it represents vulnerabilities, mitigations, alerts, patches, or firewall ruleset updates.
- Reliability Coordinator Information System (RCIS) – the RCIS tool is a bulletin-board-like system used by Reliability Coordinators (RCs) to exchange operational and reliability-based information. Currently, RCs maintain a database of postings with limited search capability.
- Power System Model Files – organizations within the electricity sector (primarily Transmission Operators [TOPs] and RC's) are required to exchange power system model files with their neighboring organizations and with their RCs. These model files can be quite large. UUDEX provides a mechanism for the exchange of either full or partial model files, allowing individual participants to retrieve model files that are of use to them.
- Phasor measurement unit (PMU) data – Institute of Electrical and Electronics Engineers (IEEE) standard C37.118.2^{TM} or IEEE Std P2664^{TM}. The current protocol for transporting PMU data is IEEE Std C37.118.2; a new standard for Streaming Telemetry Transport Protocol (STTP) is under development as IEEE Std P2664. Both of these protocols are used to gather data from PMUs in the field (which is not in scope for UUDEX), as well as transfer PMU data between phasor data concentrator nodes at control centers. UUDEX could be used to distribute PMU data snapshots and aggregated values for use across control systems.
- Market data – UUDEX can also be used for many other purposes. One potential use could be for market operators and market participants to share market data. UUDEX would allow this data to be efficiently distributed in a manner that ensures confidentiality.

1.3 UUDEX Architecture Overview

UUDEX is an architecture for the management and distribution of UUDEX Data Elements (U-Data Elements) within a closed community. A U-Data Element can be any data object, including but not limited to documents, images, UUDEX Data Sets (U-Data Sets) (e.g., power system measurements), and other forms of both structured and unstructured data. UUDEX differs from other forms of content dissemination systems in that it is built to be highly secure, with all content confidentiality and integrity protected and access to data closely controlled. UUDEX is also optimized for the types of data exchanged by the energy sector, ensuring that all content is transported with the metadata necessary for it to be understood within that context.

UUDEX employs a client-server architecture for data distribution. This architecture supports both query-response and publish-subscribe interactions. The U-Infrastructure stores information, publishes information to subscribing U-Endpoints, and responds to requests from U-Endpoints.

[3] See https://www.nerc.com/FilingsOrders/us/FERCOrdersRules/E-1_Order%20No.%20848.pdf

U-Endpoints interact with U-Infrastructure. Specifically, U-Endpoints can manage UUDEX Subscriptions (U-Subscriptions), publish data, delete data, query data, or replace data on the U-Infrastructure. U-Subscriptions can be defined that either allow for alerts of new data objects or automatic forwarding of new data objects. All of these actions are subject to security policies based on the identity associated with the U-Participant requesting the action.

1.4 UUDEX Data and Data Lifecycle

UUDEX supports a wide variety of data object types, where categorizations can include but are not limited to operational status information, incident and other reporting, security alerts and materials, and even general communications and messaging. A data object conveyed using UUDEX is called a U-Data Element.

U-Data Elements are migrated through the U-Infrastructure by publishing them to UUDEX Subjects (U-Subjects). U-Subjects are the basic unit of storage and organization in the U-Infrastructure. U-Data Elements are delivered to UUDEX Consumers (U-Consumers) by creating UUDEX Subscriptions (U-Subscriptions) to U-Subjects.

The following sequence is descriptive of the lifecycle of U-Data Elements:

1. A U-Subject is created with specifications for access control and other behavior parameters.
2. An entity (such as an RC) creates a U-Subscription and attaches one or more U-Subjects to it to be able to retrieve U-Data Elements that are published to the U-Subjects. A U-Subscription must contain one or more U-Subjects, and a U-Subject may be included in more than one U-Subscription.
3. Some other entity (such as a TOP) creates and populates the U-Data Element in a U-Data Set. U-Data Elements might be automatically generated by sensors, application programs, or other tools, or they might be manually created by parties filling out forms or writing messages. The entity that creates the data is its U-Producer.
4. The U-Producer uses their U-Endpoint to send (publish) the U-Data Set to one or more U-Subjects. In addition to the data, the U-Endpoint will add metadata that includes handling instructions and other information.
5. The U-Infrastructure receives the message from the U-Endpoint and verifies that the U-Participant associated with that U-Endpoint is allowed to add the given data to the requested U-Subject. The U-Infrastructure then responds to the U-Endpoint, noting whether the request to publish data was successful or not. Note that is there is no active U-Subscription associated with the U-Subject, the data publish request is rejected.
6. Assuming the data are accepted, the data are added to a U-Subject. A U-Subject is simply a collection of data that is of the same data type and is treated similarly with regard to access control, delivery priority, and other aspects. The U-Subject will continue to store the data until conditions are met for their deletion.
7. At the time at which the data are added, the U-Infrastructure will determine the U-Subscriptions that are associated with the U-Subject. For each matching U-Subscription, the U-Infrastructure will queue a message for delivery to the U-Endpoint that established that subscription. The queued message is one of the following:
    a. If the U-Subscriptions has specified that the U-Data Set should be delivered immediately, the message consists of the U-Data Set. This is the default behavior and is intended for real-time data delivery.

   b. If the U-Subscriptions has specified that a notification is to be delivered, the message contains metadata about the U-Data Set and a message identifier that can be used to retrieve the entire U-Data Set at a later time. This behavior is intended for large file transfer data that may or may not be of interest to the subscriber.

8. Sometime later, a different U-Endpoint might query the U-Subject to learn the contents of the persistent U-Data Sets of the U-Subject. In this case, the U-Infrastructure will send an array of unique message identifiers which can be used to obtain a data manifest, with metadata taken from the metadata of the U-Data Sets within that U-Subject.

9. At some point, a U-Endpoint might query the U-Subject for data. It could do this either by requesting the specific U-Data Set using its unique message identifier (as returned from a notification message or as might be discovered in a manifest) or it could submit a pattern the U-Infrastructure compares to its U-Subject entries. Assuming the requesting U-Endpoint has read access to the U-Subject, the U-Infrastructure will send the U-Data Set to the U-Endpoint in response to its request.

10. Eventually, the U-Data Set's U-Producer (or other authorized party) might use a U-Endpoint to instruct the U-Infrastructure to delete the U-Data Set from its U-Subject. It may also be deleted when all the pending U-Subscriptions for the U-Subject have been fulfilled. Alternately, the U-Data Set might be deleted due to other configuration details associated with the U-Subject, such as being configured to delete the oldest U-Data Sets when the U-Subject's queued message size exceeds a given limit.

2.0 Use Cases Supported by UUDEX

This section describes a set of representative use cases that are applicable to UUDEX. This in no way implies specific limitations on the use of UUDEX. Where the primary purpose of a use case is to answer the question “WHO does WHAT to WHO, WHEN, and WHY do they do it?” It is not the intent to describe “HOW” this is achieved, as this would be described by a design that depicts the underlying technical infrastructure. In the use case discussion, the term “Electric Entity” is used in a general case to represent a generic U-Participant that plays the role of an actor in the use cases. There is no explicit limitation or implication on the set of allowable actors.

In a real UUDEX environment, the U-Infrastructure consists of one or more UUDEX Servers (U-Servers). These U-Servers are what receive requests from UUDEX Clients (U-Clients), manages the U-Subjects in which U-Data Elements are stored, and which manage the U-Subscriptions that U-Clients establish to receive published U-Data Elements. U-Servers will usually be redundantly deployed to provide redundancy and load balancing. Individual U-Servers for a U-Instance would be hosted by one or more organizations that might also be U-Participants of the U-Instance that the U-Servers support. However, all of these details are mostly abstracted in the use cases provided below. As far as U-Participants are concerned, they communicate with their U-Instance's U-Infrastructure, without needing to distinguish between individual U-Servers or know where those U-Servers are hosted. For this reason, in the diagrams below, the U-Infrastructure of a U-Instance is treated as an abstract entity independent from any U-Participant other than to note the organization responsible for “hosting” the U-instance (i.e., administering the U-Instance and managing its server infrastructure). The only detail provided about the U-Infrastructure are the individual U-Subjects that are used in the use case, since U-Subjects are the units to which specific access controls are assigned, and thus can help clarify which entities have access to which U-Data Elements.

The diagrams show separate U-Producers and U-Consumers. These are different behaviors of a U-Endpoint. However, the diagrams show these behaviors as separate entities to make the data flows clearer.

The U-Infrastructure only performs actions in response to U-Endpoint direction and only the specific actions U-Endpoint's U-Participants are allowed to request. This means that the U-Infrastructure is not allowed to perform any additional processing of received data, beyond what is necessary to store it in a U-Subject and any processing necessary to efficiently serve queries and U-Subscription requests. For example, U-Servers might sort and index U-Data Elements; they will not, however, alter these elements. If there is a desire to create derived U-Data Elements from U-Data Elements (e.g., to down-sample certain U-Data Elements, or create aggregated U-Data Elements based on multiple other elements), then this will need to be done by having a U-Endpoint subscribe the relevant data from the U-Infrastructure, perform the necessary derivation, and publish the data back to the U-Infrastructure.

Figure 2-1 shows the symbols used in the use case drawings.

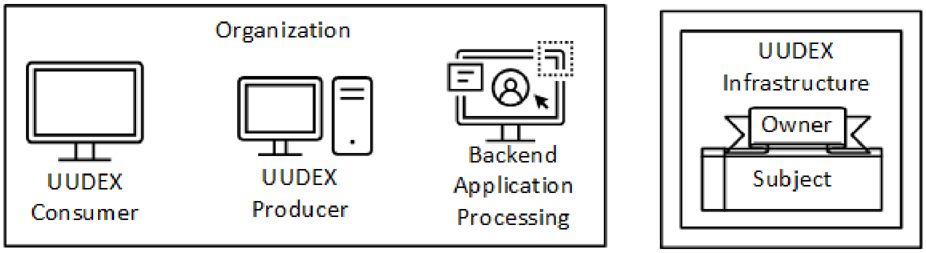

Figure 2-1: Symbols Used in Use Case Drawings

In the use case drawings, the following conventions are used:

- A single-lined box is used to designate an organization. In these examples, organization boundaries are intended just to represent grouping of ownership of the contained components. The do not necessarily represent physical co-location. The organization box contains one or more of the U-Consumer, U-Producer, or Backend Application Processing icons.
- The U-Consumer and U-Producer are defined components in UUDEX. These components interact with the U-Infrastructure using standardized interfaces. U-Consumers and U-Producers are both types of U-Endpoints. In practice, a single piece of software will likely encompass both U-Consumer and U-Producer functionality, but the diagrams below split these functions out to clarify data flows.
- The Backend Application Processing represents some component that consumes, processes, or produces data that is outside the scope of UUDEX standardization.
- A double-lined box represents a U-Instance and its U-Infrastructure. A U-Instance represents a self-contained trust environment with a common understanding of the identities of the U-Participants using that U-Instance. The U-Infrastructure comprises the U-Servers and other technical components with which U-Endpoints communicate and where U-Data Elements are stored on U-Subjects. Note that a single U-Instance might have multiple U-Infrastructures – in this case, this would be represented by multiple inner boxes. Also noted is the organization that “hosts” the U-Infrastructure, that is, the organization that provides the administrative and support functions for the U-Instance and its associated U-Infrastructure hardware and software.
- Each U-Infrastructure contains one or more named U-Subjects. U-Subjects are represented by boxes with row-column headings and appear with their name. All U-Subjects are owned by some U-Participant. The owner of the U-Subject is named in the ribbon above the U-Subject icon.

2.1 Operational Data Shared between an Electric Entity and a Reliability Coordinator

The use case drawn in Figure 2-2 represents sharing of information to a third party who is implicitly trusted to protect the data from inadvertent disclosure .

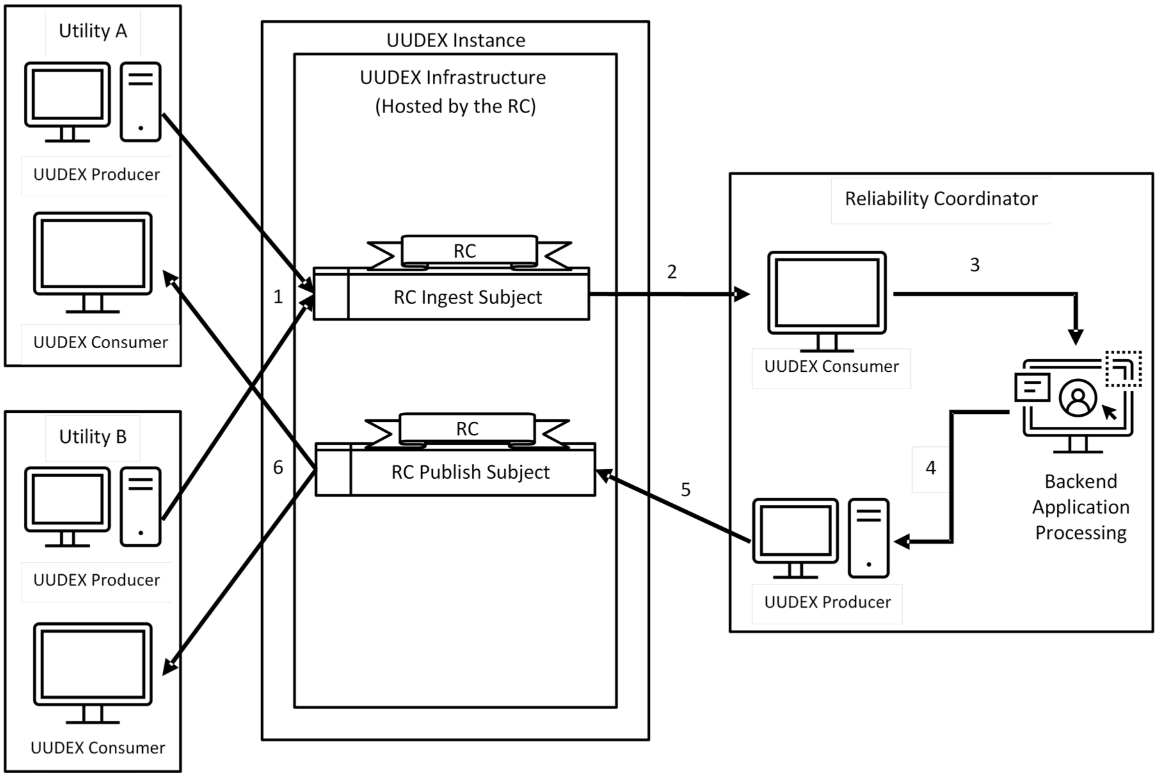

Figure 2-2: Sharing Data with a Reliability Coordinator

This use case involves the following activities:

- Electric Entities (noted in the drawing as Utility A and Utility B) will periodically provide data to the RC's U-Infrastructure using U-Producers (1). The data is stored in the RC Ingest Subject. Any utility can publish to this U-Subject, but only the RC can read its contents. The U-Subject would likely specify that the U-Data Elements in this U-Subject would likely be deleted once all the U-Subscriptions for this U-Subject are fulfilled.
- The RC’s U-Consumer will retrieve the data from the U-Infrastructure (likely a set of requests to gather data from all Electric Entities' U-Subjects for the given data type) (2).
- The RC will process the data and prepare new or modified U-Data Elements based on the provided submissions to be sent back to the Electric Entities (3, 4).
- The RC’s U-Producer will send data to the U-Infrastructure to a RC Publish Subject (5). Any utility can read from this U-Subject, but only the RC can publish to it.
- Each Electric Entity U-Consumer will request published data from the U-Infrastructure, specifically the RC Publish Subject. (6).

2.2 Operational Data Shared Between Two Electric Entities

The use case drawn in Figure 2-3 represents sharing of data between peers where limited trust exists between those peers, particularly with regard to data control .

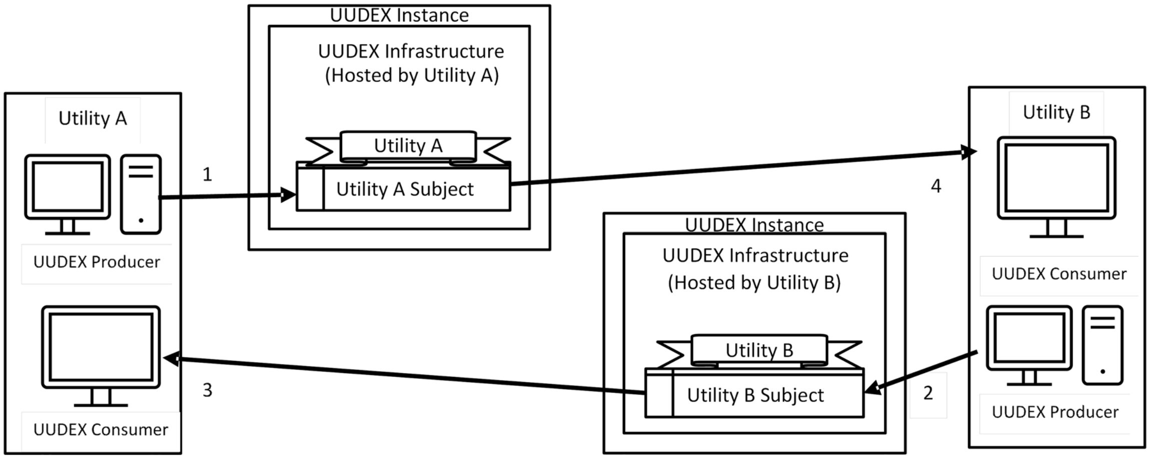

Figure 2-3: Data Sharing Between Two Electric Entities

Unlike the preceding use cases, where one party is trusted by all communicants, in this case there is only limited mutual trust between the communicating parties. Reflecting this, both Utility A and Utility B stand up their own U-Instances and U-Infrastructure. This gives each utility complete and unshared control of their U-Infrastructure elements (such as U-Servers), the U-Subjects in that U-Infrastructure, and the U-Data Elements on those U-Subjects.

- Each Electric Entity (noted in the drawing as “Utility A” and “Utility B”), for example two TOPs, will each establish a U-Subject in its own U-Infrastructure that will contain data it wishes to share with specific other Electric Entities. In this case, each Electric Entity identifies the other as having permissions to discover and read its U-Subject. (Each Electric Entity could create many such U-Subjects, each with a different set of permitted readers.) The U-Subjects would likely specify that the U-Data Elements in these U-Subjects would likely be deleted once all the U-Subscriptions for the U-Subjects are fulfilled.
- When one Electric Entity (henceforth, “the originator”) wishes to convey information to other Electric Entities (henceforth, “the recipients”), the originator will publish data to the appropriate U-Subject (based on who the recipients are) in its own U-Infrastructure using its U-Producer (1, 2).
- The U-Consumer at each recipient will request data from the appropriate U-Subject in the originator’s U-Infrastructure, only receiving data that it is authorized to receive, and store it in a local database (3, 4). Storing the data in a local database eliminates the need to go back to the data producer's U-Subject every time the data are needed in the future.
- Multiple logical point-to-point (bilateral) links from different Electric Entities can be established in this way.

The example here is the most secure as each Electric Entity has complete and un-shared control over their infrastructure, but it does require each party to manage all elements of a U-Instance. If there was a mutually trusted party to manage the U-Instance, this exchange could be accomplished within a single U-Instance and U-Infrastructure, where each Electric Entity was allowed to create and manage U-Subjects. An example of this could be the RC use case shown in Section 2.1.

2.3 Operational Data Shared between Two Control Centers of the Same Electric Entity

The use case drawn in Figure 2-4 represents sharing between locations within the same organization, for example, the primary control center and backup control center of a TOP.

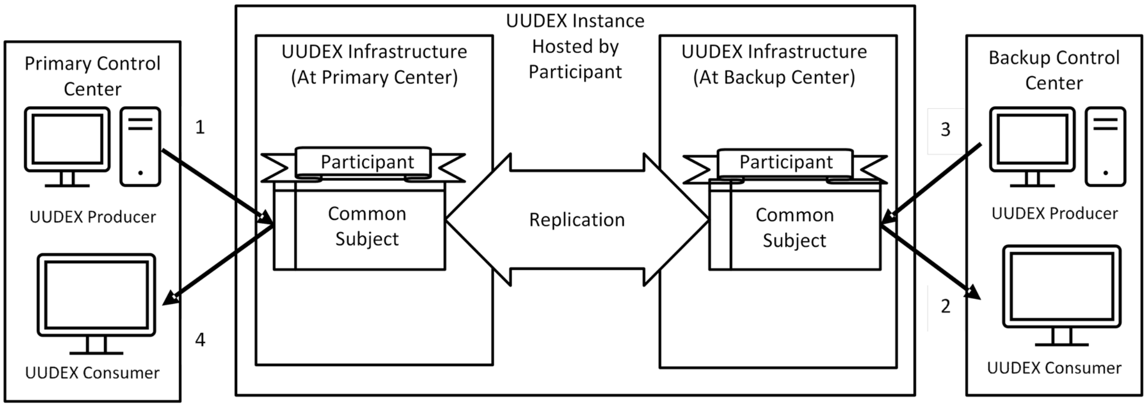

Figure 2-4: Data Sharing Between a Primary and Backup Control Center

This case highlights the fact that the U-Infrastructure within a U-Instance can be geographically distributed. This use case has a single U-Instance, but there are elements of the U-Infrastructure present at both the primary and backup control centers. Both parties interact with a common U-Subject located on their local infrastructure, and then server replication mirrors these updates to the other piece of U-Infrastructure at the other center. Only a single U-Subject is needed - since both centers represent the same U-Participant, both of them would be trusted to read and write to the same U-Subject. Should one of the Centers go offline, the other will have a local copy of the U-Subject with the necessary information to assume operational responsibilities.

- The U-Participant will create a U-Subject in the U-Infrastructure to contain data that will be shared to the other control centers of that U-Participant. (Note that this would be the same U-Participant operating both the Primary and Backup Control Center.) There is only a single U-Subject, but it is replicated[4] across U-Infrastructure servers located at both control centers in the U-Instance . The U-Subject would specify that older U-Data Elements in these U-Subjects would be deleted when new U-Data Elements are published to this U-Subjects.
- The operational primary control center will extract data from its local database and use the U-Producer to publish data to the shared U-Subject within its own U-Infrastructure in the U-Instance (1). This information is then replicated across other servers in the U-Instance, including to the U-Infrastructure in the other operations center.
- The operational backup will use the U-Consumer to request the data from the shared U-Subject in its own U-Infrastructure (2).

When the backup control center needs to send data to the primary the steps are the same:

[4] Note – the replication of data in a U-Instance is assumed to be handled by the underlying messaging transport software used to implement UUDEX. The UUDEX specification itself does not address this.

- The backup control center (now acting as a primary) will extract data from its local database and use the U-Producer to publish data to the shared U-Subject in its own U-Infrastructure (3). This is then replicated to other U-Infrastructure within the U-Instance.
- The old primary (now acting as a backup) will use the U-Consumer to request the data from the shared U-Subject in its U-Infrastructure (4).

2.4 Security Event Data Shared between an Electric Entity and the E-ISAC

The use case drawn in Figure 2-5 represents sharing of security event data from an Electric Entity to the E-ISAC and from the E-ISAC back to one or more Electric Entities noted in the figure as “Utility A and Utility B.” In addition to sharing the data with the E-ISAC, since data sent to the E-ISAC might be sensitive, there is a desire to remove the data from the UUDEX Subject as quickly as possible so that only the E-ISAC's internal copy of the data exists.

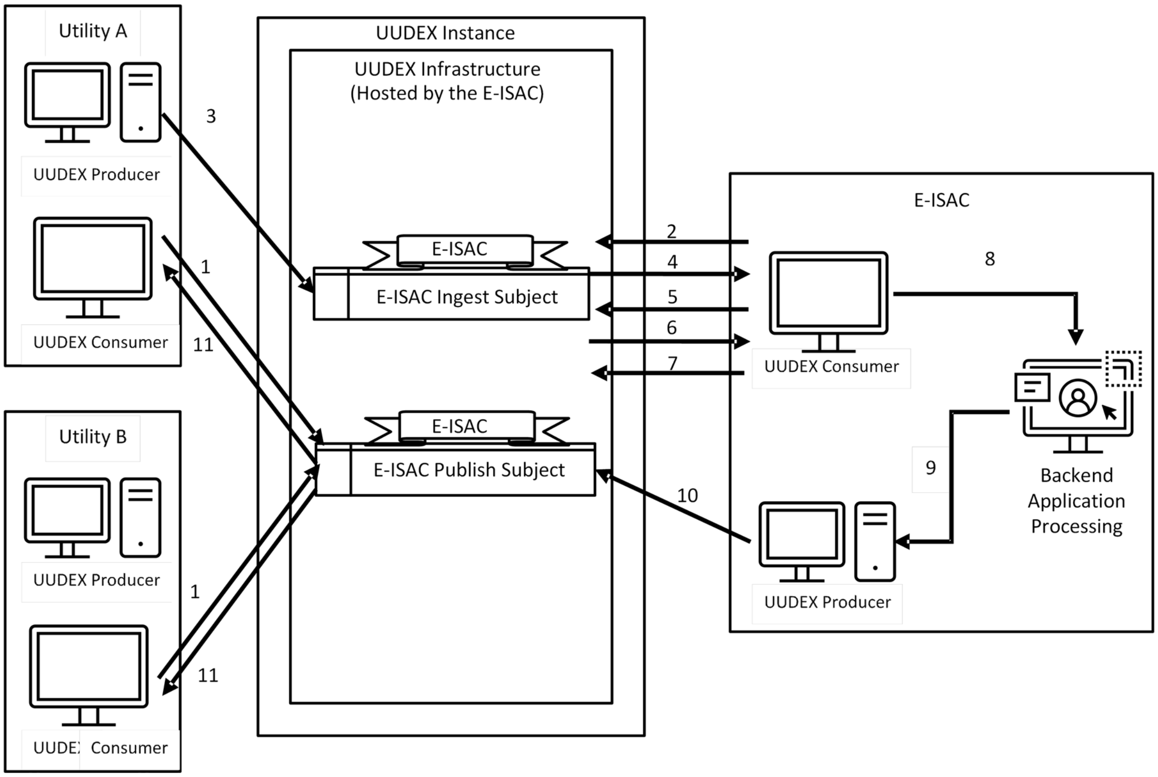

Figure 2-5: Data Sharing with the E ISAC

- The E-ISAC will establish a U-Subject for Electric Entities to publish security event data (the E-ISAC Ingest Subject). All Electric Entities will be able to publish to this U-Subject, but only the E-ISAC can read and delete content from the U-Subject.
- The E-ISAC will establish a U-Subject to publish security notices (the E-ISAC Publish Subject). Only it will be able to write to this U-Subject, but all Electric Entities are able to read from it.
- All Electric Entity participants will issue a subscribe request to the E-ISAC Publish Subject to be notified when the E-ISAC publishes a security notice (1).

- The E-ISAC U-Consumer will issue a subscribe request to its own E-ISAC Ingest Subject to be notified when security event data are published to the U-Subjects (2).
- Utility A publishes a security event to the E-ISAC's Ingest Subject (3).
- The UUDEX Infrastructure notes that the E-ISAC has a U-Subscription established on the E-ISAC Ingest Subject and sends a notification to the E-ISAC U-Consumer that new data have been added (4).
- The E-ISAC U-Consumer requests (5) and receives (6) the security event from the E-ISAC Ingest Subject in the U-Infrastructure and stores it locally.
- The E-ISAC U-Producer sends a delete data request to remove the security event from the E-ISAC Ingest Subject in the U-Infrastructure (7). (This may also be accomplished by specifying that the U-Data Element should be deleted from the U-Subject once the U-subscription processing has completed. The delete data request ensures that the U-Data Element is removed.)
- The E-ISAC develops an industry alert (8, 9) and uses its U-Producer to publish the alert to its E-ISAC Publish Subject in the U-Infrastructure (10).
- The U-Infrastructure sends a notification to each subscribing Electric Entity that an alert has been published to the E-ISAC's Subject for this data (11).
- Each utility U-Consumer requests and receives the alert if is relevant to their operations. (Not shown)

2.5 Power System Model Updates Published by an RC

The use case drawn in Figure 2-6 represents how large data sets (files) can be shared using an update notification that does not contain all the data to be shared. The use case also discusses how a U-Data Manifest listing of available large files could be made accessible in the event that obtaining a particular file is necessary.

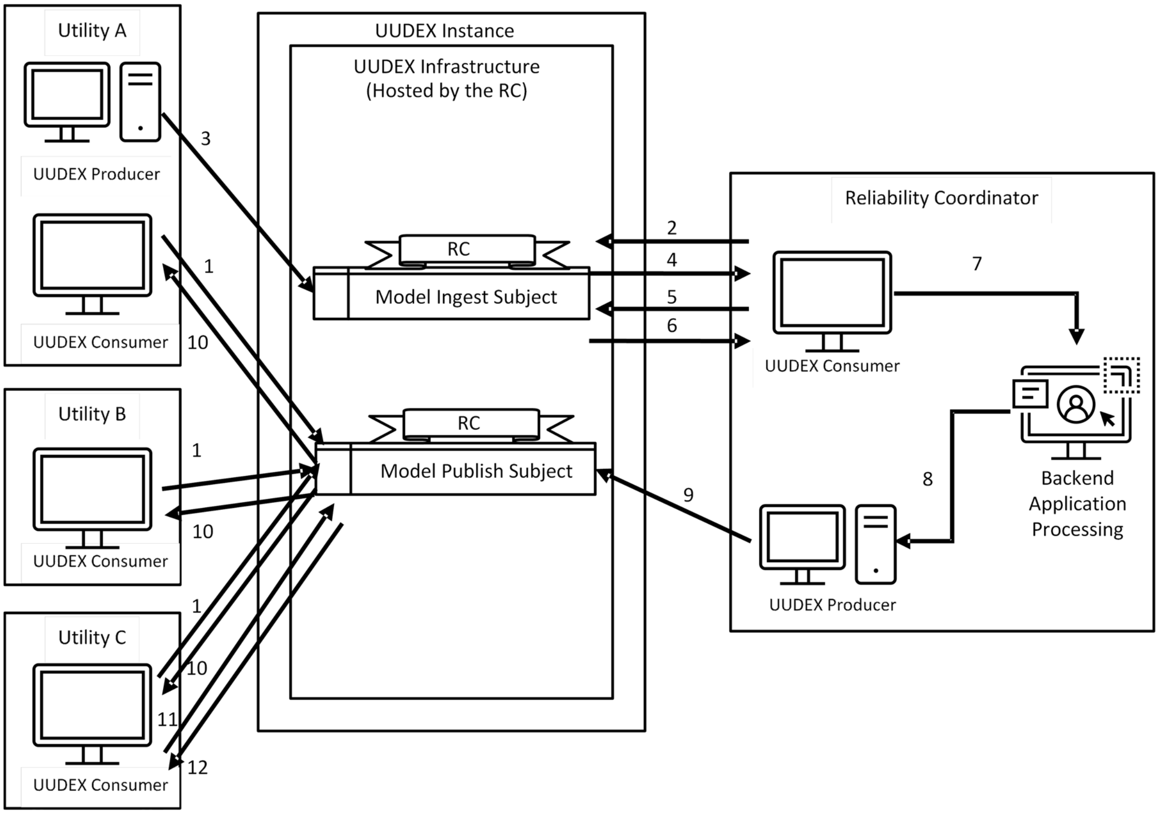

Figure 2-6: Power System Model Update Processing

- A central Electricity Entity (for example, an RC) will establish a U-Subject for sharing power system model updates (the Model Publish Subject). Multiple power system model update versions will be supported. The U-Subject will be created to maintain old version of U-Data Elements containing power system model updates.
- Other Electric Entities may subscribe to the RC's U-Subject for power system model updates they are interested in (1).
- The RC will establish a U-Subjects to which Electricity Entities will post their own model updates (the Model Ingest Subject). The RC will subscribe to each of these U-Subjects (2).
- An Electric Entity will send power system model updates to the Model Ingest Subject (3).
- The U-Infrastructure will alert the RC's U-Consumer that new data has been added to a U-Subject to which it is subscribed (4). The RC's U-Consumer will request (5) and receive (6) the power system model updates from that U-Subject.
- The RC will process the model update to produce an updated power system model (7, 8).
- The RC's U-Producer will publish power system model updates back to the Model Publish Subject (9).
- The U-Infrastructure will send an “update notification” message to each Electric Entity U-Consumer that has established a U-Subscription on the RC's U-Subject (10). The U-Subscription response contains some identifier that uniquely identifies the model file that was just published.

- Subscribing Electric Entities will receive the U-Subscription responses and determine if the power system model update should be requested.
- Electric Entities will use their U-Consumers to download power system model files applicable to them. (Not shown.)

Alternately:

- Any time an Electric Entity wishes to receive an updated power system model, the Electric Entity may query the RC's U-Subject to see what power system model updates are available (11).
- The U-Infrastructure would deliver power system model updates matching their query (12).
- In this case, the retrieving party collects new model information based on their own timeline, rather than responding to notifications from U-Subscription processing by the U-Infrastructure.

2.6 Patch Updates Published by the E-ISAC

The use case drawn in Figure 2-7 shows how UUDEX could be used by the E-ISAC, or another organization, to disseminate software patch notifications.

This use case is similar to the model updates, but patches may contain additional fields available for query (e.g., equipment type, version, patch metadata) that may also need to be communicated in the notification data. Such fields would be necessary for recipients to determine whether the patch is relevant to their local systems (i.e., whether they are running the software that the patch fixes). Fields available for query will need to be agreed upon by all U-Participants, as would the appropriate values for those fields. The latter would be necessary to avoid otherwise synonymous values producing different results (e.g., “Windows 10” vs. “Win10”).

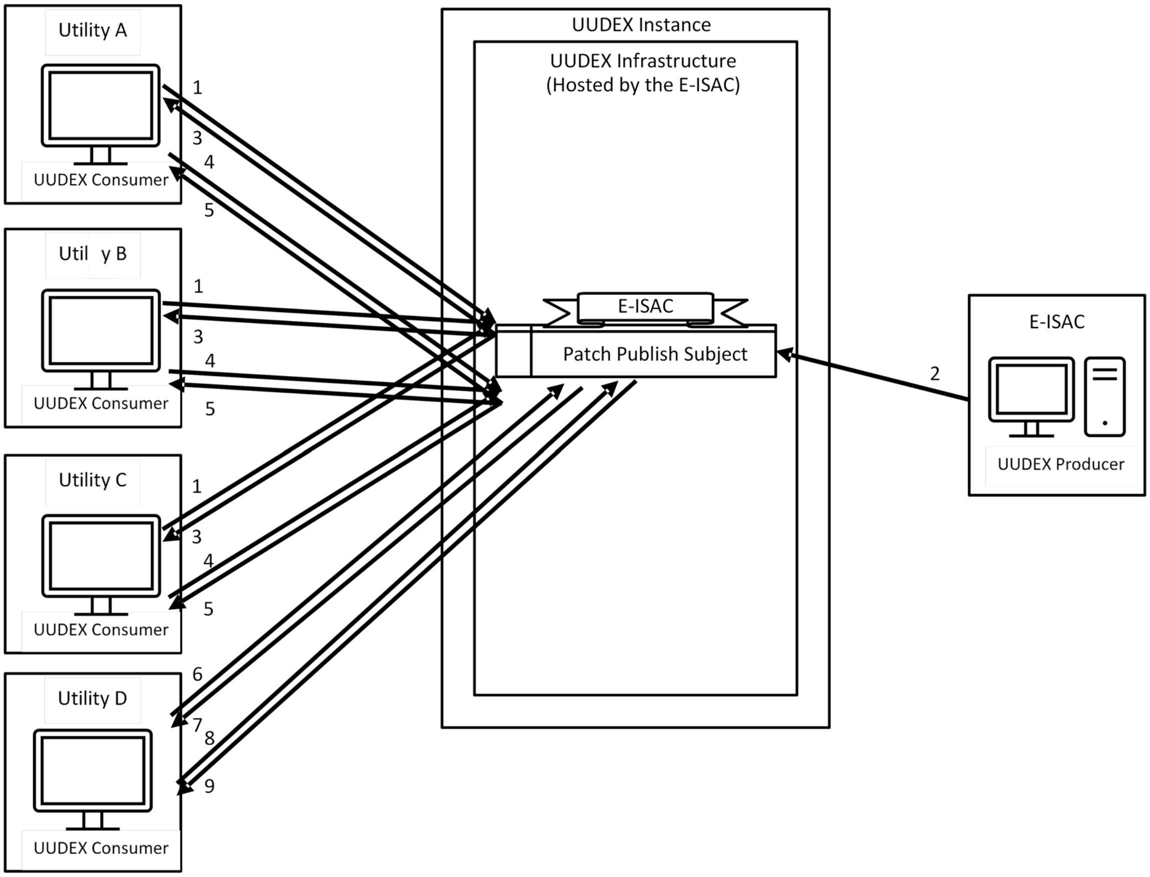

Figure 2-7: Patch Updates Published by the E ISAC

- The E-ISAC will establish a U-Subject.
  - The E-ISAC U-Subject will be configured to contain records for multiple types of patches but would all share a common U-Data Type indicating that U-Data Elements contain patch data. The U-Subject will be created to maintain old version of U-Data Elements containing patch updates.
  - The E-ISAC U-Subject will be able to be queried by Electric Utilities, but only the E-ISAC will be able to write to it.
- Organizations can use UUDEX establish U-Subscriptions to request notification whenever a new patch matching specified criteria is added to the E-ISAC's U-Subject (1).
  - Utility A indicates specifically the criteria matching the patches it is interested in.
  - Utility B and Utility C do not indicate which patches they are specifically interested in but request notification when any patch for any software is added. While this will result in many alerts about irrelevant patches, this might be desired because it means the software the utilities use is not exposed to the E-ISAC in their U-Subscriptions.
  - Utility D does not request to be notified when patches are made available.
- The E-ISAC will use the U-Producer to publish patch data to a U-Subject in the U-Infrastructure (2).
- The U-Infrastructure will send notifications to organizations that have U-Subscriptions that match the newly added patch data (3).

- Utility A receives the U-Subscription notification because the software matches the criteria in its U-Subscription. It requests (4) and receives (5) the patch file.
- Utility B and C receive the U-Subscription notification and request (4) and receive (5) the patch file, understanding that many might be discarded as irrelevant. However, the E-ISAC would be unable to determine what software Utility B or Utility C were using. (Since all communications are encrypted, an external entity wouldn't be able to tell which patches were being downloaded in any case.)
- Utility D does not subscribe to patch update notifications. Rather, it periodically queries the E-ISAC's U-Subject for a list of available patches (6). The U-Subject responds with a U-Data Manifest of patches available (7). Utility D determines which patches are necessary, and requests (8) and receives (9) the desired patches.

2.7 RCIS Messaging

The use case shown in Figure 2-8 is for a remote database update or query capability as used by the RCIS. There are many types of RCIS messages. While a given U-Instance could choose to treat all RCIS messages as a single U-Data Type and convey them over a single U-Subject, in this example, each RCIS message is given its own U-Subject. This could allow Electric Entities to select which RCIS messages they receive immediately and which ones they choose to poll for on an *ad hoc* basis.

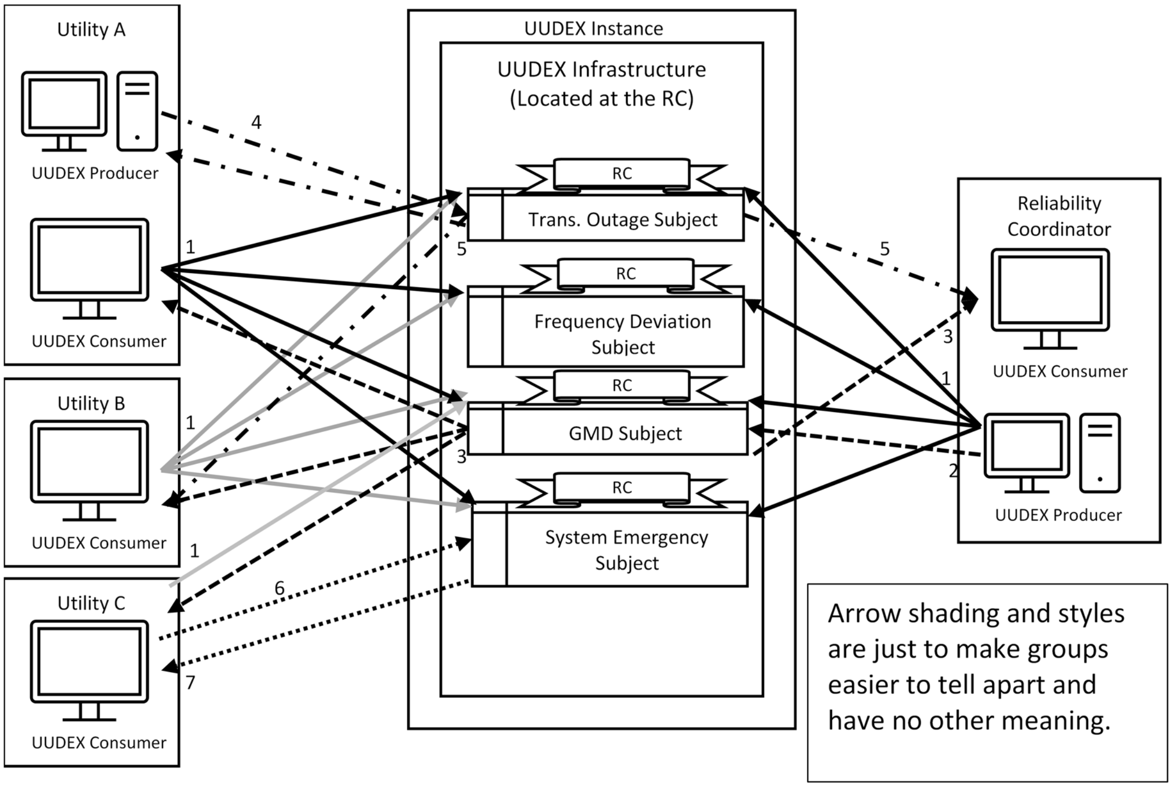

Arrow shading and styles are just to make groups easier to tell apart and have no other meaning.

Figure 2-8: RCIS Messages

- The RC will establish U-Subjects corresponding to each RCIS message type. All parties are granted rights to subscribe to and read the U-Subject. Depending on the type of message, either all parties or only the RC are granted rights to publish to the U-Subject. The U-Subject will be created to maintain old version of U-Data Elements containing RCIS messages making them available for query.

- Electric Entities and the RC will all create a U-Subscription with each of these U-Subjects for RCIS message types they are interested in (1). Note that Utility C only chooses to subscribe to the GMD subject, but not to the other subjects.
- As an example, the RC publishes a Geomagnetic Disturbance (GMD) message to the GMD subject (2). Notifications of a new GMD message are immediately pushed out to all subscribing parties (3).
- Later, Electric Entity A publishes a Transmission Outage message to the corresponding U-Subject (4). Notifications of the new Transmission Outage message are immediately pushed out to all subscribing parties (5). Not that, in this case, Utility C does not receive a notification because it did not subscribe to this U-Subject.

Alternately:

- Organizations can use the U-Consumer to query for information from the U-Subjects on an *ad hoc* basis. For example, Utility C could query the System Emergency Subject to see if there have been any updates (6). The U-Infrastructure will respond with the requested data from the U-Subject and return it to Utility C (7).

2.8 DOE OE-417 Reporting

The use case drawn in Figure 2-9 represents a mechanism for sending the same information (in this case incident reports) to multiple organizations.

The organizations receiving reports are noted as “Report Recipient Organizations” (and are the DOE, E-ISAC, NERC Bulk Power System Awareness, and an RC in this example). Each Report Recipient Organization will have their own U-Instance that it operates. Regulations limit how and when each Report Recipient Organization shares reports with the other. If a common U-Infrastructure was hosted by one of these Report Recipient Organization, having it resend messages to the other organizations could run afoul of those rules. Note that, if there was a mutually trusted third party hosting the U-Infrastructure, a single Infrastructure and U-Instance might be able to serve all parties. For this example, however, we are assuming that each Report Recipient Organization is hosting their own U-Instance and Infrastructure.

Note that the single Electric Entity has four U-Producers. This reflects the fact that a U-Participant's identity is tied to a single U-Instance. As such, the fact that each Report Recipient Organization hosts a separate U-Instance means that the Electric Entity is using a separate identity when it interacts with each U-Instance. In fact, the Electric Entity could be using the same hardware and software to interact with all four U-Instances and simply have that software authenticate using different identities in each case. However, for clarity, the example uses separate U-Producer icons for each interaction, reflecting that each U-Producer has a distinct identity.

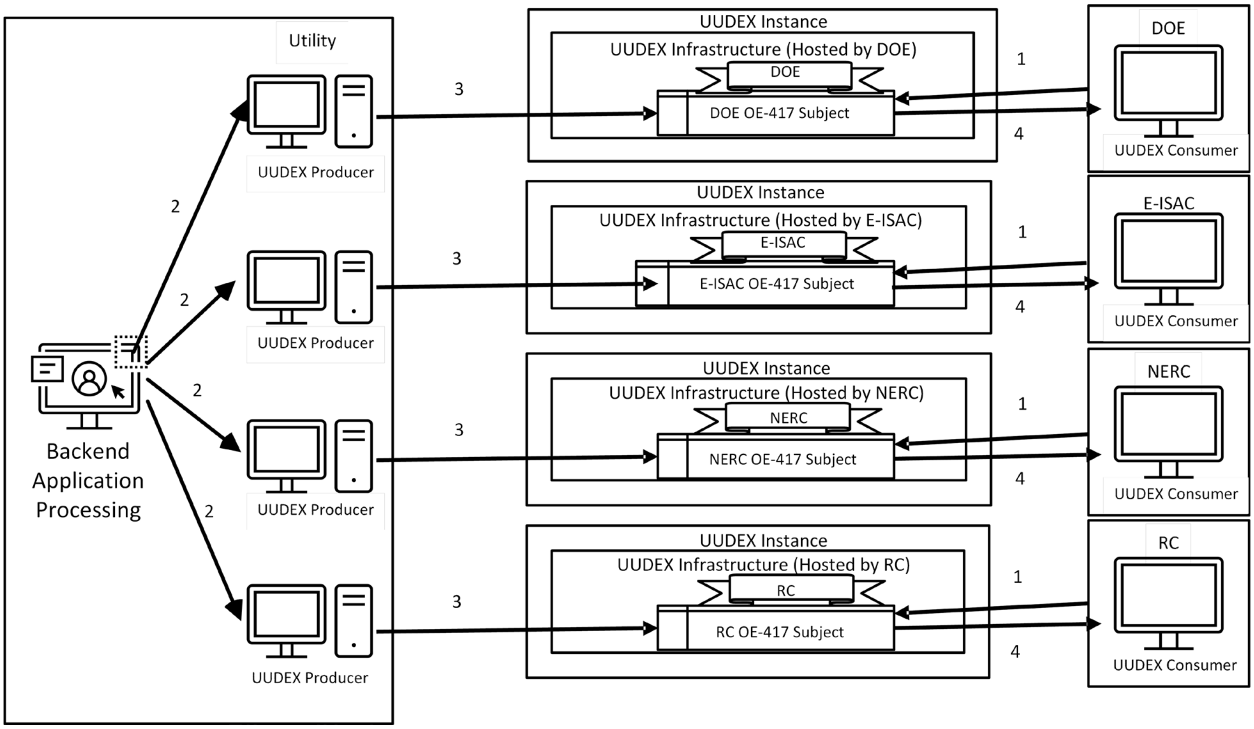

Figure 2-9: DOE OE 417 Reporting

- Within their own U-Instance and U-Infrastructure, each Report Recipient Organization establishes a U-Subject for Electric Entities to publish OE-417 reports. All Electric Entities are allowed to publish to this U-Subject, but only the Report Recipient Organization that operates the U-Instance is allowed to read the U-Subject. The U-Subject may be created to either maintain old version of U-Data Elements or delete them once all the U-Subscriptions have been fulfilled, depending on the backend processing for the specific Report Recipient Organization.
- The Report Recipient Organization establish a U-Subscription to their OE-417 U-Subject (1).
- Electric Entities will generate OE-417 report data (2) and use their U-Producer to publish the data to the U-Subject (3). Note – the same OE-417 report is sent to all Report Recipient Organizations.
- Whenever new information (i.e., a new OE-417 report) is posted to their OE-417 U-Subject, a Report Recipient Organization's U-Consumer is automatically notified and pulls the report down (4). Once retrieved by the U-Consumer, the report can be processed as necessary. In the case, if the OE-417 data were expressed using a standardized data model, the first step in this processing will likely to be to generate an actual OE-417 form from the data.

2.9 Emergency Responder Information

The use case drawn in Figure 2-10 represents how an existing U-Subject could rapidly be configured to provide specific information to a new type of organization.

The prime example of such a goal is the need to respond to a natural disaster or other crisis. In such a case, emergency responders (state, federal, local, or non-governmental organization) might be granted specific, limited information about the current and anticipated state of the power grid to plan their efforts.

For example, following a natural disaster, police, fire, and government response organizations need to know which portions of the electric grid are energized, which are de-energized, and the approximate order of restoration. Traditionally this information is available by telephone or email correspondence on an infrequent basis. By using a U-Subject within the U-Infrastructure that can be queried for current outage information and updated in near real time by an Electricity Entity's OMS, the information can be provided in near real time to the first responder with no processing or impact to the Electricity Entity other than establishing the initial link for the emergency responder.

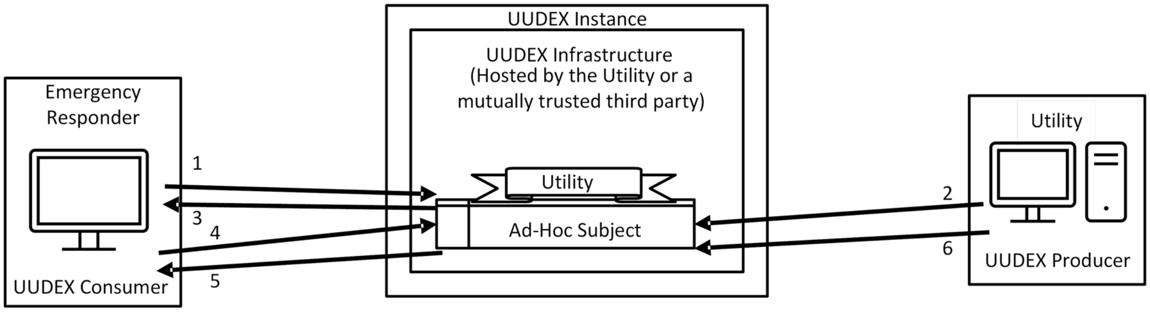

Figure 2-10: Emergency Responder

NOTE: the software used by the emergency responders to display the outage information is beyond the scope of the UUDEX project.

- The Electric Entity (for example a TOP or Distribution Operator) is already a member of some U-Instance. The Emergency Responder is onboarded to this U-Instance as a new U-Participant.
- The Electric Entity will establish a U-Subject for sharing with the Emergency Responder. (Ad-Hoc Subject) Alternately, they might grant access to outage information already present on an existing U-Subject that contains that information. The U-Subject will be created to maintain old version of U-Data Elements containing outage information to minimize the impact of adding a new Emergency Responder client to the U-Instance.
- Emergency responders can use a U-Consumer to subscribe to updates to the outage information from that U-Subject (1).
- The Electric Entity will extract information from its OMS and publish outage information and outage updates (in standardized format) to the U-Subject (2).
- The U-Infrastructure will send updates to emergency responders in fulfillment of their U-Subscriptions to the U-Subject (3). Once downloaded, an application program at the emergency responder location can process and display the outage information to facilitate emergency operations.
- Alternately, emergency responders may use the U-Consumer to request information from the U-Subject on an *ad hoc* basis rather than using a U-Subscription to receive alerts for every change to the U-Subject (4, 5).
- After a period of time, the Electric Entity may delete the outage information from the U-Subject (6).

2.10 Mass Alert

The use case drawn in Figure 2-11 represents the situation where a U-Subject's data include information that needs to be distributed to a broad audience quickly.

Examples of this could include alerts from the E-ISAC regarding an active cyberattack campaign for which a broad group of U-Participants should act. While the E-ISAC and NERC Bulk Power System Awareness are shown in this example, other organizations, such as DOE, U.S. Department of Homeland Security (DHS), or an RC, could use the same process to send alert information to a large number of U-Participants.

In this example, it is assumed that a trusted third party is managing the U-Instance and corresponding U-Infrastructure.

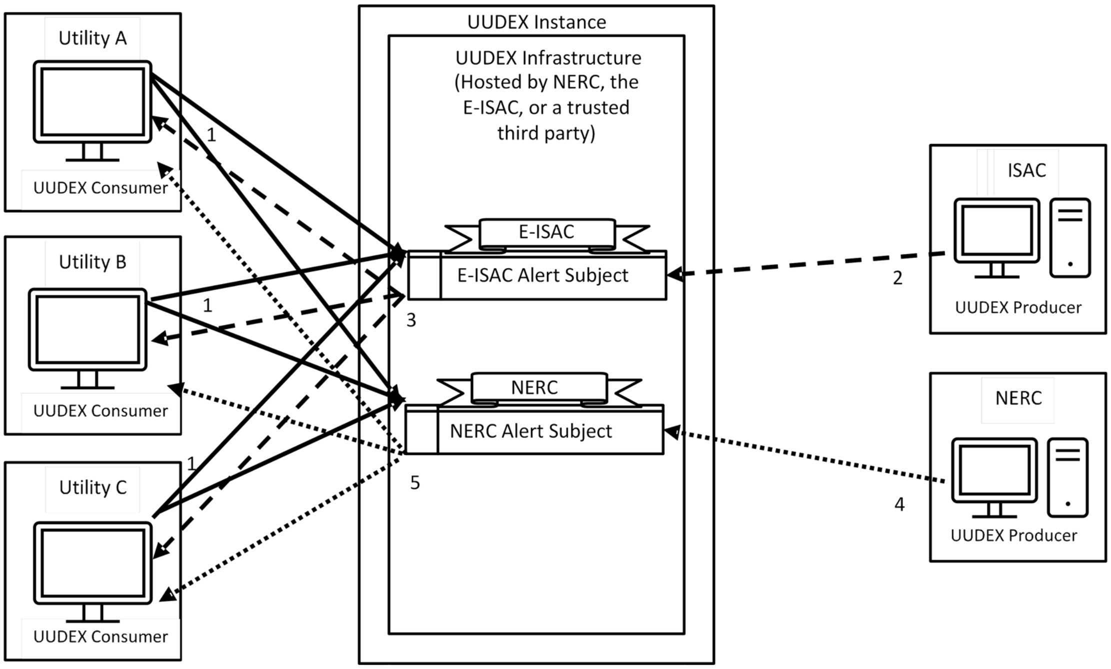

Figure 2-11: Mass Alert

- Each party originating the alerts will stand up a U-Subject for these alerts. (E-ISAC Alert Subject for the E-ISAC and NERC Alert Subject for the NERC.) The originating party will configure their U-Subject to have a high delivery priority to ensure timely delivery even when the network is degraded. Only the originating party will be allowed to publish to this U-Subject, but any party will be allowed to subscribe to the U-Subject. Both U-Subjects should be configured to delete the U-Data Elements once the U-Subscriptions have been fulfilled to minimize the possibility of inadvertent release of sensitive information.
- All parties that should receive the alert will establish U-Subscriptions with the alerting U-Subjects (1).
- When an alert is generated (2), it is immediately added to the appropriate U-Subject.
- When the alert is added to the U-Subject, the U-Infrastructure immediately compares the U-Data Element for this alert against U-Subscription criteria. This will match all

U-Subscriptions for immediate delivery of the alert, leading to fulfillment of the U-Subscriptions (3). Each utility receiving an alert would process it.

- Later, a different party might issue an alert, which passes to its corresponding U-Subject (4) and notifications are sent out to all subscribing parties (5).
- When the alert is delivered, UUDEX's reliable message delivery mechanisms will allow the U-Infrastructure to know that the given U-Consumer received the message.
- Alert messages could remain in the U-Subject. This would provide archiving of alerts and would allow parties to query for old alerts at a later date.

Variation:

- It might be desirable to have a hierarchical distribution chain, where the originator sends messages to other entities acting as relays, who then pass the message on to other parties. There might be multiple levels in this hierarchy. This might reflect operational responsibilities and thus better mesh with existing processes.
- In this case, each level of the distribution hierarchy would have their own U-Consumer, U-Producer, and U-Subject. Each distribution node's U-Consumer would subscribe for updates from the U-Subject of the next entity above them in the distribution tree. When alerts are received by this U-Consumer, they are immediately posted to that entity's U-Subject by its U-Producer.
- All children in the distribution tree would have U-Consumers who had established U-Subscriptions to this U-Subject, which would immediately be fulfilled when the new data are added. This could be repeated through any number of levels in the distribution hierarchy.

2.11 Mass Alert with Required Response

The use case drawn in Figure 2-12 is a variation of the mass alert use case in section 2.10, where a response from the recipient of an alert is required.

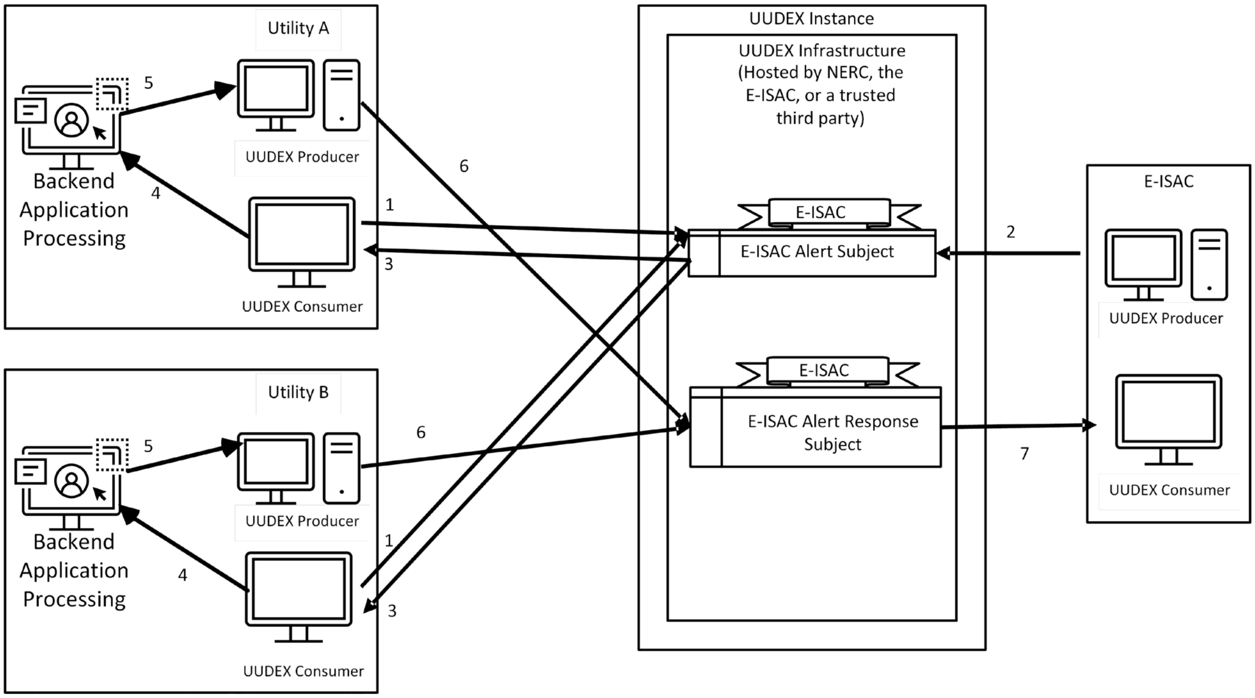
Utility A
Backend Application Processing
UUDEX Producer
UUDEX Consumer
Utility B
Backend Application Processing
UUDEX Producer
UUDEX Consumer
UUDEX Instance
UUDEX Infrastructure (Hosted by NERC, the E-ISAC, or a trusted third party)
E-ISAC
E-ISAC Alert Subject
E-ISAC
E-ISAC Alert Response Subject
E-ISAC
UUDEX Producer
UUDEX Consumer
1
2
3
4
5
6
7

Figure 2-12: Mass Alert with Required Response

In this use case, a high-priority alert that requires an asynchronous (i.e., after a period of data gathering) response is distributed.

- The first part of this use case would be identical to the mass alert use case.
  - An alert originator stands up a U-Subject. The U-Subject should be configured to delete the U-Data Elements once the U-Subscriptions have been fulfilled to minimize the possibility of inadvertent release of sensitive information.
  - All parties that should receive the alert establish U-Subscriptions for those alerts to the originator's U-Subject (1).
  - Generated alerts (2) are automatically added to the originator's U-Subject.
  - Adding to the U-Subject triggers comparisons against U-Subscription patterns.
  - The alert is sent to all subscribed U-Consumers in fulfillment of the U-Subscriptions (3).
  - When the alert is delivered, UUDEX's reliable message delivery mechanisms allow the U-Infrastructure to know that the given U-Consumer received the message.
- The UUDEX Header (U-Header) of the message will have the "Response Required" metadata field set, which will flag the message as requiring follow up. The U-Header for the alert will contain a unique message identifier (i.e., a universally unique identifier [UUID]) to verify and track submitted responses.
- Each recipient's U-Consumer will hand off the message to the organization's internal processes (4). They will handle identifying what response is necessary and crafting that response.
- The Utility will submit their response to the alert. The nature of the response will depend on the characteristics of the alert itself.

  - The alert might dictate a specific response mechanism external to UUDEX (e.g., phone call, email to a given address). In this case, UUDEX will not be employed in the response and the organization's internal processes will be responsible for ensuring the response occurs.
  - The alert might allow for responses to be delivered via UUDEX (5, 6). In this case, responses would be sent to a U-Subject established by the responding party and to which the originator of the alert is subscribed.
  - A U-Consumer at the party originating the alert will gather and process responses sent to the responder's U-Subject (7).
- As the deadline approaches, the originator could compare the list of response providers to the expected list of respondents and send reminders to those whose response is missing.

Variation:

- As with the mass alert use case, a hierarchical distribution model could be used to send out the alerts.
- In this case, the relays that delivered messages to U-Consumers would need to be the parties that collected delivery confirmation and sent out reminders to those late in responding. In the latter case, this would require coordination between the relays and the originating party to which responses need to be sent, since these might not be the same parties.

3.0 Data Exchanges supported by UUDEX

This section identifies the initial set of data exchanges UUDEX will support. Two criteria were used to create this shortlist. First, based on findings from the first set of industry interviews, data exchanges that were essential for day-to-day operation were included. Presently, these exchanges are largely performed using individual mechanisms that have several disadvantages ranging from cumbersome link setup to lack of uniform data representation. Second, data exchanges that we identify as critical to secure grid operation in the future will also be supported. These primarily include cybersecurity-specific data exchanges that will support the dissemination of threat, vulnerability, and security upgrades to enhance the overall security posture of the stakeholder.

The sections below are included as examples of some of the different types of data that UUDEX can convey. UUDEX defines the concept of a UUDEX Data Element Type (U-Data Element Type), which is a defined structure for conveying certain types of information over U-Data Elements. UUDEX will standardize certain data element types, either providing wrappers for existing data structures already in use operationally, or by defining new structures as needed. U-Instances are also free to develop their own U-Data Element Types. This section outlines some of the different classes of information that may be standardized as U-Data Element Types. Some of the types of data below might be revised and undergo changes in structure when turned into U-Data Element Types, possibly in such a way that both old and new structures are employed simultaneously within a given environment. Ultimately, UUDEX's ability to exchange data is agnostic with regard to the nature and format of that data, so UUDEX will be able to handle exchange of new or updated forms of data. Thus, the following sections are intended to call out some of the ways UUDEX might fulfill current data exchange needs, but with the recognition that UUDEX's capabilities go beyond these examples.

3.1 Grid Operational Data

The primary data exchanged by UUDEX are grid operational data. These include analog and status values telemetered by TOPs from substations and plants and shared with neighboring TOPs, RCs, and others to allow them to analyze grid conditions, and perform required functions to ensure grid reliability and maintain situational awareness.

3.1.1 ICCP Data

ICCP is used extensively by utility organizations to exchange grid operational data over wide-area communication networks. The types of information supported by ICCP include analog values, binary status data, control signals, and schedules. When analog values are conveyed, each is typically represented as a real or integer value along with data quality flags that provide further context about the data. These include flags to highlight data source (options: telemetered, calculated, estimated, and manually inputted), flags to indicate data normalcy (options: normal or abnormal), and flags to specify data validity (options: valid, not valid, held, and suspect). Typically, SCADA applications use a combination of these flags to sufficiently describe the conveyed data during a wide range of scenarios, including normal operation, telemetry failure, stale data detection, data out of range, and data conversion errors. The protocol also defines the capability to include a timestamp for the data value at the source. In the case of data values calculated from telemetry at the source, the timestamp to be included is left to the implementation of the application. Binary status data are represented by bitmasks that support per-phase representation. The data quality and timestamp attributes are used in this case to provide additional context. The control capability supported by ICCP includes device switching, raise/lower commands, set-point specification, and device turn on/off. ICCP also supports the ability to exchange scheduling information (e.g., generator schedules, interchange schedules, pricing information) between a client and server.

Each data value communicated using ICCP has a tag used by the application software at the recipient to associate the value with a specific measurement object in an application data model. It is the responsibility of applications on each end of the data link to establish any mappings that are necessary as well as data conversions to get the values to a representation where the scaling and units on both sides are well understood.

The data exchanged using ICCP are organized into "conformance blocks" representing fundamental types of service supported by the standard. These include:

Block 1 – Periodic Power System Data

The data objects categorized under this block are used for periodic exchange of field device status, analog measurements, and accumulator values between a client and server. As introduced earlier, these exchanges could be accompanied by suitable data quality flags that provide additional context about the exchanged data.

Block 2 – Extended Data Set Condition Monitoring

Also referred to as "report by exception," this block provides a client the capability to configure report generation and transmission by exception. Examples of exceptions include cases where the value or quality code of a particular data point has changed an operator-initiated push of value from server to client.

Block 3 – Block Data Transfer

This block defines an optimized method of transferring data described in Blocks 1 and 2 as groups of data, rather than individual enumerated values. The optimization is achieved by combining individual data values into blocks, removing the tags associated with each data value for mapping to suitable point at the recipient, and also removing the length fields associated each data value. In place of the variable name-based tags used for mapping in Blocks 1 and 2, this block employs an index-based tagging mechanism.

Block 4 – Informational Messages

The exchange of ASCII text and binary files is supported by this block. It is typically used for the exchange of complex information that cannot be conveyed using the other blocks (e.g., text file describing an emergency situation or system restoration summary).

Block 5 – Device Control

The device control block provides a client the capability to request to operate a remote device. As introduced earlier, the controls that can be implemented using this block range from switching actions to set-point specifications on the remote device. In addition, ICCP supports both interlocked operation, wherein select-before-operate confirmation is required, and non-interlocked operation.

Block 6 – Program Control

Block 6 introduces the capability for an ICCP client to control a program on the remote server.

Block 7 – Event Reporting

This block provides the capability for ICCP clients to receive event-specific information to which the client has subscribed.

Block 8 – Additional User Objects

This block provides the utility with the means to exchange scheduling and accounting information; examples include generator schedules, interchange schedules, and pricing information.

Block 9 – Time Series Data

This block adds the capability to exchange time series data between an ICCP client and server. This is useful to applications that do not need data sampled at a very high rate in real time (e.g., post-disturbance voltage data recorded in millisecond intervals for analysis).

The third edition of the ICCP specification (IEC 60870-6-503:2014), released in 2014, made Blocks 6, 7, 8, and 9 out of scope. UUDEX will provide a mechanism to transport data described in Blocks 1, 2, 3, and 4. It will not provide device control as specified in Block 5 or program control as described in Block 6. The flexibility of information modeling in UUDEX may allow exchange of information from Blocks 7, 8, and 9.

3.1.2 Reliability Coordinator Information System

RCIS is used primarily by RCs to post information concerning reliable operations of the Bulk Electric System (BES). The current RCIS is a web-based system that functions much like a message bulletin board, allowing users to post messages to the system and providing access for other users to view posted messages. Some messages are generated and reported autonomously by software at the RC. Users can either monitor the web interface for new activity or can sign up to receive emails when new messages are posted. Most RCIS messages are free-form text messages, although some message types have a limited message structure.

Current RCIS message types include:[5]

- **Critical Infrastructure Protection – DHS** has been supplanted by reporting through the E-ISAC portal.
- **Critical Infrastructure Protection – Free Form** has also been supplanted by reporting through the E-ISAC portal, but it is still used to alert control room staff of issues that would otherwise not be made known to them in a timely manner, especially if access to the E-ISAC portal is not provided to the control room operator 24x7.
- **Energy Emergency Alert reporting** by RCs is required by NERC Standard EOP-002 from the time such an alert is issued to the time the alert has been cancelled.
- **Frequency** is used for communicating system events that have or could result in a rapid change in frequency that significantly impacts system operation; also used to report changes in frequency for which the cause is unknown.

[5] Much of the information was extracted from minutes of the NERC Reliability Coordinator Working Group of May 3, 2011, “Exhibit C” located at: https://www.nerc.com/comm/OC/RCWG%20Agendas%20Highlights%20and%20Minutes%20DL/Agendas,%20Highlights,%20and%20Minutes%20-%202011/RCWG_Minutes_3May11.pdf (no longer available)

- **Geomagnetic Disturbance** information originates from the National Oceanic and Atmospheric Administration's Space Weather Prediction Center and is made available by designated RCs to receive and disseminate notifications of possible geomagnetic disturbances to RCs, Balancing Authorities (BAs), and TOPs.
- **System Emergency** – used to provide notification when an RC foresees transmission problems (such as a system operating limit [SOL] or interconnected reliability operating limit [IROL] violation, loss of reactive reserves, etc.), when results of operational studies for the current day or the next day indicate that there is potential for SOL or IROL violations, or when an interconnected system separation, system islanding, or blackout has occurred.
- **Transmission Outage** – messages relating to transmission line outages for facilities greater than 230 kV, automatically generated and sent to the System Data Exchange database and posted on the RCIS.
- **Generation Outage** – messages relating to generation facility outages greater than 300 MW, automatically generated and sent to the System Data Exchange database and posted on the RCIS.
- **Time Correction** – indication of the start and end of time error correction within an interconnection.
- **Transmission Line Loading Relief** – messages relating to transmission loading relief following NERC Standard IRO-006, automatically generated by the Interchange Distribution Calculator and posted to the RCIS.
- **Weather Advisory** – notifications of approaching or existing severe or extreme weather conditions that have the potential to affect system reliability. These conditions could include severe heat or cold, insulator ice bridging, large thermal generation limitations (due to fuel restrictions), tower damage (due to tornado, hurricane, or flooding), extensive ice storms, galloping on transmission circuits, forest fires, etc.
- **Free Form** – designed to capture situations that an RC determines to be appropriate to communicate with other RCs, BAs, or TOPs regarding an issue that is not directly related to any of the other message board categories available on RCIS. A number of free-form message classes have come into use, even though there is not a specific RCIS message class for them. UUDEX should consider specifically creating message classes including but not limited to the following:
  - **Test & Maintenance** – messages indicating testing of the RCIS system, such as test notification, testing of successful message submission, or maintenance of the RCIS.
  - **Drills & Exercises** – notification of operational drills, such as testing of backup control center locations, testing of new communications facilities, (parallel) testing of new platforms and applications, etc. (Note: actual evacuation should be reported as an Emergency).
  - **Software Issue** – notification that the RC is experiencing EMS or software issues, such as the State Estimator not converging, or installation or testing of major software updates.
  - **Timing Integrity Issue** – notification that an RC, BA, or TOP has detected a timing integrity issue such as that caused by spoofing the time from a Global Positioning System (GPS) receiver. The issue must be of sufficient magnitude to affect the timing alignment of SCADA, PMU, or other time-sensitive data across an RC, BA, or TOP geographic area.

- **Theft, Burglary, Vandalism** – reports of (mostly) nuisance events that do not have a direct impact on operations, like copper theft, substation break in, and bomb threats.
- **Non-Transmission Emergency** – similar to System Emergency, but not for transmission issues. Examples include control room evacuations due to fire or bomb threat, or physical damage to a transmission station (like fire or flood) that does not necessarily induce an IROL violation, islanding, or cascading (therefore not qualifying as a System Emergency).

3.1.3 Power System Models

Power system models that enable real-time and study simulation of the electricity grid are commonly exchanged between utilities and grid operators. These models describe the electrical connectivity between objects such as transformers, breakers, generators, meters, transmission lines, and distribution feeders, and then expand to include aspects such as associated measurements, electrical impedance characteristics, asset information, etc. Given the many uses of these models, profiles are defined to identify the information (e.g., classes, attributes, and relationships) needed for a given use. The models may be exchanged in forms that are suited toward a given usage. This is a realization that there is diversity in the set of information needed for use by applications such as power flow, state estimation, distribution outage management, metering, or planning.

One form used is Common Information Model/Extensible Markup Language (CIM/XML), as defined by the International Electrotechnical Commission (IEC) 61970 and IEC 61968 series of standards, where files can represent full or incremental updates to the model. These files convey the electrical characteristics of power system resources and their connectivity relationships.

There is no implicit restriction as to whether models exchanged relate to transmission, distribution, substations, or even microgrids. Models exchanged may (or may not) carry graphical relationships, which may be based on geographic (e.g., GPS) or schematic coordinate systems.

3.1.4 Phasor Measurement Unit Data

PMU data can be characterized as fined-grained, time-series voltage, current, and frequency data that are time-stamped at the source (the PMU) against GPS time. Timestamping using a GPS enables the synchronization of measurements from geographically dispersed PMUs, providing an accurate representation of the power system state at a given time. The advent of PMU technology has ushered in the development of novel control applications that offer significant benefits to improving power system reliability. These applications can be broadly classified into the following three categories: automated closed-loop control, human-in-the-loop control, and offline analyses.

Automated closed-loop controls operate very quickly from the sensing of an event or disturbance to the execution of a control action. This type of control application typically leaves the human out of the loop, although there could be use cases demanding instantaneous intervention from a system operator. Applications that fit this description include protection, fast-reactive switching, damping controls, resource integration support, and alarming and operating limit monitoring among others.

Human-in-the-loop controls involve the analysis of processed phasor measurements from an event or disturbance by a system operator, who then proceeds to implement specific control actions, if deemed necessary. These applications, implemented with the objective of providing system operators with situational awareness, have a more lenient expectation on communication latency. Applications of this type include tools for wide-area situational awareness, voltage monitoring and trending, dynamic line ratings calculation, outage restoration, and operations planning.

Offline analyses carried out using historical PMU data have the potential to enhance overall reliability of the grid. System planners benefit significantly from the additional insight received by simultaneously studying synchronized measurements from multiple locations in the grid, enabling them to improve operational strategies. Applications that fall under this category include power system baselining to support predictive tools, post-event analysis for event reconstruction and reoccurrence prevention, static and dynamic model calibration, load characterization and modeling, and design and testing of special protection systems (or remedial action schemes).

As stated earlier, UUDEX is not intended for control purposes, eliminating the need for extremely low-latency communications and the requirement to support PMU-based automated closed-loop controls. However, UUDEX is capable of supporting human-in-the-loop controls and offline analyses. In both these cases, UUDEX merely functions as a conduit for time-stamped measurements from the source to the system operator or planner for processing. It will not support the execution of any control actions identified by the applications.

UUDEX will consider protocols under development such as STTP as well as techniques that provide for down-sampling of the information in a form that might be useful to a wide array of applications.

3.1.5 Industry Incident Reports

The electricity industry has several required formats to report incidents that affect reliable operations to DOE, NERC, or the E-ISAC. Additionally, NERC developed a guideline in 2008 for Threat and Incident Reporting[6] that provides guidance for voluntary reporting.

The DOE report form (OE-417) and the NERC Standard CIP-008 refer to the reported items as “incidents,” while the NERC Standard EOP-004 refers to them as “events.” For purposes of UUDEX, incidents and events refer to the same thing.

The NERC guideline document discusses topics that should be reported, as well as suggested reporting timeframes, but does not specify a particular format or data fields that should be included in the reports.

3.1.5.1 DOE OE-417 Electric Emergency Incident and Disturbance Report

The DOE OE-417 Electric Emergency Incident and Disturbance Report is used to report certain electrical, operational, cybersecurity, and physical security incidents to the DOE, with optional

[6] Document marked as “under revision”; see https://www.nerc.com/files/Incident-Reporting.pdf (no longer available)

copies to NERC and the E-ISAC. The reports must be filed within 1 hour (emergency reports), 6 hours (normal reports), and 24 hours (system reports). The report form includes sections for:

1. Alert criteria, including whether the report is “Emergency,” “Normal,” or “System,” and providing a report status, along with information about who is filing the report (organization name and address).
2. Information about where and when the disturbance or incident occurred and whether the incident involved load or customer outage.
3. Information about the type of emergency, including its cause, impact, and any actions taken.
4. Free-text information, including contact information for the person making the report, a free-form block to describe the incident, estimated restoration time, names of any assets (electrical) that were impacted, and indication of whether the information should be shared with NERC or the E-ISAC.

For implementation within UUDEX, in addition to transporting the PDF or Word OE-417 file itself, the field identifiers and values can be extracted from the form and transported using a U-Data Element Type. Those field values can be reimported into its respective document format (PDF or Word) at the destination. This ability notwithstanding, it will always be possible to transport the whole file itself as a U-Data Element.

There are certain sub-documents of DOE OE-417 reports and NERC Standard CIP-008 reports that are assigned a 1-hour time limit by OE-417 and CIP-008 as the upper bound on reporting. The reason for the time limit is that it enables national authorities who receive these reports (DOE, E-ISAC, and the Industrial Control Systems Cyber Emergency Response Team [ICS-CERT]) to determine whether there is a coordinated cyberattack underway against the BES and to issue security guidance as quickly as possible to mitigate its further impact and geographic spread.

While a 1-hour time limit normally is not a message latency requirement that is difficult to meet, this requirement must also be met under network congestion conditions. To ensure the time limit is always met, the U-Instance operators should consider prioritizing U-Subjects that contain U-Data Elements for OE-417 and CIP-008 reports. As discussed later, UUDEX supports a way to prioritize message delivery so higher priority messages have a better chance of being delivered despite network congestion.

3.1.5.2 NERC Standard EOP-004

NERC Standard EOP-004 requires that NERC Responsible Entities must file electrical disturbance reports to NERC within 24 hours following the disturbance. NERC will accept an OE-417 report or information following a sample form provided in the standard. Copies are also to be sent to the appropriate Regional Entity, company personnel, the Responsible Entity’s RC, law enforcement, or the Applicable Governmental Authority, as described in Requirement R1 of the standard.

Information required to be submitted includes:

1. Information pertaining to the reporting entity.
2. Date and time of the event.
3. An indication of whether the event originated within the organization’s part of the electrical system.

4. Event identification and description, including selection boxes and a free-text field.

All EOP-004 field values can be mapped into OE-417 field values. For example, the EOP-004 field value "3: An indication of whether the event originated within the organizations part of the electrical system" can be mapped into the OE-417 field "13: Damage or destruction of a Facility within its Reliability Coordinator Area, Balancing Authority Area or Transmission Operator Area that results in action(s) to avoid a Bulk Electric System Emergency."

Like the OE-417 form, UUDEX can define a U-Data Element Type so that fields and values can be extracted from the EOP-004 form. These values can be transported over UUDEX and ultimately reimported into a destination form. This ability notwithstanding, it will always be possible to transport the whole file as a U-Data Element.

3.1.5.3 NERC Standard CIP-008

NERC Standard CIP-008 requires that NERC Responsible Entities file cyber incident reports with the E-ISAC for specific kinds of "Reportable Cyber Security Incidents" within 1 hour of determining a report should be made. Previous versions of the standard did not prescribe any specific format or data fields that should be included in the report.

As a result of the directives in a FERC directive issued on July 19, 2918, modifications to CIP-008-6 were made to include the following when reporting Cyber Security Incidents:

1. Functional impact that the Cyber Security Incident achieved or attempted to achieve.
2. Attack vector that was used to achieve or attempt to achieve the Cyber Security Incident.
3. Level of intrusion that was achieved or attempted as a result of the Cyber Security Incident.

These reports should be submitted to the E-ISAC and appropriate governmental authorities

UUDEX can develop a U-Data Element Type for fields associated with a CIP-008 report that encompasses the minimum information required by FERC (items 1, 2, and 3 above). If information for an OE-417 report is present, then item 1 will be satisfied by the "impact" information associated with the OE-417 reported fields.

As noted under the discussion of OE-417 reports, some CIP-008 reports might have time constraints on their delivery. As such, U-Instances may wish to ensure their U-Subject used to convey CIP-008 reports are configured to prioritize the delivery of these reports.

3.1.5.4 Physical Security Incident Reporting

While NERC Standard EOP-004 and the DOE OE-417 report both have provisions for reporting physical security incidents, there is no standardized format for these reports nor is a comprehensive list of collected information provided in either report. The UUDEX project can develop a comprehensive standardized report that captures essential information for a complete report.

3.1.6 Files (COMTRADE [IEEE Std C37.111^{TM}] and Others

Evolution of business processes and the regulatory environment can cause development of new types of information. To support this evolution, there needs to be the ability to exchange these new types of information without the necessity of software changes to the communication infrastructure or associated interfaces. This necessitates the ability to exchange files independent of type or data format.

To accommodate this, there needs to be:

- A way to indicate the type of file transmitted (such as by providing a media type, as identified in the Internet Assigned Numbers Authority Media Types registry[7]). There also needs to be a way to describe new file types on an *ad hoc* basis or refine a named type to indicate specific use of a broader media type.
- A set of data models that may be used to define the structure of some documents as a way to improve interoperability.
- UUDEX metadata fields that provide information for each file that might include, but not be limited to:
  - File type (as may relate to a specific application or defined schema)
  - File name (need not be unique, may be hierarchical)
  - File ID (unique key, e.g., a UUID message identifier)
  - File source (organization that is owner or creator of the file)
  - Created by (optional, person within the organization that created the file)
  - File format (e.g., CSV, XML, JSON, PDF, text, image, binary)
  - Schema reference (for structured documents, optional)
  - Message hash (e.g., using SHA-256)
  - File creation date (ISO 8601 timestamp, set by submitter)
  - File submission date (ISO 8601 timestamp, set by UUDEX)
  - File expiration date (after which file is no longer valid)
  - Abstract (short description of file contents)
  - Keywords (that may be useful for searches)
  - Priority
  - Status (default = ACTIVE)
  - Version (default = 1)
  - Obsoletes (optional, message identifier of file version that this replaces)
  - ObsoletedBy (optional, message identifier of file version that this is replaced by)
- In some cases, it will be necessary to encode files so that they can be transmitted in UUDEX Messages (U-Messages). For example, some file formats may need to undergo UUencoding and optional compression.

[7] Located at https://www.iana.org/assignments/media-types/media-types.xhtml

3.1.7 Operations Planning Data

In order to operate the BES reliably, RCs, BAs, and TOPs perform day-ahead and future hour power system studies to mimic future operating condition.

These studies have a dual purpose. First, they try to predict and anticipate potential operating transmission limit violations, both SOL and IROL, as well as voltage stability limits, generation shortages, or other operating condition that would threaten the reliability of the BES. In addition, they provide preventive actions to mitigate unsecure operating conditions.

To perform these studies, operating authorities (RCs, BAs, and TOPs) exchange power system models as previously indicated and it is of paramount importance to share information on load forecast or neighboring systems, generating units' operating plans, and power transfer schedules.

The following sections describe some of the operations planning data that could be transferred using UUDEX.

3.1.7.1 Load Forecast

The load forecast is normally calculated by BAs and the information that is exchanged with other neighboring entities include the following information:

- BA reporting the load
- RC the BA belongs to
- Time zone in which the load is being reported
- Period for which the load is being provided that depends on the time horizon (could be 5 minutes, 15 minutes, hourly, daily, weekly, etc.)
- Actual load forecast for the specified period

3.1.7.2 Interchange Schedules

Information on energy transfers from one balancing area to another is commonly shared between RCs, BAs, and TOPs on a periodic basis. These entities can then perform power flow and advance application studies using schedules and interchange with neighboring areas.

The schedule information that is normally exchanged for each schedule could include the following:

- Reporting Entity
- Source Point
- Source Balancing Area
- Sink Point
- Sink Balancing Area
- Schedule Start Time
- Schedule End time of the transaction

- Schedule Energy profile
- E-Tag Interchange Transaction reference, if applicable

3.1.7.3 Current Operating Plan

It is common for Load Serving Entities and Generator Operators to share the current hour, current day, and extended days operating plan for load and generating resources with their BAs, RCs, and TOPs. The information shared could be very extensive, but in most cases would include as the minimum the following:

- Delivery Date and Time
- Resource Name
- Resource Status (can indicate if the unit is on, off but available, out of service and unavailable, etc.)
- Resource Limits:
  - High Sustained Limit
  - Low Sustainable Limit
  - High Emergency Limit
  - Low Emergency Limit
- Ancillary Services that they are providing:
  - Regulation Up
  - Regulation Down
  - Responsive Reserve
  - Non-Spinning Reserve

This type of information is well established and varies in wholesale markets.

3.1.8 Asset Management

Beyond the needs to exchange power system models as described previously, the ability to exchange asset information will be a growing need over time. Information that is found in an asset model over what is required for simulation includes the following examples:

- Identification of the individual physical, serialized assets that comprise a power system resource
- Location of an asset, in terms of GPS coordinates or physical address
- Manufacturer, model, and version of a given asset
- Attributes of the asset beyond electrical characteristics, such as size, weight, height, volume, supporting structures, etc.
- References to related specifications, data sheets, etc.
- Ownership and value
- Lifecycle history of the asset

This information could be conveyed in a variety of ways, such as CIM/XML files or Shapefiles.

Other asset management data could include:

- Spare Equipment Database
- Cyber Asset Management (National institute of Standards and Technology [NIST] document under development)
- Domain Management Task Force (DMTF)

3.2 Cybersecurity Data

The information technology (IT) and IT-connected assets of the energy sector are increasingly targets of malicious activity, including criminal attacks to steal or extort money, vandalism and public reputation attacks, and state-sponsored attacks, in addition to potentially being victims of target-of-opportunity attacks, such as having assets compromised for use in botnets. For these reasons, IT operators need to ensure adequate protection of their IT assets. Key to this protection is the receipt and sharing of cybersecurity data. This includes guidance and information regarding current threats, proposed courses of action to address vulnerabilities, and sharing incident data both for reporting requirements and to use experience gained in the incident to help protect other parties. This section looks at some of the types of cybersecurity data UUDEX is able to convey.

In the case of operational cybersecurity data (i.e., data the recipient might use to alter the operations of their cyber assets), the issue of trust relationships needs to be carefully considered. Specifically, the data recipient needs to trust that the information received is accurate and provides the asserted benefits. For example, most parties trust patches released by the vendor of the product to which the patch applies. Very few would be willing to install a patch created by some unknown third party. As such, careful documentation of the provenance of cybersecurity data or the identities of parties who vouch for the accuracy of the data, is necessary.

In many cases, the use of Structured Threat Information Expression (STIX^{TM}) formatted messages can be used to transmit cybersecurity information.

3.2.1 Cyber Incident Reporting

Intrusion detection system (IDS), intrusion prevention system (IPS), and packet capture (PCAP) data types cover logs from cybersecurity tools as well as other network and endpoint monitoring tools. The data are primarily used for forensics and might also be required by certain parties as part of an incident report. The goal of sharing this type of information is to help provide comprehensive context with regard to network or other activities within a given period in time. Log files can be quite large, so will likely only be sent in response to specific incidents or needs.

Data elements of this type would need to include:

- Organization sending the log.
- Specific times associated with the log collections.
- Specific tool (including vendor, model, and version number of software or firmware) of the tool that generated the log.

- Information about the scope of the collection. For example, identification of subnets monitored, any filters applied to the data log, or other information. Because this information could vary widely in nature, it will likely need to be provided using descriptive free text.
- References to other U-Data Elements relevant to the log. This could include such things as a formal incident report submission or free-text analysis by a local operator offering their thoughts on the provided logs.

Log information is likely to be highly sensitive as it will reveal not only the identity of the security tools in use that produced the log, but often include a wealth of information about the sender's IT infrastructure. For this reason, logs will require confidentiality and integrity protection, and access to them will generally be granted to only a very small number of parties. One possible model for using this data would be to submit logs and other information to a trusted source, that then anonymizes and aggregates the data and then publishes a security advisory, informed by the submitted log information, but not traceable to the organization that provided those logs. This allows multiple parties to benefit from real incident data from a trusted source (i.e., the analyst/aggregator) while shielding the identity of the incident victim.

3.2.2 Operations Technology Cyber Incident Reporting

Due to the increasing number and types of cyberattacks on operations technology (OT) within electric power and other critical infrastructure, there is a need for better technical characterization of incidents involving these types of attacks. Leaning forward, UUDEX will support the exchange of incident reports relating to OT equipment. Fields could include the impact of an incident (both on function and on data), how long it took to recover from the attack, identities (make and model) of impacted OT equipment, vector by which the system was attacked, and how the attack was detected. Standards for encapsulating such information are under development and UUDEX will be able to support their transportation.

Data elements of this type would need to include:

- Organization sending the report.
- References to other U-Data Elements relevant to the report. This could include such things as data logs surrounding the event, change of system state information resulting from the event, or even a formal description of the indicators of compromise (IoCs).

Report information is likely to be highly sensitive as it will reveal the identity of the compromised party as well as specific OT equipment they use and how they were compromised. For this reason, reports will require confidentiality and integrity protection, and access to them will generally be granted to only a very small number of parties. As with the previous example, a trusted intermediary may serve to anonymize and aggregate submitted data prior to broader sharing.

3.2.3 Indicator of Compromise Sharing

An IoC is a concise expression of network traffic, endpoint behavior, or other patterns used to help guide the recipient in the detection of likely attacks or compromises. The goal of such U-Data Elements is that the recipient will be able to use the information to alter monitoring tools (e.g., IDS, IPS) or perform other scans to detect if the described behavior is present on their own networks. Ideally, IoCs are expressed in a way that can be directly ingested by security tools. IoCs often include additional information about the nature of the compromise being detected so that IT operators can better understand the implications of a positive or negative detection.

Data elements of this type would likely include:

- Author of the IoC.
- If the IoC can be ingested by certain security tools, the specific tools that can ingest the IoC.
- When the IoC was authored. There might also be an expiration date when the IoC will no longer be applicable.
- Severity level of the described compromise, indicating how urgently the recipient should check for the indicators.
- References to other U-Data Elements relevant to the IoC. These could include reports on certain cyber threats associated with the indicator or instructions to the recipient regarding their use of the indicator.

Some IoCs are public information and can be shared with anyone, even outside the energy sector. Others represent proprietary information and can only be distributed to parties that have purchased a license or have joined certain organizations. In other cases, it might be necessary to control release of an IoC because adversaries could be tipped off that their activities have been detected, prompting them to evolve their procedures to better avoid that detection. In all cases, IoCs need to be integrity protected so the indicator patterns cannot be corrupted, which would render the IoC useless. This is especially true of IoCs that are intended to be automatically ingested by tools.

3.2.4 Guidance

Guidance refers to any material intended to provide instruction with regard to how an enterprise is configured. Examples include Center for Internet Security Benchmarks[8] or material from NIST's National Checklist Repository.[9] Guidance comes in a variety of forms, from structured content that can be automatically ingested and used by security tools to prose descriptions of best practices. Guidance can represent general recommendations for best practices or could come with requirements to adopt the described practices as issued by a suitable authority such as NERC.

Data elements of this type could include:

- Author of the guidance.
- Date the guidance was authored.
- Applicability of the guidance (e.g., the specific operating system or software application to which the guidance applies).
- Format of the guidance. For guidance that can be automatically ingested by tools, this would identify the tools that could ingest this guidance.
- Criticality of the guidance. This could include whether some authority was mandating its adoption.

[8] Available at https://www.cisecurity.org/cis-benchmarks/

[9] Available at https://nvd.nist.gov/ncp/repository

Some guidance material is public information while others might only be releasable to parties with a certain license from the guidance author. In either case, guidance needs to be integrity protected against corruption. This is especially true of guidance that is intended to be automatically ingested by tools. Similarly, even if public guidance is being disseminated, there might be a desire not to expose the nature of the guidance being issued as it could reveal desired configurations. For this reason, confidentiality of guidance is also recommended.

3.2.5 Conformance Reports

Conformance reports describe the state of enterprise assets, usually with regard to some specific piece of guidance or patch notification. Often conformance reports are used by parties to inform an authority whether or to what extent their enterprise conforms to certain configurations, patch levels, or other standards. For example, a party might issue guidance and then expect the recipients to report where they conform or deviate from that guidance. Conformance reports are useful to gather a focused snapshot regarding certain important properties of an enterprise.

Data elements of this type could include:

- Source of the conformance report.
- Tool used to generate the report, if any.
- Reference to any guidance, patch notification, or other material that guided the generation of the conformance report.
- Date the report was generated.
- Format of the report (structured or free text).
- Any additional contextual information regarding the report (e.g., whether the report covers the whole enterprise or just specific assets).

Conformance reports will often reveal sensitive information about the party that generated the report. This could include information about software used in the enterprise, the presence of unpatched software vulnerabilities, and other network infrastructure information. For this reason, it is important that both the confidentiality and integrity of the report be protected and that the list of parties authorized to view the report be carefully controlled.

3.2.6 Patch Notification

Patch notifications are issued by software vendors to correct flaws in their software products. Sometimes patches simply add features or correct undesired user experiences with the software. Other times they are issued to fix security vulnerabilities associated with a software product. Patch notifications will either include an executable patch file that can be used to fix the described software or will include a reference (often a Uniform Resource Identifier) that can be used to retrieve this patch file from a remote source. In most cases, recipients use a patch notification to deploy the patch to applications in their enterprise, usually after first confirming the patch does not have any disruptive side effects through testing in a lab environment.

Data elements of this type could include:

- Applicability of the patch (e.g., the specific operating system or software application to which the patch applies).

- Severity of the issue the patch corrects.
- Date the patch was released (or approved).
- References to additional U-Data Elements as necessary. For example, the patch might be associated with some guidance that requires its adoption, or with a vulnerability or threat notification that the patch corrects.

In general, patches are public information. However, the integrity of patch notifications is critical as corruption of patches, or of references to patches, could at best lead to corruption of the patched software and at worst be used to introduce malware into an environment.

3.2.7 Vulnerability Notification

Vulnerability notifications are informative materials disseminated to make recipients aware of flaws in one or more software products that might be exploitable by adversaries. A notification itself is informative and intended to recommend increased vigilance against the described vulnerability. It might be accompanied by a patch notification that closes that vulnerability. In some cases, patches are not available and the recipient might respond to the notification by altering the configuration of the vulnerable software to mitigate the vulnerability (if possible), reduce access to vulnerable software to decrease exposure, or even uninstall the vulnerable software if the threat of exploitation exceeds the software's utility.

Data elements of this type could include:

- Author of the vulnerability notification.
- Applicability of the vulnerability notification (e.g., the specific operating system or software application to which the vulnerability notification applies).
- Date the vulnerability notification was authored.
- Severity of the described vulnerability. This could include whether the vulnerability is being actively exploited by malware.
- References to additional U-Data Elements as necessary. For example, the vulnerability notification might be associated with some guidance to mitigate the vulnerability or (ideally) with a patch that closes the vulnerability.

Many vulnerability notifications are public information. Some vulnerability notifications are non-public, often because the vendor has not yet developed a patch, but the notification author is sharing the vulnerability report selectively so certain parties can deploy mitigations. In the latter case, the confidentiality of the vulnerability notification is extremely important since exposure of this information to malicious parties could provide them with the information necessary to develop malware that exploits the vulnerability. Similarly, recipients of non-public vulnerability notifications need to be carefully limited to reduce the chance of exposure to malicious parties. Finally, the integrity of vulnerability notifications is important since corruption of these notifications could lead to incorrect or incomplete mitigations by the recipient.

3.2.8 Threat Notification

Threat notifications are warnings of activity by cyber adversaries, including identification of specific tools, target groups, or cyberattack campaigns. A notification itself is informative and intended to recommend increased vigilance against the described threat. A threat notification might be accompanied by an IoC to help detect the described threat. Alternately, a threat notification might simply recommend additional vigilance based on general behavior observed.

Data elements of this type could include:

- Author of the threat notification.
- Date the threat notification was authored.
- Severity of the threat.
- References to additional U-Data Elements as necessary. For example, the threat notification might be associated with guidance to mitigate the threat or with IoCs to detect the describe threat.

Some threat notifications are public information. Others are proprietary and only distributable to parties with the necessary license or group membership. Still others might be sensitive due to a desire not to tip off adversaries that their activities are being monitored. Threat notifications that have limited distribution need to be kept confidential and limited to certain recipients. All threat notifications need to be integrity protected since corruption of these notifications could mislead the recipient regarding the described threat.

4.0 Functional Requirements for Data Exchange

UUDEX will support the exchange of diverse data types that demand specific functional requirements primarily based on the nature of the data exchanged and the application it supports. This section introduces performance features that will be used to characterize the data exchange, the outcome of which is a key input to UUDEX development.

4.1 UUDEX Roles and Definitions

The following terms and associated definitions are used in the descriptions of UUDEX and related functionality.

UUDEX Administrator (U-Administrator)	Administrative users that have global responsibility for a UUDEX Framework and can authorize UUDEX Participants.
UUDEX API (U-API)	A set of parameterized instructions that UUDEX Endpoints and UUDEX Infrastructure use to interact with each other. UUDEX APIs are abstract definitions, rather than detailed functions in a particular programming language.
UUDEX Component (U-Component)	An individual hardware or software element that supports the functioning of the UUDEX Infrastructure
UUDEX Connection (U-Connection)	A communication channel between a UUDEX Participant and the UUDEX Infrastructure that conforms to all UUDEX requirements (e.g., security, performance).
UUDEX Consumer (U-Consumer)	A consumer of information, receives information from a UUDEX Subject.
UUDEX Data Element (U-Data Element)	Any data collection conveyed over UUDEX Exchanges.
UUDEX Data Element Type (U-Data Element Type)	A defined structure and format for specific classes of UUDEX Data Elements. Each UUDEX Instance defines its own set of supported UUDEX Data Element Types.
UUDEX Data Manifests (U-Data Manifest)	A UUDEX Data Element that describes the information within a UUDEX Subject, often filtered by specific search criteria, that is available to a given UUDEX Participant.
UUDEX Endpoint (U-Endpoint)	An entity that produces or consumes UUDEX Data Elements through interactions with a UUDEX Subject.
UUDEX Exchanges (U-Exchanges)	A communication UUDEX Participants and the UUDEX Infrastructure where all communicants are acting as elements of a UUDEX Framework (i.e., excludes out-of-band exchanges between UUDEX Participants). The UUDEX Exchange will involve one or more UUDEX Connections and communications will only occur over UUDEX Connections.

UUDEX Framework
(U-Framework)

Includes the totality of UUDEX, specifically UUDEX Infrastructure, UUDEX Endpoints, UUDEX APIs, UUDEX Participants, UUDEX Protocols, UUDEX Communication Fabric, and UUDEX Information Models.

UUDEX Header
(U-Header)

The portion of a UUDEX Message that contains metadata about the message exchanged between UUDEX Endpoints. The UUDEX Header controls behaviors associated with the delivery of the UUDEX Data Element. The UUDEX Header may be discarded when the UUDEX Data Element arrives at its destination, or its contents may be used to validate the UUDEX Message or other processing.

UUDEX Identity Authority
(U-Identity Authority)

An entity that creates, certifies, manages, and revokes UUDEX Identity Objects. In effect, it serves as an identity authority within a UUDEX Instance.

UUDEX Identity Objects
(U-Identity Objects)

A type of UUDEX Data Element that contains information necessary to authenticate the identity of a UUDEX Participant.

UUDEX Infrastructure
(U-Infrastructure)

The servers, communication fabric and other hardware pertaining to UUDEX.

Those UUDEX components that permit the management and flow of information to and from UUDEX Endpoints. These components provide a variety of services and are typically replicated for availability purposes.

UUDEX Instance
(U-Instance)

A collection of connected UUDEX Participants that is closed with regard to its trust environment. A UUDEX Instance is defined by a set of identities within a UUDEX Infrastructure where those identities are only valid within that UUDEX Infrastructure and no other identities are valid within that UUDEX Infrastructure.

UUDEX Message
(U-Message)

An instantiation of the data in a UUDEX Subject that is comprised of a UUDEX Header and a UUDEX Payload.

UUDEX Message Envelope
(U-Message Envelope)

A structure of a UUDEX Message that contains a UUDEX Header and a UUDEX Payload while the UUDEX Message is in transit over a UUDEX Connection..

UUDEX Notification
(U-Notification)

A response message sent by the UUDEX Infrastructure in response to a UUDEX Subscription match that indicates the presence of a UUDEX Data Element but does not contain the UUDEX Data Element's data.

UUDEX Participant
(U-Participant)

An organization that is a onboarded member of a UUDEX Instance.

UUDEX Participant Administrator (U-Participant Administrator)
An administrative user that performs activities related to the to the publication and consumption of information for a given UUDEX Participant.

UUDEX Payload (U-Payload)
The portion of a UUDEX Message conveying the information exchanged between UUDEX Endpoints.

UUDEX Producer (U-Producer)
A publisher of information, sends information to a UUDEX Subject.

UUDEX Protocol (U-Protocol)
The set of messaging patterns, message structures, UUDEX APIs, and common data structures outline in the UUDEX Protocol Design document.

UUDEX Server (U-Server)
A UUDEX Component that stores data, received data from UUDEX Producers, delivers data to UUDEX Consumers, and maintains UUDEX Subjects and associated prioritization and access control policies. In general, the UUDEX Infrastructure abstracts the concept of the UUDEX Server, allowing UUDEX Endpoints to engage with the UUDEX Infrastructure without tracking individual UUDEX Servers.

UUDEX Subject (U-Subject)
A UUDEX Subject is the basic unit of storage, access. and organization in the U Infrastructure. Data is published by a UUDEX Producer to a UUDEX Subject and delivered to a UUDEX Consumer by queueing it to a UUDEX Subscription. The ability to publish or subscribe to a UUDEX Subject is controlled by access control policies.

UUDEX Subscription (U-Subscription)
The means by which a UUDEX Consumer retrieves UUDEX Messages published to a UUDEX Subject.

4.2 Information Flow Requirements

UUDEX is responsible for facilitating the exchange of information to and from U-Participants. This section outlines the requirements that surround the general framework of data exchange supported by UUDEX.

FLOW-1 The U-Infrastructure MUST have the ability to receive requests from U-Endpoints to post U-Data Elements.

FLOW-2 When the U-Infrastructure receives U-Data Elements from a U-Endpoint, it MUST be able to store that U-Data Element and its metadata as a member of a particular U-Subject.

FLOW-3 The U-Infrastructure MUST have the ability to search their stored U-Data Elements and metadata for matches against search criteria. The U-Infrastructure MUST be able to do this efficiently so that searches complete in a reasonable amount of time. Facilitating this might require indexing data as it is added or other optimizations.

FLOW-4 The U-Infrastructure MUST be able to receive requests by U-Endpoints to retrieve stored U-Data Elements. These requests can either specify the U-Data Element’s UUID (see requirement DAT-1) or search parameters.

FLOW-5 The U-Infrastructure MUST be able to receive requests to delete a U-Data Element and its metadata from a U-Endpoint.

FLOW-6 The U-Infrastructure MUST be able to receive requests from a U-Endpoint to establish U-Subscriptions to specific U-Subjects. These U-Subscriptions specify search parameters that could match elements of U-Data Elements or its metadata, similar to U-Endpoint data retrieval requests.

FLOW-6.1 The U-Infrastructure MUST be able to store and serve U-Subscription requests until instructed to do otherwise by a U-Endpoint.

FLOW-6.2 The U-Infrastructure MUST be able to receive requests to delete, pause, resume, or replace existing s U-Subscriptions.

FLOW-6.3 When a U-Subscription is active on a given U-Subject, every time a new U-Data Element is added to that U-Subject in the U-Infrastructure, that U-Data Element and its metadata MUST be processed against the U-Subscription. Processing involves comparing the U-Data Element and its metadata against the search parameters in the U-Subscription. If the U-Data Element or its metadata matches these parameters, this is called a “match” against the U-Subscription.

FLOW-6.4 Upon detecting a U-Subscription match, the U-Infrastructure MUST immediately queue a response message to the U-Subscription holder. Depending on parameters in the U-Subscription, this response message might be a U-Notification, mentioning match and providing limited metadata about the U-Data Element, at least including its UUID. Alternately, if the U-Subscription is so configured, the queued message might contain the whole U-Data Element and its metadata. Queued messages are to be delivered to the subscriber in accordance with the U-Subscription parameters and U-Subject configuration. Prioritization (ARCH-2), supported fulfillment models (FLOW-6.6), and other configuration choices might impact the details of this delivery.

FLOW-6.5 U-Endpoints MUST have the ability to signal the U-Infrastructure to pause and resume delivery of queued messages. U-Endpoints MUST have the ability to signal the U-Infrastructure to purge messages in their delivery queue. Both of these can help the U-Endpoint recover from situations where the queue of messages to deliver is so large that it is overwhelming the U-Endpoint’s ability to receive them.

FLOW-6.6 The U-Infrastructure MUST have the ability to store messages queued for delivery in response to a U-Subscription until such time that the relevant U-Endpoints contact the U-Infrastructure and request delivery of this message queue (subject to potential size limits of this queue). This is called "deferred U-Subscription fulfillment." In addition, the U-Infrastructure MUST have the ability to immediately contact the subscribing party to deliver the queued messages. This is called "immediate U-Subscription fulfillment." Owners of a U-Subject can decide whether a given U-Subject will support deferred or immediate U-Subscription fulfillment or both.

FLOW-7 All described actions MUST be taken only after authentication of the U-Endpoint requesting the action and validation that the action is permitted by access controls associated with the relevant U-Subject.

4.3 Identity Requirements

U-Identity Authorities are standalone systems within the U-Infrastructure that perform identity proofing and provide authentication services (such as creating and delivering U-Identity Objects) to a U-Instance. As part of the role of the U-Identity Authorities the system need to provide UUDEX with a mechanism for accessing information for making access control decisions.

U-Identity Objects are data structures that associate the contact information (e.g., name, email address, physical address, and phone number) associated with an entity with a cryptographic puzzle that only the identified entity is able to solve. Thus, proof that some party is able to solve the cryptographic puzzle can serve as evidence that that party is the named entity in the identity object. Public key certificates are one example of such an identity object.

ID-1 UUDEX MUST support a distributed management model for identities used in the U-Framework. U-Instances MUST be able to avoid dependency on a single U-Identity Authority.

ID-1.1 All U-Instances MUST have at least one U-Identity Authority.

ID-2 All U-Identity Objects MUST be validated by a U-Identity Authority prior to that authority adding them to the collection of U-Identity Objects used by the U-Instance.

ID-2.1 Identity proofing goal 1: determine to a reasonable level of certainty that the entity identified in the U-Identity Object is the entity it claims to be (i.e., detect and prevent cases where party A attempts to register a U-Identity Object in party B’s name).

ID-2.2 Identity proofing goal 2: as necessary, validate the identity attributes associated with the U-Identity Object that the U-Instance uses to automatically assign to specific groups or roles. Since such identity attributes will be linked to structures that are used for access control decisions, they need to be reliable.

ID-3 U-Identity Objects MUST conform to a standard format.

ID-3.1 All U-Identity Objects will identify the U-Identity Authority that validated them.

4.4 Communications Connection Requirements

The following requirements are related to communications.

COM-1 All U-Exchanges MUST be encrypted using algorithms deemed sufficient for protecting Sensitive But Unclassified information.

COM-2 All U-Exchanges MUST be integrity protected using algorithms deemed sufficient for protecting Sensitive But Unclassified information.

COM-3 All U-Exchanges MUST be mutually authenticated.

COM-4 UUDEX MUST allow U-Participants to establish a U-Connection prior to the need to communicate data and be able to keep the U-Connection open for as long as U-Data Elements might be exchanged. Establishing the U-Connection might include activities such as contacting, encryption or integrity algorithm negotiation, agreement on encryption keys, and mutual authentication of parties.

COM-5 UUDEX MUST allow any U-Endpoint to measure the bandwidth and latency between itself and the U-Infrastructure.

COM-6 UUDEX MUST support a U-Subject delivery prioritization scheme. This scheme MUST include support for U-Subjects whose U-Data Elements MUST be delivered immediately regardless of what other elements are preempted, U-Subjects whose U-Data Elements are only to be delivered when the U-Connection in question is otherwise idle, and one or more levels of relative priority that exists between these extremes (e.g., where level 1 preempts level 2 U-Subjects, and level 2 preempts level 3 U-Subjects).

COM-6.1 Messages that control UUDEX behaviors SHOULD be assigned a priority higher than all U-Subjects, since their successful delivery will impact all other UUDEX behaviors including delivery of all priority-levels of U-Subjects.

COM-7 All U-Connections MUST use reliable message delivery. This means that the sender of a message will always know if the recipient did not receive a given message, allowing the sender to attempt to resend.

4.5 Data Storage Requirements for UUDEX Servers

The following are data storage requirements for U-Servers:

STO-1 All data stored by U-Servers, including U-Subjects, their U-Data Elements and metadata, and also including other information such as U-Subscription information, MUST be encrypted while at rest.

STO-2 All data stored by U-Servers MUST include mechanisms to detect data corruption (e.g., checksums or other integrity protections).

STO-3 U-Servers MUST support a mechanism for associating access control lists (ACLs) with U-Subjects. ACLs MUST be protected against data corruption or inadvertent modification.

STO-4 All U-Servers MUST support a mechanism for storing metadata associated with U-Data Elements, including but not limited to timestamp, source, type, retention period, and sensitivity of the data, in conjunction their associated U-Data Elements. This metadata MUST be protected against data corruption or inadvertent modification.

STO-5 All U-Servers MUST support a mechanism for performing data backup or archiving of all information they store, including U- Subjects, their associated U-Data Elements and metadata, configuration of the U-Subjects including ACLs and prioritization, and U-Subscription information. All U-Servers MUST include mechanisms to restore all this information from backup data.

STO-6 U-Servers MUST support a mechanism for controlling access to data to prevent unauthorized parties from modifying the information they store and manage.

4.6 Requirements for Testing

The following requirements relate to the capability of UUDEX to perform testing of UUDEX services without disrupting current operational UUDEX services. While testing new or existing UUDEX services UUDEX needs to have the ability to continue normal operation without compromising the capabilities or features of U-Components.

UUDEX needs to allow for non-disruptive testing of UUDEX capabilities.

TM-1 It MUST be possible to configure verbose logging messages for anomalies or exceptions.

TM-2 It MUST be possible for transportation and storage of dummy data to test the full capabilities of an operational U-Instance.

TM-3 It MUST be possible for a U-Endpoint or U-Server will be monitored to ensure it will not adversely affect the U-Instance.

TM-4 It MUST be possible to establish message exchanges between selected parties for the exchange of test messages that will not be transmitted or received from parties that are not participating in this test.

TM-5 It MUST be possible for a message to be flagged as a test message by U-Endpoints and U-Servers. Messages flagged as such will not be used for any purpose outside of the test.

TM-5.1 U-Endpoints and U-Servers receiving a test message SHOULD respond with a message indicating whether expected conditions were met from that test message.

UUDEX needs to support verbose software logging.

SD-1 Software verbose logging MUST be possible even while a U-Component is being used operationally.

SD-1.2 All logs MUST be written locally. Logs MAY also be copied to remote locations.

SD-2 When performing verbose logging, each U-Endpoint or U-Server MUST report on any anomalies detected for each software component.

UUDEX MUST support verbose U-Data Element logging.

DD-1 U-Data Element verbose logging MUST be possible while the U-Component is being used operationally.

DD-1.2 All logs MUST be written locally. Logs MAY also be copied to remote locations.

DD-2 When performing verbose logging, each U-Endpoint or U-Server MUST report on any anomalous structure or inappropriate values detected for each U-Data Element.

UUDEX needs to support verbose message logging.

MD-1 Message verbose logging MUST be possible while the U-Component is being used operationally.

MD-1.2 All logs MUST be written locally. Logs MAY also be copied to remote locations.

MD-2 When performing verbose logging, each U-Endpoint or U-Server MUST report on any anomalous structure or inappropriate values detected for each U-Exchange.

4.7 Architectural Requirements

These requirements deal with architectural aspects of UUDEX that need to be supported in the design. They focus on capabilities that UUDEX elements need to be able to support; actual utilization of many of these features within deployed environments would remain at the discretion of operators.

ARCH-1 U-Servers MUST be capable of supporting redundant deployment, where multiple U-Server implementations (i.e., nodes hosting U-Servers) support a single U-Server role. These redundant U-Servers will all be capable of serving the same U-Endpoint requests. These U-Servers will have the same U-Subjects with the same U-Data Elements and metadata and governed by the same ACL and prioritizations policies and U-Subscriptions replicated between them. This provides both service redundancy (allowing other, redundant U-Servers to field requests if one U-Server becomes unavailable) and data redundancy (ensuring that other U-Servers are effectively providing data backup in the case of corruption of one U-Server's data).

ARCH-1.1 U-Servers SHOULD support distributed redundant deployment, where a set of redundant U-Servers are not geographically co-located.

ARCH-2 All U-Producers MUST be able to assign a priority value to a U-Subject. This priority value is used to prioritize activities within a U-Instance in the case where demand for UUDEX capabilities is outstripping the U-Instance's ability to supply those capabilities.

ARCH-3 The U-Infrastructure and U-Endpoints MUST have the ability to detect network connectivity status of a U-Connection by which they communicate.

ARCH-3.1 U-Endpoints MUST be able to detect and alarm when they are not able to establish a U-Connection to the U-Infrastructure and vice versa.

ARCH-3.2 The U-Infrastructure and U-Endpoints MUST be able to detect and alarm whenever the bandwidth of the U-Connection is less than the bandwidth required to deliver the data within the latency requirements of that data (e.g., channel congestion detection).

ARCH-4 The U-Infrastructure MUST be configurable with a prioritization policy. The prioritization policy MUST allow the U-Infrastructure to limit query responses, U-Subscription fulfillment, and possibly other U-Endpoint interactions at certain times based on each U-Subject's priority level. Specifically, the U-Server can choose to only send U-Data Elements from U-Subjects with higher priorities when the policy is executed. This policy MAY contain more than one priority level, allowing the U-Infrastructure to scale its restrictions based on appropriate factors (e.g., a minor network constrain might cause the U-Server to only deliver U-Data Elements from medium or high priority U-Subjects, while a more significant network constrain might cause the U-Server to only deliver U-Data Elements from high priority U-Subjects).

ARCH-4.1 The U-Infrastructure MUST be able to execute its prioritization policy on a U-Connection by U-Connection basis. Thus, at any given time, some U-Connections might be constrained by the prioritization policy while others might not.

ARCH-4.2 The U-Infrastructure MAY include the ability to automatically engage its prioritization policy based on the state of network connectivity or other factors. Whether or not the U-Infrastructure has the ability to automatically set the prioritization level, the U-Infrastructure MUST have the ability to manually set the prioritization level, and the manual level MUST take precedence over the automatic level.

ARCH-5 U-Endpoints SHOULD have the ability to constrain their behavior based on a prioritization policy. This could include deferring or dropping data publication or requests for data that would be beneath the priority threshold in effect between that U-Endpoint and the U-Infrastructure.

ARCH-6 U-Infrastructures MUST be able to support fielding multiple simultaneous U-Connections between a given U-Endpoint and the U-Infrastructure in order to provide redundancy in the communications fabric. It MUST be possible for these U-Connections to be over different communications media, including mixes of TCP/IP and non-TCP/IP networks.

4.8 Data Lifecycle Requirements

All data that are exchanged using UUDEX undergo a common lifecycle of creation, distribution, and ultimately deletion. The following requirements govern the treatment of data throughout this lifecycle.

DAT-1 All U-Data Elements that are added to a U-Subject MUST be assigned a unique message identifier (i.e., a UUID).

DAT-1.1 UUID structure MUST conform to the format and creation rules outlined in RFC 4122 for “name-based” UUIDs.

DAT-1.2 UUIDs are only necessary when data are added to a U-Subject. It is possible that the U-Data Element itself might have been created significantly before being added to a U-Subject, in which case the UUID is only required to be added at the time the U-Data Element is added to the U-Subject. That said, for processing reasons, content creators MAY assign UUIDs to U-Data Elements before they are added to a U-Subject, even if there is no guarantee that the U-Data Element would ultimately be added to a U-Subject.

DAT-1.3 UUIDs MUST NOT be reused. Even if a U-Data Element is deleted, the UUID associated with the deleted U-Data Element remains forever bound to that U-Data Element and cannot be reassigned.

DAT-2 When a U-Endpoint requests a U-Data Element or deletion of a U-Data Element from a U-Subject, the U-Infrastructure MUST first verify that the U-Identity Element associated with this U-Endpoint request is valid. The U-Infrastructure MUST then compare the identity, roles, or other attributes of the U-Identity Element against the relevant U-Subject's access controls to determine whether the action is allowed. Only if the action is permitted by the ACL associated with the U-Data Element's U-Subject will the action proceed.

DAT-2.1 If a U-Endpoint action request (e.g., read, delete, or post) is denied, the message returned by the U-Infrastructure MUST NOT leak information to the U-Endpoint. For example, if a U-Subject's access controls prevent a particular U-Endpoint from knowing of the existence of that U-Subject, and the U-Endpoint requests that U-Subject, the U-Infrastructure's response MUST be the same as if the U-Subject did not exist, rather than reporting that access to the U-Subject was denied.

DAT-3 If a U-Data Element is deleted from a U-Subject, the U-Infrastructure MUST immediately make the U-Data Element unavailable to all U-Endpoints and treat any requests for the U-Data Element as being made to non-existent data. The U-Infrastructure MUST also remove the U-Data Element from memory at its earliest convenience. The latter might be later than the former due to needs to ensure the deletion is replicated across all U-Servers in a redundant deployment.

DAT-3.1 U-Servers MAY employ backup systems to protect against data loss in the case of server failure or storage corruption. If such a backup exists, deleted U-Data Elements SHOULD be purged from the backup if and when doing so is possible.

DAT-3.2 It is generally not the intended role of the U-Infrastructure to provide a long-term, historical archive of U-Data Elements. A historical archive might retain deleted U-Data Elements for historical reasons, but that would violate the previous requirements about removing deleted U-Data Elements from memory. If an organization wishes to maintain a historical archive of U-Data Elements, this needs to be done by downloading the U-Data Element from the U-Subject to a U-Endpoint, which can then use the downloaded data to create and maintain such an archive.

DAT-4 When a U-Data Element is published to a U-Subject, the U-Infrastructure MUST perform any preprocessing necessary to make the U-Data Element available to the appropriate U-Endpoint requests. This might include indexing the U-Data Element, if possible, to enable the U-Data Element to be matched against search and U-Subscription requests.

DAT-5 When a U-Subject is created it MAY specify performance constraints to minimize the impact of large amounts of data stored in the U-Server for extended periods of time. These performance constraints may be specified either as the number of U-Data Elements queued in the U-Subject, or as the total size of queued U-Data Elements in the U-Subject. Only one performance constraint may be specified for each U-Subject. If no performance constraints are specified, no performance constraint processing is performed.

DAT-5.1 If the performance constraint is specified to delete old entries, the oldest U-Data Elements in the U-Subscription queues are deleted until the new U-Data Element can be processed without violating the specified performance constraint.

DAT-5.2 If the performance constraint is specified to block new entries, the U-Data Element is not stored, and an error is returned to the U-Publisher. The U-Publisher is responsible for retrying the publish action at a later time.

DAT-6 When a U-Subject is created it MAY specify that the U-Data Element be either retained once all pending U-Subscriptions for the U-Subject have been filled, or it MAY specify that the U-Data Element be deleted from the U-Subject's persistent storage once all pending U-Subscriptions for the U-Subject have been filled.

5.0 UUDEX Messages and Data

The purpose of this section is to provide a very high-level overview of the metadata and message structure used by UUDEX. The intent is to describe the interfaces and infrastructure, while remaining largely agnostic with respect to the data models of U-Data Elements being exchanged. The overall UUDEX message structure is shown in Figure 5-1.

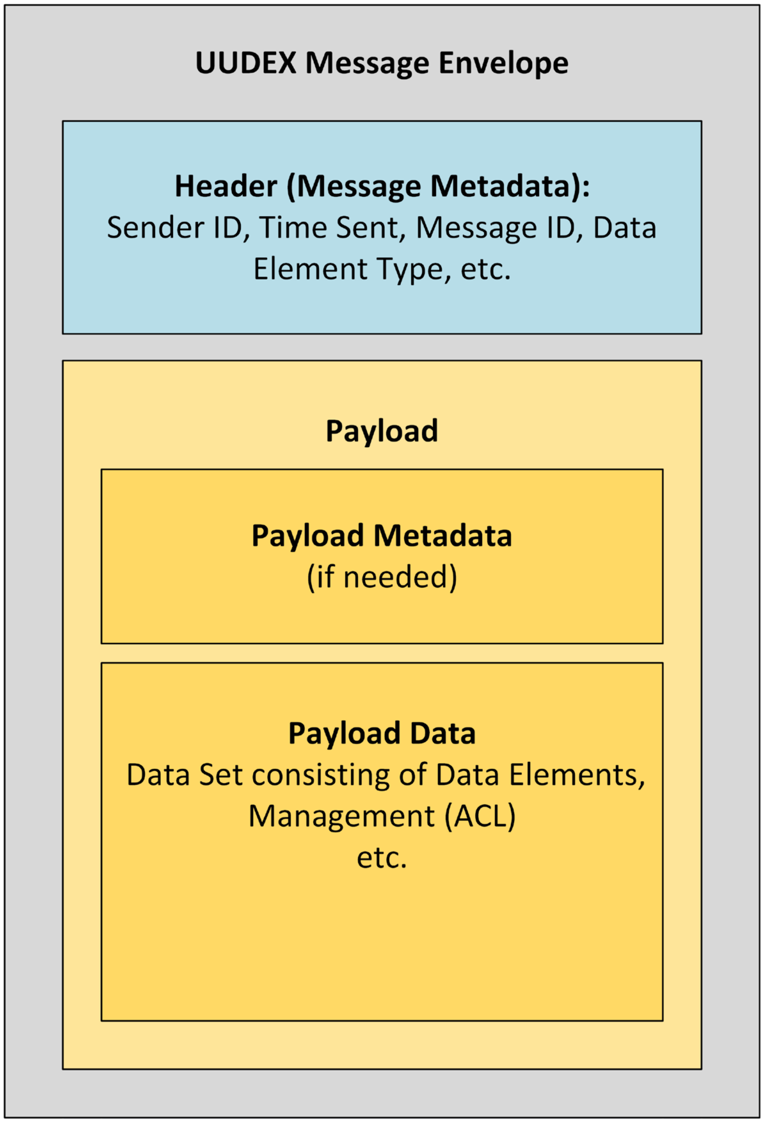
UUDEX Message Envelope
Header (Message Metadata):
Sender ID, Time Sent, Message ID, Data Element Type, etc.
Payload
Payload Metadata
(if needed)
Payload Data
Data Set consisting of Data Elements,
Management (ACL)
etc.

Figure 5-1: UUDEX Message Structure

5.1 The UUDEX Message Envelope and UUDEX Header

The U-Message Envelope shown in Figure 5-1 is a JSON structure used to wrap a U-Header and a U-Payload (generally consisting of a set of U-Data Elements) and associated metadata (if needed) while it is in transit to and from U-Endpoints across a U-Connection.

Once the message has been received and processed by the U-Consumer, the U-Header information can be discarded (although some fields might be used for validating the U-Message or for other processing).

The U-Header is used to assist with the delivery and processing of the U-Data Element and its metadata. It contains the message metadata that is related to the flow of the message.

The U-Header minimally includes:

- Sender identifier
- Time sent
- Data element type

Other fields might be included in the U-Header, such as:

- Special handling instructions (e.g., acknowledgement required)
- Additional security related information
- Additional information used by the U-Server for routing the U-Data Element to the appropriate U-Subject.

The U-Header and other metadata are designed to be easily extensible, so additional fields might be specific to a given U-Instance.

5.2 UUDEX Payload

All U-Messages stored in the UUDEX system may have additional metadata as shown in Figure 5-1. This metadata contains key information needed to support handling of the U-Message. Currently, there is no metadata defined at the U-Payload level; all defined metadata is associated with a U-Data Element as defined by the UUDEX Data Element Type (U-Data Element Type) definitions.

5.2.1 UUDEX Data Element Types

The U-Data Element Type defines components that can be considered metadata as well as components that can be considered data.

The metadata component of the U-Data Element Type specifies information about the U-Data Elements and applies to all the U-Data Elements contained in the U-Message. Each U-Data Element Type defines the metadata that is important for that specific U-Data Element Type. Common information could include:

- A reference to the U-Data Element Type schema
- The version of the U-Data Element Type schema
- A name for the U-Data Element
- The unique message identifier for the U-Data Element
- A description of the U-Data Element
- Keywords (tags) for searching the U-Data Element

- Encoding and compressing methods used in the U-Data Element
- Comments

The data component of the U-Data Element Type specifies the content and format of the U-Data Element. Some U-Data Element Types, such as the Power System Information U-Data Element Type, are complicated, while others such as the Electrical Disturbance Reporting U-Data Element Type are simpler.

U-Exchanges involve the use of messages. There are two categories of messages:

1. Messages for the management of U-Endpoints and U-Subscriptions
2. Messages that convey all other types of U-Data Elements. Details on the types of U-Data Elements are provided below.

UUDEX can be configured to enable the exchange of an extensible set of U-Data Types. The types of U-Data Elements would fall into one of a number of categories:

- Time series data snapshots, as are commonly conveyed using ICCP, where data points are defined for the capture of values and changes over time.
- Structured documents that convey data using a common format (e.g., JSON, XML, CSV) based on some information model such as (but not limited to) the IEC Common Information Model and that may be parsed by applications.
- Unstructured documents, which are conveyed using formats such as PDF, JPEG, or binary executable and are not typically parsed by applications.
- Power system network and asset models, which represent the descriptions and relationships of objects that comprise the portions of or changes to the electricity grid.

Given that UUDEX tries to be largely agnostic to the data models used by U-Data Elements, when applicable, different U-Data Types could be based on different logical information models. The IEC CIM is one example of these models.

There is also the issue of granularity. A simple example is where a U-Producer may publish a set of generator measurements. It will be possible to impose access controls via a ACL at the U-Subject level such that a U-Consumer may subscribe to or access only messages in a specific U-Subject that conveys the granularity the U-Endpoint requires.

The diagram shown in Figure 5-2 is a high-level class hierarchy for U-Data Elements. This is used to categorize the different types of U-Data Elements. This hierarchy would be extended by adding U-Data Element Types for structures such as OE-417, RCIS, COMTRADE (IEEE Std C37.111), etc. as needed for U-Exchanges. Note that this hierarchy is intended to be informative and is not complete, especially at the lower layers.

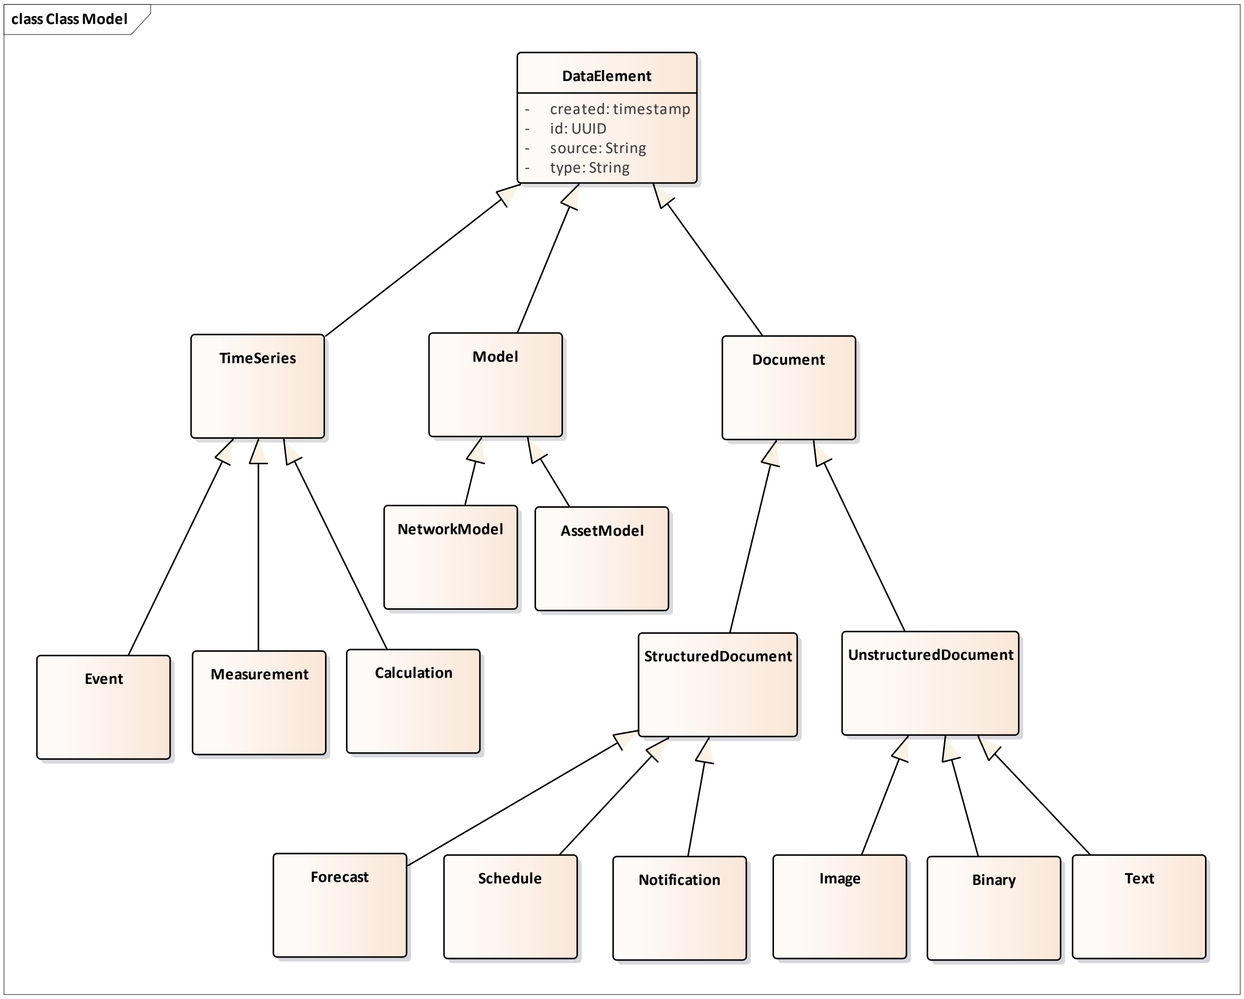

Figure 5-2: Data Element Hierarchy

6.0 Data Exchange Architectures

This section describes high-level data exchange architectures and data flows.

6.1 Publish-Store-Forward

Figure 6-1 illustrates the basic sequence of a publish –subscribe information exchange using UUDEX.

In this example, UUDEX Subscriber 1 and UUDEX Subscriber 2 have already established U-Subscriptions to a particular U-Subject on the U-Server. The ACL associated with this U-Subject allows both subscribers to consume data from this U-Subject. As soon as the producer adds the U-Data Element, the U-Infrastructure checks for U-Subscriptions to the subject, notices that both UUDEX Subscriber 1 and UUDEX Subscriber 2 have active U-Subscriptions, and sends each a message containing the newly published U-Data Element.

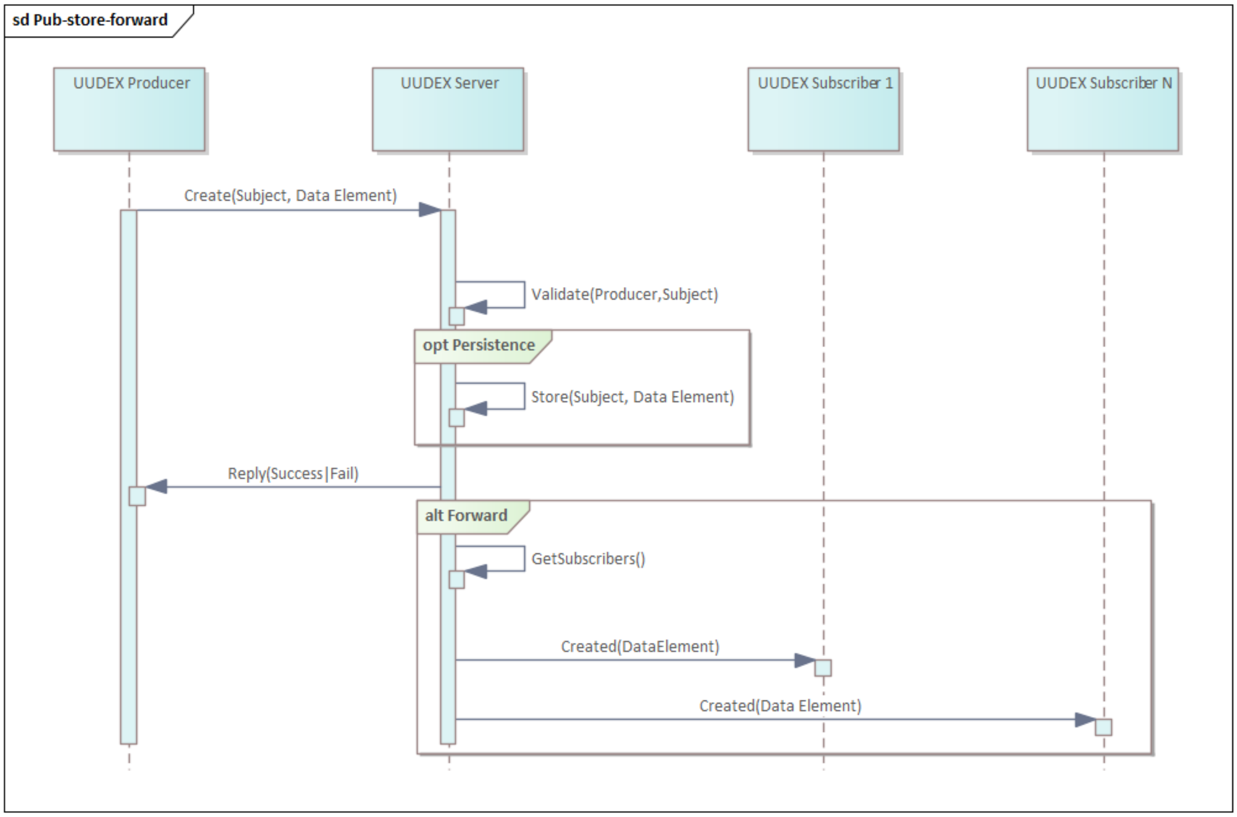

Figure 6-1: Publish-Subscribe Store-forward

6.2 Publish-Store-Notify

Figure 6-2 describes the sequence of a UUDEX publish-and-notify information exchange. In this pattern, the published U-Data Elements are stored in a U-Subject and a notification is issued to potentially interested U-Consumers who have subscribed to the subject. Upon receipt of the notification, the U-Consumer can then decide to retrieve the information from the U-Subject at a convenient time.

In this example, both UUDEX Subscriber 1 and UUDEX Subscriber 2 have established U-Subscriptions to the relevant U-Subject. However, instead of sending the U-Data Element itself in fulfillment of the U-Subscriptions, the U-Server sends a U-Data Manifest that identifies the new U-Data Element to both U-Subscribers. Such U-Data Manifests are smaller than the U-Data Elements they identify. In this case, both U-Subscribers receive the U-Data Manifest alerting them to the new U-Data Element. UUDEX Subscriber 1 decides not to retrieve the U-Data Element. UUDEX Subscriber 2 decides to retrieve the U-Data Element and sends a query to the U-Subject that contains the message identifier for that U-Data Element as specified in the U-Data Manifest it received. The U-Infrastructure processes this request and returns the requested U-Data Element.

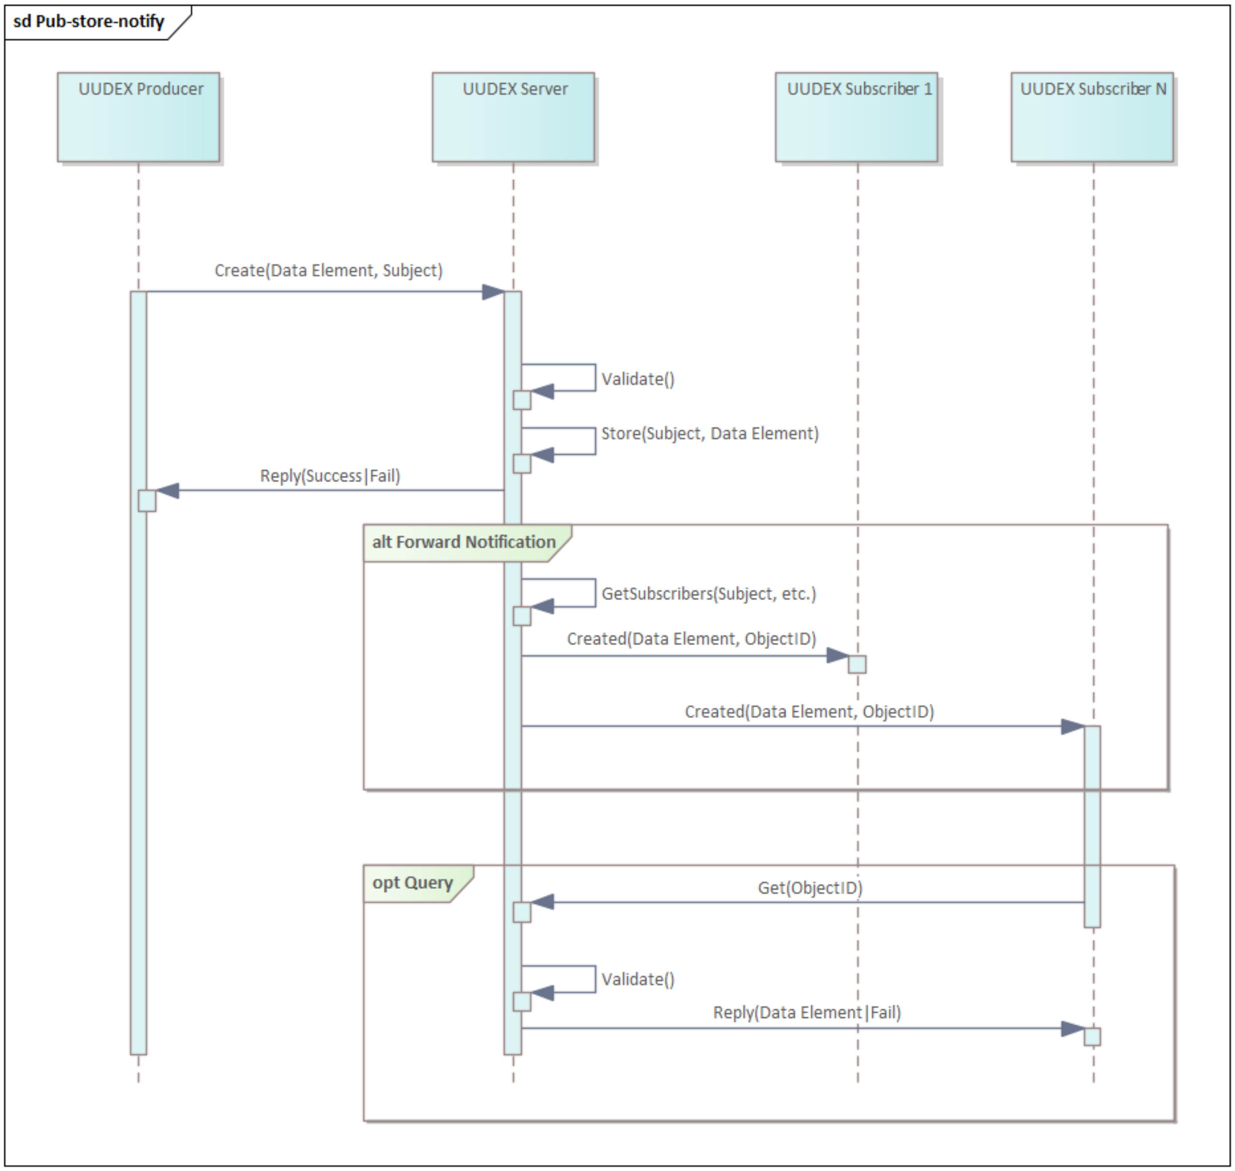

Figure 6-2: Publish Store Notify

6.3 Query

The sequence diagram shown in Figure 6-3 describes a U-Endpoint querying the U-Infrastructure for information from a given U-Subject. Provided that the U-Endpoint is authorized for the specific type of information as defined by the U-Subject's ACL, the query can be honored by the U-Infrastructure.

In this case, the U-Endpoint sends a query to the U-Infrastructure asking for U-Data Elements that match a certain set of criteria. Criteria could include data element type, time the data were submitted, source of the U-Data Element, etc. The request is validated, checked against the U-Subject's ACL and, if correct and permitted, the collection of matching U-Data Elements are compiled and returned to the U-Endpoint.

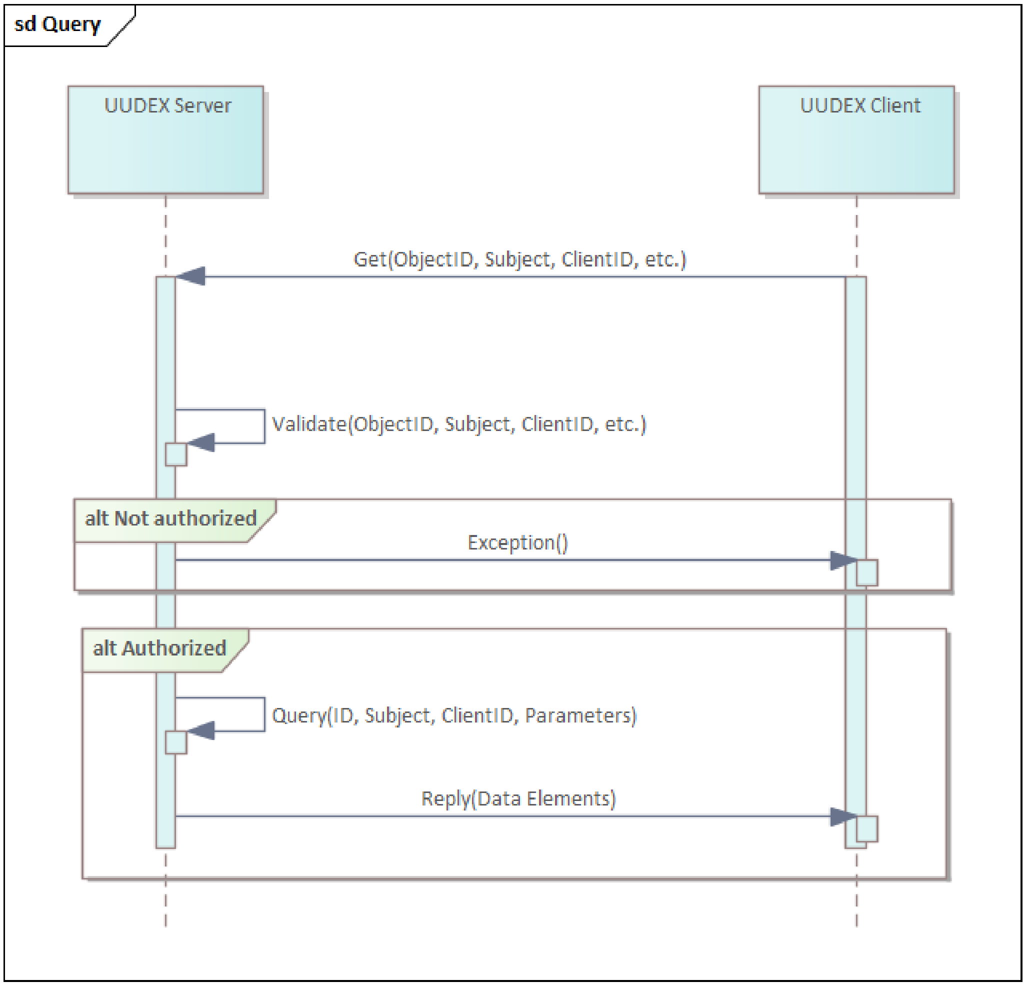
sd Query
UUDEX Server
UUDEX Client
Get(ObjectID, Subject, ClientID, etc.)
Validate(ObjectID, Subject, ClientID, etc.)
alt Not authorized
Exception()
alt Authorized
Query(ID, Subject, ClientID, Parameters)
Reply(Data Elements)

Figure 6-3: Query

6.4 Set Access Control List

The sequence diagram in Figure 6-4 shows an authorized U-Producer creating or updating an ACL for a given U-Subject. The ACL is persisted by the U-Server and used to validate requests involving that U-Subject. The ACL can be defined to permit or prohibit specific actions on the U-Subject by different U-Participants.

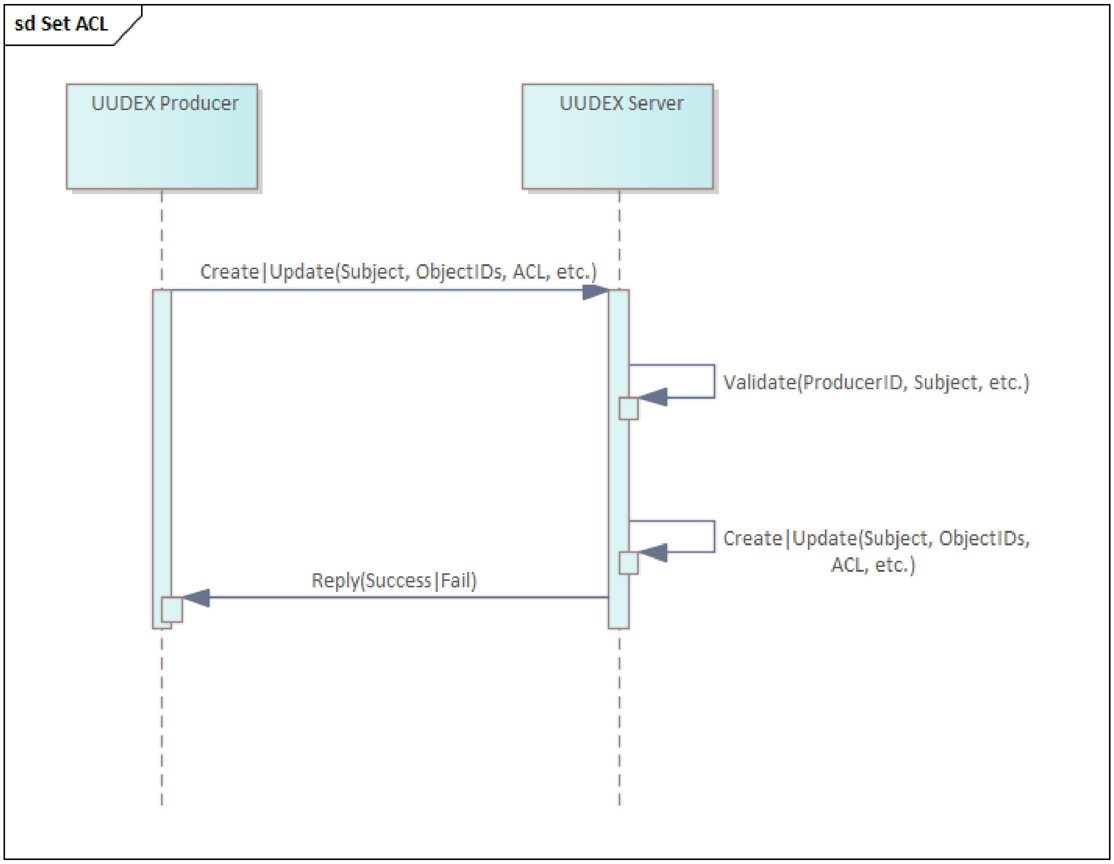
sd Set ACL
UUDEX Producer
UUDEX Server
Create|Update(Subject, ObjectIDs, ACL, etc.)
Validate(ProducerID, Subject, etc.)
Create|Update(Subject, ObjectIDs, ACL, etc.)
Reply(Success|Fail)

Figure 6-4: Set Access Control List

6.5 Subscribe

The sequence diagram shown in Figure 6-5 illustrates a U-Endpoint subscribing to a given U-Subject. The U-Subscription is managed by the U-Infrastructure, which acts as an intermediary for information exchanges.

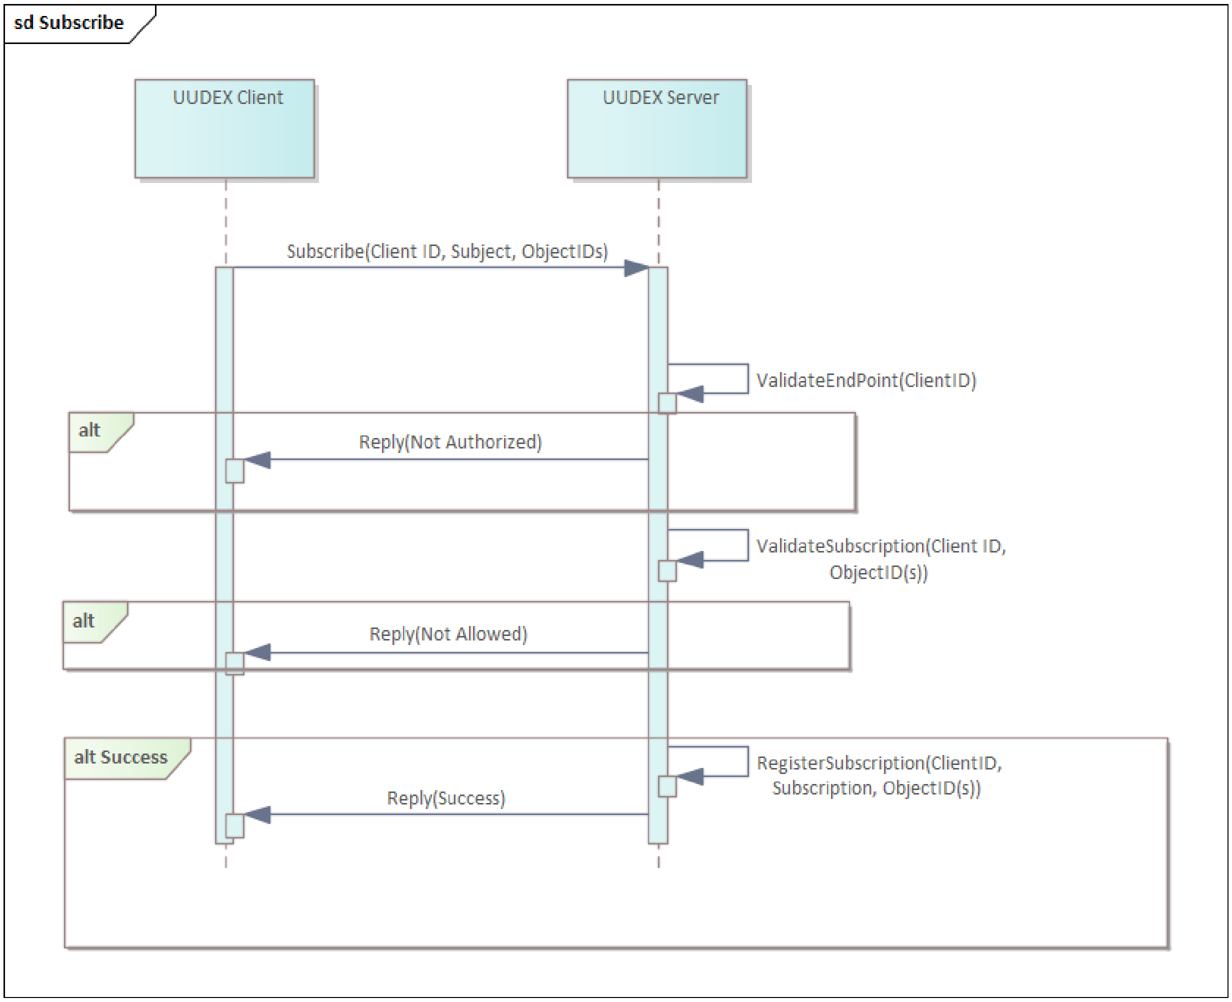
sd Subscribe
UUDEX Client
UUDEX Server
Subscribe(Client ID, Subject, ObjectIDs)
ValidateEndPoint(ClientID)
alt
Reply(Not Authorized)
ValidateSubscription(Client ID, ObjectID(s))
alt
Reply(Not Allowed)
alt Success
RegisterSubscription(ClientID, Subscription, ObjectID(s))
Reply(Success)

Figure 6-5: Subscribe

7.0 UUDEX Security Considerations

This section considers the security aspects of the U-Framework. In particular, it identifies security risks and security trust relationships within UUDEX.

7.1 Security Risks

This section identifies key security risks within the U-Framework. A security risk includes both risks imposed by security threats as well as concerns with security implications regardless of whether or not they result from malicious actors. This list is not exhaustive and other risks might exist within specific contexts. However, the list below identifies common, significant risks that any UUDEX deployment will face.

7.1.1 Information Disclosure

Many types of U-Data Elements exchanged over UUDEX are sensitive, and disclosure to unauthorized parties can result in damage to an organization or disruption of services. Examples of sensitive data include:

- Details about the configuration of IT or OT assets, particularly details about IT or OT security systems that could be used by malicious parties to plan cyberattacks.
- Disclosures of compromise of IT or OT assets that could impact an organization’s reputation and be used by malicious parties to plan cyberattacks.
- Market-sensitive data that could give competitors, customers, or suppliers an unfair market advantage.
- Data from third parties to whom the recipient has a contractual obligation to secure the data from disclosure. Possible consequences of the disclosure of these data include loss of future access to the data or financial penalties.

To guard against unauthorized disclosure, data must be protected both when at rest and at transit. In addition, all interactions with the data need to be constrained by ACLs to ensure that only authorized parties are able to view the data.

Note that it is not only the data that require protection from disclosure. Metadata linked to a U-Data Element must also be carefully managed to prevent unauthorized disclosure. For example:

- The fact that a given entity created a cyber incident report would indicate that entity experienced and detected a cyberattack, even if the incident report could not be read. This could have reputational impact on the report creator and could also tip off the attacker that their efforts had been discovered.
- The ACL of a U-Subject that contains market-sensitive data could reveal an entity’s commercial partners in an economic transaction, even if the data were not readable. This could give a competitor unfair leverage in the market.
- In some cases, the mere existence of some types of U-Data Elements might reveal sensitive information about an entity or the state of the grid. For example, certain classes of alerts might be sensitive because they direct operational changes and the issuance of such directives reveal a lot about the overall state of the power grid in ways that adversaries could use to their advantage.

For these reasons, interactions with data cannot leak information about associated metadata. In some cases, there needs to be no way to distinguish between a request that fails because access to the U-Data Element was denied and a request that fails because the requested U-Data Element does not exist. These and other error conditions need to be reviewed to ensure that they do not expose information to unauthorized parties.

7.1.2 Information Corruption

UUDEX does not support the exchange of control instructions, such as commands that directly manipulate the behavior of energy infrastructure devices. However, system operators and components will still be using information transported over UUDEX to make critical decisions that impact the functioning of their software and physical assets. For these reasons, it is critical that information stored by U-Servers and transported over U-Connections be protected against corruption, both deliberate and accidental.

Information corruption covers both changes to a U-Data Element’s meaning (such as changing an "off" indicator to an "on" indicator) as well as changes that render the data unintelligible. These different situations are likely to have different causes (e.g., the first is likely due to deliberate malicious activity, while the latter might be accidental), but both types of corruption can be damaging and must be addressed.

Data need to be protected from corruption both when at rest (i.e., when stored on a U-Server) and when in motion (i.e., when transmitted using a U-Connection). Ideally, such protections will prevent the corruption from occurring in the first place. At the very least, protections need to exist such that any data corruption will be detectable. This is necessary because small data corruptions, which could have significant operational impact, might not be readily detectable by all U-Participants (e.g., a false 10-degree voltage phase angle shift reported by PMU data).

Protection of the data includes protection of the U-Data Element's metadata. The metadata associated with a U-Data Element impacts how data are discovered, organized, identified, and accessed. As such, corruption of the metadata could result in multiple issues, including denial of data (if corrupted information means the data can no longer be found). For these reasons, the U-Data Element's metadata needs to receive the same protections against corruption as does its associated U-Data Element.

7.1.3 Denial of Service

Entities that use UUDEX services will rely on it for critical information and communications central to their operations. As such, loss of these services, through accident or malice, could disrupt those operations. Hence, UUDEX will need to include mechanisms to reduce the chance that it can be used to deny U-Participants necessary services.

One of the key services that UUDEX provides is support for U-Endpoints retrieving information from the U-Infrastructure. Denial of these services could come in many ways including, but not limited to:

- Loss of or congestion in the U-Connection between the U-Endpoint and U-Infrastructure.
- Rendering the U-Infrastructure, or portions thereof, unavailable or unable to adequately respond to U-Endpoint requests.

- Corrupting or deleting the U-Data Element the U-Endpoint is requesting, either on the U-Server or when it is in transit to the U-Endpoint.
- Altering access rights on the U-Subject so the U-Infrastructure's access control mechanisms prevent access.
- Corruption or deletion of U-Endpoint's U-Subscriptions established on a U-Subject such that the U-Endpoint is not alerted to the publication of relevant U-Data Elements.
- Corruption of the U-Infrastructure's indexing or search functionality, causing requests by the U-Endpoint to fail to find necessary U-Data Element.

Implementation of software products will need mechanisms to mitigate the chance that access to necessary U-Data Elements will be denied due to any of these circumstances.

Similarly, U-Producers will depend on the U-Infrastructure to deliver U-Data Elements to the appropriate U-Consumers. Specifically, the message needs to be sent to a U-Subject from a U-Producer and then made available to the appropriate U-Consumers. This service could be denied by events similar to those listed above:

- Loss of or congestion in the U-Connection between the U-Producer and U-Infrastructure or loss of connection between the U-Infrastructure and one or more valid U-Consumers for the given data.
- Rendering the U-Infrastructure, or portions thereof, unavailable or unable to adequately receive new U-Data Elements or to respond to U-Endpoint requests.
- Corruption or deletion of the U-Data Elements, either on a U-Server or in transit from the U-Producer or to the U-Consumer.
- Altering access rights on the U-Subject so the U-Infrastructure's access control mechanisms prevent access by legitimate U-Producers or U-Consumers.
- Corruption or deletion of U-Subscriptions from U-Consumers on the U-Subject such that those U-Consumers are not alerted to the publication of the U-Data Element.
- Corruption of the U-Infrastructure's indexing or search functionality, causing requests by the U-Consumers to fail to find the posted U-Data Element.

Implementations of U-Components will need to be able to mitigate these types of threats to minimize the chance that necessary services are lost.

7.1.4 Identity Spoofing

Key to any access control strategy is the ability to authenticate actors so that requested actions can be compared against controls. If one entity can masquerade as another then access controls can be circumvented. This, in turn, could lead to many of the aforementioned issues, including, but not limited to, information disclosure, unauthorized data manipulation, insertion of false data under the identities of trusted parties or services, and denial of service.

There are several ways identity information could be spoofed. These include, but are not limited to:

- Inadequate vetting of parties requesting U-Identity Objects. This occurs when a party provides false information when it requests a U-Identity Object. It might claim to be a different party or might claim roles or attributes that it should not be granted in its U-Identity Object.

- Falsifying identity evidence in messages. Some types of network attacks might allow valid credentials from one entity to be copied and used by another entity to pose as the former.
- Stolen private identity information. Identity evidence is often supported by calculations that only the valid holder of that identity could perform. However, if attackers were able to steal the secrets that allowed those calculations, they could pose as the entity.

Depending on the specific mechanisms UUDEX uses for managing identities, other attacks might also be possible. U-implementations will need to include mechanisms to protect against these threats.

There may be situations in which suppliers of UUDEX information need to remain anonymous. For example, reporters of attack details often wish not to be identified. The need to support anonymity, at least anonymity relative to certain other U-Participants, does not negate the need for strong identity controls. Mechanisms will be needed so participants are anonymous with regard to certain U-Participants, but can still be identified by other, trusted parties.

7.2 UUDEX Trust Relationships

Trust relationships exist where parties must rely on others to conform to certain behaviors in the absence of any way to enforce those behaviors. This section outlines the key trust relationships in UUDEX.

7.2.1 UUDEX Servers

U-Servers receive and act upon instructions from U-Endpoints. They also are responsible for storing U-Data Elements and delivering them as appropriate. U-Servers are trusted as follows:

- To enforce security policies on U-Subjects with regard to requests to post (publish), read (subscribe), and delete U-Data Elements.
- To accurately process queries for U-Data Elements. This means they must correctly identify matching U-Data Elements within a U-Subject to which a requesting party has access and accurately respond to the requestor based on this information.
- To accurately maintain and serve U-Subscriptions established by U-Endpoints.
- Not to add, modify, or delete U-Data Elements except at the direct instruction of an authorized U-Endpoint.
- To execute commands from authorized U-Endpoints (e.g., if a U-Server is instructed to delete a U-Data Element by an authorized U-Endpoint, the server is trusted to perform that action).
- To accurately report its status (e.g., whether its services are currently degraded).
- To conform to behaviors dictated by prioritization policies.

7.2.2 UUDEX Endpoints

U-Endpoints issue commands to the U-Infrastructure to post (publish), retrieve (subscribe), and delete U-Data Elements. U-Endpoints are trusted as follows:

- To adequately protect U-Data Elements they retrieve from the U-Infrastructure. In particular, they are trusted not to disclose the U-Data Element (intentionally or unintentionally) to parties that are not authorized to view the U-Data Element.

- Not to send false information in U-Data Elements.
- Not to create undue communications load by sending excessively large amounts of U-Data Elements to the U-Infrastructure.
- Not to create undue processing loads on the U-Infrastructure by making excessive U-Query or U-Subscribe requests.

7.3 UUDEX Access Control Overview

This section is intended to provide an overview of the role and scope of access control within the U-Framework.

- Establishing U-Connections – All U-Connections are required to be mutually authenticated. The U-Infrastructure controls access at this stage. As shown in Figure 7-1, only U-Participants (organizations that have been vetted and onboarded to the U-Instance) are allowed to establish communications. Onboarding involves providing the U-Participant with identification tokens that allow it to prove its identity to the U-Infrastructure.

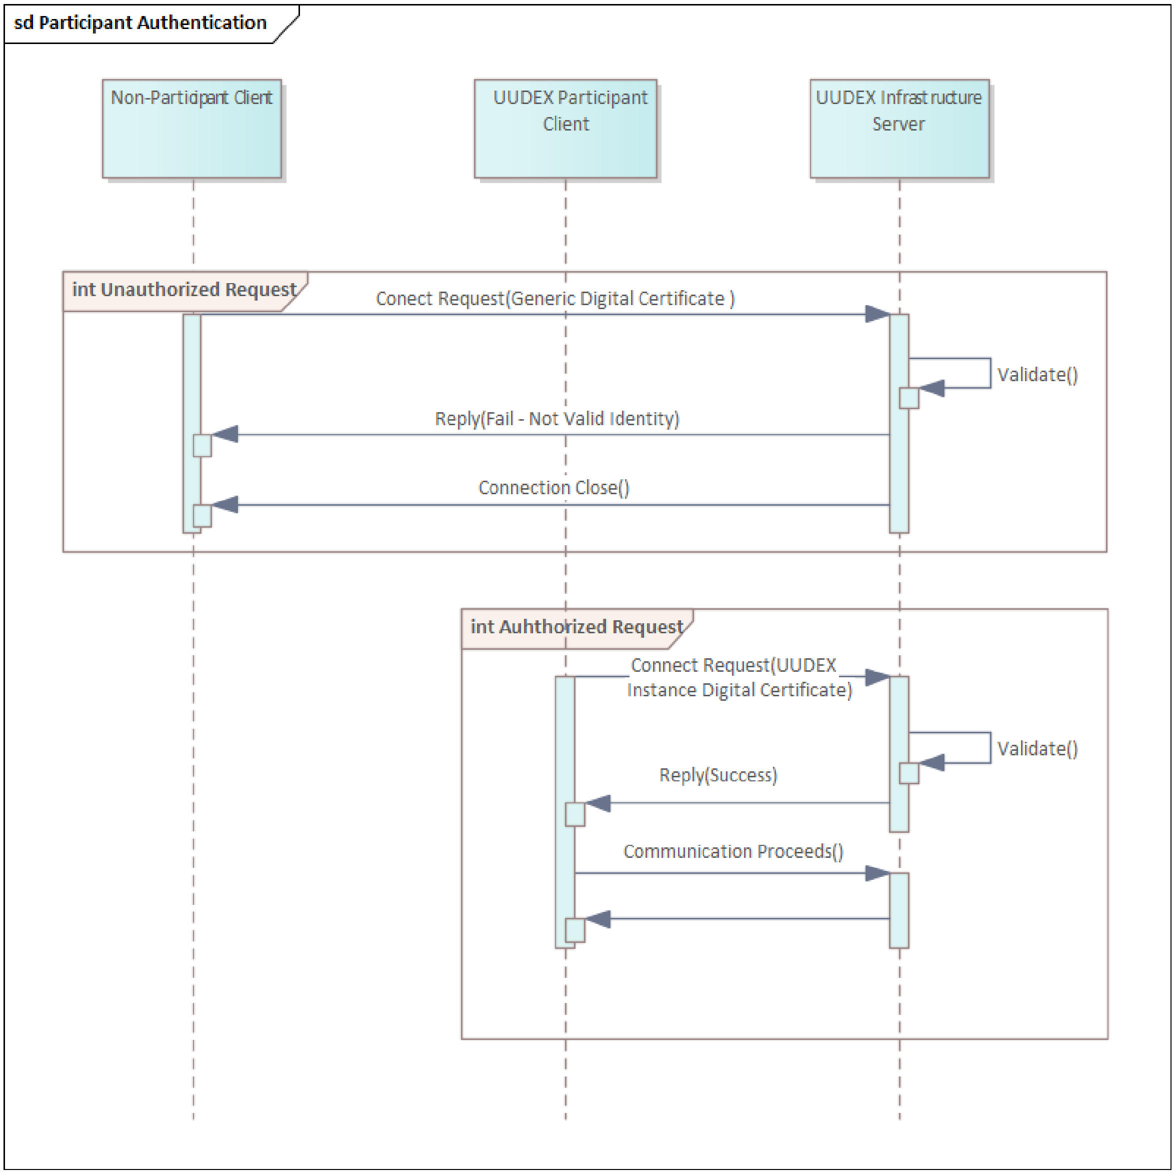
sd Participant Authentication
Non-Participant Client
UUDEX Participant Client
UUDEX Infrastructure Server
int Unauthorized Request
Conect Request(Generic Digital Certificate)
Validate()
Reply(Fail - Not Valid Identity)
Connection Close()
int Auhthorized Request
Connect Request(UUDEX Instance Digital Certificate)
Validate()
Reply(Success)
Communication Proceeds()

Figure 7-1: Participant Authentication

- When a U-Producer sends a U-Data Element to the U-Infrastructure, they identify the U-Subject to which it will be added. This U-Subject will have an associated ACL that governs subsequent access to this and all other U-Data Elements within that U-Subject.
- Access to UUDEX Subjects can be by individual U-Endpoint, by U-Participant, or by other roles or attributes. Boolean combinations of identities, attributes, and Boolean expressions determine the final list of parties allowed to access information.

8.0 References

Internet Engineering Task Force Requests for Comment:

RFC 3986, “Uniform Resource Identifier: Generic Syntax”

RFC 4180, “Common Format and MIME Type for Comma-Separated Values Files”

RFC 4122, “A Universally Unique Identifier URN Namespace”

RFC 7158, “The JavaScript Object Notation (JSON) Data Interchange Format”

International Electrotechnical Commission documents

IEC 61970, “Common Information Model / Energy Management”

IEC 61968, “Common Information Model / Distribution Management”

IEC 60870-6, “Telecontrol equipment and systems - Part 6: Telecontrol protocols compatible with ISO standards and ITU-T recommendations”. Also known as Inter-Control Center Communications Protocol or Telecontrol Application Service Element 2 (TASE.2).

Institute of Electrical and Electronics Engineers

IEEE Std C37.111, “IEEE Standard Common Format for Transient Data Exchange (COMTRADE) for Power Systems

IEEE Std C37.118.2, “IEEE Standard for Synchrophasor Data Transfer for Power Systems”

IEEE Std P2664, Proposed “IEEE Standard for Streaming Telemetry Transport Protocol”

North American Electric Reliability Corporation Standards

EOP-004, “Event Reporting”

CIP-008, “Cyber Security — Incident Reporting and Response Planning”

International Standards Organization Standards

ISO 8601, “Data elements and interchange formats – Information interchange – Representation of dates and times”

ISO/IEC 10918, “Information technology -- Digital compression and coding of continuous-tone still images (JPEG)”

World Wide Web Consortium

“eXtensible Markup Language (XML)”

European Computer Manufacturers Association

ECMA 404, “The JSON (JavaScript Object Notation) Data Interchange Syntax”

U.S. Department of Energy

“Electric Emergency Incident and Disturbance Report (Form OE-417)”

OTHER

“Structured Threat Information eXpression (STIX™)”

“Trusted Automated eXchange of Indicator Information (TAXII™)”

“Packet Capture” (PCAP)

-

Appendix A – Data Characteristics

This Appendix contains a set of notional data element characteristics considered during the development of the functional specification as a table. It shows various types of data planned for UUDEX as columns, and data attributes as rows.

This information will be used to derive the requirements and implementation considerations for UUDEX.

	ICCP Data	RCIS	Power System Models	PMU	OE-417 Report	Files (COMTRADE, others)	Market Data	Asset Management	Cyber Incident Reporting (IDS/IPS logs, PCAP, etc.)	Indicator of Compromise sharing	Guidance (firewall, configuration, etc.)	Patch notification	Vulnerability or Threat notification (STIX, DOE, NERC, DHS) (Non-public)	Vulnerability or Threat notification (CVE, STIX, etc.) (public)
Additional security	Org role, ACL	Org role, ACL	Org role, ACL			File type, org role, ACL	Yes	org role, ACL	Yes – confidentiality, destination authentication	Yes – Integrity, confidentiality, mutual authentication	Yes – integrity, source authentication	Yes – integrity, source authentication	Yes – confidentiality, mutual authentication, integrity	Yes – integrity
Permissible delay between transmission and reception of the data set	Near R/T	Near R/T	minutes	Near R/T	minutes	Seconds-minutes	Near R/T	Minutes	Minutes	Minutes	Minutes	Minutes	Minutes	Minutes
Durability		Yes	Yes		Yes	Yes	Yes	Yes						
Retention or expiration	Ephemeral	Persistent	Persistent	Ephemeral except post event	Persistent	Persistent	Varied	Persistent	Persistent	Persistent	Persistent (short term)	Temporary	Persistent (short term)	Persistent (short term)
Core, desired or optional?	Core	Core	Core	Desired	Core	Optional	Optional	Optional	Desired	Desired	Desired	Desired	Core	Desired

	ICCP Data	RCIS	Power System Models	PMU	OE-417 Report	Files (COMTRADE, others)	Market Data	Asset Management	Cyber Incident Reporting (IDS/IPS logs, PCAP, etc.)	Indicator of Compromise sharing	Guidance (firewall, configuration, etc.)	Patch notification	Vulnerability or Threat notification (STIX, DOE, NERC, DHS) (Non-public)	Vulnerability or Threat notification (CVE, STIX, etc.) (public)
Frequency of transmission	2 seconds+	*Ad hoc*	*Ad hoc*, periodic	stream	*Ad hoc*	*Ad hoc*, daily+	hourly, daily, transaction	*Ad hoc*	*Ad hoc* (weekly?)	Periodic (daily); *ad hoc* for high- priority	Periodic (weekly); *ad hoc* for high- priority	Periodic (weekly); *ad hoc* for high- priority	*Ad hoc*	Periodic (weekly); *ad hoc*
Subscribable?	By point, point type, publisher	Yes	By model profile, by publisher	By point	Yes	By file type, by publisher	By object type, by publisher	By asset type, by publisher	Reported to designated authority	Yes (by multiple parameters)	Yes – by device covered	Yes – by device covered	Yes (by multiple parameters)	Yes (by multiple parameters)
Persistent publisher connection required?	Yes	No?	No	Yes	No	No	Yes?	No	No	No	No	No	No	No
Persistent subscriber connection required?	Yes	No?	No	Yes	No	No	Yes?	No	No	No	No	No	No	No
Stored for future use?	Recent history	Yes	Yes	Selectively; Recent snapshots	Yes	Yes	Recent history	Yes	Yes	Yes	Yes (short term)	Unlikely	Possibly	Possibly
Intermediate processing (e.g., down sample)	Down sample, Periodic snapshot, significant change, event detect	Store and forward	Store and notify	Down sample, Periodic snapshot, significant event detect	Store and forward	Store and notify	Store and notify	Store and notify	Probably: store, augment, anonymize, redistribute	Possibly: store, augment	Unlikely	Unlikely	Possibly – store, augment, anonymize, redistribute	Possibly – store, augment, anonymize, redistribute

	ICCP Data	RCIS	Power System Models	PMU	OE-417 Report	Files (COMTRADE, others)	Market Data	Asset Management	Cyber Incident Reporting (IDS/IPS logs, PCAP, etc.)	Indicator of Compromise sharing	Guidance (firewall, configuration, etc.)	Patch notification	Vulnerability or Threat notification (STIX, DOE, NERC, DHS) (Non-public)	Vulnerability or Threat notification (CVE, STIX, etc.) (public)
Sender down sample (e.g., for degraded link performance)	Yes	No	Directory, change log	Yes	Directory, log	Directory, log	Directory, log	Change log	Unlikely	Unlikely	No	No	No	No
Priority of Message				Event-dependent										
QOS (Traffic Prioritization)	High (?)	High		Event-dependent	High				Med	Med (periodic) or high (*ad hoc*)	Med	Med	High	Med
Logical Topography	P2P, Pub/Sub	Pub/Sub	Pub/Sub	Peer-to-Peer, Pub/Sub	Pub/Sub	Pub/Sub	Pub/Sub	Pub/Sub	Hub-spoke	Pub/Sub	Pub/Sub	Pub/Sub	Hub-spoke	Pub/Sub
Minimum expected bandwidth (?)									MBs	MBs	MBs	MBs	MBs	MBs
Size of logical message / “data set”	Continuous, MBs	<1k	Large	Continuous	Small / medium	Large	Varied	Varied	10K-100MB	10K-100MB	10K-10MB	1MB-100MB	10K-100MB	10K-100MB
Originator Control required?	Yes			Yes					Yes – sensitive info included	Yes	Maybe	No	Yes	No

	ICCP Data	RCIS	Power System Models	PMU	OE-417 Report	Files (COMTRADE, others)	Market Data	Asset Management	Cyber Incident Reporting (IDS/IPS logs, PCAP, etc.)	Indicator of Compromise sharing	Guidance (firewall, configuration, etc.)	Patch notification	Vulnerability or Threat notification (STIX, DOE, NERC, DHS) (Non-public)	Vulnerability or Threat notification (CVE, STIX, etc.) (public)
Recipient designations	Subscriptions as permitted by role or ACL	Subscriptions as permitted by role or ACL	Subscriptions as permitted by role or ACL	Subscriptions as permitted by role or ACL; BAs, RCs	Subscriptions as permitted by role or ACL	Subscriptions as permitted by role or ACL	Subscriptions as permitted by role or ACL	Subscriptions as permitted by role or ACL	Yes (to authority)	Subscriptions as permitted by role or ACL	Subscriptions	Subscription	Subscriptions as permitted by role or ACL	Subscription
Special Handling Instructions					Yes				Yes	Yes	No	No	Yes	No
Sensitivity marking	Yes?	Yes	Yes		Yes	By file type	By object type		Yes	Yes	Maybe (unlikely)	No	Yes	No
(Maximum) Sensitivity level	Different point types could have different sensitivity					By file type	By object type		Major org/ op impact	Private	Private	Public	Major org/ op impact	Public
Integrity sensitivity	High	Moderate	Moderate (can be independently validated before use)	Could have minor organizational/ operational impact; Timing data must be intact	Moderate	High	High	Moderate	Moderate to low	Moderate	High	High	High	Moderate
Reasonable number of recipients	0-20	>100	>100	0-10	>100	>100	>100	>100	0-20 (at least initially)	>>100	>>100	>>100	>>100	>>100

	ICCP Data	RCIS	Power System Models	PMU	OE-417 Report	Files (COMTRADE, others)	Market Data	Asset Management	Cyber Incident Reporting (IDS/IPS logs, PCAP, etc.)	Indicator of Compromise sharing	Guidance (firewall, configuration, etc.)	Patch notification	Vulnerability or Threat notification (STIX, DOE, NERC, DHS) (Non-public)	Vulnerability or Threat notification (CVE, STIX, etc.) (public)
Trigger for sending	Periodic, on change	On create	On change or run	Periodic, stream; upon request	On create	On create, on update	On create	On create, on update	Automated or manual creation	All	All	All	All	All
Trigger for requesting			Notification receipt	A disturbance or event		Notification receipt	Notification receipt	Notification receipt	N/A (not requested from source)	Manual, periodic, or subscription	Manual, periodic, or subscription	Manual, periodic, or subscription	Manual, periodic, or subscription	All

Appendix B Changes from November 2018 Version

Changes Made May 2021

- Deleted the terms UUDEX Auditor, UUDEX Bridge, UUDEX Cloud, UUDEX Configuration Object, UUDEX Directory, UUDEX Monitor, and UUDEX Workflow from the set of defined terms. Some of these concepts have been merged with other terms (e.g., documents now use UUDEX Infrastructure to cover uses formerly tied to the UUDEX Director term). Other terms only ended up being used once or twice across the documentation, and thus are best served by explaining their meaning at those locations rather than defining a global term.
- Updated requirements to reflect that UUDEX Instance is being used to define an isolated trust domain. This is not a change in this term's meaning as much as it is a reflection how this particular aspect of UUDEX Instances has been amplified in subsequent documentation.
- Removed references to the modification of UUDED Data Elements after their storage in a UUDEX Subject. Subsequent design decisions have removed this behavior as adding unnecessary complexity to the architecture. This involved deleting functional requirements previously numbered as FLOW-5, DAT-1.4, DAT-1.5, DAT-2, DAT-2.1, and DAT-2.2. Where necessary, the remaining functional requirements have been re-numbered to remove gaps in numbering.
- Removed references to an explicit "Test Mode" in UUDEX. The current design allows test data to be sent over the UUDEX Infrastructure with appropriate markings, but this is done alongside regular data rather than using a separate "mode". Message prioritization as well as marking UUDEX Data Elements as test messages ensures these tests are not disruptive. This edit involved deleting functional requirement TM-4.1.
- Removed references to a UUDEX Participant owning a UUDEX Subject. Participants will still own Subjects, but the original wording implied an exclusivity of activities conferred by ownership that subsequent revisions decided to change. Subject ownership is explained in more detail in other documents, so references to subject ownership were removed to avoid the incorrect impression they were originally giving.
- Because of multiple changes in how components interact in a UUDEX Infrastructure, the use cases have been completely refactored. The same set of use cases are presented, but the means by which UUDEX supports those use cases has been revised to reflect current designs.
- Changed “UUDEX xxx” references to “U-xxx” for readability

Changes Made November 2019

- Replaced UUDEX Tunnel with UUDEX Connection. Tunnels imply a point-to-point connection, but the UUDEX Directory concept implies support for brokered connections.
- The term UUDEX Data Envelope has been removed. Instead, where appropriate the document talks about "metadata associated with a UUDEX Data Element".
- Prioritization levels are now assigned to UUDEX Subjects rather than individual UUDEX Data Elements.
- ACLs are now assigned to UUDEX Subjects rather than individual UUDEX Data Elements.
- Deleted Functional Requirement ARCH-2.1: "UUDEX Servers MUST be configurable with a policy that compares the assigned priority to other fields of a UUDEX Data Element or

UUDEX Data Element Envelope and rejects or downgrades the priorities of UUDEX Data Elements that have been assigned an inappropriate priority. For example, a UUDEX Server could define a policy where certain classes of routine messages are not allowed to be given a high priority, and automatically reject or reduce the priority of such routine messages whose priority has been set too high by their UUDEX Producer. Those who deploy UUDEX Instances are not required to make use of this capability, but UUDEX Server implementations must be capable of letting operators define and enforce such a policy. (This document refers to priorities as high, medium, or low for the sake of examples, but any number or type of levels of prioritization might be employed so long as they are strictly ordered.)" Because prioritization is now assigned to UUDEX Subjects rather than UUDEX Data Elements, it makes less sense for a UUDEX Server to dynamically analyze UUDEX Data Elements and reassign their priorities individually.

- Deleted Functional Requirement ARCH 2.2: "UUDEX Server MUST be able to assign priorities to UUDEX Data Elements based on fields of a UUDEX Data Element or UUDEX Data Element Envelope in the case that the priority is not assigned by the UUDEX Producer. This allows UUDEX Instances to delegate assignment of UUDEX Data Element priorities to UUDEX Servers if the operators wish to do so." Because prioritization is now assigned to UUDEX Subjects rather than UUDEX Data Elements, scanning fields and reassigning priorities based on field values is likely to be impractical.
- Deleted Functional Requirement ARCH 6.1: "UUDEX Participants MUST have the ability to automatically switch between different UUDEX Connections based on the connectivity status of any link." This made sense when talking about distinct UUDEX Tunnels, but not with more general UUDEX Connections.
- Deleted Functional Requirement ARCH-6.2: "UUDEX Participants MUST have the ability to manually switch between different UUDEX Connections." This made sense when talking about distinct UUDEX Tunnels, but not with more general UUDEX Connections.
- Deleted Functional Requirement DAT-1.5: "In some cases, the producer of a UUDEX Data Element might wish not to be associated with the UUDEX Data Element. For example, entities wish to submit cyber threat intelligence data regarding a detected intrusion anonymously, so they do not reveal they were able to detect the attack. For this reason, the UUID associated with the UUDEX Data Elements does not need to come from the party that produced it. One option to accomplish this includes having a service to which UUDEX Data Element producers can submit data, which will assign a UUID to that UUDEX Data Element and submit it to a UUDEX Server on behalf of the original UUDEX Producer without assigning attribution to the original producer. Another option would be to set up a service by which a UUDEX Producer could request a UUID generated by a third party. Either of these would allow the UUDEX Data Element to enter a UUDEX Server with a UUID that is not associated with the original data producer." All UUDEX Subjects explicitly identify the source of the UUDEX Data Elements they publish, which precludes anonymous publication. For data where broad publication is needed without disclosing the identity of the source (such as for certain classes of cyber incident reporting) the source would publish to a UUDEX Subject that only had trusted subscribers and the trusted subscriber would republish the data under their own UUDEX Subject (after any necessary data sanitization).
- Deleted Functional Requirement DAT-4: "UUDEX Servers MAY establish access controls over UUDEX Repositories. This is in addition to access controls that individual UUDEX Data Elements in a UUDEX Repository might include. UUDEX Repository access controls could limit which entities were permitted to post data to that UUDEX Repository. For other actions (e.g., read, modify, delete), a UUDEX Client would need to be granted access both to perform

the given action by the UUDEX Repository and by the UUDEX Data Element for the action to proceed." Given that the UUDEX Directory abstracts the UUDEX Server concept, applying controls on a Server-by-Server basis no longer makes sense.

- Deleted Functional Requirement DAT-6.1: "All UUDEX Data Elements published to a UUDEX Server MUST (implicitly or explicitly) automatically grant the UUDEX Server Infrastructure the rights to be aware of the UUDEX Data Elements. The UUDEX Server MUST reject requests to publish UUDEX Data Elements that do not grant this right." Now that access control is managed by Subject, it makes little sense for a Subject to be established in the UUDEX Infrastructure that the UUDEX Infrastructure cannot be aware of.
- Deleted Functional Requirement DAT-6.2: "A UUDEX Server MAY receive UUDEX Data Elements to which it is not granted read access. In this case, the server r MUST only use the UUDEX Data Element in the UUDEX Data Envelope for the purpose of storing the UUDEX Data Element in the appropriate UUDEX Repository and for matching against search and UUDEX Subscription requests. If the UUDEX Data Envelope does not have the necessary information to allow the UUDEX Server to do these tasks, the UUDEX Server MUST reject the request to publish the UUDEX Data Element." Since the location UUDEX Subjects within the UUDEX Infrastructure is abstracted, talking about the access of individual UUDEX Servers does not make sense.
- Deleted Functional Requirement DAT-6.3: "A UUDEX Server does not require access rights to a UUDEX Data Element regarding a particular action in order to undertake the action as instructed by an authorized UUDEX Client. For example, if an authorized UUDEX Client requests the deletion of a UUDEX Data Element, the UUDEX Server can fulfill that request even if the UUDEX Server does not have access rights to delete the UUDEX Data Element. Thus, with regard to the UUDEX Server and access control of UUDEX Data Elements, those controls describe UUDEX Server processes rather than strictly enforced controls. The UUDEX Server must be trusted not to violate the terms of the access controls, despite the fact that, in practice, it will have the ability to do so." Since the location UUDEX Subjects within the UUDEX Infrastructure is abstracted, talking about the access of individual UUDEX Servers does not make sense.
- Deleted the UUDEX ACL Object and replaced with ACL.
- Removed the term UUDEX Client in favor of UUDEX Endpoint. The terms started switching in the Protocol Design document and the Workflow Design document exclusively uses UUDEX Endpoint.
- Removed the term UUDEX Repository as the concept is subsumed by the concept of the UUDEX Directory. The terms started switching in the Protocol Design document and the Workflow Design document exclusively uses UUDEX Endpoint.
- Removed the term UUDEX Role as it is both obsoleted and becomes confused with the roles of persons that get defined in the Workflow Design. The term is not critical and removing it removes this confusion. When appropriate, this was replaced with UUDEX Component.
- Clarified that UUDEX Participants to not create UUDEX Connections between each other. Instead, connections are between UUDEX Participants and the UUDEX Infrastructure, as noted in both the Workflow and Protocol Designs.

Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)

www.pnnl.gov