MHHS PROGRAMME
Industry-led, Elexon facilitated

MHHS Data Integration Platform - Functional Specification

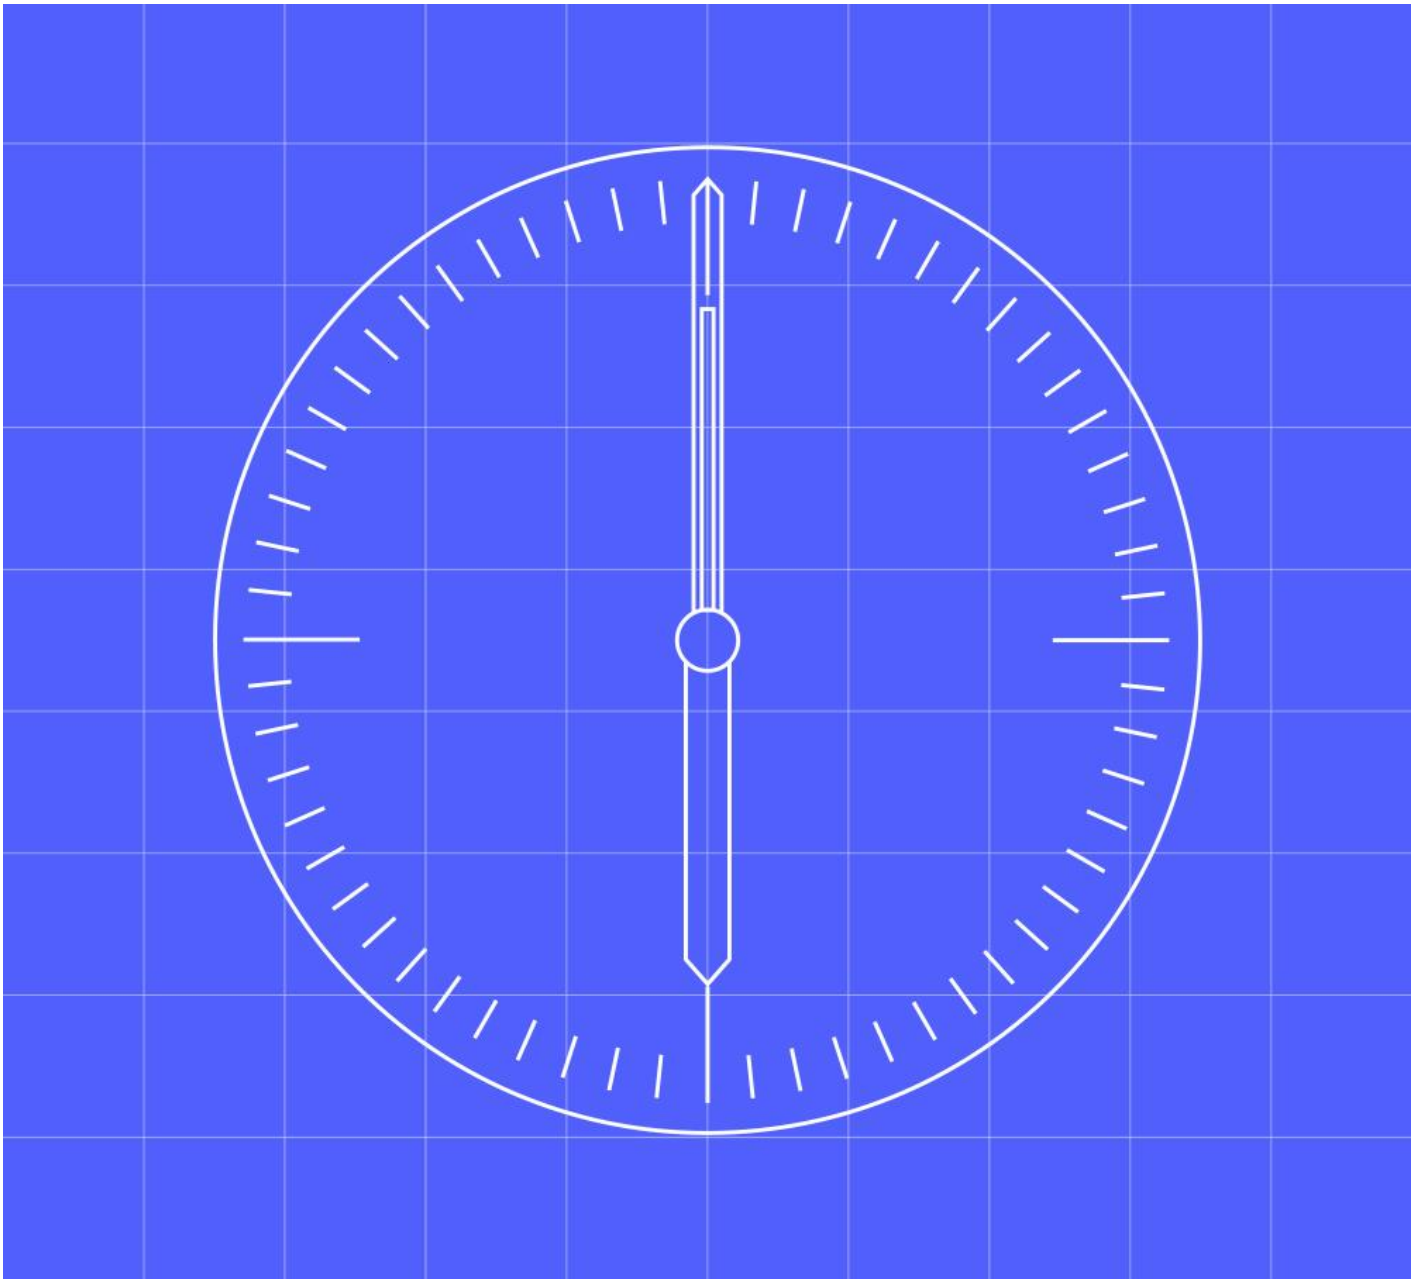

Document owner	Document number	Version	Status:	Date
MHHS DAG	**MHHS-DIP001**	**Version 2.0**	**For Sign Off**	**27 April 2022**

1 Contents

1.1 Change Record

Date	Author	Version	Change Detail
November 2021	RG	0.1	Initial draft; sent for informal review to TDWG
04 February 2022	RG	0.2	The second draft was sent to TDWG for formal review prior to the RFP release
25 February 2022	RG	0.3	Final draft prior to RFP release
03 March 2022	RG	1.0	First Issue
28 March 2022	RG	1.1	Updates following RFP comments
27 April 2022	RG	2.0	Sent to DAG for sign-off

1.2 Reviewers

Reviewer	Role
TDWG	Review
DAG	For information and sign off

1.3 References

Document/Link	Publisher	Vers	Date	Additional Information
MHHS Programme Governance Framework - (Strawman)	MHHS	v1.0		
The Target Operating Model for Market-wide Half Hourly Settlement - Design Working Group's Recommendation to Ofgem	Elexon	V1.1	12 Feb 2019	
MHHS AWG Recommendation	Elexon	v1.0	22 Apr 2021	
MHHS Risk and Management Process	MHHS	v0.1	August 2021	In preparation
MHHS Architecture Principles	MHHS	V0.1	March 2022	In preparation
MHHS End-to-End Security Architecture	MHHS	V1.0	March 2022	
MP162 'SEC changes required to deliver MHHS' Business requirements – version 0.5	SEC	V0.5		December 2021 https://smartenergycodecompany.co.uk/document-download-centre/download-info/mp162-business-requirements-v0-2/
MHHS-DIP002-Non-Functional Requirements	MHHS	V1.0		March 2022

1.4 Terminology

Term	Description
API	Application Programmable Interface
ARP	Advanced Retrieval and Processing Service
AWG	Architectural Working Group
BRP	Balancing Responsible Party
BSCCo	Balancing and Settlement Code Company (Elexon Limited)
COS	Change of Supplier

DIP	Data Integration Platform
DIPSP	DIP Service Provider
DCC	Data Communications Company
DNO	Distribution Network Operator
DPIA	Data Protection Impact Assessment
DTC	Data Transfer Communications
DTN	Data Transfer Network
DWG	Design Working Group
ECOES	Electricity Central Online Enquiry Service
ESO	Enduring Service Owner, i.e. the party with overall responsibility for the DIP
EDA	Event-Driven Architecture
HHR	Half Hour
HHS	Half-Hourly Settlement
HTTP	Hypertext Transfer Protocol
iDNO	Independent Distribution Network Operator
IP	Internet Protocol
JSON	JavaScript Object Notation
LDSO	Licensed Distribution System Operator
LSS	Load Shaping Service
MDR	Meter Data Retrieval Service
MDS	Market-wide Data Service
MHHS	Market-Wide Half-Hourly Settlement
MPAN	Metering Point Administration Number
MRS	Meter Reading Service
MSA	Metering Service (Advanced)
MSS	Metering Service (Smart)
mTLS	Mutual Transport Layer Security
PII	Personal Identifiable Information
PSS	Processing Service (Smart)
RBAC	Role Based Access Control
RECCo	Retail Energy Code Company (Retail Energy Code Limited)
RR	Register Reading
RFP	Request for Proposal
SD	Settlement Day
SDS	Smart Data Services
SMRS	Registration Service
SP	Settlement Period
SUP	Supplier
SVA	Supplier Volume Allocation
SWIKI	Switching Public Key infrastructure
TLS	Transport Layer Security
TOM	Target Operating Model
UMDS	Unmetered Supplies Data Service
UMSO	Unmetered Supplies Operator Service
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
VAS	Volume Allocation Service

2 Introduction

2.1 Background

In its 2017 Significant Code review Ofgem identified that Half-Hourly Settlement (HHS) on a market-wide basis would be introduced into the UK electricity market. A cross-industry Design Working Group (DWG) was established to understand the feasibility of HHS and how it could be delivered. The DWG produced a Target Operating Model (TOM) that outlines the new ways of working which could deliver HHS into the market (*Reference Target Operating Model for Market-wide Half Hourly Settlement).*

In conjunction with the DWG, an Architectural Working Group (AWG) was established to propose an IT architecture that could support the business process outlined in the TOM. The AWG recommendation was that an Event-Driven Architecture (EDA) be implemented (*Reference MHHS AWG Recommendation*). Hence, a new message orientated/event-driven middleware component is required – the Data Integration Platform (DIP) - to support the flow of events/messages between industry participants proposed by the EDA.

The MHHS Programme was set up to continue the preparatory work undertaken by the DWG & AWG, refine the TOM further, and then oversee its delivery into the industry.

2.2 Objective

This functional specification provides an initial view of the end-to-end solution architecture that will support the proposed MHHS TOM and the requirements for the DIP. It expands on the good work produced by the AWG to inform industry participants how the new business process operates, describes the supporting IT architecture, and provides a level of design detail whereby all industry parties understand the scope of the changes affecting them. It should also assist parties preparing to bid for the provision of the DIP under the forthcoming Request for Proposal (RFP).

Many solution architecture principles have been taken from the AWG work and are documented *(Reference MHHS Architecture Principles, March 2022, in draft).*

2.3 Document Scope

The scope of the document includes

- An overview of the different MHHS Actors and Roles
- An overview of the MHHS business data flow – this document does not describe the business processes (these will be made available in later programme documents)
- A view of the functional and non-functional requirements for the DIP
- Describes the design principles and assumptions that will underpin the DIP
- Describes a generic messaging architecture that can be used to describe the required capabilities of the DIP
- Describes the service management functions required to support the DIP and its users

The document provides a narrative for the detailed requirements for the DIP that can be found in the accompanying spreadsheet (*Reference MHHSP-DIP002-Non-Functional Requirements, March 2022).*

2.4 "Event-Driven Architecture"

The terms **message** and **event** are used synonymously throughout this document. The document takes an agnostic approach to the underlying architecture platform required to implement the DIP and hence presents the DIP requirements without referencing an underlying platform/architecture. The AWG recommendation is that an Event-Driven Architecture should be implemented. It references a Gartner Report that describes Event-Driven Architectures, which in turn defines three basic types of events brokers:

- Queue-oriented (like Solace PubSub+, RabbitMQ, Azure Service Bus, etc.)
- Log-oriented (like Apache Kafka, Amazon Kinesis)
- Subscription-oriented (such as Amazon EventBridge and Azure Event Grid).

Hence, the term "message" is more suited to queue orientated brokers, whilst "event" is more suited to log and subscription orientated brokers. The DIP design is broken down into "message/event channels", each implemented by

a specific pattern via a template. Each template pattern implements a set of common requirements; the RFP bidder should consider the best technology platform for matching the requirements.

Please note that when the various capabilities are described, they are presented in a messaging-based system rather than an event-based system; this is convenient as many of the requirements are easier to explain using a messaging system over an event system.

3 MHHS TOM - Functional Architecture

3.1 Objectives

The objective of the MHHS Programme is to create a durable, faster, more accurate settlement process for all market participants, enabling broad change across the electricity industry.

The functional architecture will define a set of services required to deliver Settlement Period (SP) data from a Meter to a central Settlement body to enable the calculation of the amount of energy that the electricity Supplier's customers have consumed (or exported) in each Settlement Period for each Settlement Day (SD). This calculation is then used in the Imbalance Settlement process, which compares the Supplier's contracted purchases of energy to the amounts deemed to have been consumed (sales) by each of the Supplier's customers (and recognises any amounts of energy contracted by National Grid under the Balancing Mechanism). Settlement Data is also provided for network charging.

In addition to these core services, a number of supporting services need to be established to ensure the smooth running of the electricity market with the move to market-wide half-hour metering.

3.2 Key Assumptions/Principles

The following are the key design assumptions when establishing the message flows within the MHHS TOM:

1. The DIP will broker new message flows between Market Participants supporting the business process underpinning the MHHS TOM, i.e. there are no direct point-point interfaces.
2. Some of the reworked existing business processes falling under the programme's scope have existing Data Transfer Catalogue (DTC) flows that use the DTN (Data Transfer Network), and these will be retained where there is no apparent change to that interface. Any interfaces that require change will be re-implemented in the DIP.
3. The DIP will be a 'stateless' messaging system meaning that it will not be responsible for orchestrating the underlying business process, i.e. each message/event is considered distinct and has no dependency or interaction with any other message. Although from the wider MHHS TOM context the DIP itself will be stateless, the status of individual messages/events as they progress through the DIP will be stateful, i.e. persist and survive service restarts.
4. In terms of business logic, the remit for the DIP will be:
    - It will
        i. Validate event/message headers
        ii. Validate the event/message structure
        iii. Address and route messages based on content
    - It will NOT:
        i. Orchestrate the business processes, e.g. synchronise message handling
        ii. Validate message body content

3.3 MHHS - Target Operating Model

The scope of the MMHS work covers the “Meter to Bank” process for all Supplier Volume Allocation (SVA) Settlement Meters – i.e., all Settlement Meters connected to distribution networks. This includes:

- Meter Registration - the recording of information pertinent to Settlement Metering Systems;
- Meter Operations - fitting and maintaining Settlement Meters;
- Data Retrieval - getting information from Settlement Meters;
- Data Processing – validating and estimating Settlement Meter data;
- Data Aggregation - summing Settlement Meter data to required granularity; and
- Volume Allocation – allocating Meter volumes to Trading Parties' signatories to the BSC.

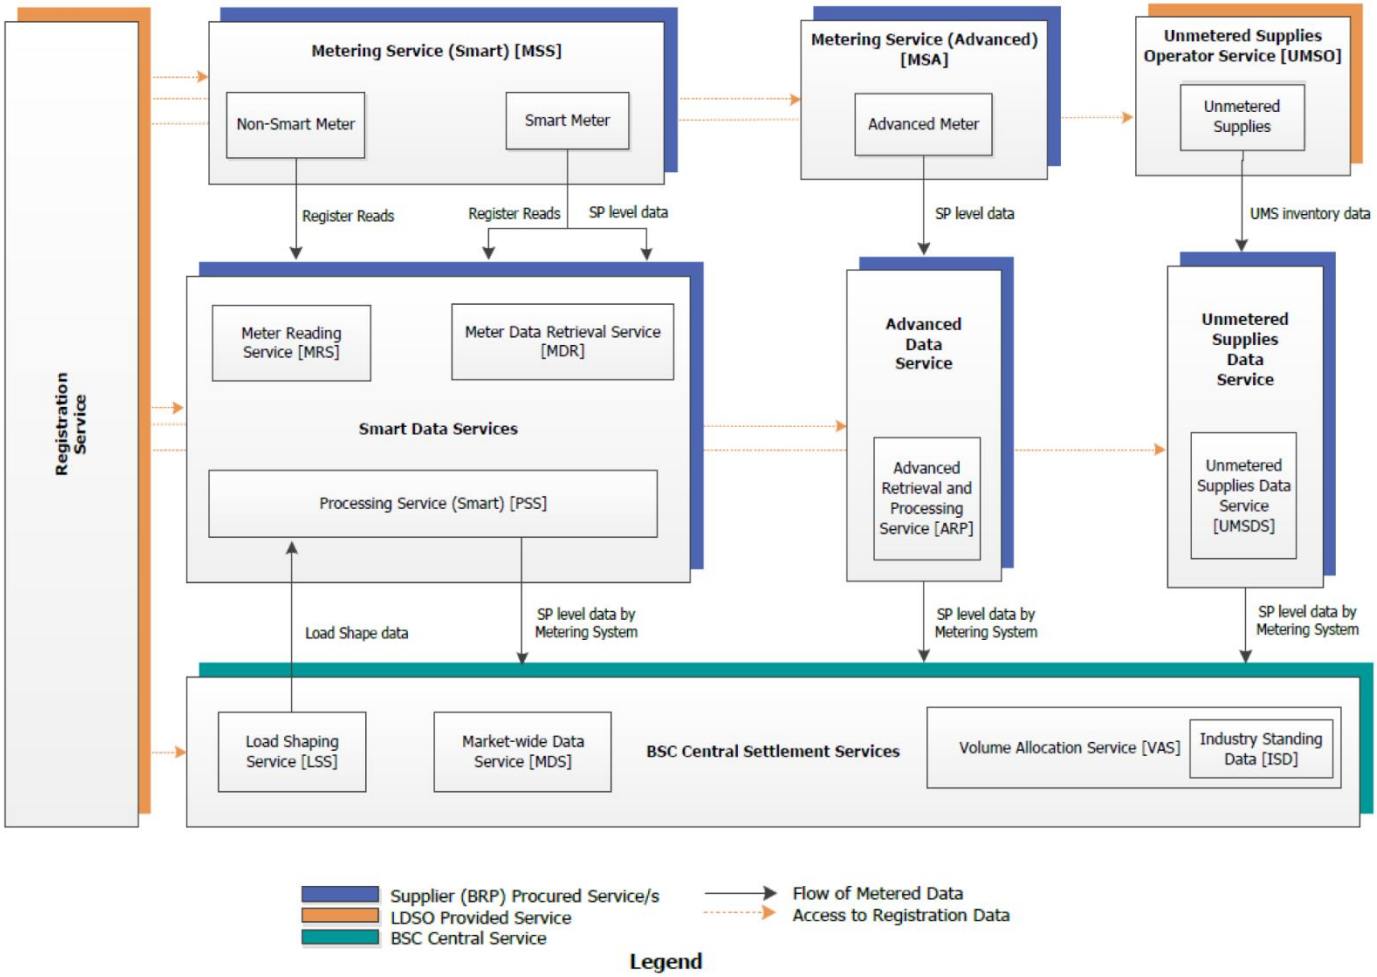

Figure 1 - Scope of the MHHS TOM

The business architecture is described as a set of services to be undertaken. A service is a set of requirements and processes required to deliver one function of MHHS. A service is agnostic of current organisation roles (such as Supplier, Meter Operator, Data Collector, Data Aggregator).

3.3.1 Roles

A number of different roles and Users (Market Participants) have been identified:

Role	Market Segment	Estimated Number of Users
Data Services	Smart	21
	Advanced	7
	Unmetered	7
Metering Service	Advanced	50
	Unmetered	35
Registration Services		27

Licensed Distribution System Operator	27
Supplier	60
DNO	6
iDNO	13

The following roles and services are described:

Service ID	Service Name	Market Segment/Role
MSA	Metering Service (Advanced)	Advanced Market Segment and Advanced Data Service
ARP	Advanced Retrieval and Processing Service	
MSS	Metering Service (Smart)	Smart and non-smart Market Segments and Smart Data Services (SDS)
MDR	Meter Data Retrieval Service	
MRS	Meter Reading Service	
PSS	Processing Service (Smart)	
UMSO	Unmetered Supplies Operator Service	Unmetered Supplies Market Segment and Unmetered Data Service
UMDS	Unmetered Supplies Data Service	
MDS	Market-wide Data Service	BSC Central Settlement Services (CSS)
LSS	Load Shaping Service	
VAS	Volume Allocation Service	
SMRS	Registration Service	Registration
SUP	Supplier	Supplier

3.3.2 Metering Service (MSA & MSS)

The principal functions of a Metering Service (Smart (MSS) and Advanced (MSA)) are to install, commission, test, maintain, rectify, energise and remove faults in respect of Metering Equipment (including, where applicable, associated Communications Equipment). The MSS and MSA will also maintain and make available Meter asset information and, where required, Meter configuration information.

3.3.3 Advanced Retrieval and Processing Service (ARP)

The Advanced Retrieval and Processing Service (ARP) is responsible for obtaining raw meter readings (SP level and Register Reads), validating and estimating (where needed) for Advanced Meters. On an optional basis, this service will also be responsible for complying with the shared metering arrangements - which allocate Metering System data between one or more BRPs.

3.3.4 Meter Data Retrieval Service (MDR)

The Meter Data Retrieval (MDR) Service is the service that submits Service Requests for settlement data via the Data and Communications Company (DCC). The Service Request type and schedule are provided by the Processing Service (Smart) (PSS) for each Metering Point Administration Number (MPAN) responsible for the PSS.

3.3.5 Meter Reading Service (MRS)

The Meter Reading Service (MRS) is the service that provides Register Readings (RRs) for Meters where Settlement Period Level data is not available or cannot be accessed from the Meter by the MDR. The MRS will operate on a transactional basis providing Register Readings (RRs) to the Processing Service (Smart) (PSS). The MRS can obtain RRs by making a physical site visit and providing a service to collect remote readings where appropriate communications are available.

3.3.6 Processing Service (Smart) (PSS)

The Processing Service (Smart) is responsible for collecting, validating and estimating Settlement Period level data from smart and non-smart Meters. It receives remotely received data from the Meter Data Retrieval (MDR) Service (Smart Meters), and meter reader reads from the Meter Reading Service (MRS) (non-smart Meters) then passes validated Settlement Period (SP) level data to the Load Shaping Service (LSS) and the Market-wide Data Service (MDS).

3.3.7 Unmetered Supplies Operator Service (UMSO)

The Unmetered Supplies Operator (UMSO) is responsible for validating the detailed unmetered supplies inventory data for equipment attached to its distribution network and providing information to other industry stakeholders. It interfaces with customers who own/operate the unmetered equipment (referred to as the Unmetered Supplies customer).

3.3.8 Unmetered Supplies Data Service (UMDS)

The Unmetered Supplies Data Service (UMDS) is responsible for calculating Settlement Period (SP) level consumption data for unmetered equipment, for example, streetlights and traffic signals.

3.3.9 Market-wide Data Service (MDS)

The Market-wide Data Service (MDS) is responsible for processing Settlement Period level data from the PSS for smart and non-smart Meters; and Advanced Retrieval and Processing Services (ARP) for Advanced Meters and UMDS for unmetered equipment. The MDS will provide data aggregations for Imbalance Settlement and other purposes (such as network charges and flexibility offerings (if required).

3.3.10 Load Shaping Service

The Load Shaping Service (LSS) is responsible for calculating energy consumption (import and export) Load Shapes for a number of defined categories of Metering Systems. The LSS uses validated actual Settlement Period (SP) level data accessed from the PSS. The PSS will then use the Load Shape data to convert RRs or daily consumption values into SP level data. The Load Shape data will also be used to estimate invalid SP level data for smart Meters and default where data is missing or unavailable.

3.3.11 Volume Allocation Service

The Volume Allocation Service (VAS) is responsible for accessing aggregated SP level data from the MDS; and SP level data (Grid Supply Point Group Takes) from the Central Data Collection Agent (CDCA). The VAS calculates SP level energy volumes for Balancing Mechanism Units (BMUs) using these two datasets. The data is processed for each Settlement Day in a scheduled run called a Volume Allocation Run (VAR). The processed BMU data is used in the Imbalance Settlement calculations. The VAS will also allocate or aggregate data for other purposes and provide a wide range of data reporting.

3.4 Supported Business Process

The initial list of business processes the DIP must support is given below. The DIP will not be responsible for orchestrating the business process but for exchanging information between the different actors. The DIP will need to be designed so that new and existing message/event channels required to support new and modified business processes can easily be added and modified.

3.4.1 List of Business Processes and Interfaces

Each message channel will support an interface. The current list of business processes and interfaces is provided below. Please note that the current business process analysis work is in progress, and this list is being refined. By the time the contract DIPSP is awarded, this list will be under configuration control as well as the details of each of the interfaces.

Business Processes	MHHSP Interface ID	Interface Name
BP001 - COS	MHHSP-IF-001	Notification of Change of Supplier
BP002 - COS	MHHSP-IF-002	Registration Update to Supplier (COS Gain)
BP09 - Change of Meter	MHHSP-IF-005	Metering Service MTD Updates to Registration
BP09 - Change of Meter	MHHSP-IF-006	Registration Service Notification of MTD Updates
BP008 - Chg En Status	MHHSP-IF-007	Change of Energisation Status Outcome
BP008 - Chg En Status	MHHSP-IF-008	Registration Service Notification of Change of Energisation Status
BP07 - Disconnection	MHHSP-IF-009	Registration Service Notification of Disconnection
BP05 - Data Processing	MHHSP-IF-013	MDS Defaults Applied
BP010	MHHSP-IF-018	Notification of Registration Data Item Changes
BP005 - Data Processing	MHHSP-IF-021	UTC Settlement Period Consumption Data
BP005 - Data Processing	MHHSP-IF-022	LSS Period Data
BP005 - Data Processing	MHHSP-IF-023	LSS Totals Data
BP004 - Data Collection	MHHSP-IF-024	Supplier Advisory Notifications
BP010 - Registration Upd's	MHHSP-IF-025	Supplier Updates to Registration
BP010 - Registration Upd's	MHHSP-IF-026	Registration Service Notification of Supplier Data Chg

BP02 - Chg. Of Serv	MHHSP-IF-31	Supplier Service Appointment Request
BP02 - Chg. Of Serv	MHHSP-IF-32	Registration Service Response to Supplier Service App Request
BP02 - Chg. Of Serv	MHHSP-IF-33	Registration Service Request for Service Appointment
BP02 - Chg. Of Serv	MHHSP-IF-34	Service Response to Appointment Request
BP02 - Chg. Of Serv	MHHSP-IF-36	Registration Service Notification of Service of Appointment & Supporting Info
BP02 - Chg. Of Serv	MHHSP-IF-37	Registration Service Notification of Service De-Appointment
BP02 - Chg. Of Serv	MHHSP-IF-38	Customer Direct Contract Advisory
BP02 - Chg. Of Serv	MHHSP-IF-39	Customer Direct Contract Advisory Response
BP02 - Chg. Of Serv	MHHSP-IF-40	Correlated MPAN Activity
BP02, BP08, BP09,	MHHSP-IF-041	Cumulative Meter Reading
BP02	MHHSP-IF-042	Customer Name Information
BP11	MHHSP-IF-045	Request Change of Segment
BP11	MHHSP-IF-046	Registration Service Notification of Change of Segment
	MHHSP-IF-47	Publish ISD
	MHHSP-IF-48	Publish Transitional MDD
	MHHSP-IF-49	Publish UMS Standing Data

4 High-Level Design

The high-level design looks at the overall MHHS TOM landscape and provides a view of the end-to-end technical architecture. The high-level design concentrates on the functionality of the DIP.

The End-to-End solution architecture will be covered in more detail later in a separate document.

4.1 End-to-End Architecture

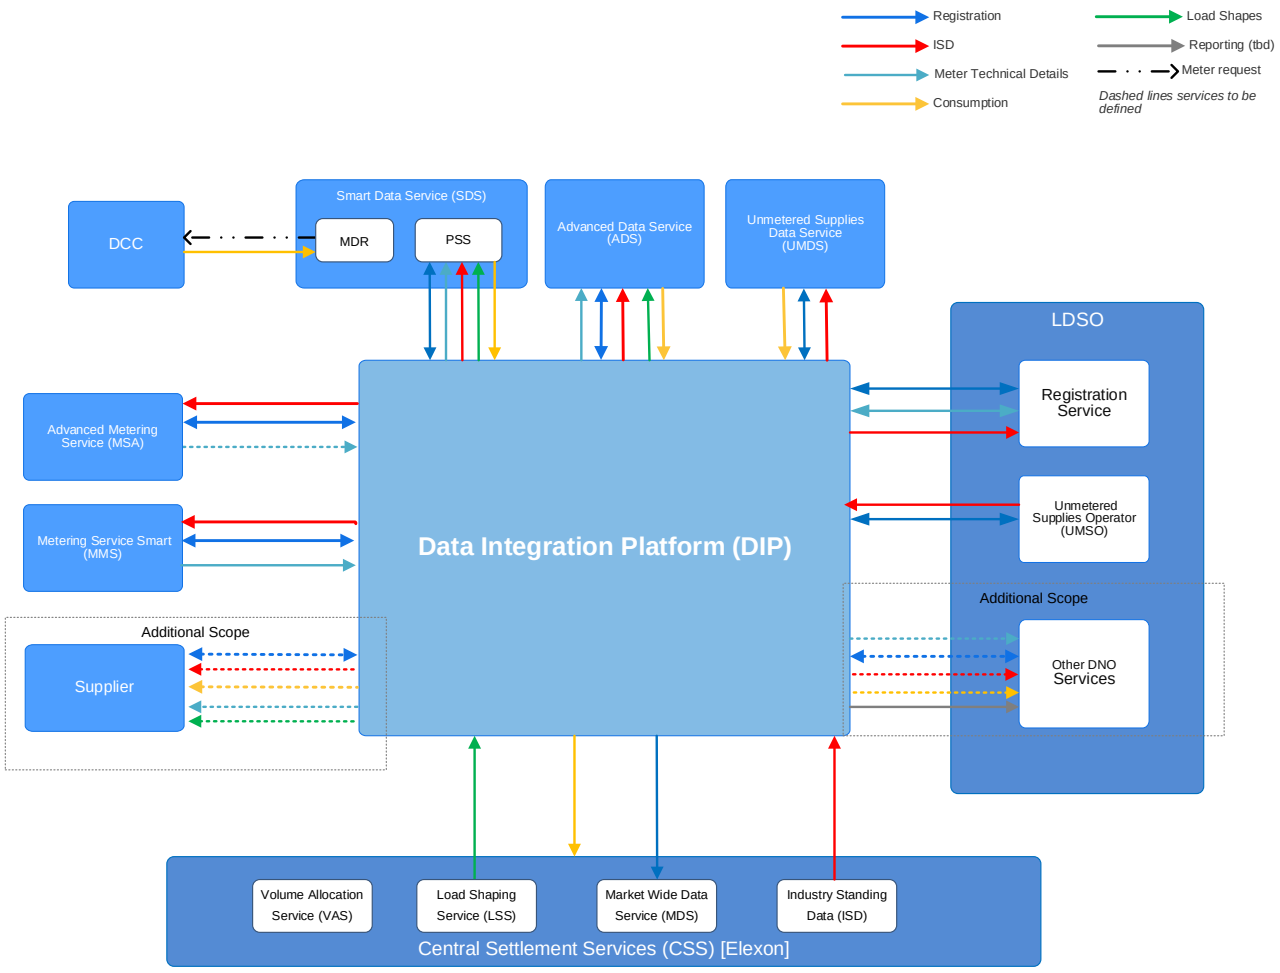

Figure 2 - End-to-End High-Level Solution Architecture

MHHS TOM landscape will be a distributed network of services and roles requiring constant data communication for operational purposes. There are approximately 27 Registration Services that will need to maintain operational data integrity and consistency across approximately 35 Data Services, 85 Metering Services, 17 DNO's and 60 Suppliers. In addition, the Data Services must provide an approximate total of 32 million daily (15 billion annual) consumption events to Central Settlement.

The AWG proposed an Event-Driven Architecture (EDA), an architectural pattern used to produce, manage, and consume data messages/events that enables the creation of a responsive/reactive, asynchronous, non-blocking/concurrent and de-coupled systems topology. This proposed EDA has evolved into the Data Integration Platform (DIP).

5 Data Integration Platform

The Data Integration Platform (DIP) is at the core of this architecture, responsible for brokering the communication between all industry participants operating under the TOM.

A service orientated view of the DIP is presented below:

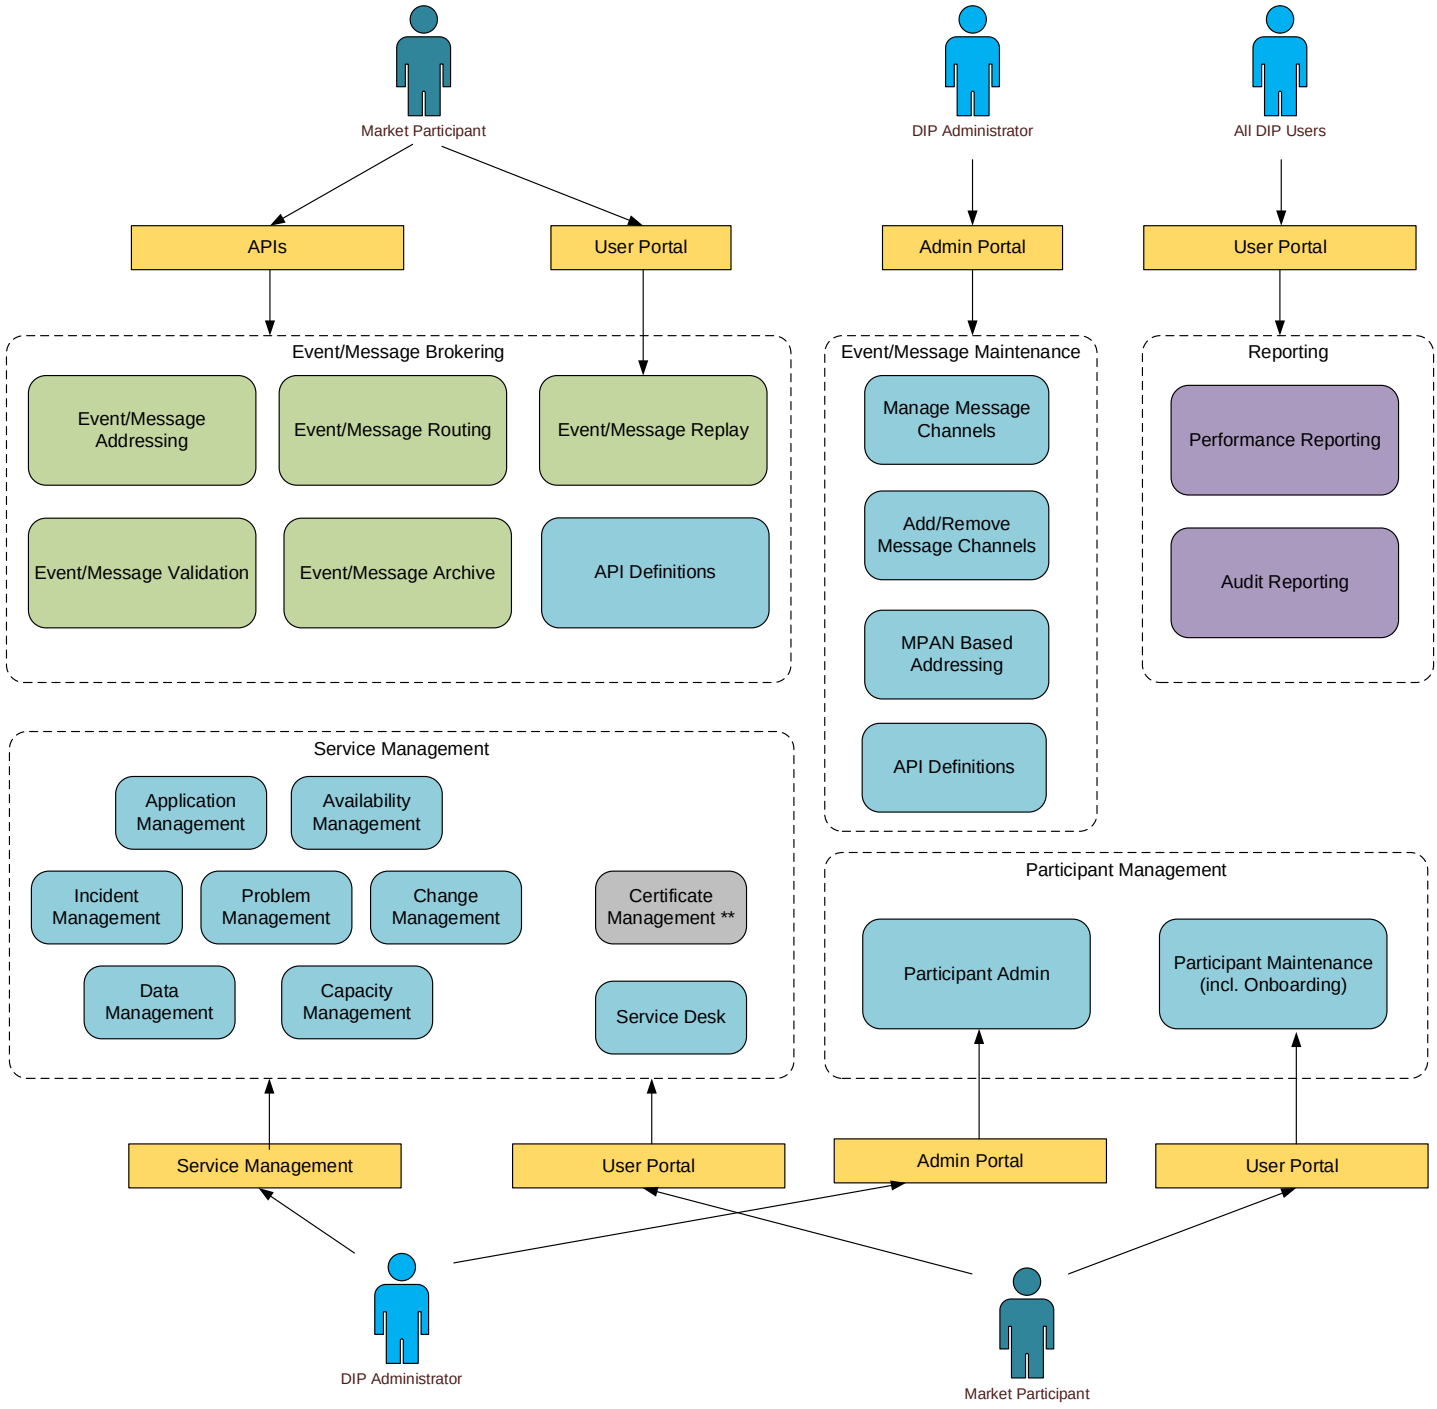

Figure 3 - Service Orientated View of the DIP

The core service areas are:

- Message Brokering

  The core component of the DIP responsible for relaying messages between Market Participants includes the following functions:

  - Message Validation, Security and Repudiation

    - Message Archiving
    - Message Replay
- Message Brokerage Management

  This includes the following functions:
    - Manage Message Channels
    - Add/Remove Message Channels
    - MPAN based addressing
- Participant Management
    - Onboarding and offboarding
    - Participant Admin including Role maintenance
- Service Management
- Reporting
    - Tracking workflows across MHHS business process

Three distinct users are recognised:

- Market Participants – individual companies, or their agents, involved in the orchestration of the business processes underpinning the TOM, covering the all the roles identified in section 3.3.1.
- DIP Administrator – the DIP service provider responsible for managing the DIP (the lower-level design should identify a dissection into sub-roles)
- Read Only Users – access the DIP for reporting purposes
- All DIP Users – Users from all the above groups

The solution requirements for the message brokering are written generically that endeavour to be agnostic to the choice of the target platform.

The working assumption is that the DIP will be a cloud-hosted, serverless/containerised, compute/messaging system that will leverage the benefits of distributed cloud architectures to achieve the resilience, availability, and scalability required.

5.1 DIP Logical Design

This section provides a logical view of the Data Integration Platform, and it does not provide a view of the MHHS TOM, just the technical architecture of the DIP.

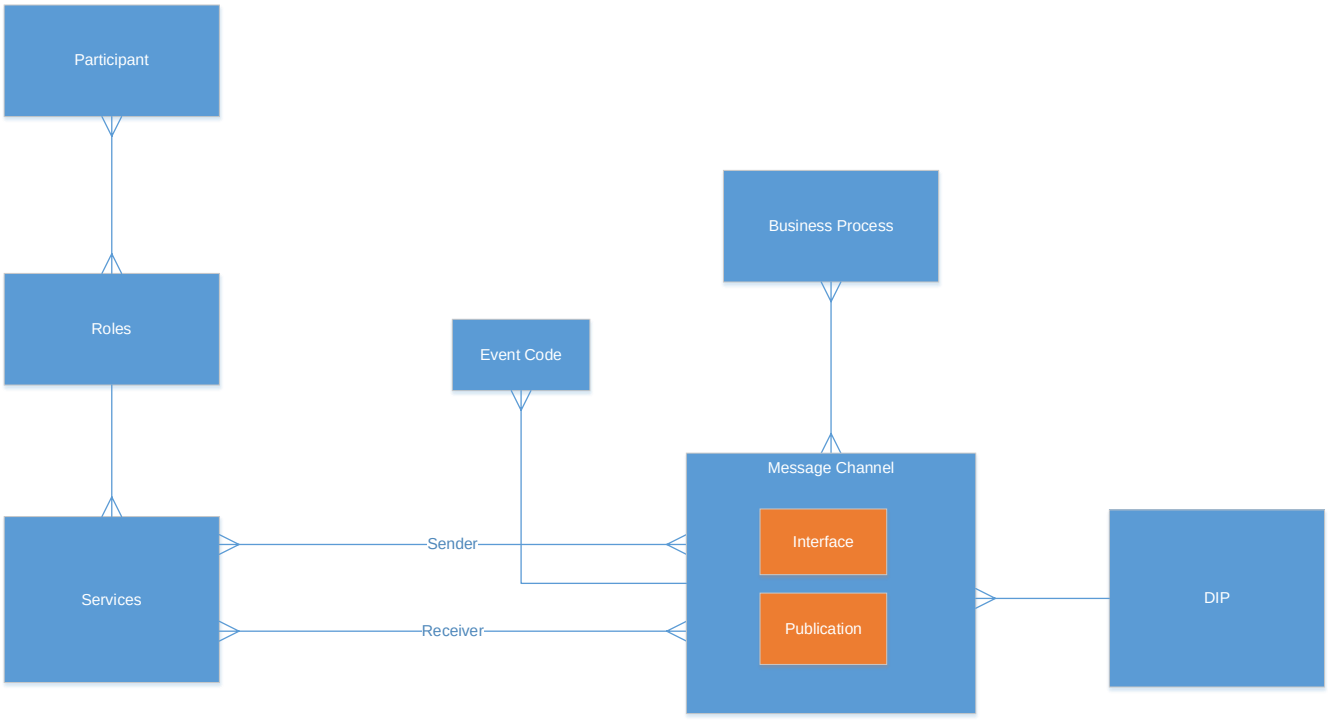

Figure 4 - DIP Logical Design

The ***DIP*** hosts multiple ***Message Channels***. Each ***Message Channel*** represents the logical and physical transfer of messages. ***The business process*** will be orchestrated by ***Message/Event channels*** where messages are exchanged between ***Market Participants***. As well as the business process and data payload, each channel will be defined by a set of ***services***, or possibly only a single service, that can send or receive the messages on that channel. Their inherent roles govern the access to each ***service*** for each ***Market Participant***.

Each ***Message Channel*** is defined by a specific ***interface*** definition, effectively the data payload, a ***publication***, the outgoing message and also an ***event code***. In addition, a set of rules will exist for each channel in terms of the permissions on the channel, data obfuscation, archiving, message addressing/routing etc. These will be set up when the channel is deployed.

5.2 Message Brokering

The message brokering function covers the exchanges of messages/events between Market Participants, as is the DIP's primary function.

There is a set of common requirements that fit across all message channels within the DIP:

5.2.1 Connection Patterns

The requirement will be to have standardised connection patterns across all services. All services will be expected to present as minimum API (inbound)/webhook (outbound) HTTPS interfaces with JSON payload and encrypted in transit with mTLS. This is the minimum requirement for all services and should not rule out the possibility of having other connections on specific services where considered appropriate, for example, the use of proprietary cloud connectors should be considered for the high-volume interfaces, i.e., half-hourly consumption data if the DIP and the source systems are located within the same cloud platform.

5.2.2 Message Privacy & Repudiation

Some of the data flowing through the DIP will include sharing personal data and fall within the remit of the UK GDPR or confidential or sensitive information. Hence, there will be a requirement to ensure the security of the data, which should be achieved by encrypting payloads and possibly headers. There is also a requirement for message repudiation, i.e., ensuring a message is sent from only the expected party.

To help simplify the message privacy classification, initial analysis has identified four different security categories, and these, along with the encryption/signing requirements for each, are:

Category	Description	Signing/Encryption requirement
1	Public Data	Digitally signed/No action
2	MPAN	Digitally Signed/Encrypted
3	MPAN + PII	Digitally Signed/Encrypted
4	MPAN + Consumption data	Digitally Signed/Encrypted

Each of the MHHS interfaces will be assigned one of these categories.

As the data are classified, four different security patterns have been identified that will ensure the integrity of the data as it transitions the DIP. With all the patterns, end-to-end security will need to be ensured with messages encrypted in transit with mTLS. The security patterns are:

Security Pattern	Description
0	No security action
1	The message is digitally signed
2	End-to-End Encryption with DIP decrypting and Re-encrypting messages

The requirement for full end-to-end encryption, which was in earlier drafts of this documents, has been dropped and hence removed from the document. Only patterns 1 & 2 are required.

Security Pattern 0 – No Security Action

No security action is taken on the data.

Security Pattern 1 – Message is digitally signed

The messages will be digitally signed, but no encryption will be applied. With this pattern, the following actions are taken during the transition of the message from sender to receiver:

- The Publisher digitally signs the message using private keys provided by MHHS PKI.
- The DIP verifies the message from the Publisher.
- The MHHS DIP signs the message using its private key for onward transmission.
- The recipient verifies the message data using the public key of the MHHS DIP.

Security Pattern 2 - End-to-End Encryption with DIP decrypting and re-encrypting messages

The following security steps are followed in this pattern to ensure data privacy:

- The Publisher encrypts the payload using certificates provided by the MHHS PKI before the message is sent to the DIP
- The DIP will decrypt the encrypted data using the private key of the Publisher's MHHS certificate
- The DIP will archive a decrypted copy of the Publisher's message using the platform's own encryption (FIPS 140-2 compliant)
- Every message forwarded to a recipient will be sent encrypted using the recipient's MHHS public certificate.

- The recipient decrypts the message data using the private key of the MHHS Certificate.

With this pattern, Market Participants will never interact with another Market Participant to exchange certificate(s) / keys.

- All certificates and keys will be managed individually between the MHHS DIP and the Market Participant.
- There will be no shared keys using this pattern.

The proposal is for the defined categories to adopt the preferred security patterns:

Category	Description	Payload Configuration	Preferred option
1	Public Data	No action	Security Pattern 1
2	MPAN	Encrypted	Security pattern 2
3	MPAN + PII	Encrypted	Security pattern 2
4	MPAN + Consumption data	Encrypted	Security pattern 2

There is a requirement for the potential to define new security categories and patterns within the DIP. The requirements of currently undefined message flows could easily be added in the context of these new categories and patterns.

5.3 Event/Message Channel Maintenance

The first draft of the end-to-end business process described briefly in section 3 has identified the need for the DIP to initially support 12 business processes and 24 message/event channels. The expectation is that the number of business processes and message channels will increase during the course of the programme and in the future to support other industry initiatives. In addition, a number of the reporting requirements are still to be defined.

The DIP Administrator is responsible for undertaking event/message channel maintenance.

A number of message/event channels will exchange messages/events between participants. A specific design pattern will implement each channel.

The design should be "templated" so that multiple message channels can adopt the same messaging pattern, albeit with different configuration information.

Once designed and developed, the use of templates will reduce system development times, thereby reducing the time and costs for the introduction of new channels,

The expectation is that the following activities can be automated and be under the control of the DIP administrator via an admin portal:

a) Add new Message/Event Channel;

b) Remove Message/Event Channel;

c) Manage Message/Event Channel:

   i. Add/remove Sender to Message/Event Channel

   ii. Add/remove Recipient to Message/Event Channel

d) Add New participant to all Message/Event Channels (based on roles);

e) Remove the Participant from all Message/Event Channels.

The DIP is a platform for the future and should be designed such that additional message channels can easily be added without requiring any system outage. This would make the system easier to deliver and maintain, and where an outage to a specific service for maintenance will not mean an outage to the whole system.

5.4 Message/Event Channel Instances

At present, two different distinct message patterns: Message Pattern 'A' and Message Pattern 'B', have been identified to implement the 24 message/event channels; the requirements are described below. It is expected that the reporting requirements to be defined will adopt Message Pattern ‘B’.

The proposed patterns only implement the initial set of business processes identified by the MHHS TOM. However, the DIP will need to support new messaging patterns as different industry initiatives may require different message/event exchange patterns.

The DIP will need to support different integration patterns with the following characteristics:

- Synchronous (pub/sub; request/reply) and asynchronous (push/pull; stream).
- One-to-many, many-to-one and one-to-one (pub/sub) messaging patterns.
- Exactly Once Processing
    - Idempotent Writes - a message/event is written once, and only once, by a sender.
    - Idempotent Read - a message/event is delivered only once to a receiver (and/or each message/event will need to be uniquely identified, and receiver processes can detect duplicates and gaps in message sequences)
- The requirement for FIFO has been dropped (21/03/2021) ~~First-In-First-Out (FIFO) Delivery - the order in which messages are sent and received is strictly preserved.~~
- Market Participants receiving messages will be able to process messages/events at their own speed (“message throttling”) independently to other Market Participants.
- Durable subscriptions - when a Market Participant is off-line due to a planned or unplanned outage, the DIP must retain any messages/events for the Market Participant for up to 14 days for standard consumption. After this time, the Market Participants may use the Event Replay facility (see the section 5.5.11). However, this is not considered to be standard consumption.
- Message exchange traceability – ensure message delivery with repudiation for both receipt and delivery built-in.
- Message archive and replay

A mixed product platform solution is a possibility. Some channels are implemented by one template on one messaging/event platform and other channels on a different platform by a different pattern. Even though there may be a mixed platform architecture, the requirement to have all message/event channels within the same cloud tenancy still endures.

At present, two distinct patterns are recognised, Message Pattern A and Message Pattern B, and each pattern supports a different set of capabilities based on the specific requirements of the supported business process:

5.5 Message/Event Channel – Pattern ‘A’

Message Pattern ‘A’ is a generalised publish and subscribe pattern that allows both many-to-many, one-to-many and many-to-one message exchange and has targeted message recipients.

5.5.1 Message Traceability

Messages exchanged over Message Pattern 'A' requires a technical acknowledgement for the receipt and the subsequent sending of the message/event by the DIP, i.e. the API HTTP return codes and response bodies to/from Participants. The requirement to log all API activity is deemed sufficient to meet any traceability/non-repudiation requirement and provide the acknowledgement back to the message sender.

Audit reporting facilities will be made available to allow participants to check on the progression of individual message exchanges.

5.5.2 Message Validation

The DIP will only undertake 'schema' validation rather than content validation. The following principles will be adopted:

- The message header is syntactically correct, and all the necessary data required for the DIP to process the message is present.
- Message payload ensures the message schema, mandatory parameters, and data types are correct.

Messages failing validation will be reported back to the Market Participant that sent the message.

The message recipients carry out detailed validation of the message payload containing the pertinent business data.

5.5.3 Message/Event Routing

Message Pattern 'A' has a requirement for three types of message addressing/routing:

- Targeted (primary) Routing – where the message sender knows the intended recipient, i.e., Participant, and the DIP will use the routing information in the message header.
- MPAN Based Lookup (secondary) Routing – the DIP will need to route a message based on the MPAN within the message. The requirements around this service are described below in section 5.5.14.
- Always – where the DIP will always send a message to either a named participant or all participants assigned to a designated role (and the role is assigned to the message channel).

The DIP can apply a single or all types of routing to a single message channel.

5.5.4 Message Formats

Messages will have three primary components and the principles of each of the components:

1. The message header provides uniqueness for the data exchange and should enable duplicate transports of the same data flow to be identified and allow for messages to be correctly routed. The following data will be included (i.e. set by the message sender):
    a. Message Channel Id/ Interface ID
    b. Message sent date/time
    c. Sender Participant ID
    d. Sender Participant Role – capacity in which the message was sent
    e. Sender Unique Reference – Sender participant Unique Reference used by the Participant to uniquely identify a message, for example, a date with a long sequence number. Provides a key that enables idempotent data exchange, i.e. if a recipient receives two messages with the same Sender Unique Reference, then the second message can be considered a duplicate.
    f. Sender Envelope Reference – provides a facility for a Sender to logically group a set of messages.
    g. Recipient Participant Id – used to route messages to specific parties (Primary Addressing) (optional)
    h. Routing MPAN (Primary Addressing) (optional)
2. The payload provides the required data items for a business process to be completed. Payloads may need to be encrypted, dependent on the classification of the data in transport.
3. On receipt of an event/message, the DIP will create a transaction wrapper with the following details:
    a. Transaction Id - uniquely identify the event/message within the DIP
    b. Timestamp the receipt of the message
    c. Environment indicator – confirming the receiving environment (Production, Pre-production, etc.)
    d. Replay Indicator - to indicate this is not a replayed message
    e. Workflow Correlation ID to the message (optional – configurable on each channel) (see section 5.5.5).

    These attributes will be appended to the message so that the receiver gets this information. This information is also written to the message archive.

The requirement is for all message payloads to be JSON.

5.5.5 Workflow Correlation Id

For some MHHS business processes, there is a requirement for each specific workflow instance to be uniquely identified and then for all messages in that workflow to reference the unique ID. This will be the Workflow Correlation ID. (A similar requirement exists in Faster Switching). This will operate as follows:

1. The Market Participant initiating the workflow will send the first message in the sequence to the DIP.
2. The DIP will recognise events/messages on certain channels require a Workflow Correlation ID to be generated.
3. The DIP will generate the Workflow Correlation ID, written to the message for onward processing and returned to the message sender as part of the HTTP response body.

Section 5.5.13 describes the generation of the Transaction ID in the broader context of the end-to-end choreography.

5.5.6 Message/Event Obfuscation

The requirement is for the DIP to obfuscate the contents of a message based on the recipient's role within a particular message channel. This requirement is implemented so that the message sender does not have to send multiple messages to different recipients containing a subset of data; instead, a single message is sent to multiple recipients. The DIP will remove the required fields, and the default position will not apply a filter.

5.5.7 Message/Event Archiving

There is a requirement for messages to be archived, and each message channel will have its own archiving requirements. The message archive will be used as the source for the message/event replay capability and event/message audit reporting.

5.5.8 Dead Letter Handling

A facility to move ‘stale’ messages, i.e. those that have not been processed within a certain time (configured on each channel, defaulted to 14 days) by the intended recipient, is required to a dead letter queue (DLQ). Messages on the DLQ will be reported back to the message's sender.

5.5.9 Bad Message/Event Handling

An error reporting capability to return errors to the message sender is required. For example, if the DIP cannot read the message header, including any routing information, then this will need to be communicated back to the message sender. In addition, any format issues with the message that the recipient encounters will also need to be relayed back to the Sender.

This capability is separate from the main message flows that orchestrate the business process and should not be used for business process orchestration. It only reports when a processing thread terminates due to data or processing problems.

5.5.10 Connection Patterns

The requirement is to have standardised connection patterns across all services. Each messaging/event channel will primarily access via a pair of API HTTP(s) (incoming) / webhook (outgoing) endpoints over mTLS.

Although these will be the primary connection pattern for accessing the DIP, this does not preclude using other connection mechanisms, for example, proprietary cloud connectors.

Inbound API

There will be a single API endpoint for all Pattern 'A' message channels; the API will receive a message comprising a standard header and a mixed payload (encrypted). Each connection will be able to transmit a number of different messages during each connection (limits will need to be specified during the design phase).

Each connection will result in a HTTP return code that will indicate the success or otherwise of the complete transaction.

Response	Meaning
202 – Created	All messages created
207 – Multi status	Some messages created; see response body for details
4xx	Bad request – no messages sent

The response body of the HTTP call will deliver the Sender a transaction ID and optionally a correlation ID against the Sender's Unique Reference for each message.

Outbound Webhook

There will also be a single webhook for all Pattern 'A' message channels. The expectation is that each receiving Participant will be served from an outgoing message/event queue, and messages will be queued in the order received. The response body of the HTTP call will record the transaction ID against the success/failure of the call.

Some participants may need to logically split the data received from the DIP before sending it, as their applications may be hosted in different locations. Hence, participants will be able to define a number of logical queues/endpoints in order to group the message/event channels they choose. The Participant will be able to assign each message/event channel to a specific endpoint. These logical queues/endpoints may relate to individual services defined in the TOM.

5.5.11 Message/Event Replay Facility

A basic query facility is required for message/event replay for each message channel for all participants. In the production environment, a request for message/event replay can be used to assist participants' downstream systems is suffering from an unrecoverable loss of data. The request for replay will consist of the following:

- Message Channel ID/Event Code
- Start Message Transaction ID – the message from which to start the replay sequence from, or
- Date/time from – transaction time from which the first message needs to be replayed.
- End Message Transaction ID (optional) – the last message from which to start the replay sequence from, or
- Date/time to (optional) – transaction time message to be replayed

If End Message Transaction ID and Date/time to are not specified, the request will send all archived messages up to real-time.

If Start Message Transaction ID and Date/time from are both specified, then the earliest message of the two will be used, and End Message Transaction ID and Date/time to are specified, then the latest message will be used

There will be two methods by which the Event Replay facility can be initiated:

- Through the UI interface on the DIP
- Through an API Call

As well as being delivered via a separate API, the event replay will deliver the message under a new transaction wrapper with a Transaction ID and Correlation ID (with a reply prefix) so that the event/message are identified as a replayed, and hence downstream processes are not triggered. The initial message transaction wrapper will also be sent.

Only messages the Participant is entitled to view will be sent.

5.5.12 Worked Example

An example of a Messaging Pattern ‘A’ Template using a message queue-based architecture:

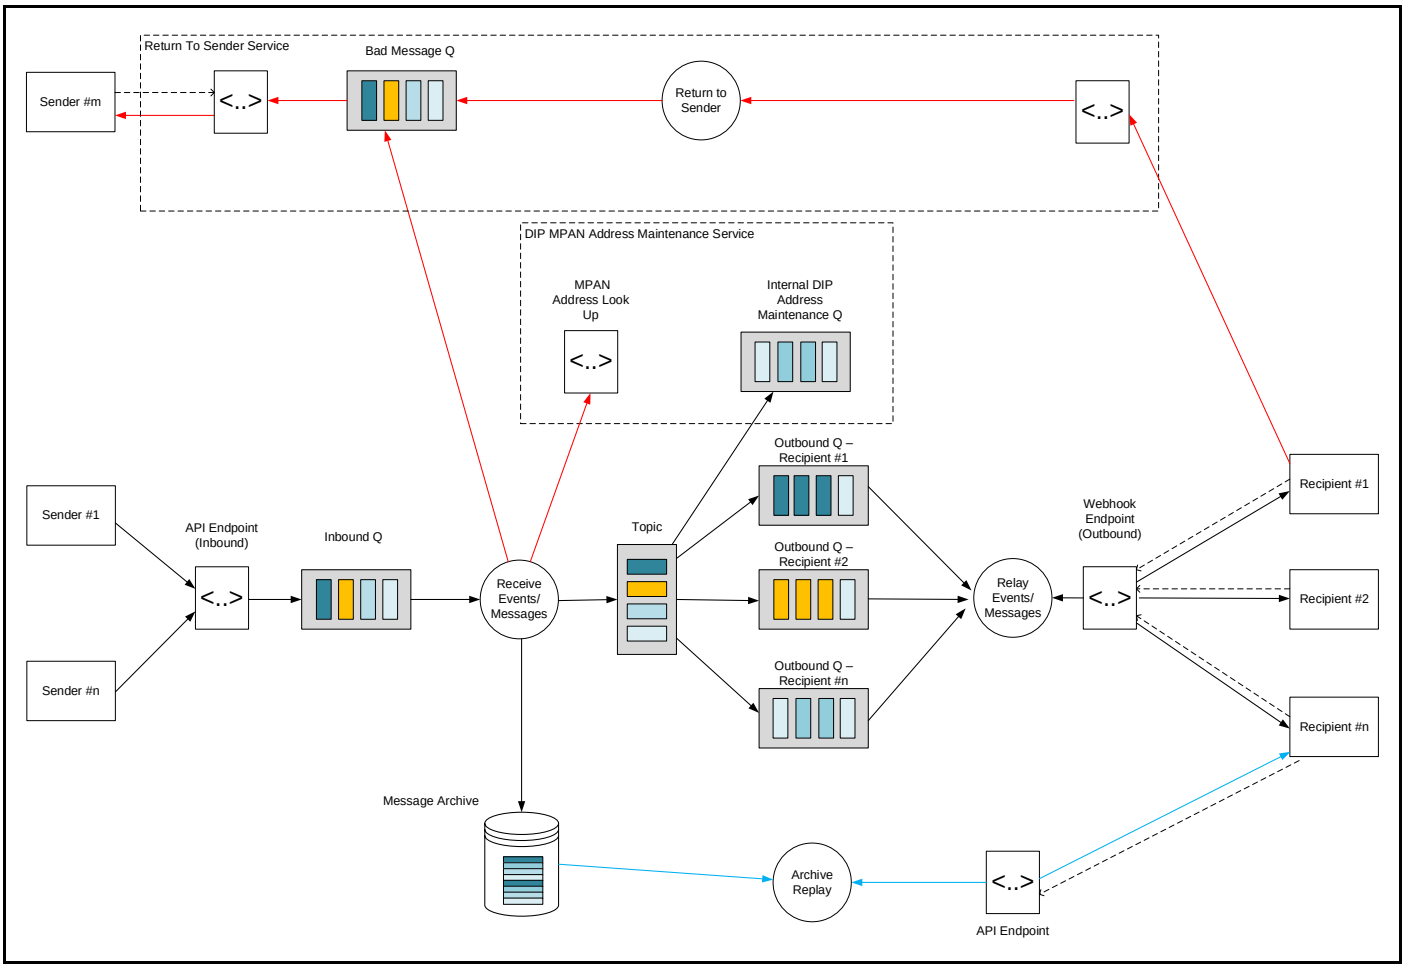

Figure 5 -Message Pattern 'A' worked example

The actual details in the data described below are only shown for representational purposes and **not** for accuracy.

Message Channel Id	#909
Interface Id	INTF-909
Message Flow	MHHS- Notification of Data Service Appointment
Senders	MSS, MSA
Targeted Receivers	SMRS
Default Receivers	SUP

Payload

			SMRS	SUP
DI-063	MPAN Core	Integer 13		
DI-020	Contract Reference Metering Service	String 10		
DI-087	Supplier ID	String 4		

DI-062	Metering Service ID	String 4		X
DI-061	Metering Service Effective From Date	Date/time		X
DI-010	Appointment Code	String 1		X

When the message channel is generated, access to the inbound endpoint is granted to participants with the MSS and MSA roles. Access to the outbound endpoint is granted to participants with the SMRS and SUP roles. The inbound API specification is extended to include the new payload (if the payload is encrypted, nothing will need to be done).

Messages to the SUP role participants are missing the three last items: DI-062, DI-061 & DI-010. Messages sent to the SUP roles contain all items.

In this example, a targeted registration system identified in the message header is sent the full contents of the message, copies of the message are sent to all participants with the SUP role. The DIP implements this default routing, i.e., the Sender does not need to undertake the addressing; if addressing is required to be undertaken by MPAN, then this will use the MPAN Address Lookup Service, which is described below.

'Bad messages' are returned to the Sender via the 'Return to Sender' service.

Dead-letter queue handling is not shown in the example.

5.5.13 Message Choreography

The following section describes a few of the expected message exchanges required between participants and the DIP. The list is only meant to be a representative sample to provide bidders of the type of interactions to expect and is by no means a complete list. An End-to-End Solution Architecture and an Operational Choreography Document will be produced in the next phase of the MHHS programme that will provide information on all possible exchanges.

The following diagrams describe the expected message exchange between the participant services and the DIP for Pattern 'A':

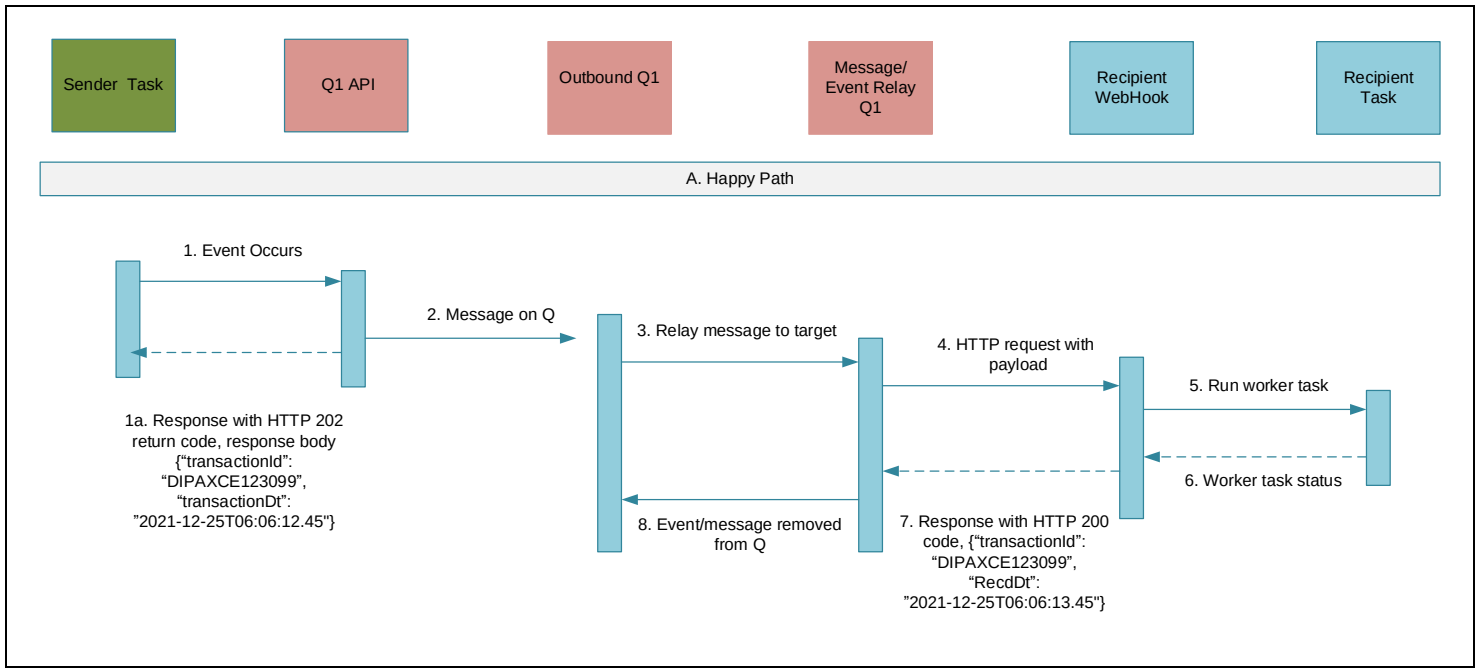

Figure 6 - Pattern A - Happy Path

A. Happy Path – the diagram above shows the normal sequence of actions in a single message exchange through the DIP.

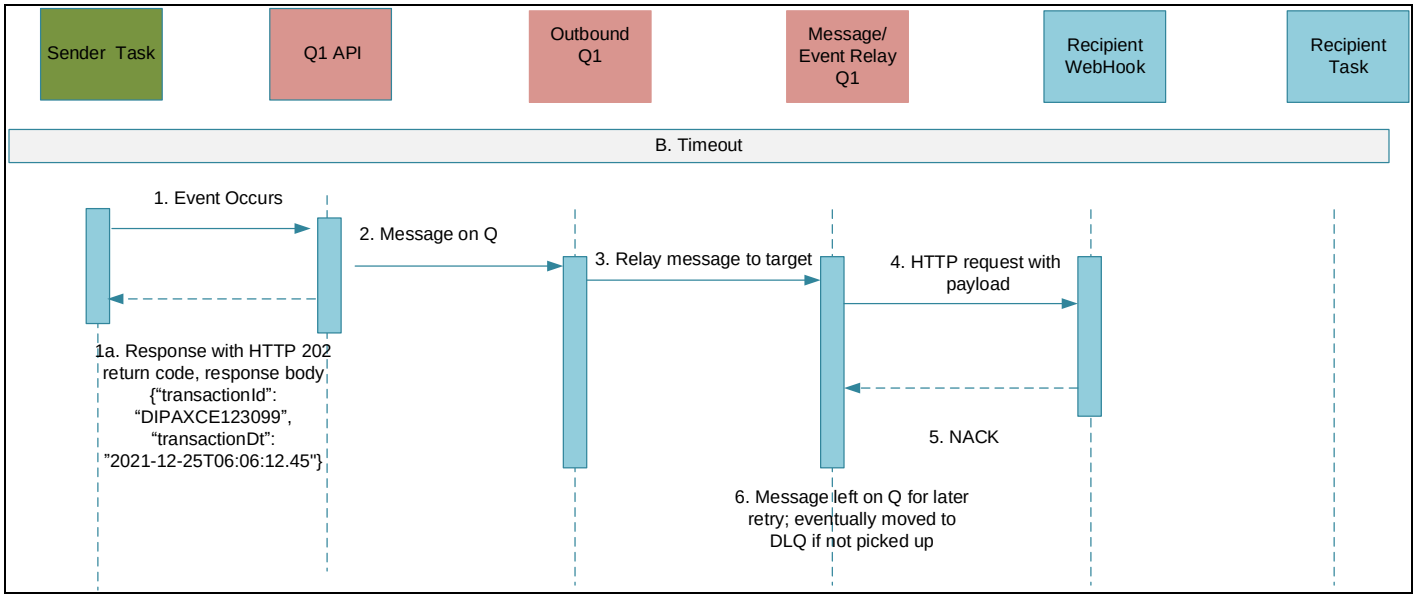

Figure 7 - Pattern A -Timeout

B. Timeout – above shows the case where there has been no response from the recipient webhook at the call back times out. The message is retained in the queue for later processing. The DIP will have the logic to attempt to resend the message after a timeout period.

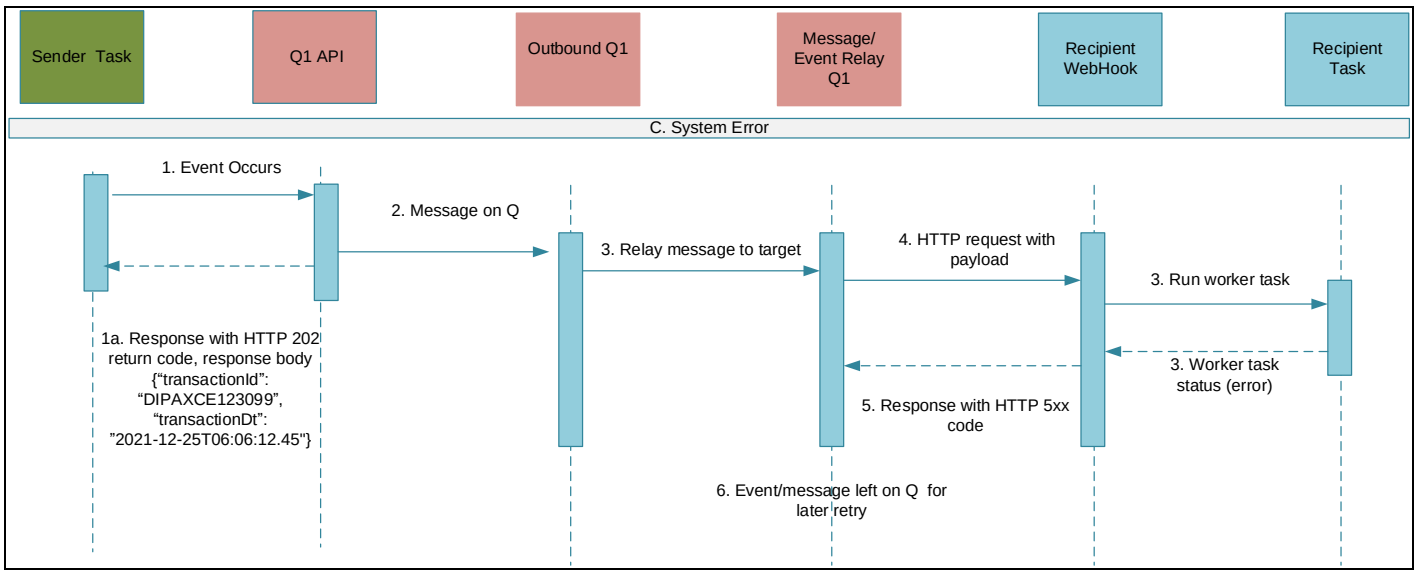

Figure 8 -Pattern A- Recipient Error

C. Recipient System Error – above shows the case where a response from the recipient webhook indicates a system error. The message is retained in the queue for later processing. The DIP will have the logic to attempt to resend the message after a timeout period.

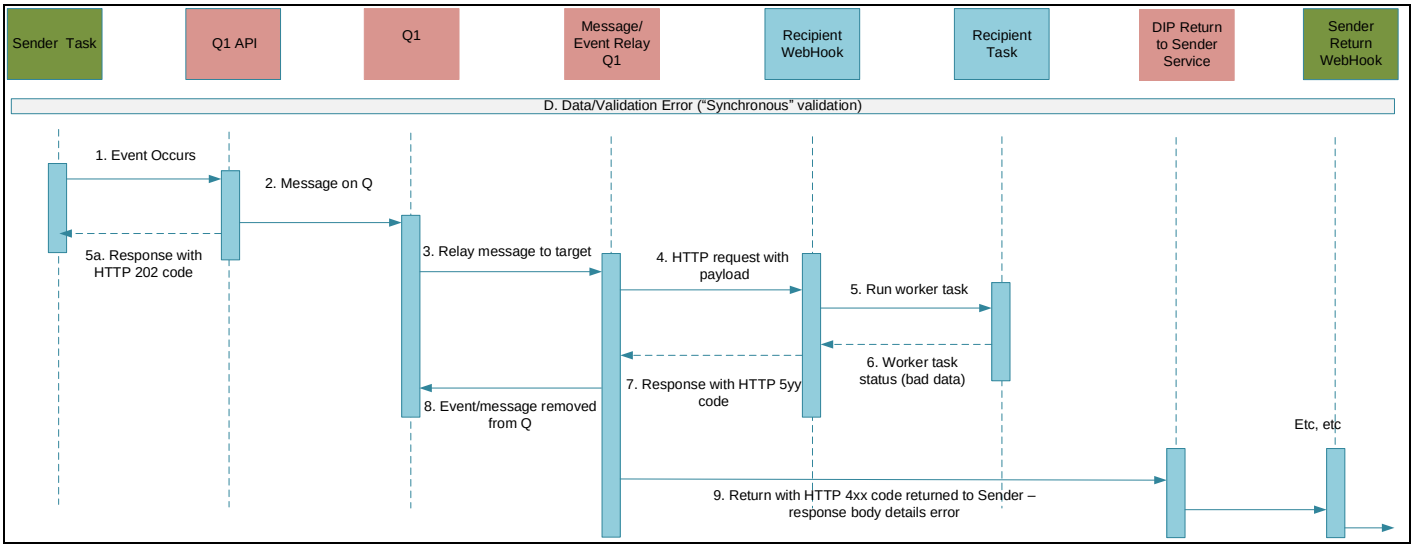

Figure 9 - Recipient Data Error - synchronous reporting

D. Recipient Synchronous Error Condition– above shows the case where the recipient’s system rejects the message via the webhook call. In this scenario, the webhook reports a mixed/error return, and the 'Return to Sender' service relays the error to the originating party.

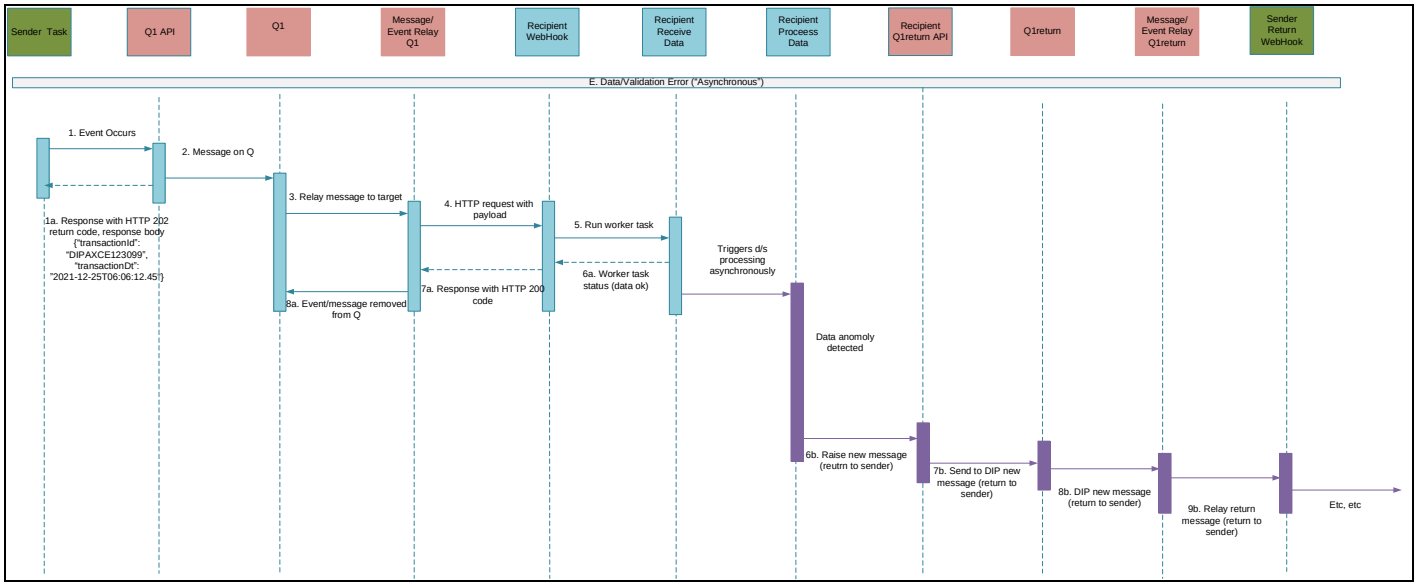

Figure 10 - E - Recipient Data Error - Asynchronous reporting

E. Recipient Asynchronous Error condition– above shows the case where the recipient’s system can consume the data; however, there is an inconsistency with the data, but it is not reported on the initial call back. In this scenario, the webhook reports a normal success return. A second message thread is then initiated to report the data inconsistency to the originating party via the DIP ‘Return to Sender’ service.

5.5.14 MPAN Lookup Addressing Service

To facilitate the addressing and routing of messages, a MPAN addressing service is required. The MPAN lookup addressing service is responsible for maintaining a routing table that provides the messaging services with an instant

address lookup for incoming messages based on MPAN. Each message/event channel will have a set of distinct roles that each message needs to be addressed.

The lookup table will be based on MPAN, Message Channel, intended recipient roles and date/time.

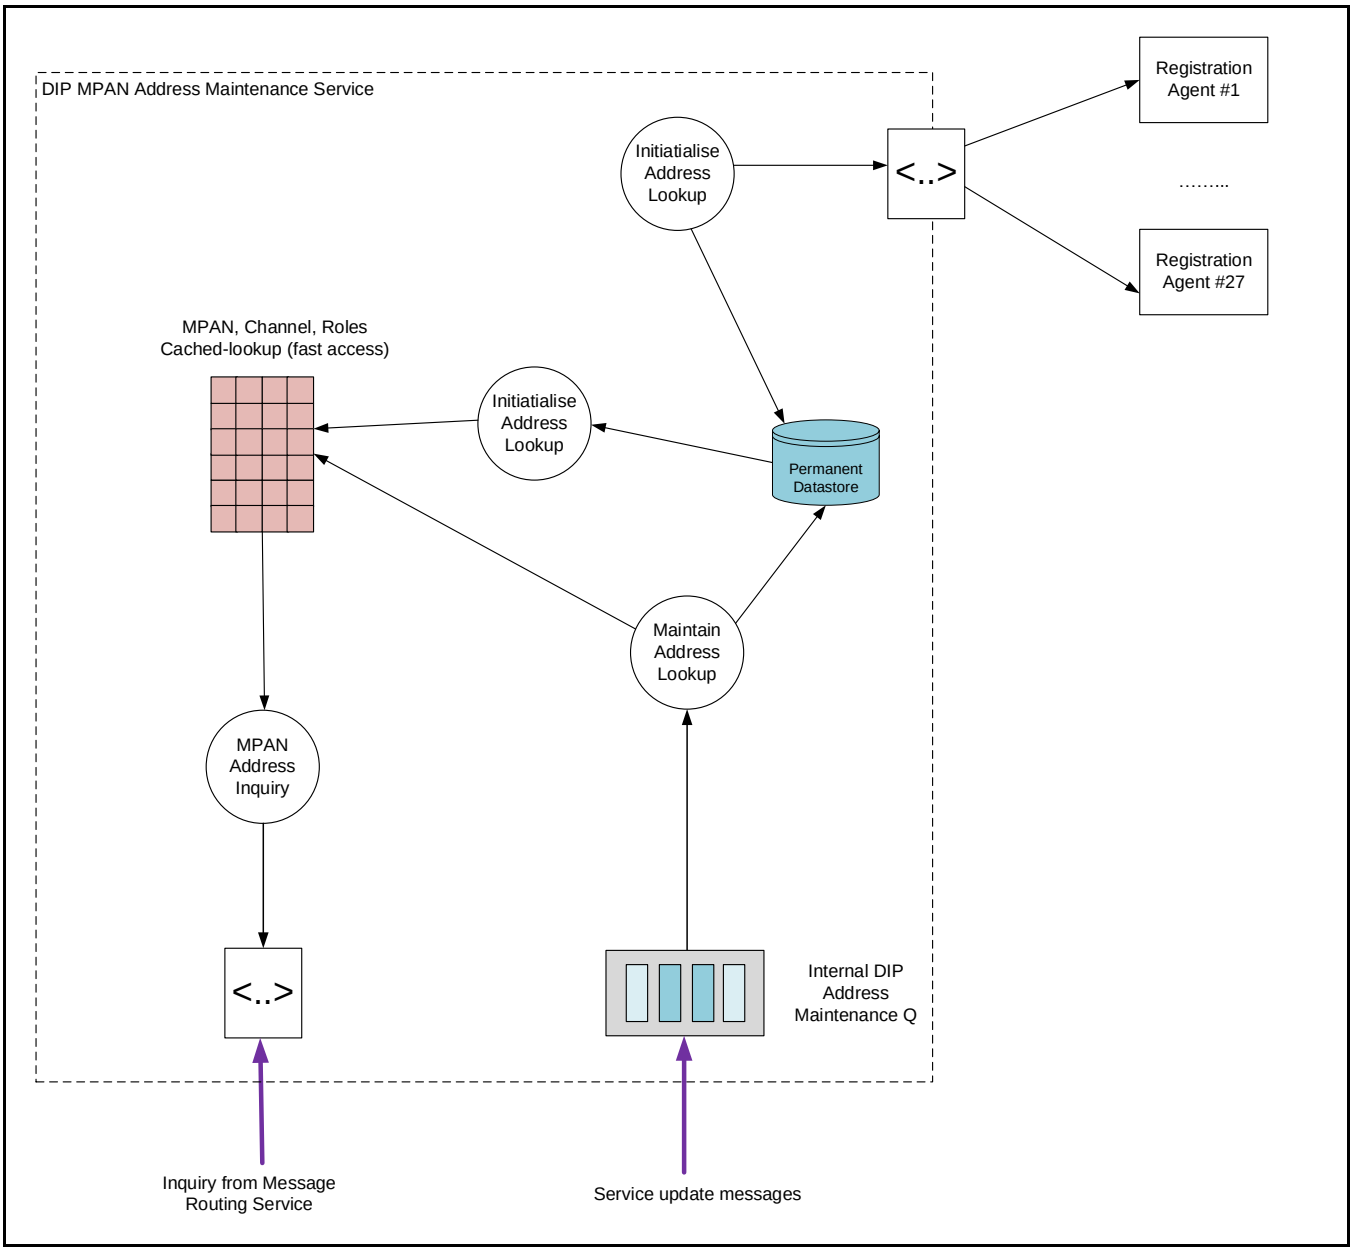

Figure 11 - MPAN Address Maintenance Service

Requirements for this process:

- Each Registration Service must present an API that the DIP can query to initialise its permanent data store.
- The working assumption is that the data will need to be kept in two places: a permanent and a data cache. The corresponding use of these is self-evident – the cache is populated from the permanent store and provides a fast lookup; the permanent store provides the permanent store.
- The MPAN inquiry service will receive an input MPAN and message channel and return a list of downstream recipients.
- After initialisation, the DIP is responsible for maintaining the Address Lookup data in real time. All message channels pertinent to the registration data flow will send a message to an internal queue within the DIP. This queue provides the Maintain Address Lookup function information: both permanent and cache are updated with the new details.
- There is a requirement for the lookup service to cope with a change of details and be retained for the historical lookup. 'Old' messages may arrive that need to be sent to a previously responsible party; for example, late consumption data (INTF-011) would need to be sent to the set of secondary parties that are pertinent for the

specific day for the meter read. In this scenario, the header would define an applicable date/time field used for addressing.

5.6 Message/Event Channel - Pattern 'B'

Pattern B is a generalised publish and subscribe pattern with a one-to-many message exchange. The pattern allows for a single message with a large payload (+10Mb) to be distributed to multiple participants. As the payload is large, and there is no requirement for dedicated routing, a notification message is sent out to registered participants with a locator for the message payload. Also, the design allows for multiple payloads to be sent together, i.e., it supports the output from a batch type program. The recipient message contains the URI for the payload(s) and the other meta-data, i.e., date/time, version. The payloads are separated from the messages as they are not deleted once the message has been read; instead, they are available for later consumption.

Market Participants will also query the data store to see current and past payloads. This will also enable anonymous access to the data store, and a facility to download current and past payloads is also available.

5.6.1 Archiving

There is no requirement to archive the messages sent on this pattern, and the message payload store will act as an online archive. Each separate channel employing this pattern will have specific retention requirements for the payload data.

5.6.2 Message Security

At present, there are no requirements for encrypting the payload data for any of the channels employing this pattern as it is primarily meant for those distributing publicly available data.

5.6.3 Connection Patterns

As per Pattern A, the requirement will be to have a single standardised connection pattern across all services. Therefore, each messaging/event channel will primarily be accessed via a pair of API HTTP(s)/webhook endpoints. The method for securing the service endpoints is covered below.

5.6.4 Worked Example

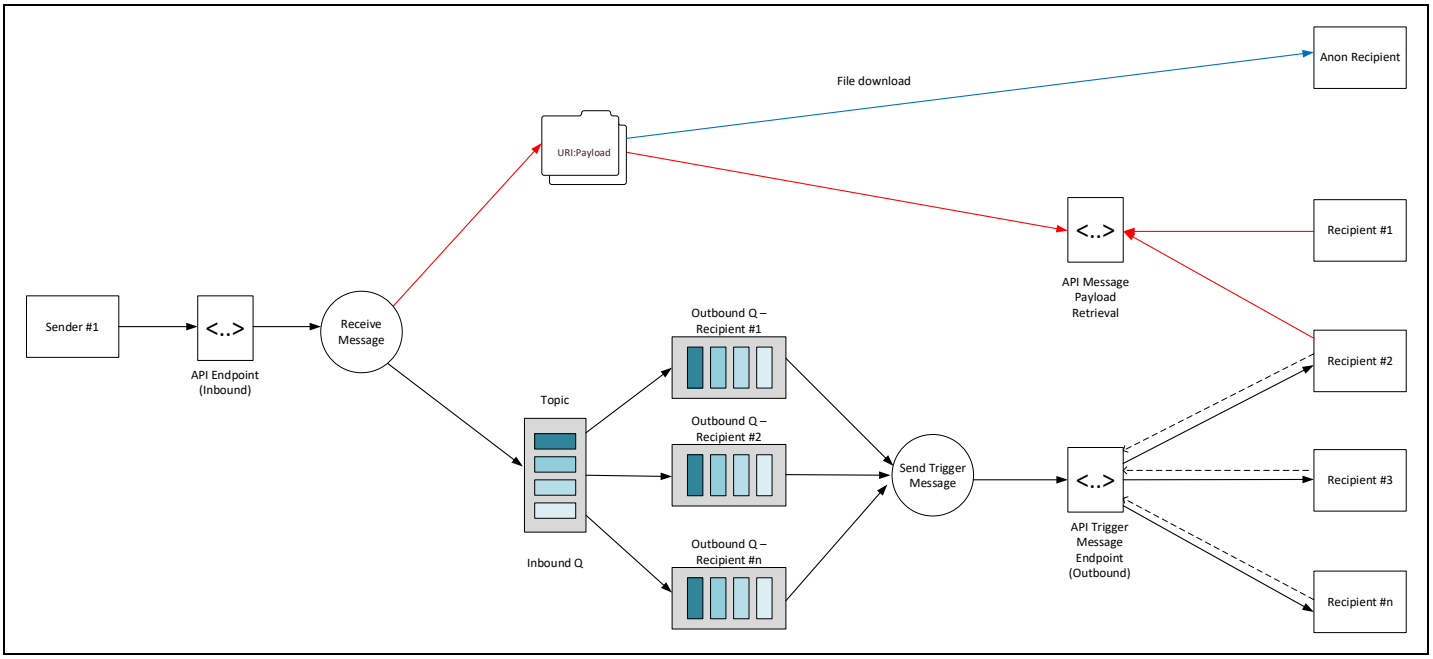

Figure 12 – Message/Event Pattern 'B'

1. Sender#1 produces a set of data, and they are written to a known location (URI: payload)
2. Sender#1 publishes a trigger message to the DIP providing the new payload(s) details (this could be delivered via the Pattern 'A' method)
3. The DIP will relay the message to all registered recipients.
4. Recipients determine the contents of the new payloads by interrogating the message and can access the payload by either:
    a. An API – the API will be able to provide simple filter capabilities so that only selected elements of the payload are retrieved.
    b. A file download.
5. Anonymous recipients can access and download data from the URI: payload at any time.

5.6.5 Message Choreography

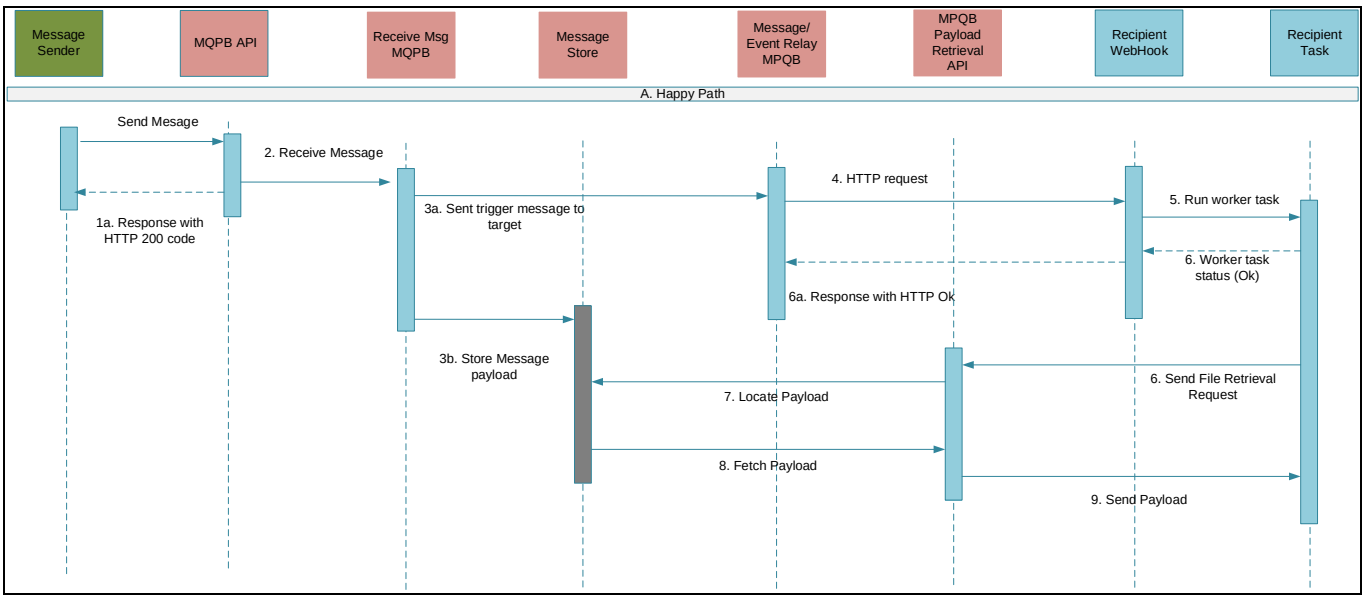

Figure 13 - Message Pattern 'B' Orchestration

5.7 Participant Management

5.7.1 Participant On-boarding

When a participant is enrolled with the DIP, they will need to be automatically assigned the roles they are qualified for when undertaking their BSC Registration.[1]The working assumption is that the details of qualified participants will be sourced from Elexon central systems, which is described in the onboarding process. Each role will have a number of services that a participant can undertake, and a participant can have multiple roles and hence access to multiple services.

Participants may want an agent to undertake the responsibility of interacting with the DIP on their behalf; in this scenario, the responsible Participant would need to transfer the necessary roles to the appointed agent.

The current working assumption is that an API will be provided to the DIP to retrieve the initial onboarding information for the Elexon systems. The expectation is that this information will include Market Participant ID (MPID) and associated market role codes for the Market Participant. As well as the initial set of data, any changes to existing participants, where roles are added or revoked, or participants are fully rescinded, will be published over the API.

The market roles inherited from the Elexon feed will determine the basis of the RBAC (Role Based Access Control) for a Market Participant and the message channels they can subscribe to. In addition to this base set of roles, the Participant will be allowed to sign up for optional roles that will give them access to other channels that allow optional access.

When a new participant is onboarded into the DIP, as part of this process, they need to be automatically added to the message/event channels applicable for their role(s). Similarly, when a participant is off-boarded, the Participant is removed from the current message channels they are assigned.

In addition, the onboarding process will need to cater for third-party agents who provide messaging services for participants. The agents will interact with the DIP on behalf of the participants they represent. A single agent can represent more than one market participant. A process allowing market participants to assign responsibility to agents needs to be included in the onboarding process.

5.7.2 Role Management

[1]The onboarding requirements may eventually depend on whether BSCCo or RECCo are deemed the ESO to the Data Integration Platform. The assumption here is that it will be BSCCo.

Facilities will exist for both the Participant and the DIP Administrator to undertake role management for participants that have already on-boarded. At present is not yet known whether the request for roles changes will be managed within the DIP or with the

5.8 API Management

The requirement is that the solution provider will provide an Open API 3.0 Specification (Swagger) platform to allow industry participants to visualise and interact with the DIP API's resources without having any of the implementation logic in place. (This should be automatically generated from the OpenAPI Specification).

5.8.1 API activity archiving

There is a requirement for all API activity to be logged and made available for reporting. The activity log must include the following characteristics:

- All API requests, including User Id, URL and response time, along with HTTP status codes & response bodies, must be logged
- Each API request needs to be uniquely identifiable
- Use a standard open framework
- Use a structured framework and format (JSON)
- Any Personal Identifiable Information (PII)/sensitive data must be obfuscated in the activity log

Audit logs should be stored in a "read-only" form that cannot be accidentally or maliciously purged.

5.9 Data Management

The API definitions alone will not provide sufficient information for DIP Participants to develop their interfaces for successful message exchange with the DIP, as many payloads are encrypted. Hence, there is a requirement to stand up a portal for participant access that describes the following:

- Data Dictionary - defines all the individual items used in the interfaces. Provides a code and data type.
- Interface definition – provides both a logical and physical view of all interfaces. The logical interface defines the interface in terms of the items from the data dictionary, while the physical interface defines the JSON.
- High-level business flows – a description of the workflow and the ordering of the publications required for each business process. This information helps understanding of the DIP audit reports.

Whilst the programme is inflight, it will provide this information; however, an enduring repository is required once the programme completes.

5.10 Physical Characteristics

This section details the physical dimensions and performance requirements.

5.10.1 Dimensions

The DIP will need to accommodate the following estimates of the number of message/event channels and DIP users:

Quantity	Initial Numbers	Annual Increase
Number of Message Channels	50	5-10
Number of Senders	200	20

Number of Recipients	200	20

5.10.2 Data Volumes

Data Entity	Interface	Target	Annual Volumes		Annual Data Size	
			Typical	Maximum	Typical	Maximum
Registration	Appointment	Metering Service	6,202,000	31,020,000	1.9 GB	9.7 GB
		Data Service	6,202,000	31,020,000	1.9 GB	9.7 GB
	De-Appointment	Metering Service	6,202,000	31,020,000	402 MB	2 GB
		Data Service	6,202,000	31,020,000	402 MB	2GB
	Accept/Reject Appointment	Registration Service	12,404,000	62,040,000	800 MB	3.8 GB
	Updates	Metering Service	1,551,000	31,020,000	365 MB	7.2 GB
		Data Service	5,271,000	31,020,000	1.6 GB	9.7 GB
		Central Settlement	6,000,000	30,000,000	840 MB	4.2 GB
Consumption	Daily SP level data	Central Settlement	15,309,000,000	n/a	30 TB	n/a
	Load Shapes	Data Service	297,410	n/a	74 MB	n/a
ISD	All Entities	n/a	2,520	n/a	33.8 GB	n/a

The data volumes defined above are preliminary estimates, and further work is being undertaken to provide more refined values. This analysis will be available in a separate document *(Reference MHHS Data Integration Platform - Functional Specification – Appendix A – Transaction Volumes, in preparation)*.

5.10.3 Performance - Message Latency

The DIP requires near real-time message delivery, with 90% of all messages needing to be delivered within 3 seconds of receipt and 100% of messages within 30 seconds. These requirements are for all currently identified business processes that the DIP must support.

In order to future-proof the DIP is recognised that the platform will need to be able to support near-real-time, i.e. <1 second, message flows.

5.11 Reporting

The DIP will present a number of different reports to allow Market Participants and the DIP Administrator to view and track the performance of business processes and message flows. DIP Users will have access to these reports and real-time alerting via a dashboard. The Users will be able to configure their individual reporting and alerting requirements from the dashboard, e.g. via email.

The expectation is that the messages will have a number of key attributes (tags) available in the message header, which can be commonly linked across related message channels, such as a correlation ID, transaction ID and MPAN. These can be used to create reports that link business processes across the various channels.

Primary access to reports for DIP users will be achieved through RBAC. In addition to the primary access control, the second level of control is required: DIP Users will only be allowed to see reports for items they are authorised to view; for example, they will only view individual MPANs that they are responsible for. Reports that do not show specific MPANs, say for performance throughput where totals are reported, will not have this restriction.

In addition API

5.11.1 Message Channel Throughput

A series of reports is required to measure each channel's message/event throughput per time interval.

- Volumes - both incoming and outgoing message/event volumes and error messages. The outgoing volume will be greater on most channels as messages are sent to multiple recipients. Volumes can be reported by each sender on the channel, and likewise, outgoing volumes can be reported by individual recipients.
- Data Latency – end-to-end transaction times between the time messages arrive via the incoming API call, then written to the incoming queue to the time the recipient reads the corresponding message.
- Data latency can be measured on a maximum, minimum, and average basis for each reported period, and these can be measured against the end-to-end times for all recipients.
- Both volume and data latency metrics can be filtered and reported on either a sender/receiver basis or across multiple and individual message channels.
- Successful/failed Message selection
- The time interval reported needs to be flexible.

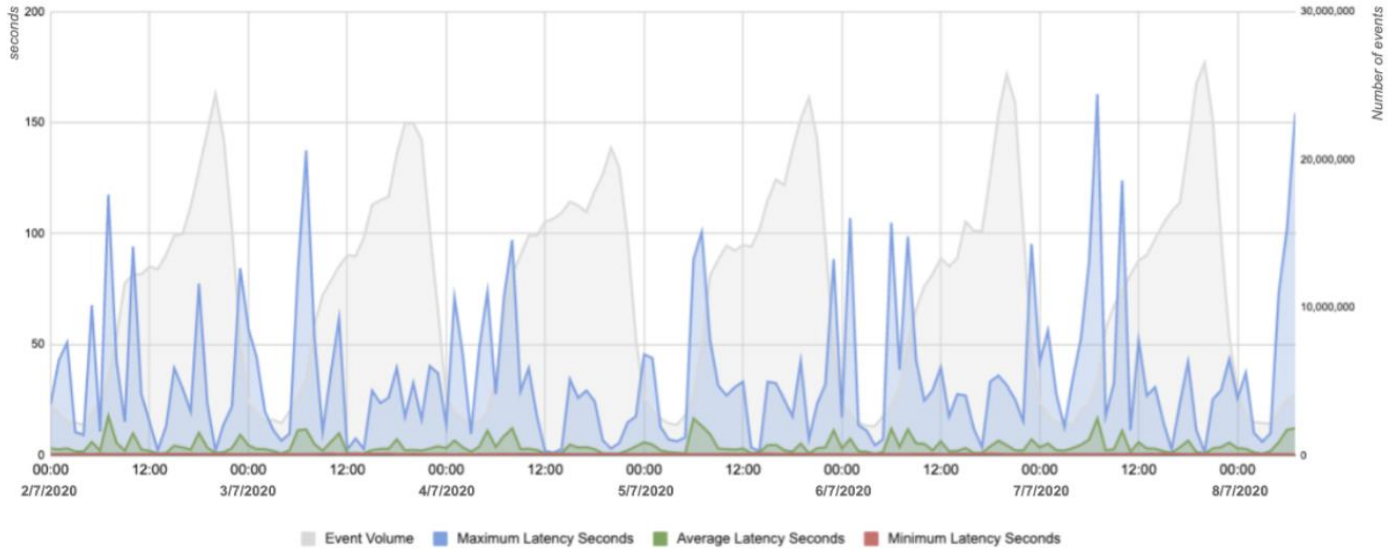

Figure 14 – Example of Data Volume and Latency Report

5.11.2 DIP Internal Performance Report

Like the Message Channel Throughput report, the DIP's internal performance needs to be reported on, i.e. the time taken to process a message from receipt to send. As above:

- The time interval reported on needs to be flexible.
- Volumes - both incoming and outgoing volumes and error messages on each channel need to be reported.
- Data Latency – end-to-end transaction times between message receipt and the outgoing topic/queue message.
- Data latency can be measured on a maximum, minimum and average basis for each reported period.

5.11.3 Reporting Message/Event Flow Activity

A report is required to provide an audit trail for each of the business processes undertaken within the DIP. The DIP cannot provide a detailed report on business process workflow outcome; however, it can provide a report based on the publication message channels used with each business process or group of business processes linked together.

In order to provide DIP Users real-time access to an audit trail of their transactions as they pass through the DIP, a report on each message channel is proposed. To provide flexibility, users will be allowed to configure reports based on a series of these 'base' reports as they are being linked together. The linking of reports will provide a view of an end-to-end business process. The linking of different reports can be achieved based on correlation ID, a Transaction ID or an MPAN (dependent on the final design of the message header and the tags that can be extracted). A set of predefined linked reports representing the orchestrated business process will be made available.

The following table provides an indication of the business process that need to be linked together along with the corresponding interfaces (IF), i.e. Message Channels:

Aquistion Process

		MPAN								
Tracking Linkage		Corrleation ID	Correlation ID					MPAN	MPAN	
Process Transitions		BP01- COS	BP02 - Service Appointment (x2)					BP04 - Data Collection	BP05 - Data Processing	
Acquisition Stages	Gain Confirmed	IF-001								
	Reg Info Provided		IF-02							
	Service App Requested			IF-031						
	Service App Rejected			IF-031	IF-032 (RJ)					
	Service App Accepted			IF-031	IF-032 (AC)					
	Service App Lapsed			IF-031	IF-032 (AC)	IF-032 (LP)				
	Service App with Agent			IF-031	IF-032 (AC)	IF-033				
	Rejected Appointment Req			IF-031	IF-032 (AC)	IF-033	IF-034 (RJ)			
	Approved Appointment			IF-031	IF-032 (AC)	IF-033	IF-034 (AC)			
	Sucessful Appointment			IF-031	IF-032 (AC)	IF-033	IF-034 (AC)	IF-036		
	Sucessful Appointment + Read			IF-031	IF-032 (AC)	IF-033	IF-034 (AC)	IF-036	IF-041 (Opt)	
	First HH Submission			IF-031	IF-032 (AC)	IF-033	IF-034 (AC)	IF-036	IF-041 (Opt)	IF-021
	Failed HH Submission			IF-031	IF-032 (AC)	IF-033	IF-034 (AC)	IF-036	IF-041 (Opt)	IF-021	IF-013

De-Acquisition Process

		MPAN		
Tracking Linkage		MPAN	MPAN	
Process Transitions		BP01- COS	BP02 - Service Appointment (x2)	
De-Acquisition Stages	Loss Confirmed	IF-001		
	De-Appointed	IF-001	IF-0037	
	De-Appointed + Read	IF-001	IF-0037	IF-041 (Opt)

Change of Service

		MPAN							
Tracking Linkage		Correlation ID					MPAN	MPAN	
Process Transitions		BP02 - Service Appointment (x1)					BP04 - Data Collection	BP05 - Data Processing	
	Service App Requested	IF-031							
	Service App Req Rejected	IF-031	IF-032 (RJ)						
	Service App Req Lapsed	IF-031	IF-032 (LP)						
	Service App Accpeted	IF-031	IF-032 (AC)						
	Service App Sent to Agent	IF-031	IF-032 (AC)	IF-033					
	Rejected Appointment Req	IF-031	IF-032 (AC)	IF-033	IF-034 (RJ)				
	Approved Appointment	IF-031	IF-032 (AC)	IF-033	IF-034 (AC)				
	Sucessful Appointment	IF-031	IF-032 (AC)	IF-033	IF-034 (AC)	IF-036			
	Sucessful Appt + Read	IF-031	IF-032 (AC)	IF-033	IF-034 (AC)	IF-036	IF-041 (Opt)		
	First HH Submission	IF-031	IF-032 (AC)	IF-033	IF-034 (AC)	IF-036	IF-041 (Opt)	IF-021	
	Failed HH Submission	IF-031	IF-032 (AC)	IF-033	IF-034 (AC)	IF-036	IF-041 (Opt)	IF-021	IF-013

Metering Activity

		MPAN								
		Correlation ID (Opt)					Correlation ID (Opt)			MPAN
Tracking Linkage		Envelope ID (opt)								
Process Transitions		BP09 - Change of Meter					BP08 - DeEnergisation			BP07 - Discon
Metering Actvitity Stages	Mtr Serv - Updates Required	IF-005								
		IF-005	IF-005 (Opt)							
	Registration Rejects Updt	IF-005	IF-005 (Opt)	IF-006 (RJ)						
	Registration Confirms Works			IF-006 (AC)						
	Metering Activity Comepleted	IF-005	IF-005 (Opt)	IF-006 (AC)	IF-041					
		IF-005	IF-005 (Opt)	IF-006 (AC)	IF-041	IF-041 (Opt)				
	Meter Works + DeEn	IF-005	IF-005 (Opt)	IF-006 (AC)	IF-041	IF-041 (Opt)	IF-007			
	DeEn Rejected	IF-005	IF-005 (Opt)	IF-006 (AC)	IF-041	IF-041 (Opt)	IF-007	IF-008 (RJ)		
	DeEnergised	IF-005	IF-005 (Opt)	IF-006 (AC)	IF-041	IF-041 (Opt)	IF-007	IF-008 (AC)		
	DeEnergised+Read	IF-005	IF-005 (Opt)	IF-006 (AC)	IF-041	IF-041 (Opt)	IF-007	IF-008 (AC)	IF-041 (Opt)	
	Disconnected	IF-005	IF-005 (Opt)	IF-006 (AC)	IF-041	IF-041 (Opt)	IF-007	IF-008 (AC)	IF-041 (Opt)	IF-009
	DeEnergised Requested	IF-007								
	DeEnergised Rejected	IF-007	IF-008 (RJ)							
	DeEnergised	IF-007	IF-008 (AC)							
	DeEnergised+Read	IF-007	IF-008 (AC)	IF-041 (Opt)						
	Disconnected	IF-007	IF-008 (AC)	IF-041 (Opt)	IF-009					

AC- Accept
Opt - Optional
RJ - Reject
LP - Lapsed

Figure 15 - Cross Business Process Tracking

The following conceptual diagrams provide a 'mock-up' of the types of reports required and hence should only be seen as an indicative design. RFP bidders have the flexibility to propose their own designs, and the following only suggest how the requirements could be met. Bidders need to provide a design that meets the requirements in the most cost-effective approach. It is recognised that this is an area that requires much further work in the design phase.

General Audit Report

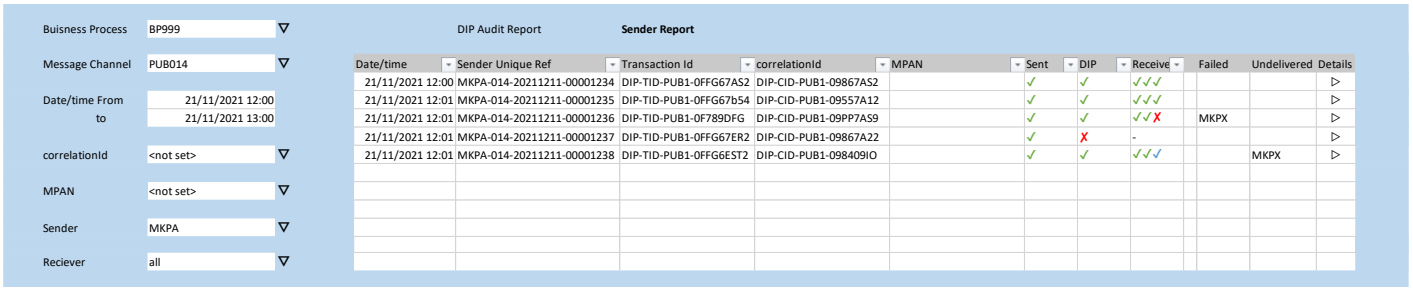

Figure 16 - Audit Report - example 1 – Sender Report

In the first example, the focus is on the message's Sender. The DIP user is allowed to select a business process, message channel and a date/time range. If a business process, then the message channels within the BP are selected, and the Sender automatically defaults to the DIP User.

The columns in the report describe whether the message was successfully received (API call) [Sent], processed by the DIP [DIP] and Received by the intended recipients (webhook call) [Received] – a green tick for each successful message received, a red cross indicates an error condition.

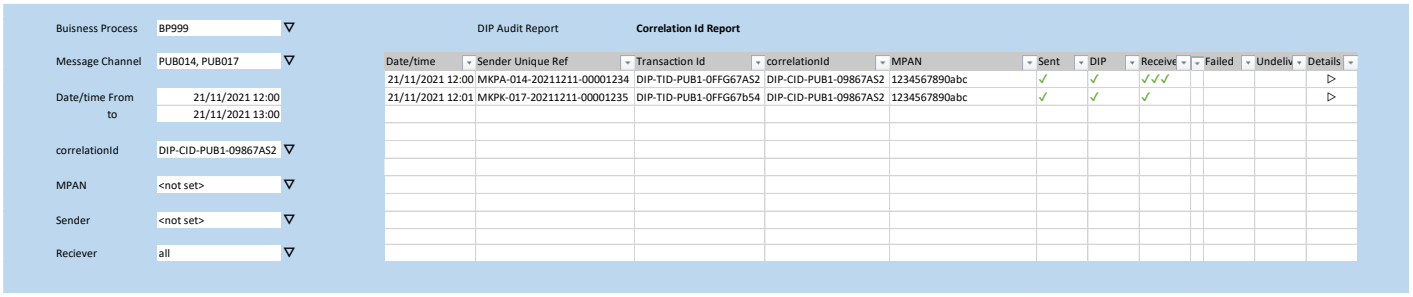

Figure 17 - Audit Report - example 2- Transaction Id Report

The second example of the audit report focuses on the Correlation ID. A business process is selected; hence all the pertinent messages/channels are automatically selected. The User can then select a Transaction ID and trace it through the various message channels.

In addition, the User will require a facility to track an individual MPAN across multiple message/event channels:

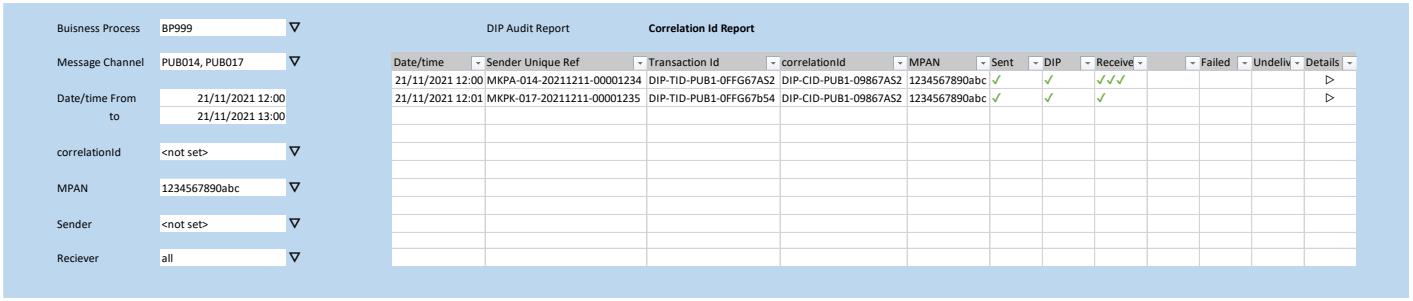

Figure 18 - Audit Report - example 3 - MPAN Report

Detailed Audit Report

The detailed audit is standalone or accessed from the General Audit report. It details the workflow of a selected message, including the API returns codes and responses from initial message ingestion to the webhook response when the message is retrieved by the recipient(s). The message timings are also included from inception to delivery. Any errors trapped by the DIP, sender or recipients are also shown.

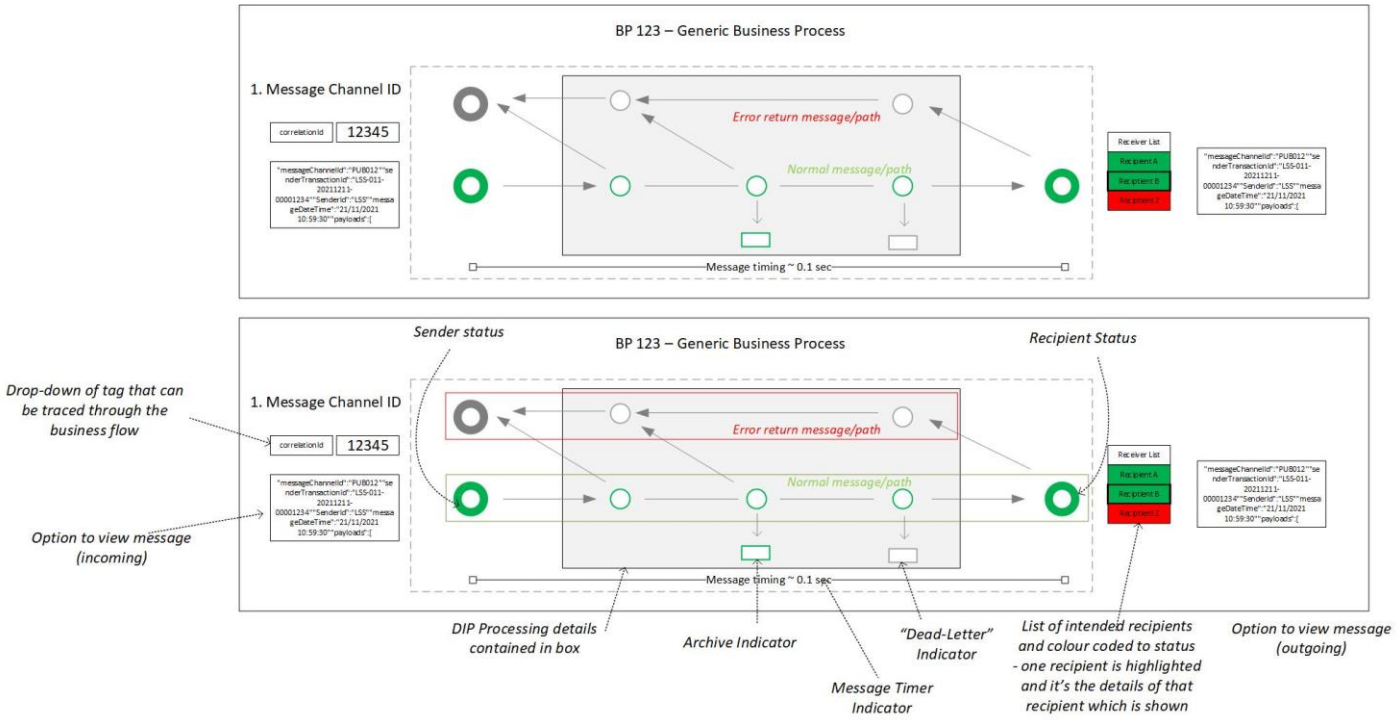

Figure 19 - Detailed Audit Report

- Green circles at the beginning and end of workflow denote successful message exchange.
- At the beginning or end, a red circle denotes a problem with the API call and the rejected message (not shown on above report).
- The timing indicator shows the elapsed time from the initial message receipt to the message delivered to the recipient.
- Different receivers can be selected at the endpoint as they may have different outcomes and timings.
- Green circles in the DIP processing box denote the flow and success of the message transitioning through the DIP, with indicators showing message written to archive and message written to DLQ.
- A green circle with a red dot on the error return path denotes an error message sent back successfully.
- A grey circle denotes step not activated.
- Option to view the message header and payload.

5.11.4 Alarms

Alarms are required on message/event channels where the latency has extended beyond a predefined threshold and messages have been moved to 'dead letter' queues. This would indicate that messages are not being processed promptly, and hence an investigation would need to be initiated.

5.12 User Interfaces

5.12.1 User Portal

The DIP will present a portal (webpage) that will facilitate user management and provide Users with some essential services:

User Onboarding	Provide a facility for a company to onboard as a DIP user.

User Roles Request	Registered Users will view the role(s) they have currently been assigned and the business services they can interact with. They will have the ability to sign up for access to optional message channels.
Agent Management	Market Participants will be able to nominate Agents that submit data on their behalf.
Certificate Management	Registered Users will manage their X509 certificates for accessing the system and encrypting their data. {*may not be available here – yet to be decided as certificate management may be delivered separately*}
Payload download	A web page will be available for interrogation and download for all message channels adopting Message Pattern B.
DIP Reporting	Access the DIP Performance, Audit Reports and alerting
DIP Message Replay	Access DIP message replay facility
Data Portal	Access Data management facilities

5.12.2 DIP Administration Portal

The DIP will present a portal (webpage) that will allow the management of DIP Users:

Create User	Provide a facility where the DIP admin will create a DIP user and assign them specific roles based on their industry role.
Manage User	Provide a facility where the DIP admin can manage a DIP user and assign/remove specific roles based on their industry role. Also remove a user.
Certificate Management	Provide a facility for the Administrator to issue and manage X509 certificates to Users. {again, may not be available here – yet to be decided as certificate management may be delivered separately}

5.13 Hosting

The current assumption is that a single cloud provider will host the solution (as the benefits provided by a multi-cloud solution do not outweigh the added complexity and costs).

To satisfy the availability requirements, the working assumption is that an online DIP service will be located in at least two availability zones/regions where the underlying cloud services are replicated between sites. The method of replicating services and maintaining system availability will depend on platform choice and progressed in the detailed design phase.

In addition, a disaster recovery (DR) capability is required.

5.13.1 Cloud Hosting Pattern

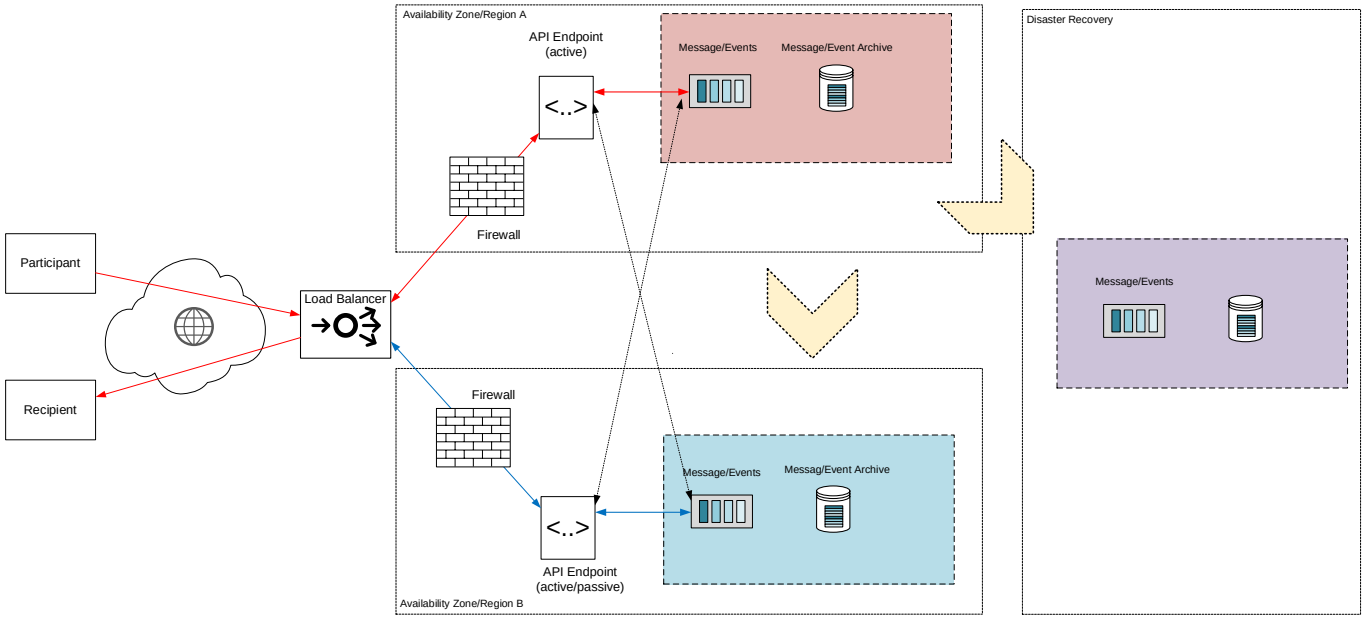

Figure 20 - Proposed Cloud Hosting Pattern for DIP

5.13.2 System Availability

The platform will have the following availability requirements:

Percentage of Uptime	99.95% (unplanned)
Mean Time to Recovery (MTTR)	60 mins
Mean Time between Failures (MTBR)	-
Recovery Time Objective (RTO)	60 mins
Recovery Point Objective (RPO)	0

5.13.3 Environments

Multiple instances of the DIP will need to be provisioned in order to support the different phases of development of the MHHS Programme:

a) Development environment set up and configured for the use by developers in the development (build) of the DIP (DEV)

b) Test environment set up and configured for Unit Test (UIT)

c) Systems Integration Test, consisting of all participants (with test tools to simulate participant messages /responses and other interactions) (SIT)

d) Pre-Production Environment (PRE-PROD)

e) Production Environment (PROD)

These environments will need to operate concurrently; the deployment capabilities described above need to be tested and trialled in the non-production environments before being rolled out to Production. The design needs to cope with

different sized environments; the Production and Non-Production will have different sizing/capacity requirements. Different certificates will be required to access Production and Non-Production environments.

5.13.4 Cloud Edge Services

To support the security and availability requirements for the DIP, the design assumes the need for the following cloud services: load balancer, firewall and application gateway.

- Load Balancer – provides the primary method for routing participant traffic to the active DIP services
- Firewall – protects the back-end services from malicious
- Application Gateway – provides the following capabilities:
  - API endpoint and routes request accordingly to the back-end services.
  - Access control from connections to the API endpoint based on certificates and keys.
  - API logging of requests and responses.
  - Supports API version control

5.13.5 Maintenance

The expectation is that the system will require maintenance from time to time; details of forthcoming maintenance windows will be published on the system website and communicated to users via email.

The maintenance of an individual message channel should not impact the rest of the system.

5.13.6 Ownership

The current assumption is that the solution deployed will exist as a standalone domain and will not be integrated into the delivery partner's or the ESO's IT estate, i.e., a new cloud-based solution. This means a single standalone management group in Azure or AWS, a separate organisation. The solution must adopt a suitable domain name and public IP addresses which can be transferred to a new service manager in the event of a change in the service contract.

6 Change Management

The DIP Service Provider will need to demonstrate best practice approaches to development and testing during the complete development and maintenance cycle for the DIP. The use of best practice documentation standards such as Unified Modelling Language (UML) etc., will be required. Adoption of a recognised development methodology, e.g. Agile, is essential. The DIP Service Provider must be open and willing to share information with industry participants through online collaborations tools.

7 Service Management

The DIP Service Provider will be responsible for the service management functions that facilitate the day-to-day maintenance and support of the DIP. The expectation is that the service management function should adopt a recognised industry framework such as ITIL.

The provision of service management capabilities for the DIP does not need to be standalone; i.e. it is expected that any potential DIP Service Provider would leverage existing service capabilities they already provide so long as the SLAs required for the DIP are not compromised in such a shared service.

The DIP Service Provider will need to provide a Service Desk for in-hour working day support (core hours) and Out-of-hours support for emergency calls outside the core support hours where DIP users can raise an issue with the DIP's performance. DIP users will be encouraged to raise all non-urgent issues via the User Portal, including standard user requests.

The service windows will be:

	Hours of Operation
Service Desk	Business Day 09:00 to 17:00
Out-of-Hours	Business Day 17:00 to 09:00, all day weekends and Bank Holidays

The Business Day and Business Hours definition will follow the definition determined by the ESO.

7.1 Service Desk

The primary responsibilities of the Service Desk will be:

- Management and ownership of incidents throughout their lifecycle
- Providing a professional interface between customers/users and the IT Service Provider
- Providing first level IT Service support
- Providing management information on IT service provision and producing associated reports

The Service Desk provides a single contact point for Customers and Users and is there to log, track and manage any issues brought to its attention. DIP users will call (or email) the Service Desk when an Incident occurs or has a query or issue. A standard set of Service Management Tools will be used to support these activities, and the knowledge base/record or events built up from these tools must be transferrable in the event of a change of DIP Service Provider.

7.2 Out of Hours Support

The out of hours support function will only be utilised to communicate P1 incidents outside standard service desk hours. Phone calls within this period should be answered immediately (no response to emails) or within 60 minutes if no personnel are immediately available (via voicemail).

7.3 Incident Management Process

The expectation is that the incident management process and the categorisation and prioritisation of incidents will be integrated into any prospective DIP Service Provider service management function. However, all incidents/service requests must be recorded, classified and prioritised according to a defined procedure that considers the impact and urgency of the incidents.

As part of the incident process, the following must be followed:

- Incident Owner – assigning an incident owner is essential to ensure that all activities occur promptly.
- Incident Tickets – all contacts and interactions with the customer must be documented into an incident ticket.
- Incident Priority – incident priority or severity should be set using an incident priority matrix.
- New Incidents – if the customer contacts the Help Desk about a new issue, the Help Desk Agent will create a new incident ticket.

- Existing Incidents – If the customer is contacting the Help Desk about an existing issue, the Help Desk Agent will search for existing tickets and provide the User with a status update. The incident ticket must be updated with a summary of the interaction.
- Escalation Queue Management – If the Help Desk cannot resolve an incident, the Help Desk Agent will assign the incident ticket to the appropriate Escalation Queue for the escalated work team.
- Incident resolution – The incident ticket should be resolved when the service has been restored to standard operation, a permanent fix or a temporary workaround. Incidents should not be moved to a status of "resolved" until service has been restored.
- Incident closure – Incidents should not be moved to a “closed " status until the incident resolution has been confirmed with the customer.
- Incident reopen – An incident in a "closed" status should never be reopened. If the incident was not resolved, a new incident ticket should be opened, which will be related to the previous incident.
- Root cause analysis to understand the cause of the event.

7.4 Major Incident Management

A major incident ("P1") will be defined as a loss or complete loss of the DIP service or a partial loss of one or more message channels where participants are unable to send or receive messages.

The DIP Service Provider must have an Incident Management Communication Plan to follow when an outage to a service occurs. The Incident Management Communication Plan detailing how people will be initially notified, what information they need, when status updates will be communicated, and what resolution steps occur when a service has been restored.

7.4.1 Post-Incident Review

A post-incident review (PIR) process needs to be followed to evaluate the incident response for major, critical, and high priority incidents. The post-incident review will be initiated once the incident has been resolved and will review the incident from start to finish. The goal will be to determine if the incident could have been handled better.

8 Security Architecture

The security architecture of the DIP is covered in greater depth in the *MHHS End-to-End Security Architecture document*. This document provides a high-level summary of the important security topics concerning the DIP.

8.1 Data Security

The DIP project is mandating an ISO27001 accreditation for data security and management. The requirement is that the solution developed must adhere to this standard for the duration of the contract term.

Some of the payloads will be personally identifiable data under UK GDPR, and the expectation is that data must be encrypted in transit, during processing and at rest.

8.2 Connection Security

The expectation is that most users' connection to the new platform will be over the public internet, and authentication will be via mTLS.

8.3 User Management

The current assumption is that the DIP will support users based on a role-based access control (RBAC) approach. Users will be given access to the different services based on their role, i.e., read /write to different service endpoints. There will be the need for a DIP Administrator responsible for creating and managing Users.

8.4 Certificate Management

A dedicated robust certificate management process would need to be implemented to maintain certificate integrity.

The DIP service provider shall implement and operate a dedicated PKI infrastructure to be used for all parties connecting to the DIP

The current assumption is that the DIP may leverage the certificate services implemented for the faster switching programme rather than burden the industry with the cost of additional certificate services. However, should the DCC solution prove to be unworkable for the MHHS implementation, MHHS would need to establish certificate services of their own to meet the security requirements of the MHHS TOM.

The expectation is that DIP users will require different certificates for Production and Non-Production environments.

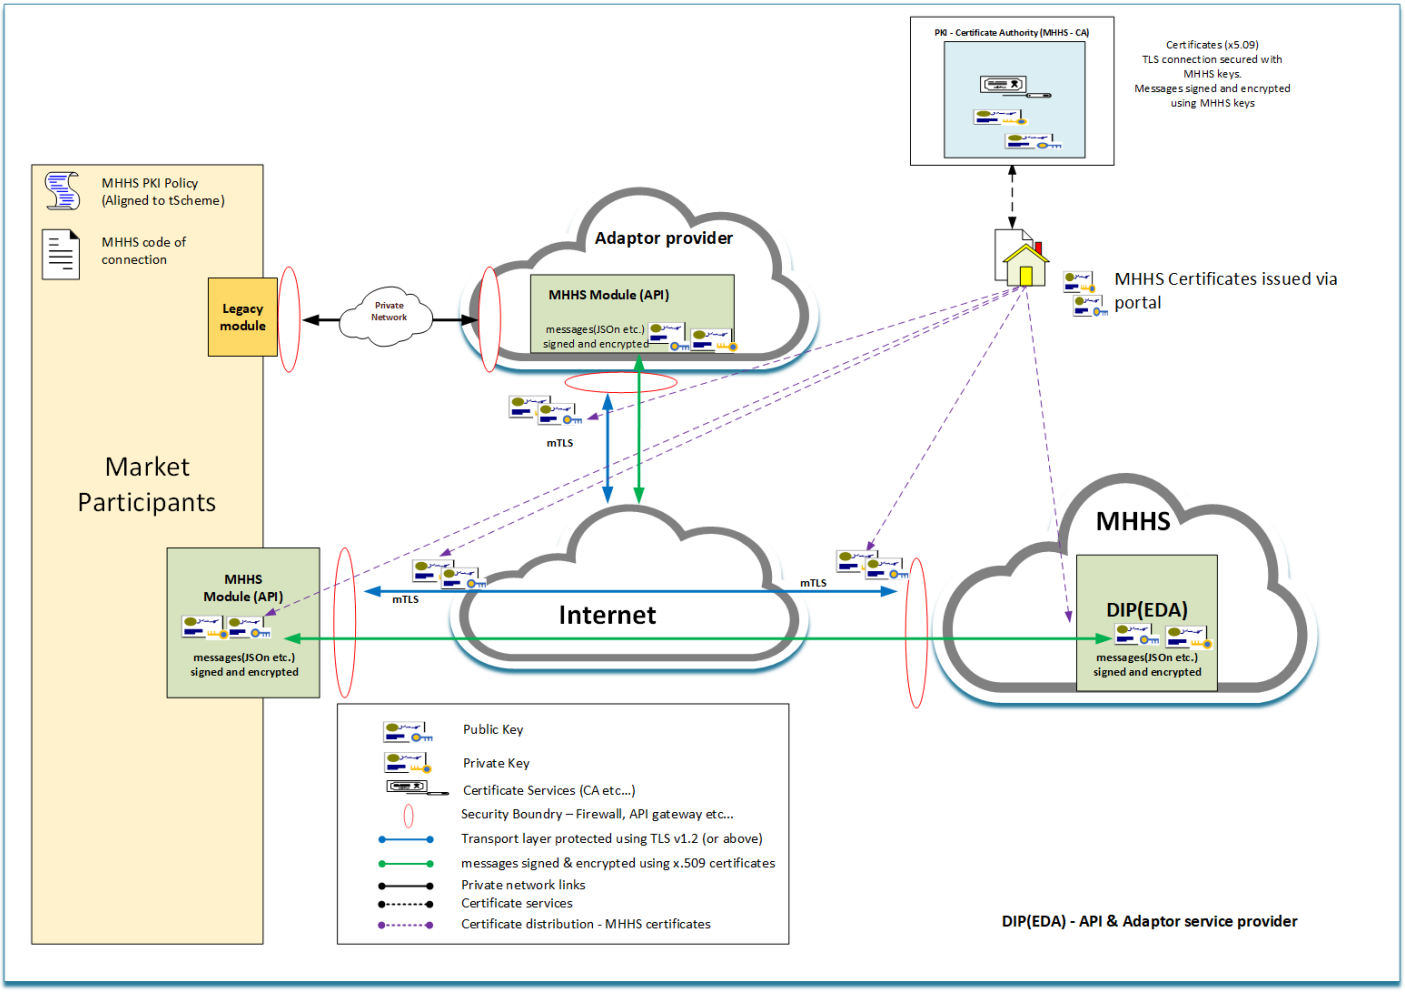

Figure 21 - DIP Certificate Services (proposed)

8.5 Operational security

The DIP will be managed according to industry best practices regarding security incidents and investigation (CREST, NIST) and ISO 18788:2015.

All assets within the DIP will be built/configured in line with industry best practices following secure configuration standards such as CIS, NIST, and NCSC.

The DIP Service Provider should focus on identifying seemingly innocuous actions that could inadvertently reveal critical or sensitive data to a cyber-criminal. OPSEC is both a process and a strategy; the DIP Service Provider should view their operations and systems from the perspective of a potential attacker.

- Analytical activities and processes
- Behavioural monitoring
- Social media monitoring
- Security best practice.

The DIP Service Provider should use their risk management processes to discover potential threats and vulnerabilities in the DIP processes, the way they operate, and the software and hardware, both physical and logical. Looking at systems and operations from a third party's point of view enables the DIP Service Provider's security teams to discover issues they may have overlooked and can be crucial to implementing the appropriate countermeasures that will keep their most sensitive data secure.

The DIP will be regularly scanned and tested for vulnerabilities, and vulnerability analysis will be undertaken.

- Countermeasures based on the analysis of any threats and vulnerabilities detected will be applied to the DIP regularly.

https://www.ncsc.gov.uk/collection/cloud-security/implementing-the-cloud-security-principles/operational-security

8.6 Secure Software Development

To provide premium security while also applying faster process speed, accessibility, and scalability and finding creative solutions by breaking down barriers between development teams.

Establish responsibility for security activities and expectations for behaviour, and set mandatory expectations related to software security.

The code developed for the DIP will be developed using secure coding standards (OWASP, CREST) and static code scanning tools (SAST) and software composition analysis (SCA) to understand better the impact of code on risks related to security.

All software development (application and cloud infrastructure) should follow a recognised software development lifecycle such as OWASP / Microsoft SDL.

Figure 22 - Diagram for illustration purposes only.

8.6.1 SDLC phases:

- Planning and requirements
- Architecture and design
- Test planning
- Coding
- Testing and results
- Release and maintenance

8.6.2 Define the key security policies

1. Software security. Build security into product requirements, implementation, procurement, deployment, and operations:
    - Secure SDLC. Use is not optional.
    - Application risk ranking. Identify where the most significant technical risk lies.
    - Application design. Define security controls built into the DIP based on this document and the detailed security requirements.
    - Application development. Require specific technology stacks and mandatory coding standards. Adhere to software secure-by-design principles such as OWASP / Microsoft SDL.
    - Application testing. Defined schedules and testing intervals for:
        - Static and dynamic code evaluation.
        - White box / black box testing.
    - Defect severity and remediation. Establish rules for setting bug and flaw severities and timelines for fixing coding bugs and design flaws.

2. Network security. Determine protocols and authorisation levels to help define the DIP security.
3. Data security. Identify and classify sensitive data (MPAN, PII, Consumption data) apply the correct security features based on the data privacy classifications aligned to SPaR.
4. Virtual infrastructure security. Govern access control to secure the virtual infrastructure of the DIP.
5. Disaster recovery. Determine steps to take in the event of an attack, including reporting, recording, and resolution for attacks against applications

8.7 Governance

A governance framework will be implemented ensuring procedures, personnel, physical and technical controls continue to work through the lifetime of the DIP. It should also respond to changes in the service, technological developments and the appearance of new threats. CSA CCM v3.0 ISO/IEC 27001)

https://www.ncsc.gov.uk/collection/cloud-security/implementing-the-cloud-security-principles/governance-framework

8.7.1 Risk Management

A robust risk process will be established to identify any risks posed to the DIP when connecting Market participants to the DIP. Risk assessments of the DIP will be undertaken in line with industry best practices (NIST, ISO 27001).

Guidance will be provided as to the minimum security requirements the market participant must meet to be connected to the DIP.

Where possible, MHHS intends to leverage the investment already made by the Faster [& More Reliable] Switching Programme (FSP) concerning risk assessing the Market Participants' security posture in meeting the minimum security requirements set out in the DCC Code of Connection. Therefore a gap analysis will be undertaken to identify market participants that have already met the minimum security requirements when connecting to faster switching.

9 User Requirements

DIP project requirements defined will be based on the characteristics as defined in the quality framework ISO/IEC 25010. https://www.iso.org/standard/35733.html. The framework adopts the following characteristics:

Functional Requirements can be measured against:

- Functional suitability
  - ***Functional Completeness*** - the degree to which the set of functions covers all the specified tasks and user objectives.
  - ***Functional Correctness*** - the degree to which the functions provides the correct results with the needed degree of precision.
  - ***Functional Appropriateness*** - the degree to which the functions facilitate accomplishing specified tasks and objectives.

Non-Functional Requirements:

- Performance Efficiency
  - ***Time behaviour*** - the degree to which a product or system's response and processing times and throughput rates meet requirements when performing its functions.
  - ***Resource utilisation*** - the degree to which the amounts and types of resources used by a product or system meet requirements when performing its functions.
  - **Capacity** - the degree to which the maximum limits of the product or system parameter meet requirements.
- Reliability
  - ***Maturity*** - the degree to which a system, product or component meets needs for reliability under normal operation.
  - ***Availability*** - the degree to which a product or system is operational and accessible when required for use.
  - ***Fault tolerance*** - the degree to which a system, product or component operates as intended despite the presence of hardware or software faults.
  - ***Recoverability*** - the degree to which a product or system can recover the data directly affected and re-establish the system's desired state in the event of an interruption or a failure.
- Maintainability
  - ***Modularity*** - the degree to which a system or computer program is composed of discrete components such that a change to one component has minimal impact on other components.
  - ***Reusability*** - the degree to which an asset can be used in more than one system or in building other assets.
  - ***Analysability*** - degree of effectiveness and efficiency with which it is possible to assess the impact of an intended change on a product or system to one or more of its parts, diagnose a product for deficiencies or causes of failures, or identify parts to be modified.
  - ***Modifiability*** - the degree to which a product or system can be effectively and efficiently modified without introducing defects or degrading existing product quality.
  - ***Testability*** - degree of effectiveness and efficiency with which test criteria can be established for a system, product or component and tests can be performed to determine whether those criteria have been met.
- Usability
  - ***Appropriateness recognizability*** - the degree to which users can recognise whether a product or system is appropriate for their needs
  - ***Learnability*** - the degree to which a product or system enables the User to learn to use it effectively in emergencies.
  - ***Operability*** - the degree to which a product or system is easy to operate, control and appropriate to use.

- ***User error protection*** - the degree to which a product or system protects users against making errors.
- ***User interface aesthetics*** - the degree to which a user interface enables pleasing and satisfying interaction for the User.
- ***Accessibility*** - the degree to which a product or system can be used by people with the broadest range of characteristics and capabilities to achieve a specified goal in a specified context of use.

- Portability
  - ***Adaptability*** - the degree to which a product or system can effectively and efficiently be adapted for different or evolving hardware, software or other operational or usage environments.
  - ***Installability*** - degree of effectiveness and efficiency in which a product or system can be successfully installed and/or uninstalled in a specified environment.
  - ***Replaceability*** is the degree to which a product can replace another specified software product for the same purpose in the same environment.
- Compatibility
  - ***Coexistence*** - the degree to which a product can perform its required functions efficiently while sharing a common environment and resources with other products, without detrimental impact on any other product.
  - ***Interoperability*** - the degree to which two or more systems, products or components can exchange information and use the information that has been exchanged.
- Security
  - ***Confidentiality*** - the degree to which that data is accessible only to those authorised to have access.
  - ***Integrity*** - the degree to which a system, product or component prevents unauthorised access to, or modification of, computer programs or data.
  - ***Non-repudiation*** - the degree to which actions or events can be proven to have taken place so that the events or actions cannot be repudiated later.
  - ***Accountability*** - the degree to which the actions of an entity can be traced uniquely to the entity.
  - ***Authenticity*** - the degree to which the identity of a subject or resource can be proved to be the one claimed.

The complete list of Non-Functional Requirements for the MHHS Programme is available in the spreadsheet (MHHS-DIP002). At this stage, the expectation is that the functional requirements of the MHHS programme will not significantly impact the architecture of the DIP. The business process is being defined, from which the user requirements are being driven out, making use of the DIP solely for the exchange of messages.