Market Interface Technology Enhancements Working Group (MITE WG)

Wednesday 2 October 2024

(1:00pm to 5:00pm)

This meeting will be recorded for minute taking purposes.

We acknowledge the Traditional Custodians of the land, seas and waters across Australia. We honour the wisdom of Aboriginal and Torres Strait Islander Elders past and present and embrace future generations.

We acknowledge that, wherever we work, we do so on Aboriginal and Torres Strait Islander lands. We pay respect to the world's oldest continuing culture and First Nations peoples' deep and continuing connection to Country; and hope that our work can benefit both people and Country.

'Journey of unity: AEMO's Reconciliation Path' by Lani Balzan

AEMO Group is proud to have delivered its first Reconciliation Action Plan in May 2024. 'Journey of unity: AEMO's Reconciliation Path' was created by Wiradjuri artist Lani Balzan to visually narrate our ongoing journey towards reconciliation - a collaborative endeavour that honours First Nations cultures, fosters mutual understanding, and paves the way for a brighter, more inclusive future.

Read our RAP

Housekeeping

1. This meeting will be recorded for minute taking purposes
2. Please mute your microphone, this helps with audio quality as background noises distract from the conversation.
3. Use the ‘Raise hand’ function should you wish to speak to an item.
4. Use the ‘Chat’ function for any other questions or comments you may have.
5. In attending this meeting, you are expected to:
    - Not only represent your organisation’s interests but also the interests of Industry and its customers
    - Have an open mindset
    - Contribute constructively
    - Be respectful, both on the call and in the chat

1. Welcome

Blaine Miner

Objective of today's session

The MITE WG has been established to define and develop Technical Procedures/guides for IDAM, IDX and Portal Consolidation. These initiatives seek to deliver foundational capability supporting interactions between participants and AEMO and based on the agreed scope to transition or enable decisions on transitioning of existing business services

This workshop aims to cover:

- The forward plan
- Recent focus / working group sessions:
  - IDAM API Authentication and authorisation
  - IDX Decision Tree
  - IDX Business Function Endpoints
- IDX Archiving

The ask of participants:

- Invite and share this pack with your technical experts who will support the MITE WG / FG process to provide context and background
- Engage in the workshop – questions are welcome

Link to the target state pack established in consultation with the industry stakeholders

Agenda

#	Indicative Timings	Topic	Presenter
1	1:00pm-1:05pm	Welcome	Blaine Miner
2	1:05pm-1:10pm	Forward Plan	Blaine Miner
3	1:10pm-2:20pm	IDAM API Authentication and authorisation – Playback	Sivaraj Ganesan
4	2:20pm-2:35pm	IDAM API Authentication and authorisation FG Actions	Sivaraj Ganesan
5	2:35pm-2:40pm	IDAM Future Topics	Phil Hayes
6	2:40pm-2:45pm	IDX Actions	Blaine Miner
7	2:45pm-3:40pm	IDX Decision Tree – Playback	Sri Gundu
8	3:40pm-4:10pm	IDX Business Function End Points – Playback	Sri Gundu
9	4:10pm-4:40pm	IDX Archiving – Working Group Session	Selwyn Sequeira
10	4:40pm-4:55pm	IDX Future Topics	Andrew Bell
11	4:55pm-5:00pm	General Business and Next Steps	Blaine Miner
	Appendix	Appendix A: AEMO Competition Law Meeting Protocol, Appendix B: IDAM End-End Worked Example Appendix C: IDAM Steps Appendix D: Message Threshold Options Appendix E: Current Archiving use cases	

2. Forward Plan

Blaine Miner

Consultation Workshop Structure

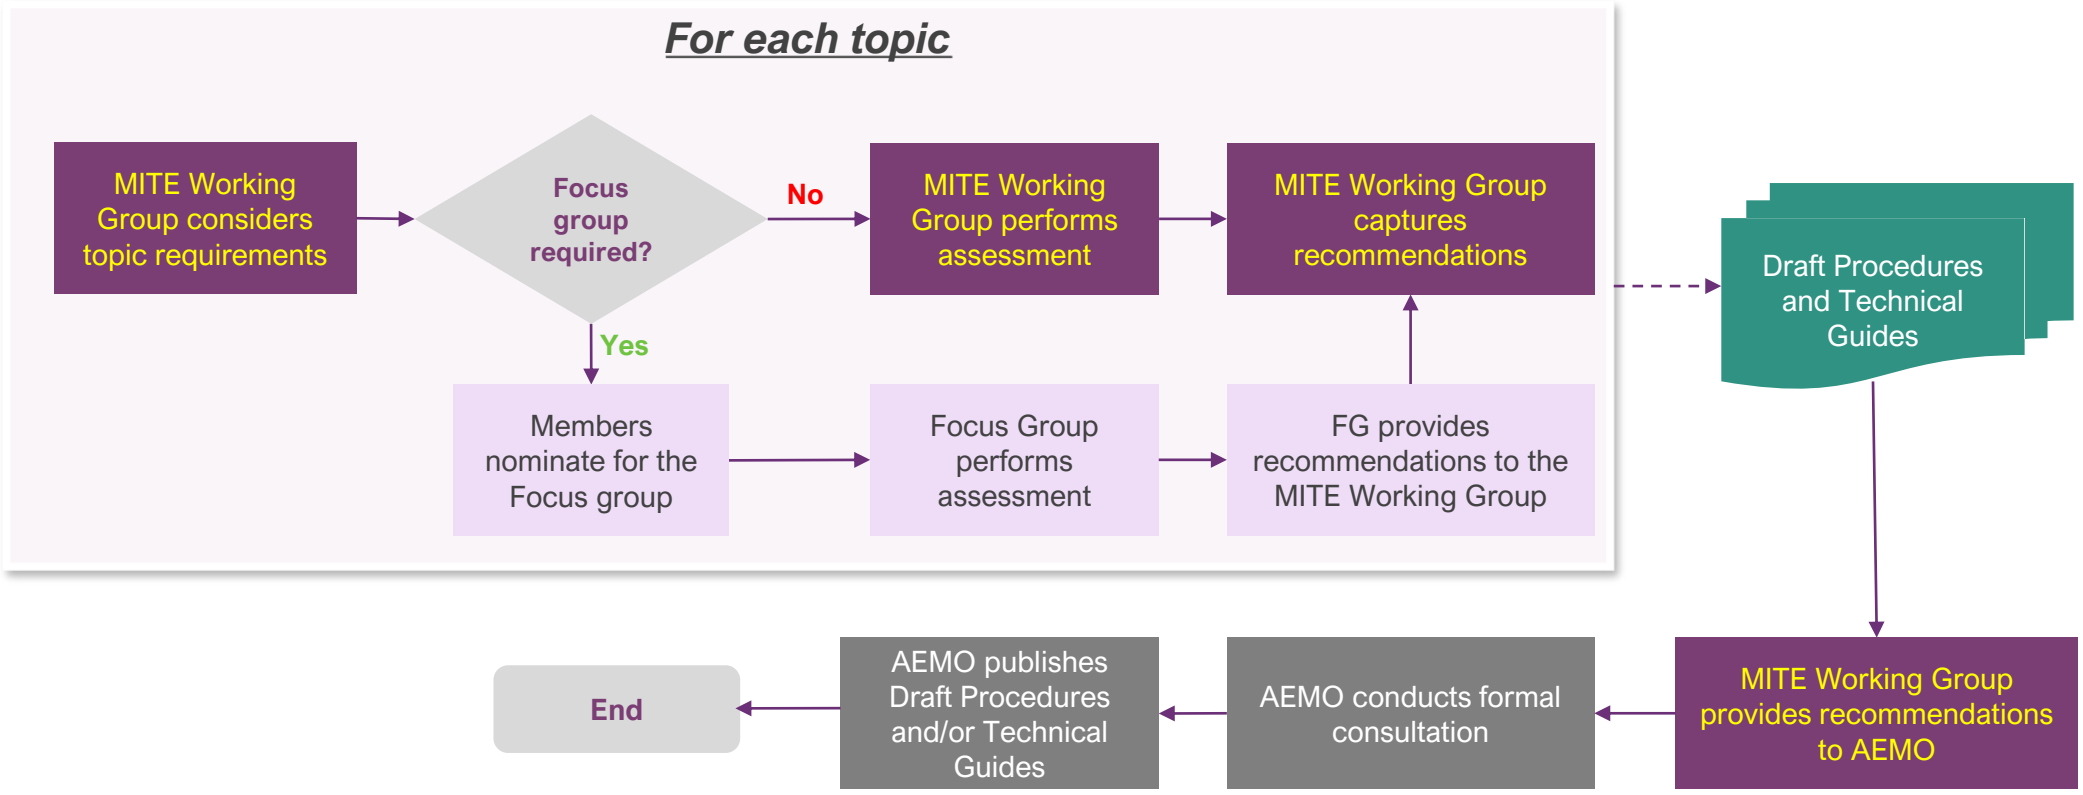

MITE Working Group

- **Actively participate** in highly technical workshop discussions to assess options, co-design draft deliverables.
- **Review key drafts** of documentation prepared by the Focus Group.
- **Consult** internally within own organisation to test, socialise and ultimately champion.

Focus Group (as required)

- **Co-design** draft deliverables for consultation with working group members
- **Actively participate** in the Focus Group workshops and activities
- **Participate** in **highly technical discussions**, including engaging within their business prior, to provide detailed responses to matters under discussion
- **Champion** technical discussions with their peers and within own organisations

Forward Plan

- 3 WG meetings left for 2024 (30 Oct, 27 Nov (IDX only) and 12 Dec)

- 2 Focus Group sessions have been conducted, outcomes/recommendations to be shared today
  - IDX: Decision Tree Focus Group and IDAM: API Authentication & Authorisation Patterns

- 8 focus groups scheduled for 2024
  - IDX
    - ~~9 October –~~ New date TBD - Asynchronous and Event Notification Group
    - WC 4 November - Large Share File
    - WC 11 November - Payloads
    - WC 18 November - Low Volume Interface
    - WC 2 December - AEMO Gateway Software
    - WC 9 Dec - Inquiry Platform
  - IDAM
    - 18 October - Organisation Hierarchy
    - 14 November - Security Compliance Obligations

- WG and FG learnings to-date?

- Call for nominations for November and December FGs will be sent out shortly

Strawman Timeline for upcoming IDAM / IDX Sessions

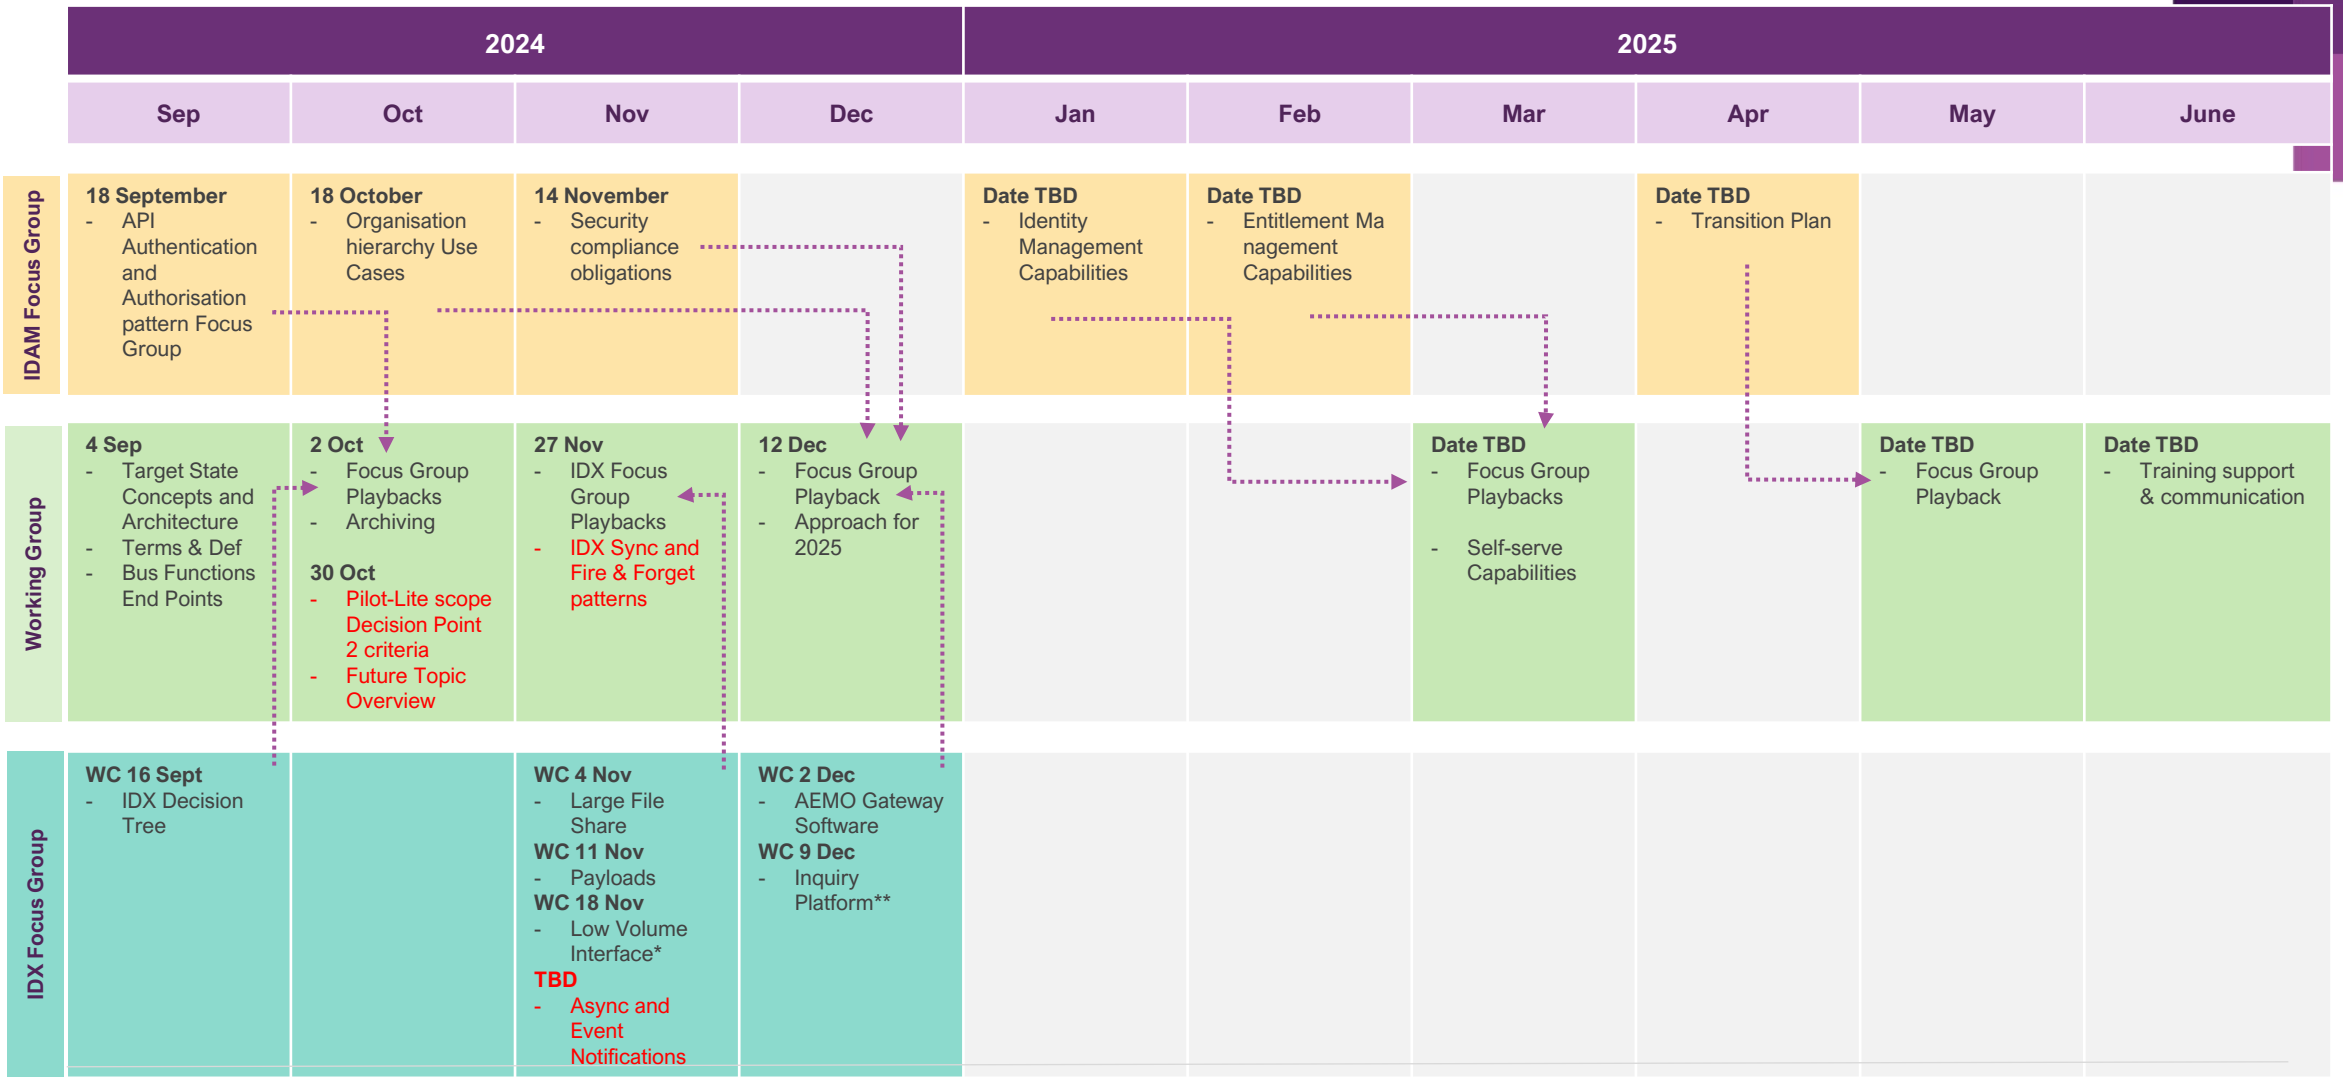

*These proposed dates are indicative dates

** Potential to be deferred dependent upon Focus Group content and feedback

Notes

- Blaine Miner provided the WG an update regarding the MITE's forward plan

- Blaine mentioned that the Asynchronous and Event Notification Group focus group meeting, previously scheduled for the 9 Oct, has been deferred to allow for additional time to consider and prepare the required content
    - The new date and time is still to be determined

- As a result of the Asynchronous and Event Notification Group focus group meeting being deferred, the WG scheduled for:
    - 30 Oct will now focus on the IDX Pilot-Lite scope and the IDX Decision Point 2 criteria
    - 27 Nov will now focus on the IDX Sync and Fire & Forget patterns

- Blaine also mentioned that call for nominations for the November and December FGs will be sent out shortly
    - Note, the call for nominations was circulated to the MITE coordinators on Monday 7 Oct, with nominations closing Monday 14 Oct

Identity and Access Management

3. Recap of IDAM API Authentication and authorisation Pattern

API Authentication and Authorisation Focus Group - Recap

This workshop covered:

- An overview of Service Accounts Authentication and Authorisation using OAuth 2 Client Credentials Flow
- Discussed proposed API Authentication and Authorisation Flows

This workshop did not cover:

- Usage of Federated Credentials for Service Accounts, this is to be covered as part of the Jan ‘25 FG Session.
- Authentication and Authorisation for other IDX channels (eg: Large File Share). These will be covered in future FG sessions.
- Finalising the approach for using Client Credential flow for API Authentication and Authorisation. This pattern will be finalised in the Feb '25 FG session.

Current Pain Points: API Authentication and Authorisation

IDAM for Market Interfaces refers to a technical framework and a set of processes that govern the management of digital identities and the control of access to AEMO's market systems. This IDAM platform is also a step towards compliance with the SOCI Act, ensuring that the access to critical energy infrastructure is tightly controlled and monitored.

- Multiple access controls to access AEMO systems
- Multiple user credentials are required to access AEMO systems
- Multiple Authentication patterns e.g., API keys, Basic Auth and OAuth used
- Inadequate capabilities for managing password rotation
- Lack of designation of account to a specific AEMO environment such as pre-production or production

API Authentication Options considered

AEMO had assessed four potential options for API Authentication.

It is assumed that mTLS using the Participant certificate is used for Transport Layer Security for all of the options below .

1. API Keys	API keys are a simple method for authenticating API access, where a client includes a key in requests. This option would be used with the Participant Certificate for Transport layer security.
2. Basic Auth	It is a method of HTTP authentication where the client sends a username and password encoded in Base64 as part of the request's Authorisation header, which the server verifies and either grant or deny access to the requested resource based on its validity. This option would be used with the Participant Certificate for Transport layer security.
3. mTLS (Mutual TLS)	Provides client authentication and secure communication between services by exchanging client and server certificates.
4. OAuth 2.x	Oauth 2.x allows applications (clients) to securely access resources, typically using access tokens which are short-lived minimising security risks and ensures that the credentials of the resource owner remain safe and protected. This option would be used with the Participant Certificate for Transport layer security.

API Authentication Options considered

Based on feedback from the focus group, AEMO has assessed the feasibility of using HMAC-based authentication, which uses a cryptographic hash function and a shared secret key to verify message integrity, in place of the currently proposed OAuth based approach.

Criteria	Current Patterns			Proposed Pattern	FG Feedback
	mTLS + API Keys	mTLS + Basic Authentication	mTLS	mTLS + OAuth	**HMAC**
Security	More secure than just mTLS due to the usage of two factors for Authentication (Participant Certificate and API Keys)	More secure than just mTLS due to the usage of two factors for Authentication (Participant Certificate and Service Account credentials)	Less secure due to only using one factor (Client Certificate) for Authentication	Highly secure, token-based with fine-grained control and short-lived tokens.	Provides integrity but vulnerable if secrets are leaked
Token Expiration & Revocation	Manual only	Manual or TTL based	Revocation of the Certificate and issuing a new one requires manual steps	Tokens are short-lived and revocable. Easy to manage lifecycles.	No built-in support for expiration or revocation
Short-lived Granular Access Control	Not supported	Not supported	Not supported	Allows precise short-lived scoped access for client apps, supports least privilege flows.	Not supported natively, requires external system
Ability to support increasing number of market participants	Limited	Scalable	Limited Scalability due to high PKI Infrastructure costs	Highly scalable in dynamic environments, integrates with identity providers.	Becomes complex as the number of secrets increases
Auditing & Monitoring	Limited monitoring infrastructure	Good, but limited monitoring infrastructure	Good, but limited monitoring infrastructure	Comprehensive logging, easy to monitor and auditing capabilities.	Basic logging, lacks detailed auditing features

While HMAC offers simplicity and message integrity, it lacks the comprehensive features required to handle the complexity and security needs of modern B2B interactions. For this specific use case, where secure, scalable, and fine-grained access control is essential, OAuth 2.x Client Credential Flow is the more appropriate and suitable choice.

Proposed Recommendation – OAuth 2.x

Focus Group Recommendation

OAuth is ideal for API security because it provides token-based authentication, enabling fine-grained access control with scopes for specific permissions. It supports short-lived tokens, reducing the risk of long-term credential exposure.

It allows:

- **External Participant Access to APIs**: OAuth facilitates granting external participant limited access to APIs, without sharing the user's credentials.
- **Token-based Authentication**: Rather than relying on traditional credentials (username/password), OAuth uses tokens that the API can validate.
- **Secure Authorisation Flows**: OAuth provides different authorisation flows based on the client's nature (web, mobile, or machine-to-machine), ensuring secure API access for each use case.
- **Fine-grained Access Control**: API providers can control what resources are accessible and for how long by issuing tokens with specific scopes and lifetimes.

OAuth flows and their suitability for API Security

Authorisation Code Flow (PKCE)	The Authorization Code Grant with PKCE is a more secure way to handle OAuth 2.0 authorization for public clients that cannot store client secrets securely. This is relevant for interactive sessions involving users.
Client Credentials Flow	The Client Credential Flow in OAuth is specifically designed for Machine-to-Machine (M2M) applications where a system or application needs to authenticate and access resources without any user involvement. *AEMO recommends the use of Client Credential Flow (detailed in subsequent slides).*
Resource Owner Password Flow	The Resource Owner Password Credentials (ROPC) Grant is used by highly trusted clients, where a client (application) can directly obtain an access token by collecting the user's credentials (username and password) and sending them to the authorization server. This is deprecated in OAuth2.1 due to the inherent risk associated with credential security
Implicit Flow	The Implicit Flow is designed primarily for public clients, such as single-page applications (SPAs) or mobile apps, that cannot securely store a client secret. This is deprecated in OAuth 2.1 due to the inherent risk associated with access token exposure and its vulnerability to potential attacks like cross site scripting (XSS)

Focus Group Recommendation

Proposed OAuth Credential Types

Aspect	ClientID + Secret	ClientID + Certificate	ClientID + mTLS
Authentication Method	Shared secret	Signed certificate	Mutual TLS
Credential strength	Low (shared secret)	High (asymmetric keys)	Very High (mutual authentication)
Resistance to Credential Theft	Low (passwords can be stolen)	Medium (certificates are more secure)	High (private keys are harder to compromise)
Proof of Possession	No	Yes	Yes
Data Protection	In transit only	In transit and at rest	In transit and at rest
Scalability	Limited	Good (with automation)	Excellent (with automation)
Implementation Cost	Low (no PKI needed)	Medium (requires certificate management)	High (requires full mTLS infrastructure)
Implementation Complexity	Low	Medium	High
Rotation Complexity	Low	Medium	High
Revocation Complexity	Simple (invalidate secret)	Complex (CRL/OCSP)	Very Complex (CRL/OCSP + token binding)
Administrative Overhead	High (secret rotation and management)	Medium (certificate lifecycle management)	High (client/server cert management)

Focus Group Assessment

The focus group provides an interim recommendation to use of Client ID and Secret as the preferred credential type. This approach strikes an optimal balance between security, cost, and simplicity, ensuring secure authentication for applications while maintaining manageable implementation complexity. AEMO is currently exploring a more secure version of the Client Credentials flow, in alignment with RFC 8705, through a series of Proof of Concepts. The findings from this exploration will be presented to the February 2025 Focus Group for further discussion and finalisation.

Overview: API Authentication and Authorisation Pattern

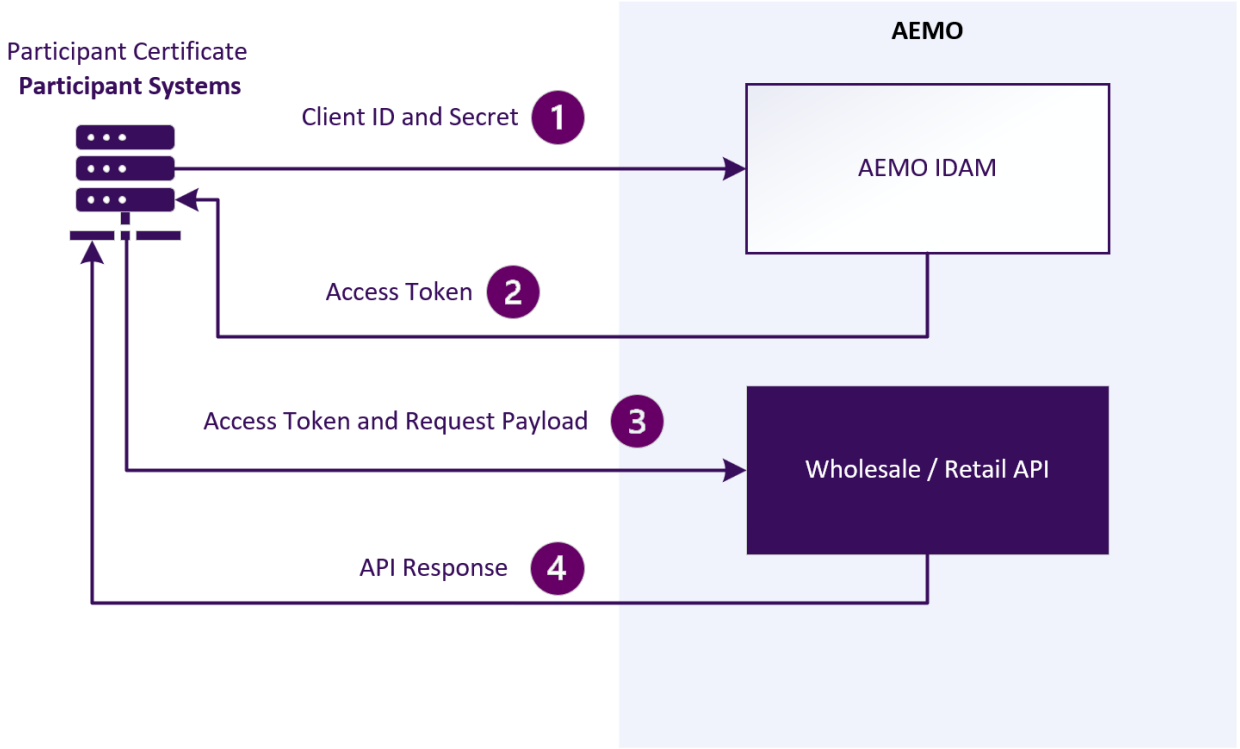

*** Client Secrets based flow shown here is only shown for illustrative purposes, AEMO is looking into the feasibility of implementing RFC 8705 based Client Authentication*

Client Authentication (With Participant Certificate):

In addition to using the client ID and client secret, the client (e.g., a Participant System) must present a **Participant certificate** during the token request. The communication between the client and the IDAM system occurs over a **mutual TLS (mTLS)** connection.

Token Request (Using mTLS):

1. The client initiates a POST request to AEMO IDAM's token endpoint over a mutual TLS connection using the Participant Certificate.
   The client presents the **client ID**, **client secret** (using the Basic Authorisation Scheme).

2. Clients specifies the API **scopes** that it need access to (in the message body). Clients will be able to request scopes or scope bundles which are mapped to a specific collection of Rights.

3. The **Participant's access token** is used to verify the client's identity.

4. The API returns the response back to the Participant.

Optional parameters for filtering scopes

For organizations with multiple participant identifiers, a potential use case for IDX is a common Service Account with access to multiple Participant IDs requesting permission to perform a specific action on an IDX business function, such as "Bidding" or "Metering."

By using optional scope parameters to reduce / filter the scope of the feature, Participants would have the flexibility to request access to an IDX business function, either for a specific Participant ID, or a combination of the Participant IDs that the account has been granted access to.

For instance, say, a particular Service Account has been granted access to the "Bidding" business function for **three** Participant IDs in NEM - pid_123, pid_234, and pid_345, by the PA. When requesting access to "Read bids“, the Participant would have the flexibility to specify that only **two** pids - pid_123 and pid_234 need to be included in the request, as illustrated in the example on the right.

Similarly, other filter criteria can be defined using additional optional parameters. **The full list of these parameters, along with guidance on their use, will be published by AEMO**. These optional parameters can be included either in the request body, as demonstrated in the example, or as part of the request query string.

Token Request Format

POST /oauth2/token HTTP/1.1

Headers
Host: identity.aemo.com.au (TBD)
Content-Type: application/x-www-form-urlencoded
Authorisation: Basic base64(client_id:client_secret)

Body (Form Parameters):
grant_type=client_credentials
scope=bids:r (TBD)
market=nem (TBD)
pids=pid_123 pid_234 (Optional parameter)

In the proposed approach, the use of optional parameters to limit or filter the scope of the feature will allow participants to apply various criteria to refine access levels to the resources they wish to request, tailored to their specific use cases.

AEMO is seeking participant feedback on the access token scoping capabilities needed to meet participant use cases specific to both functional scope and ParticipantID context.

Proposed Access Tokens Types: Opaque

For Step 2 (refer to previous slide)

	Transparent Tokens	Opaque Tokens
Token Structure, Readability & size	• The payload includes user data, roles, and claims, making it accessible to clients and servers without needing additional validation. • The payload is readable and contains three parts: Header, Payload, and Signature. • The token size is typically larger due to the included claims and signature	• The payload is a random string that holds no inherent meaning to the client. • The payload will not expose any server-side information to the client. • The token size is typically smaller compared to JWT tokens.
Revocation Handling	Transparent Tokens once issued are stateless, meaning they cannot be easily revoked unless additional mechanisms like a revocation list or short expiry times are implemented.	Easier to revoke in real-time since each use requires checking with the authorisation server, which can immediately invalidate the token if necessary.
Security	While the contents are signed and encrypted, they are often exposed to the client, making them more vulnerable to potential leaks or misuse if improperly secured (e.g., stored unsafely). However, cryptographic signatures ensure data integrity.	Since they reveal no internal data, even if an opaque token is intercepted, it is of little use without validation by the authorisation server. This makes it less prone to misuse if exposed to a client or third party.
Future extensibility	Transparent tokens (like JWT) expose the token's structure to clients. This can lead to client applications building dependencies on the token format. As a result, if AEMO modifies the token format in the future, it may impact market participants' systems that rely on specific details of the token's structure.	Opaque tokens prevents the client applications from building a dependency on the token format. This would allow AEMO to modify the token format in the future without affecting market participants, as participant systems will not be reliant on the token's format.

Focus Group Assessment

Focus group members agreed that **opaque access tokens** should be used for the production runtime environment to enhance security and prevent the exposure of sensitive information. Additionally, the group recommended that AEMO consider using transparent tokens in development environments, where they could offer easier debugging and monitoring capabilities without compromising the integrity of the production system. AEMO will evaluate the feasibility of providing this capability or an equivalent function in Development / Sandbox environments.

Key Recommendations: Voting Poll

- AEMO seeks in-principle agreement with the working group on the recommended Pattern
- A consensus position is sought with the industry. Where there are diverging views, these will be acknowledged.

Key Recommendation
OAuth Client Credential Flow: OAuth client credential flow is the recommended method for IDX API authentication and authorisation, with additional consideration given to HMAC-based authentication.

AEMO welcomes any feedback from the MITE Working Group to NEMReform@aemo.com.au.

Notes

Sivaraj Ganesan spoke to the section and did a playback of the recommendation from the Focus group regarding API Authentication and Authorisation Patterns.

The following questions were raised:

1. **Slide 20**: Is there a Secret Rotation Policy in place? Is that Mandatory Requirement? **AEMO Response** Security Best practice is to rotate only the Secret. If there is an automated way that this can be done via API then we can expose the API through a mechanism which can rotate the credentials. AEMO will provide guidelines around these in the future sessions.

2. **Slide 23:** Are there other better ways which could offer easier debugging rather than providing the Transparent Tokens in the non-production environment and then different token in the production environment which is without modifying AuthN and AuthZ. **AEMO Response** AEMO will assess other options and will evaluate the feasibility of the providing the capability of easier debugging and monitoring or an equivalent function in Development/Sandbox environment.

3. **Slide 23**: Opaque tokens in production and capability around debugging? **AEMO Response**: One of the options which we also discussed in the Focus group was that we could provide an API endpoint which you could call with your opaque token and the response on that API endpoint would give you the more detailed breakdown of the set of scopes and claims that were embedded in that token as a facility to do debugging in production without necessarily changing the overall configuration of production. So that is one of the suggestions that AEMO will be looking at that as a potential capability.

4. Can an organisation have distinct service accounts? **AEMO Response** So the definition of service accounts is 100% in the hands of participants, you will have the choice if you're an organisation that has multiple participant IDs or has a unique organisational structure. You can create one master service account that does as much as you'd like. You can create an individual service account for every single organisational element for every single business function. If that's the way you choose to operate, AEMO deliberately want to provide that flexibility to participants, particularly because there are very significant ranges of participant size involved in this process.

Key Recommendation:
OAuth 2.X Client Credential Flow is the recommended method for IDX API authentication and authorisation.

4.IDAM API Authentication and Authorisation Focus Group Actions

API Authentication and Authorisation Focus Group Feedback and Actions

AEMO consulted focus group members in regard to the API Authentication and Authorisation Pattern for machine-to-machine interaction . The following is a summary of the discussion

1	In line with the focus group's recommendations, AEMO will re-evaluate the process for **Certificate Validation,** including exploring the use of third-party certificate validation.	Third party validation of certificates are already performed by an organisation named DigiCert
2	Based on feedback from the focus group, AEMO will assess the feasibility of using **HMAC-based authentication**, to verify message integrity, as an alternative to the current OAuth recommendation.	AEMO has assessed the feasibility of using HMAC based authentication. The assessment has concluded that using other OAuth 2 Credentials would be a better fit for this use case.
3	The focus group recommended providing participants with the flexibility to configure the Client ID across three options—common credentials across all markets and functions, credentials grouped by energy markets and functions, or credentials specific to each business function—without enforced validation by AEMO.	The focus group recommended allowing participants to configure the client ID, without enforced validations by AEMO. AEMO is assessing the security implications of one of the options, with outcomes to be discussed and finalised at the February 2025 Focus Group.
4	AEMO to overlay the **CDR API authentication** pattern onto the current authentication patterns assessment included in the pack and assess any differences.	The AEMO has conducted a review of the Consumer Data Right (CDR) API authentication patterns and has determined that the proposed IDX IDAM patterns will enhance the future-proofing of API authentication and authorisation strategies.

Action 1: Certificate Attestation Process

- AEMO is using DigiCert as its certificate issuing authority for issuing Participant certificates.
- The root certificates being used are issued by DigiCert as OV certificates, meaning they undergo third-party validations to ensure their authenticity.
- These private certificates are part of a private chain, which isn't automatically pushed out to browsers. As a result, for these certificates to be trusted, the chain of trust needs to be distributed in advance to ensure the certificates are recognised and trusted by external organisations. While the chain of trust needs pre-distribution, the Online Certificate Status Protocol (OCSP) and Certificate Revocation Lists (CRL) remain fully verifiable externally.
- While there is an inherent trust in AEMO, third-party validation adds an additional layer of assurance.
- Having an external authority like DigiCert verify the certificates, ensures that the certificates are not only from AEMO but have also undergone a rigorous validation process.
- This independent verification increases the level of trust and confidence for organisations relying on AEMO's certificates, providing greater assurance of their authenticity and integrity.

Action 2: HMAC Based Authentication

Based on feedback from the focus group, AEMO has assessed the feasibility of using HMAC-based authentication, which uses a cryptographic hash function and a shared secret key to verify message integrity, in place of the currently proposed OAuth based approach.

Criteria	Current Patterns			Proposed Pattern	FG Feedback
	mTLS + API Keys	mTLS + Basic Authentication	mTLS	mTLS + OAuth	**HMAC**
Security	More secure than just mTLS due to the usage of two factors for Authentication (Participant Certificate and API Keys)	More secure than just mTLS due to the usage of two factors for Authentication (Participant Certificate and Service Account credentials)	Less secure due to only using one factor (Client Certificate) for Authentication	Highly secure, token-based with fine-grained control and short-lived tokens.	Provides integrity but vulnerable if secrets are leaked
Token Expiration & Revocation	Manual only	Manual or TTL based	Revocation of the Certificate and issuing a new one requires manual steps	Tokens are short-lived and revocable. Easy to manage lifecycles.	No built-in support for expiration or revocation
Short-lived Granular Access Control	Not supported	Not supported	Not supported	Allows precise short-lived scoped access for client apps, supports least privilege flows.	Not supported natively, requires external system
Ability to support increasing number of market participants	Limited	Scalable	Limited Scalability due to high PKI Infrastructure costs	Highly scalable in dynamic environments, integrates with identity providers.	Becomes complex as the number of secrets increases
Auditing & Monitoring	Limited monitoring infrastructure	Good, but limited monitoring infrastructure	Good, but limited monitoring infrastructure	Comprehensive logging, easy to monitor and auditing capabilities.	Basic logging, lacks detailed auditing features

While HMAC offers simplicity and message integrity, it lacks the comprehensive features required to handle the complexity and security needs of modern B2B interactions. For this specific use case, where secure, scalable, and fine-grained access control is essential, OAuth 2.x Client Credential Flow is the more appropriate and suitable choice.

Action 3: Guidelines for Configuring Entitlements for Client Credentials

Dimensions	Option 1 - Common client credential across all markets, participants IDs and business functions	Option 2 - Client credentials grouped around energy markets, and related business functions	Option 3 - Client credential per business function
Security risk	High – client credentials do not provide any restriction across energy markets, participant IDs and unrelated business functions	Low – client credentials cannot be exploited across markets nor unrelated business functions	Very Low – credentials limited to discrete functions and scopes
Operational impact	Low – single common (shared) client credential	Moderate – more client credentials, logically grouping related functions	High – significant number of client credentials to be deployed and maintained
Assessment	Risk outweighs operational impact	Most balanced approach, contains risk without requiring excess proliferation of client credentials	Option remains available, however may not be justified for all participants given operational impact vs limited improvement in security risk

AEMO Assessment

The focus group has recommended offering participants the flexibility to configure the client ID using one of the three options: common credentials across all markets and functions, credentials grouped by energy markets and functions, or credentials specific to each business function, without enforced validation by AEMO. AEMO is currently assessing the security implications of Option 1 and any necessary mitigations. The results of these explorations will be presented at the February 2025 Focus Group for further discussion and finalisation.

Action 4: CDR API Authentication Alignment

Dimensions	OAuth Client Credentials Flow with mTLS	CDR
Authentication Method	Uses mutual TLS (mTLS) with Participant Certificates for both client and server authentication along with Oauth 2.	Relies on private_key_jwt for client authentication. Uses signed JWT assertions.
Security Model	Ensures mutual authentication between client and server via certificates, reducing risk of man-in-the-middle attacks.	Secures API interactions using OAuth assertions. No mTLS required.
Complexity	Complexity arises from managing certificates (issuance, rotation, validation), but the token exchange is simple once the connection is secured.	More complex due to managing client metadata, OAuth assertions, scopes, and consent. Registration involves more configuration.
Dynamic Client Registration	Not typically designed for dynamic client registration. Clients are pre-registered, and tokens are exchanged securely via mTLS.	Supports dynamic registration of new clients, with detailed metadata such as client URIs, software roles, scopes, and other parameters.
Primary Use Case	Secure communication between client and server for API authentication where both need to verify each other.	Enables dynamic client registration and management of clients in a regulated ecosystem like open banking.

Conclusion

The OAuth Client Credentials Flow with mTLS is stronger for mutual authentication using certificates, which is crucial in environments needing enhanced security, such as energy markets. On the other hand, CDR approach is designed specifically for the Australian open data environment, focusing more on dynamic registration, metadata management, and integrating OAuth into broader compliance frameworks like open banking.

IDAM API Authentication and Authorisation Summary

Key Recommendation	Outcome
OAuth Client Credential Flow: OAuth client credential flow is the recommended method for IDX API authentication and authorisation, with additional consideration given to HMAC-based authentication.	Agreed
Client Credential Type: The preferred credential type is a combination of Client ID and Secret.	In Review
Opaque Access Tokens: Opaque access tokens will be used for production to improve security.	Preferred
Client ID Flexibility: Participants prefer flexibility in configuring Client IDs across three options (markets and functions) without enforced validation by AEMO.	In Review
Optional Parameters for filtering scope for refining access levels to requested resources	In Review

Notes

Sivaraj Ganesan addressed the team and provided an explanation of the actions and the reasoning:

1. **Action 1: Slide 28**: AEMO is using DigiCert as its certificate issuing authority for issuing Participant certificates. (Action is closed)

2. **Action 2: Slide 29:** AEMO has assessed the feasibility of using HMAC-based authentication, which uses a cryptographic hash function and a shared secret key to verify message integrity, in place of the currently proposed OAuth based approach. While HMAC offers simplicity and message integrity, it lacks the comprehensive features required to handle the complexity and security needs of modern B2B interactions. For this specific use case, where secure, scalable, and fine-grained access control is essential, OAuth 2.x Client Credential Flow is the more appropriate and suitable choice. (Action is closed).

3. **Action 3: Slide 30:** AEMO is currently assessing the security implications of Option 1(Common client credential across all markets, participants IDs and business functions) and any necessary mitigations. The results of these explorations will be presented at the February 2025 Focus Group for further discussion and finalisation. (Action remains Open).

4. **Action 4: Slide 31** The OAuth Client Credentials Flow with mTLS is stronger for mutual authentication using certificates, which is crucial in environments needing enhanced security, such as energy markets. On the other hand, CDR approach is designed specifically for the Australian open data environment, focusing more on dynamic registration, metadata management, and integrating OAuth into broader compliance frameworks like open banking (Action is closed).

5. IDAM Future Topics

IDAM Focus Group: Organisation Hierarchies and Enhanced Data Sharing

18th October

The objective of this focus group meeting is to thoroughly explore the technical capabilities introduced as part of the new Organisation Hierarchy and Enhanced Data Sharing Capabilities.

In the target state for Entitlements Management AEMO proposed a new set of capabilities known as Organisation Hierarchy and Enhanced Data Sharing, which will enable participants to more accurately model their corporate group structures within AEMO systems. These capabilities will facilitate the management of complex data sharing and user entitlements scenarios in a structured manner, including those involving third-party service providers.

The focus group will discuss the application flows and would seek inputs on design decisions regarding the Organisational Hierarchies and Data Sharing Framework.

Audience Skill Set for Focus Group Discussion

- Business leads who coordinate multiple PIDs arrangements
- Service Provider leads who support data sharing configurations
- Technical leads / Architects supporting multiple PID solutions and data sharing solutions

Topics for discussion

- New and enhanced Entitlements Management capabilities due to the Organisation Hierarchies
- Enhanced Data Sharing Framework.

Note: This focus group discussion will be relevant to the organisations with multiple participants Ids and all stakeholders who leverage data sharing capabilities.

6. Industry Data Exchange Actions

IDX Actions

Description	Responsible	Status	Comments
IDX Business Function End Points: AEMO to create a slide(s) to demonstrate the NEM Retail use case for both Option A and Option B business Endpoint Options.	AEMO	Closed	Slides added into this pack for review.
IDX: Participants asked if AEMO can provide a link to the AEMO Gateway Software	AEMO	Closed	This link was provided in the NOTE from the 4th Sept IDX Working group Target State Proposal for Technical Focus Groups (aemo.com.au)

Notes

Blaine spoke to the IDX actions, no comments or questions were raised.

7. IDX Decision Tree Focus Group Playback

Sri Gundu

Focus Group held on the 17th September

Objective of the Focus Group

The MITE FG met to assess the Decision tree topic.

This focus group aims to..

- Re-Cap the pain points the decision tree is designed to alleviate / address
- Review and discuss drafted Definitions and Objectives
- Review and discuss drafted Decision Tree principles
- Review and discuss drafted Decision Tree process
- Review and discuss drafted material provided by AEMO that will feed into the Decision Tree Drafting (including use cases)

The ask of participants...

- **Participate** in highly technical discussions, including engaging within their business prior, to provide detailed responses to matters under discussion
- **Champion** technical discussions with their peers and within own organisations.
- **Review** draft documentation prepared by the Focus Group and provide input.

Decision Tree Pain Points

Pain Points	Proposed Principle(s)	Target State Concept
Industry raised pain-point: • Cost and complexity. • Lack of alternative data exchange mechanisms *AEMO's reading of Industry pain points:* • AEMO offers multiple patterns for the same regulated transactions, each with different infrastructure requirements. This creates unnecessary complexity. • Management multiple patterns, most of which have had **zero** uptake (B2BMessagingSync, B2BMessagingPull, B2MMessagingPull), has high ongoing operational and implementation costs for AEMO and, in turn, industry.	• For each use case, **a single channel and protocol** is to be offered.	• The IDX platform will offer multiple channels and protocols. However, for each specific use case, **an industry-agreed-upon decision tree** for data exchange will lead to the selection of a single channel and protocol.

Decision Tree– Definition and Objectives

Concept

- A **decision tree** is a decision support hierarchical model that uses a tree-like model of decisions and their possible consequences.
- The selection of a channel for an Interface will be facilitated by the decision tree based on the Interface attributes

Objectives

- All new use cases get assessed against the corresponding decision tree and will use the available channels. New channels are added only when the current channels in a decision tree do not cater for the requirements.
- For any given use case there will be only a single channel supported (outside of transitional arrangements).

Decision Tree– Principles

Principles

- All AEMO External Interfaces deployed in IDX Foundation are in scope and are governed by decision trees
- When using the decision tree for a particular use case, there should only be one outcome (the last hop in the decision tree)
- The decision tree must logically cover all identified use cases.
- The decision tree will provide visibility of capabilities available and not yet available for IDX Foundation.
- The protocol will be defined based on the channel.
- Payload will be determined by payload decision tree.*
- The decision tree will be maintained through the IDX Framework governance process.
- The decision tree only applies to system-to-system interactions
- Legacy Interface channels not aligned with decision tree will be assessed during Transition .

* Payload decision tree will be presented in the payload focus group

Decision tree process

Pattern definitions

Synchronous	An immediate business response is required to confirm a request has been received and processed.
Asynchronous	Additional processes or validations are required to enable a response to the request. For each interface required to support the process, decision tree is applied to determine a channel. It follows a multi-legged approach to deliver the business response.
Fire & Forget	Request message is sent with no expectation of receiving a business or technical (e.g. MACK) response message.
Inquiry-Flexible	The Interface data exchange consists of an on-demand service that allows for a dynamic subset of data elements to be returned.
Inquiry-Flat	The Interface data exchange consists of an on-demand service that allows for a fixed set of data elements to be returned.

Step 1: Determine message pattern

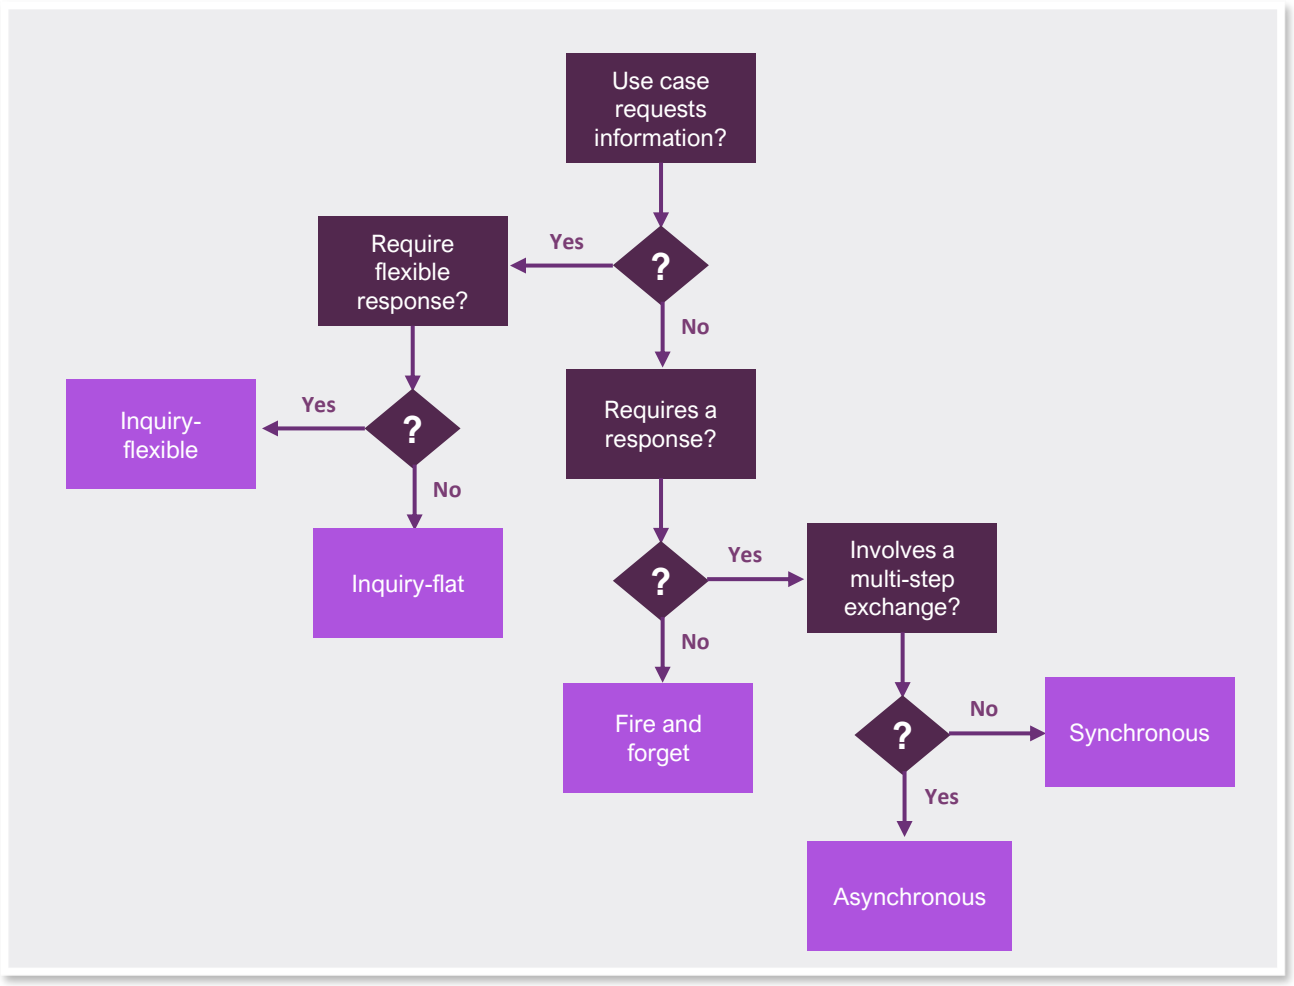

- **Decision Tree to illustrate the scenarios how a use case is assessed to determine which pattern it should follow**
- **For the decision tree to function this process must result in a consistent determination of pattern for like use cases**

Are there other patterns we should consider?
Are these the right questions to determine the patterns?

Message pattern worked example: Submit Bids

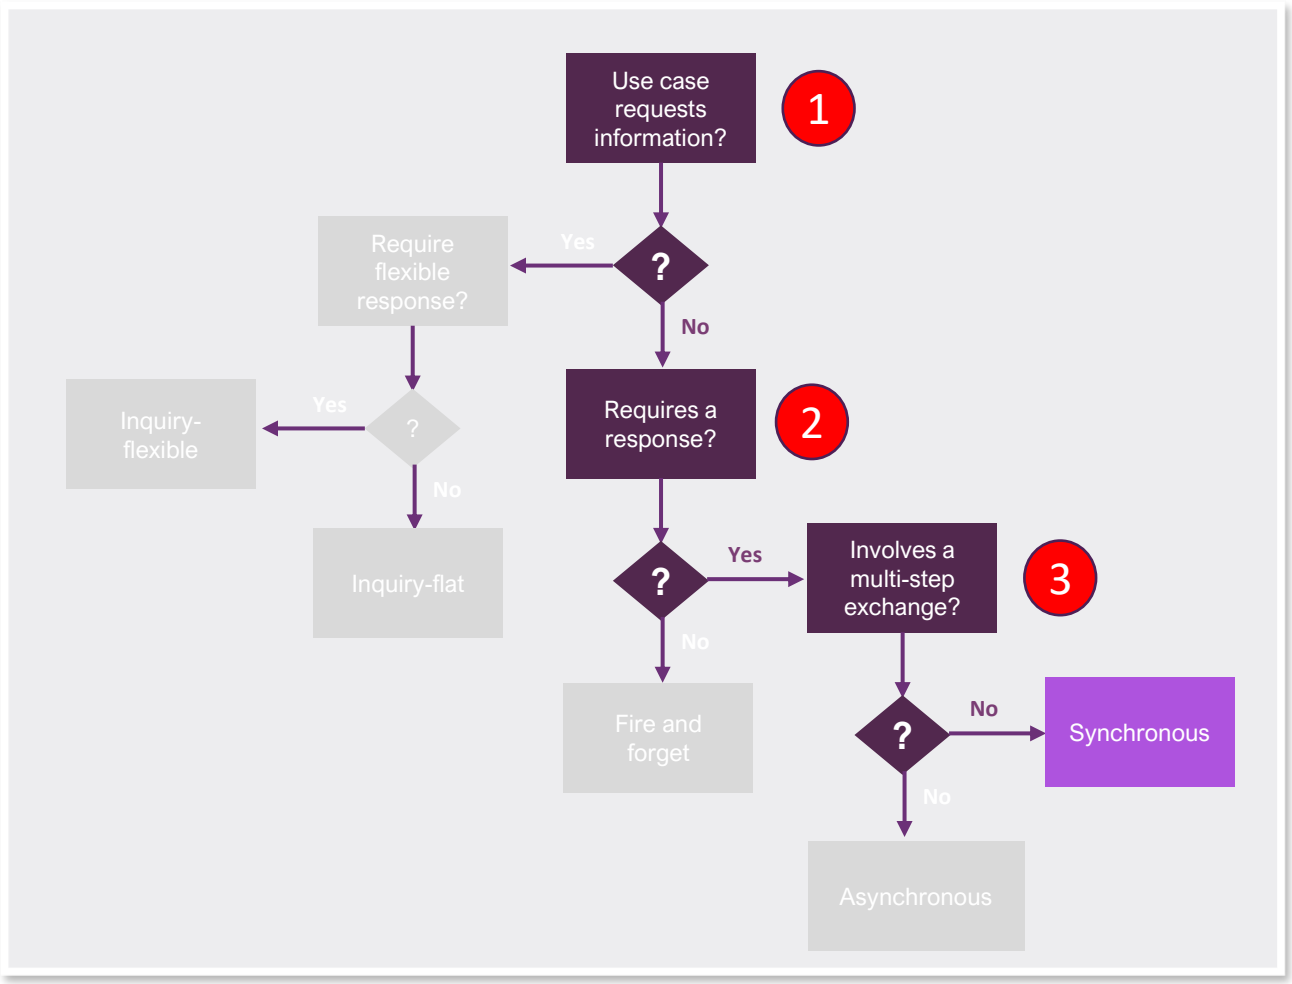

Use Case: Participant submitting Bids to AEMO

Decision tree applied criteria

1. Data (bids) is being submitted, not requested
2. A business response is required post processing the bid
3. The business response for bid submission is required immediately

Message pattern worked example: Retail B2B Service Orders

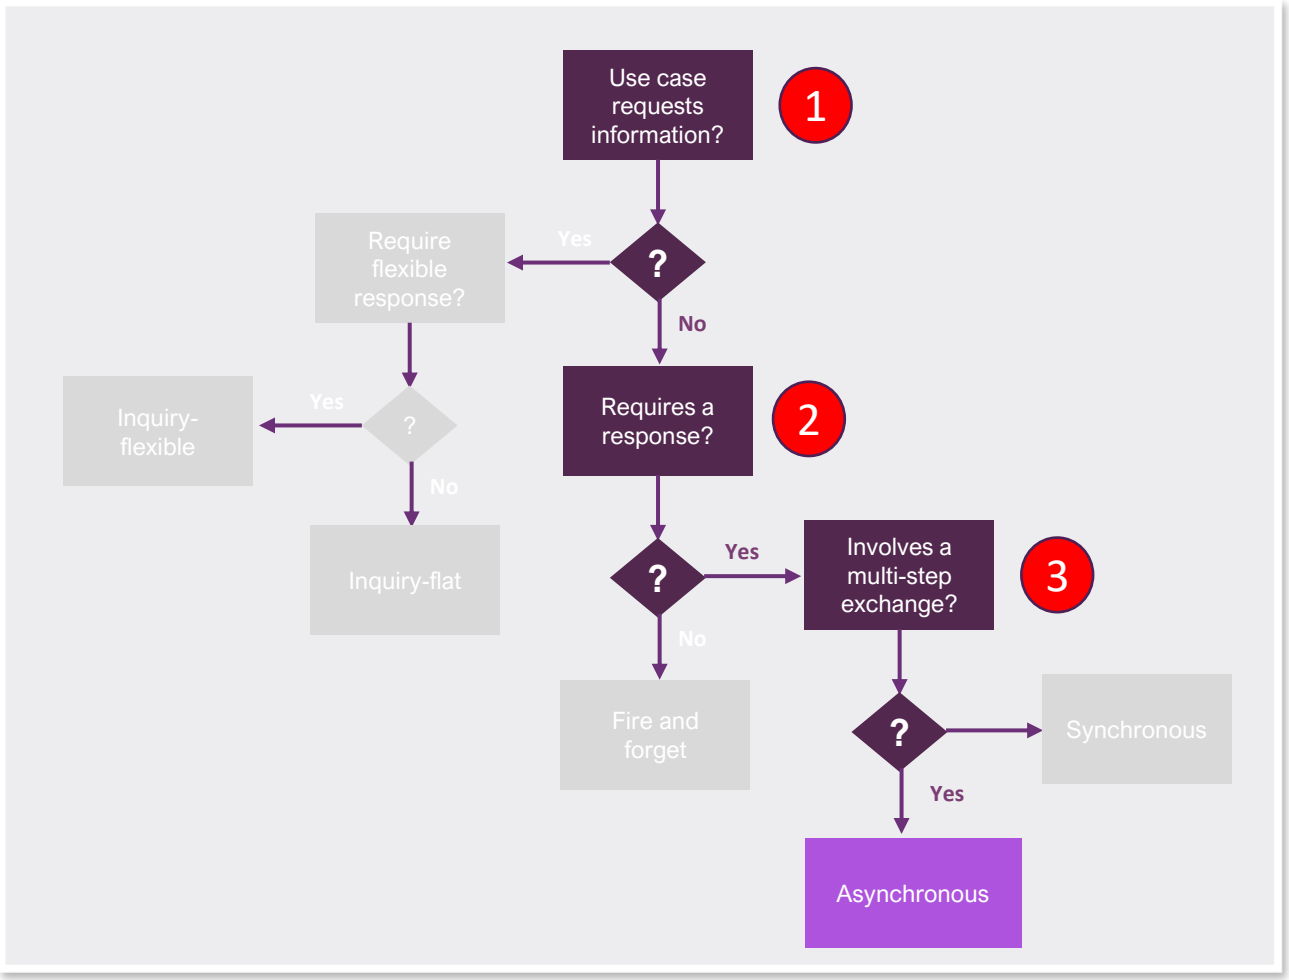

Use Case: Participant Service Order Request to AEMO

Decision tree applied criteria

1. Data (Service Order Request) is being submitted, not requested
2. A business response (Service Order Response) is required
3. The business response is not passed immediately, as it has further steps in the process to be completed i.e. the sequence of steps to complete Service Orders is Service Order Request -> Hub MACK -> Participant MACK -> TACK for Service Order Request -> Service Order Response -> Hub MACK -> Participant MACK -> TACK for Service Order Response. In summary, multiple data exchanges are required to close the business process

Message pattern worked example: Inbound Weather Data

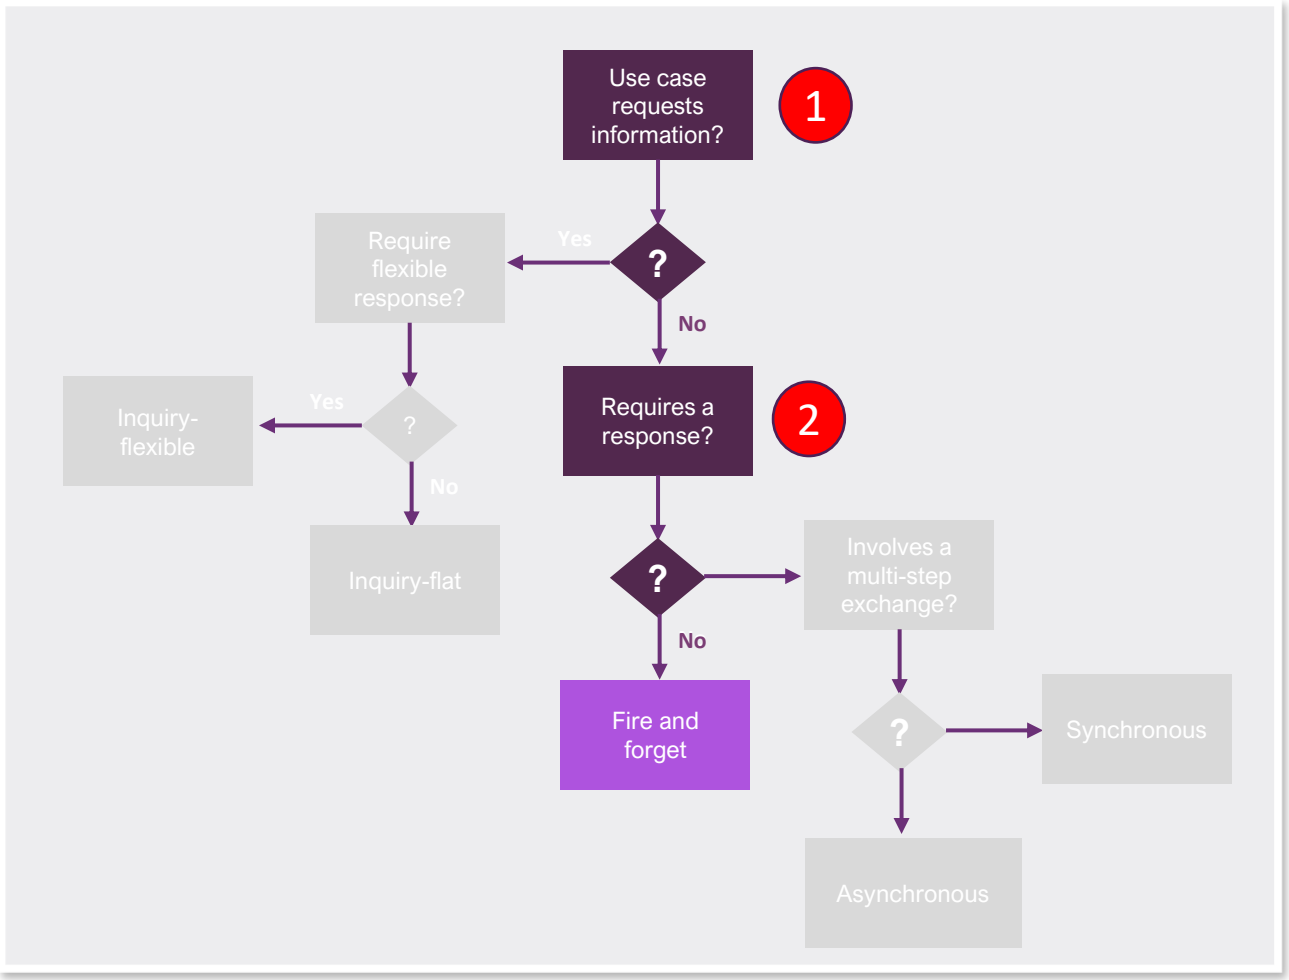

Use Case: BOM weather data inbound to AEMO

Decision tree applied criteria

1. **Data is being submitted, not requested**
2. **A response is not required, as there is no expectation of validation or acknowledgement for receiving the weather data .**

Message pattern worked example: Get Bids

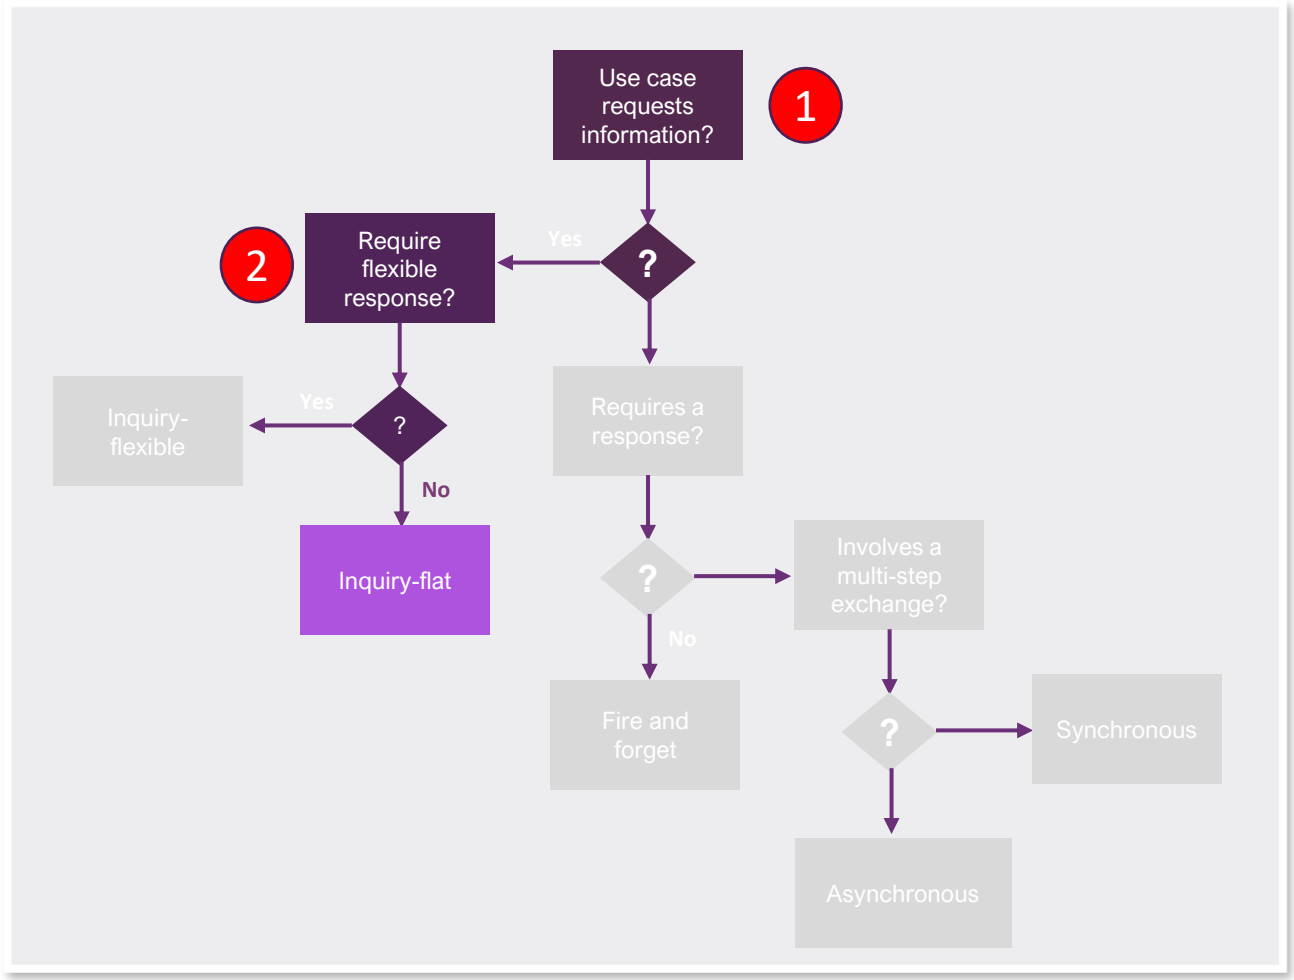

Use Case: Participant retrieves the list of submitted bids from AEMO (or) retrieves the details of a submitted bid.

Decision tree applied criteria

1. **Data is being requested i.e. retrieving the details of the submitted bid**
2. **The response structure is defined so it is not flexible for the consumer to define what attributes from the response schema are available in the response payload**

Message pattern worked example: NMI Discovery

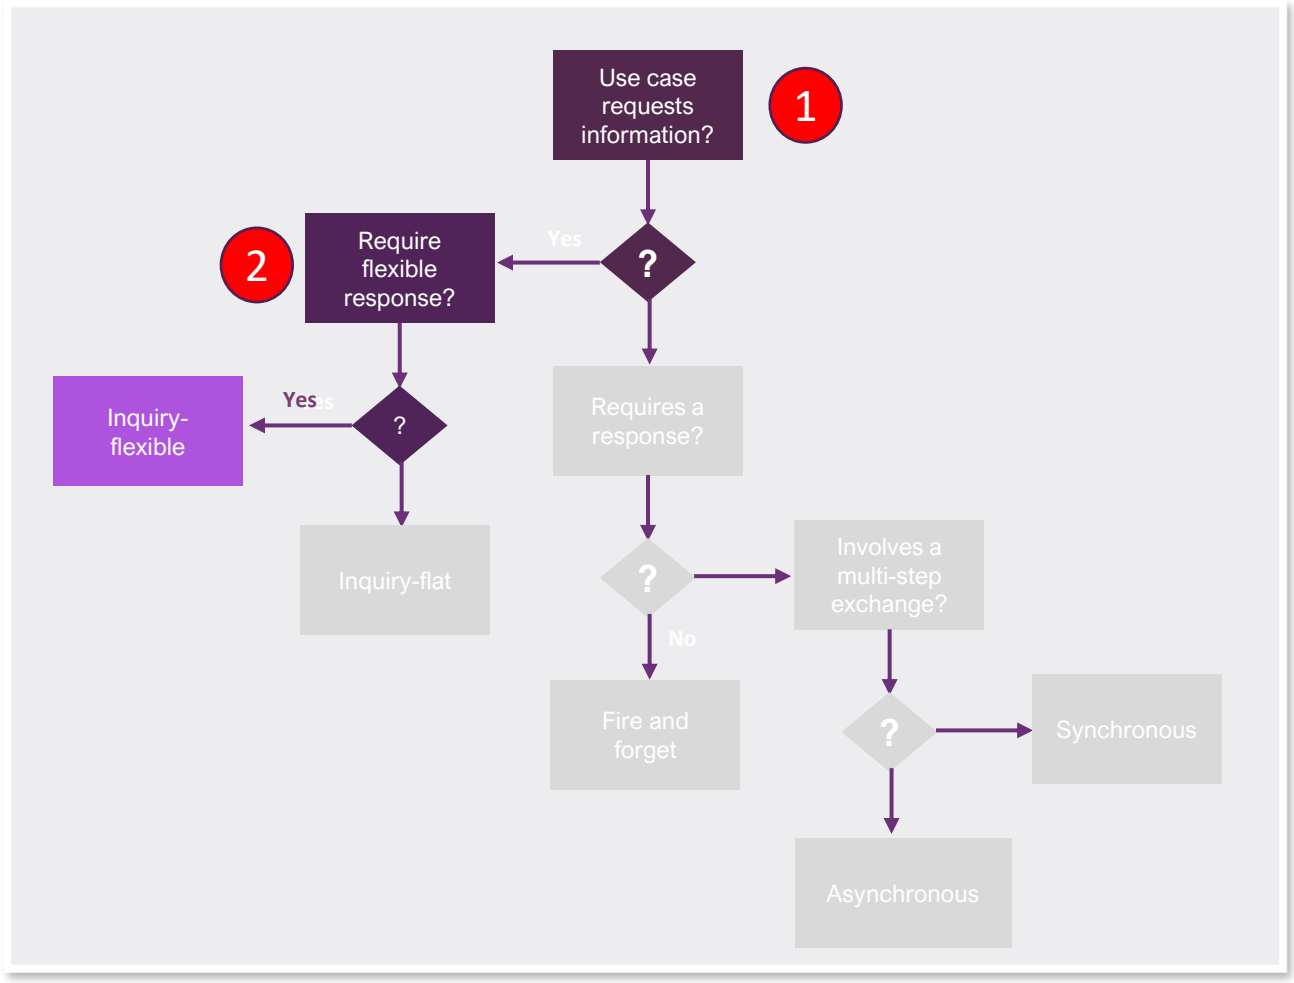

Use Case: NMI Discovery request to AEMO

Decision tree applied criteria

1. **Data is being requested in the form of 'NMI Discovery Request'**
2. **Participant defines the list of attributes from the response schema that must be sent in the response document i.e. defines the list of attributes that are to be sent in 'NMI Discovery Response'**

Terminology and Definitions

Message: A message is an exchange of data between sender(s) and recipient(s) that may contain one or more distinct transactions as shown in the example below
e.g. B2B SORD aseXML message

Transaction: A transaction is the technical realisation of a specific Business Document
e.g. <ServiceOrderRequest>, <ServiceOrderResponse>

Message & Transaction: 1:1

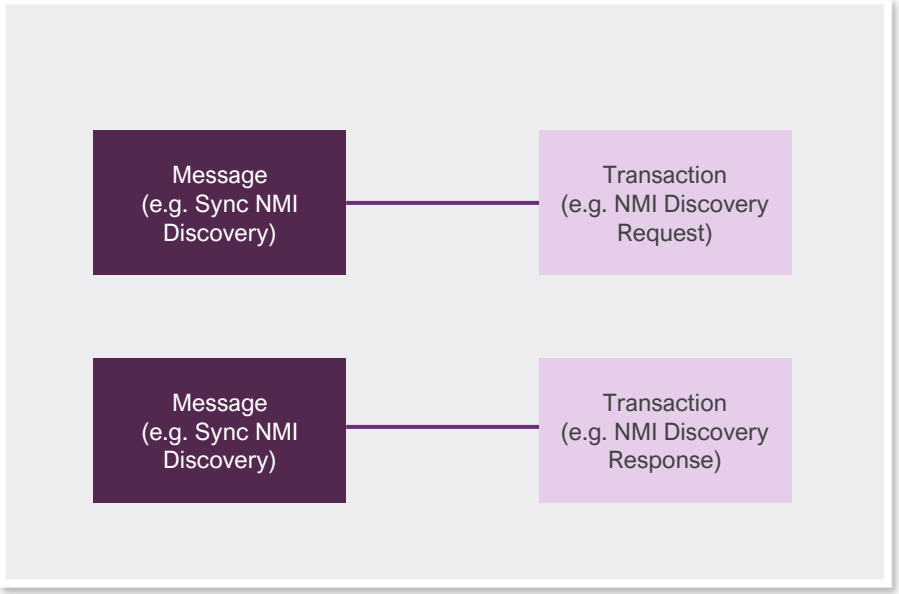

Message & Transaction: 1:Many

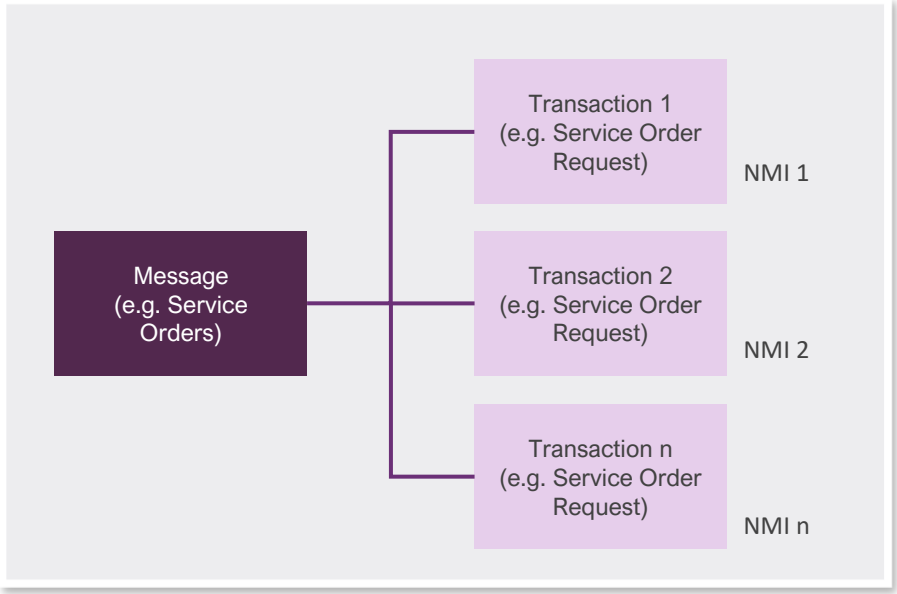

Message Compression & Bundling

The following guard-rails apply to inbound & outbound data exchange and request & response payloads. Compression could be enabled using 'Content-Encoding' & ' Accept-Encoding' header parameters via the RESTful API channel (or) 'zips' via the large file transfers.

Default: Compressed
e.g., Retail B2B messages, Retail B2M messages, Wholesale NEMReports, NEM Wholesale Bids

Exception: Consuming application does not support compressed payloads

In the rest of the section, decision tree(s) refer to 'Message Size'. The following guard-rails must be used when determining the message size
a. If the message is compressed; message size = compressed message size
b. If the message is not compressed; message size = uncompressed message size

Note: recommendation on compression algorithms (e.g. gzip, deflate) will be discussed in the upcoming focus groups (Async & Payload Focus Group)

The following table illustrates if message exchange patterns support bundling of transactions/records in the request & response payload

Pattern	Supports Transaction Bundling?	Example Use Case
Synchronous	✓	Bid submission (Inbound) Retrieve metadata of outbound hub queue (outbound)
Asynchronous	✓	Retail B2B Service Orders (Inbound & Outbound)
Fire & Forget	✓	Inbound Weather Data for forecasting (Inbound) Gas INT Reports (Outbound)
Inquiry with flexible payload formats (Sync)	✗	NMI Discovery (TBC with industry during DP2)

Message Bundling & Unbundling

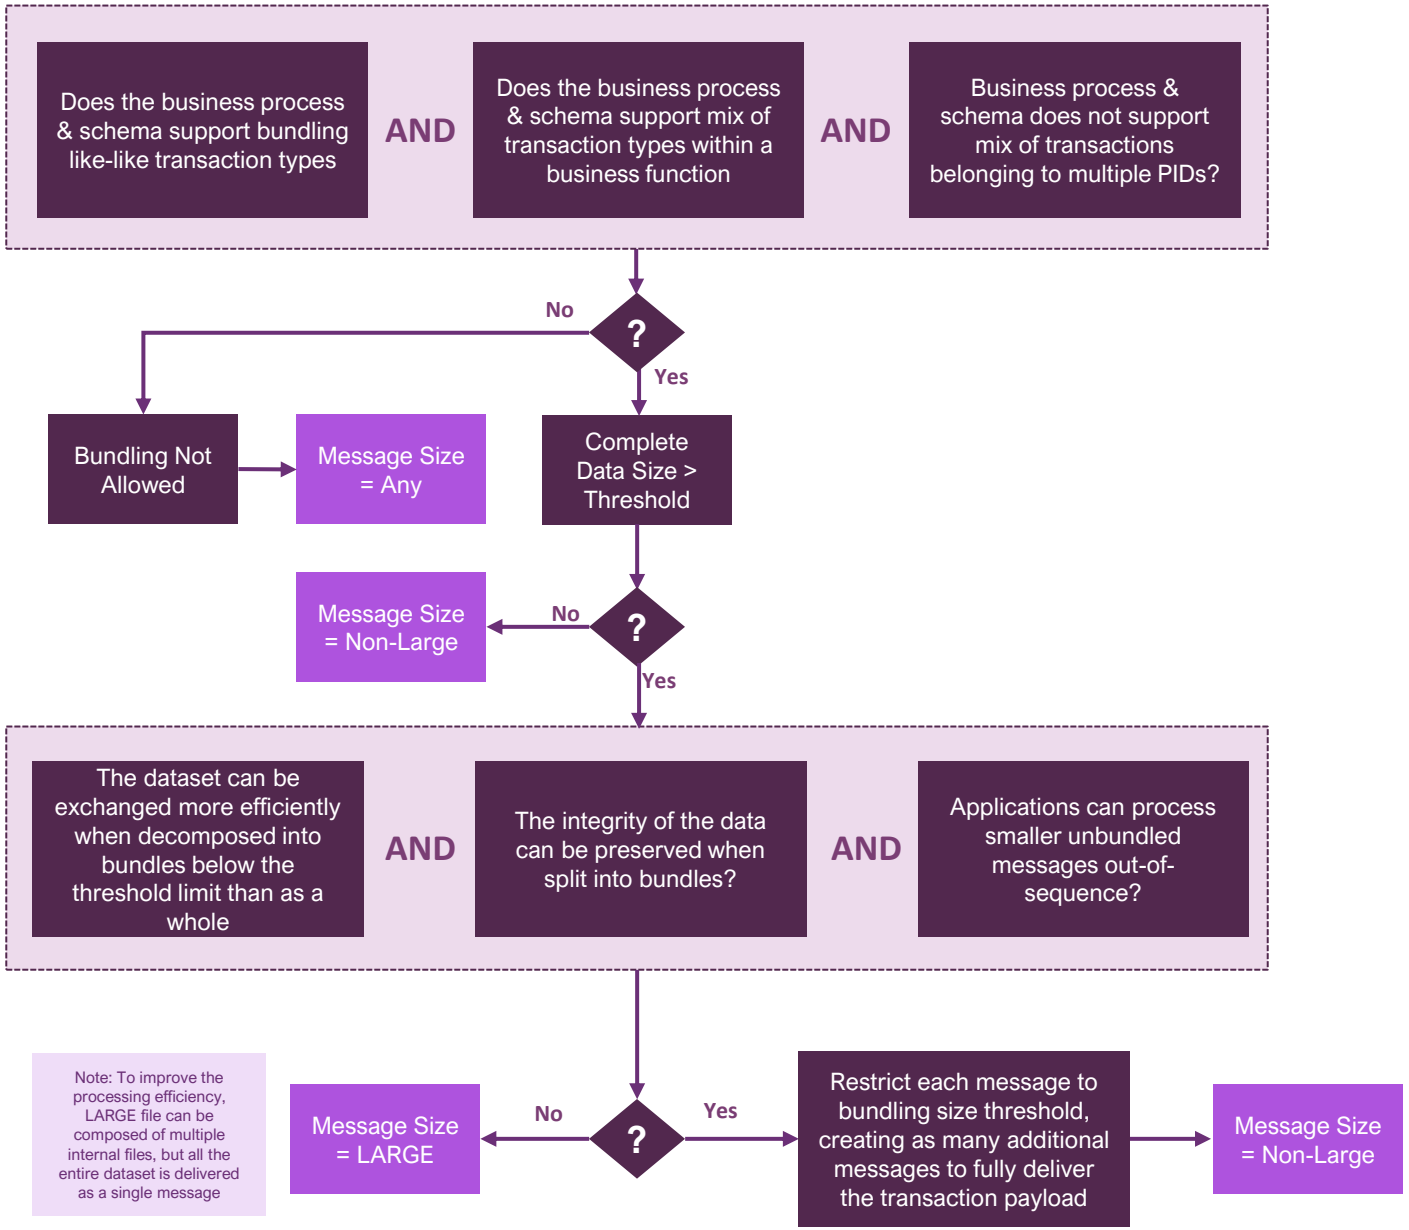

- **Decision Tree to illustrate the scenarios how bundling and un-bundling of transactions in a message are to be managed**
- **Decision Tree to illustrate the determination of message size – 'Large' vs 'Non-Large'. Refer Slide 23 on the definition of 'Threshold' value**
- **Best Practice: Bundling eligible transactions are time bound. Systems do not wait for a longer period to bundle the transactions**

Are there any other parameters that AEMO should consider when determining the decision of 'when to bundle and when to un-bundle'?

Message Bundling Worked Example: NMI Discovery Request-Response

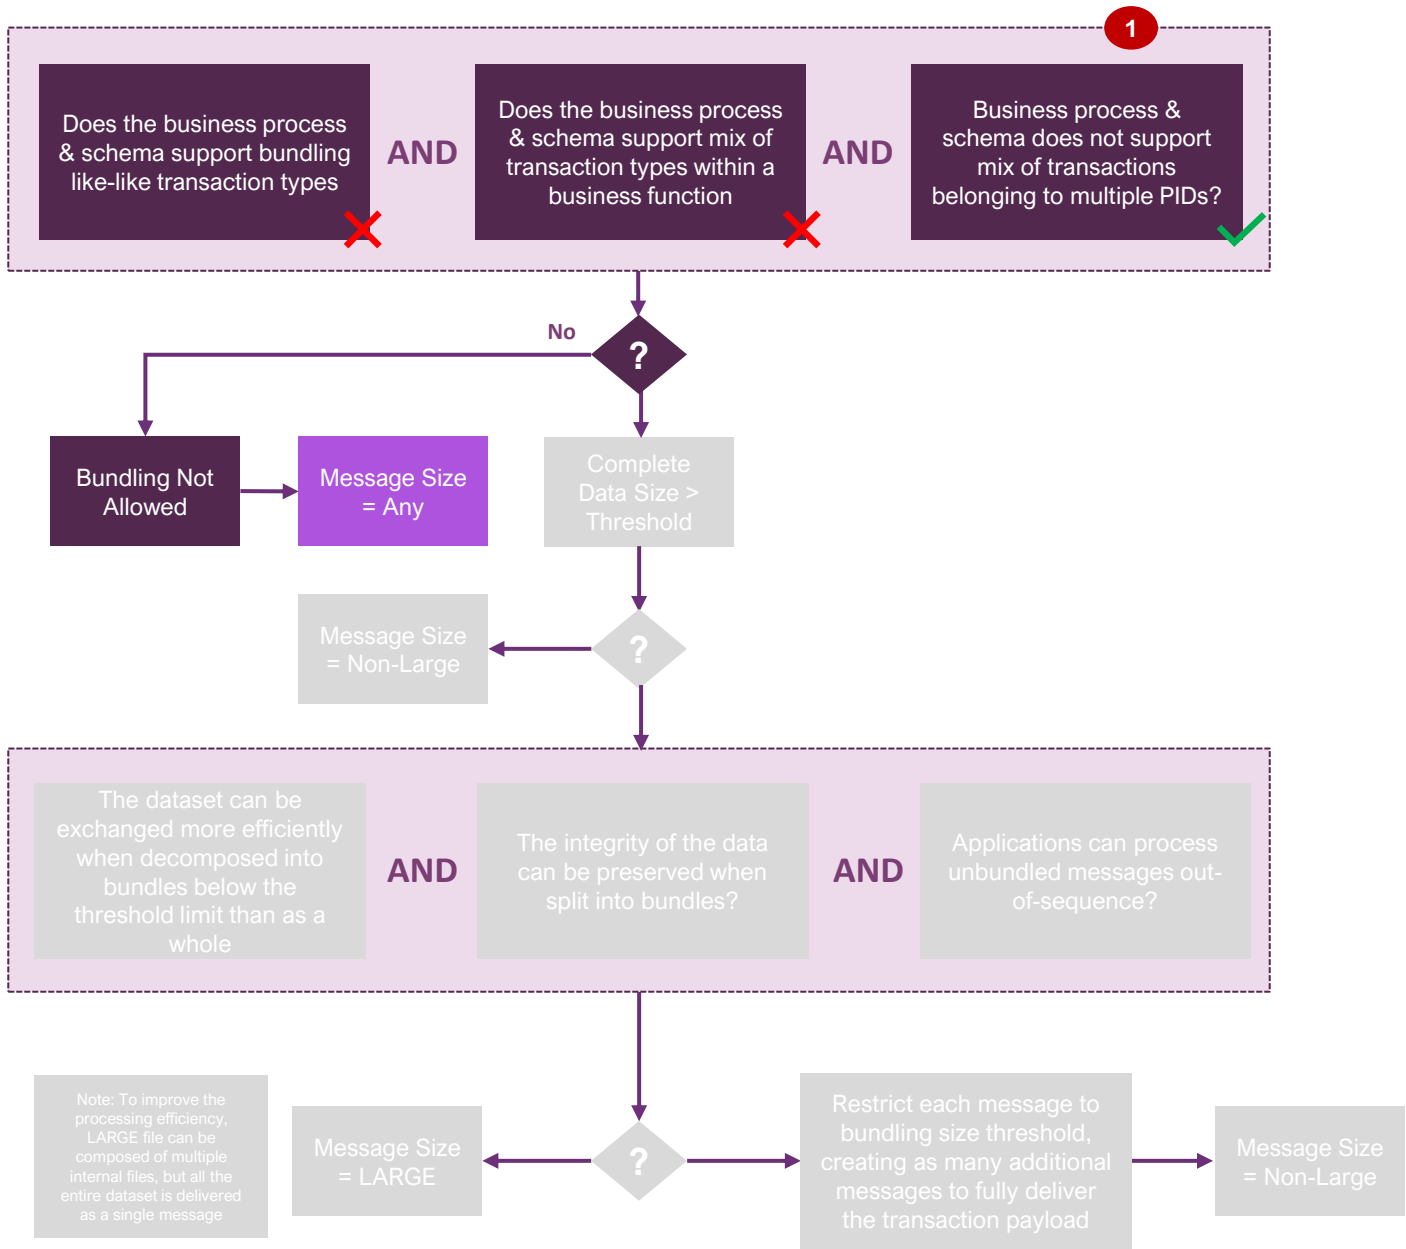

Use Case Description

NMI Discovery on Sync Mode: Participants invoke Sync NMI Discovery service wherein the NMI Discovery Response is provided in the blocking thread of the invocation

For illustration: Message size threshold = 1MB

Average Compressed Message size for NMI Discovery Response: ~100 KBs

Request Transaction: NMI Discovery Request
Response Transaction: NMI Discovery Response

Applying Decision Tree for the Use Case

1a. This is a sync inquiry service that only supports one transaction per message
1b. This is a sync inquiry service that doesn't support multiple transactions in a message
1c. A single discovery request is limited to a ParticipantID

Message Bundling Worked Example: NEMReports Next Day Public Reports

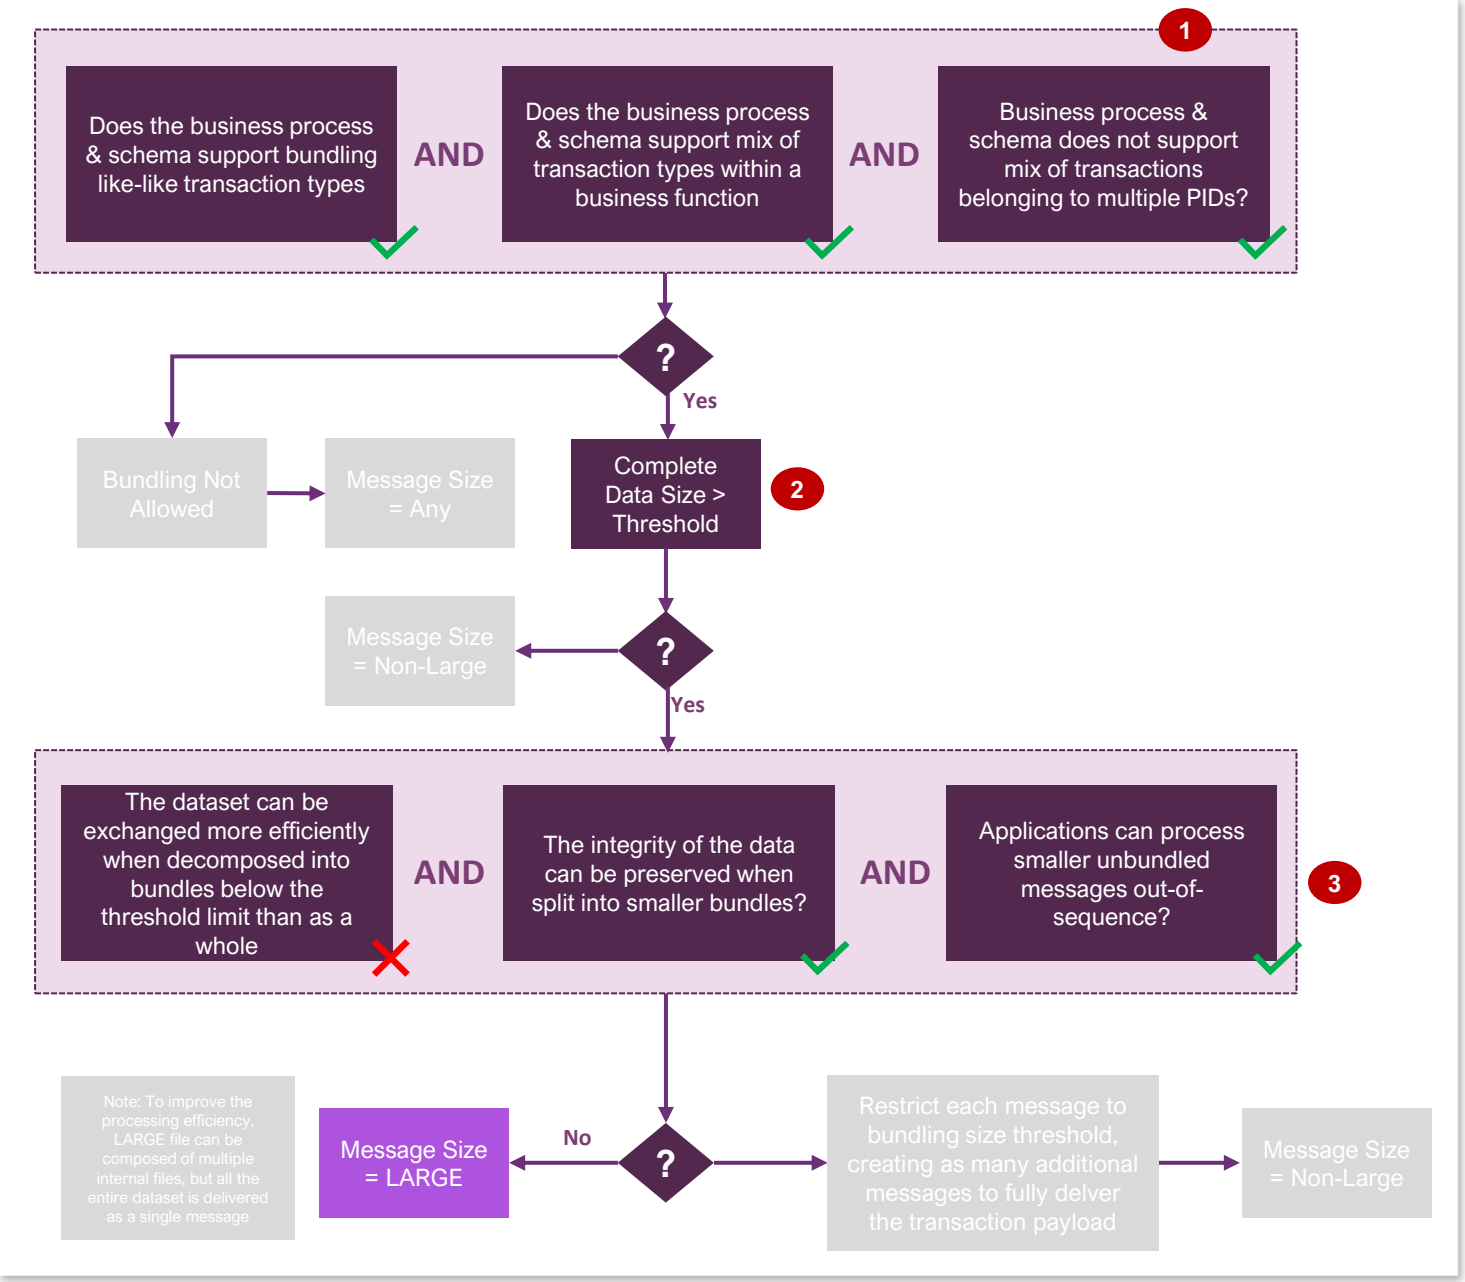

Use Case Description

Use Case: AEMO delivering Next Day Public Reports (Dispatch outputs) to Participants

For illustration: Message size threshold = 1MB

Report Name: Next_Day_Dispatch

Average Zipped Report Size: ~6MB

Next_Day_Dispatch has mix of multiple transaction types such as 'UNIT_SOLUTION', 'LOCAL_PRICE', 'OFFERTRK', 'CONSTRAINT", "MNSPBIDTRK'

Applying Decision Tree for the Use Case

1a. Next day public reports allow bundling of like-like transactions
1b. Next day public reports allow mix of transaction types (e.g. Unit_Solution & Local_Price in a single message)
1c. Public reports do not carry private ParticipantID specific confidential data

2. Complete Data Size (6MB) > Threshold (1MB for this example)

3a. Dataset is efficiently processed by Participants or pdrBatcher when the entire dataset is sent as one message
3b. Integrity of the data can be preserved when split into smaller bundles
3c. Applications can process smaller unbundled messages out-of-sync

Message Bundling Worked Example: Retail Snapshot Reports

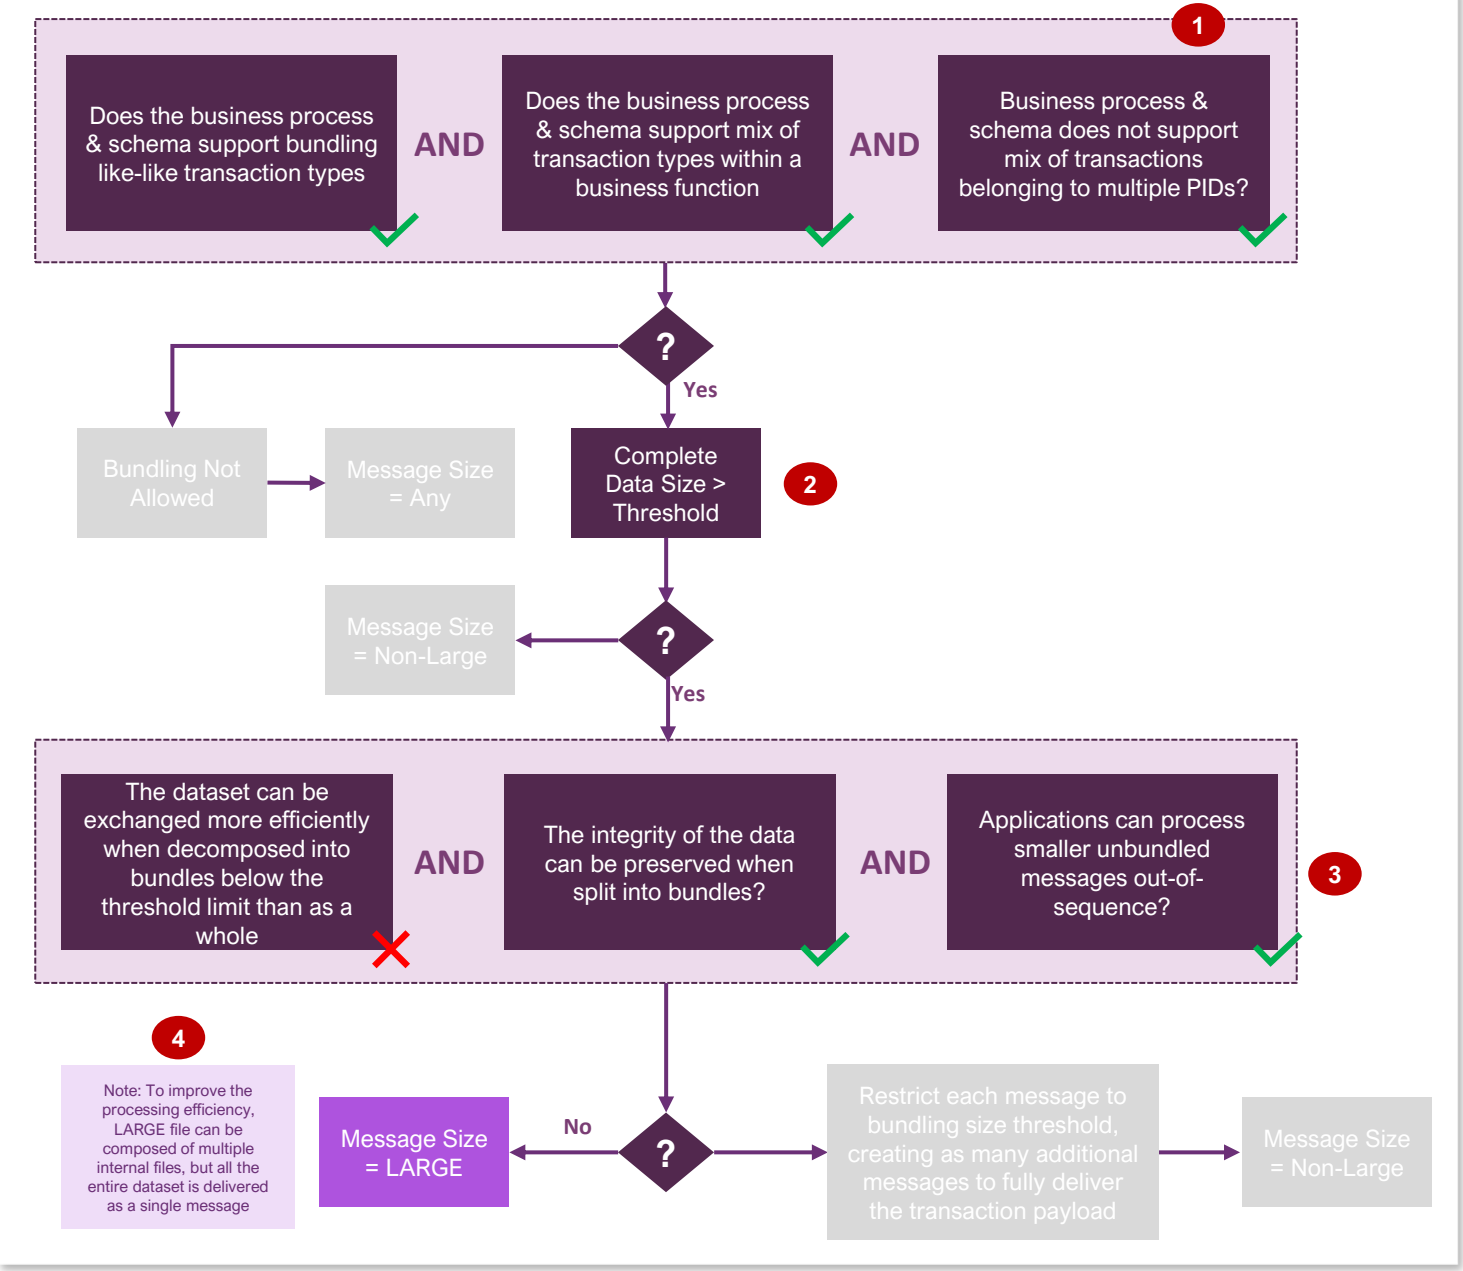

Use Case Description

Use Case: Monthly Report Snapshot Reports

For illustration: Message size threshold = 1MB

Report Name: Daily Snapshot Report

Average Zipped Report Size: ~100MB – 1GB

Payload has data from multiple Standing Data tables – CATS_NMI_DATA, CATS_NMI_PARTICIPANT_RELATIONS, CATS_METER_REGISTER, CATS_REGISTER_IDENTIFIER, CATS_NMI_DATA_STREAM

Applying Decision Tree for the Use Case

1a. Snapshot reports support bundling of like-like transactions
1b. Snapshot reports allow mix of transaction types (e.g. CND, CRI, CMR, CNDS & CNPR data) within a business function
1c. Snapshot reports are run for a ParticipantID

2. Complete Data Size (~100MB to 1GB) > Threshold (1MB for this example)

3a. Dataset is efficiently processed by Participants when the entire dataset is sent as one message
3b. Integrity of the data can be preserved when split into smaller bundles
3c. Applications can process smaller unbundled messages out-of-sync

4. The Large file can be composed of multiple internal files (say 25MB each) to improve the efficiency of processing these large files

Message Bundling Worked Example: Delivery of Meter Reads

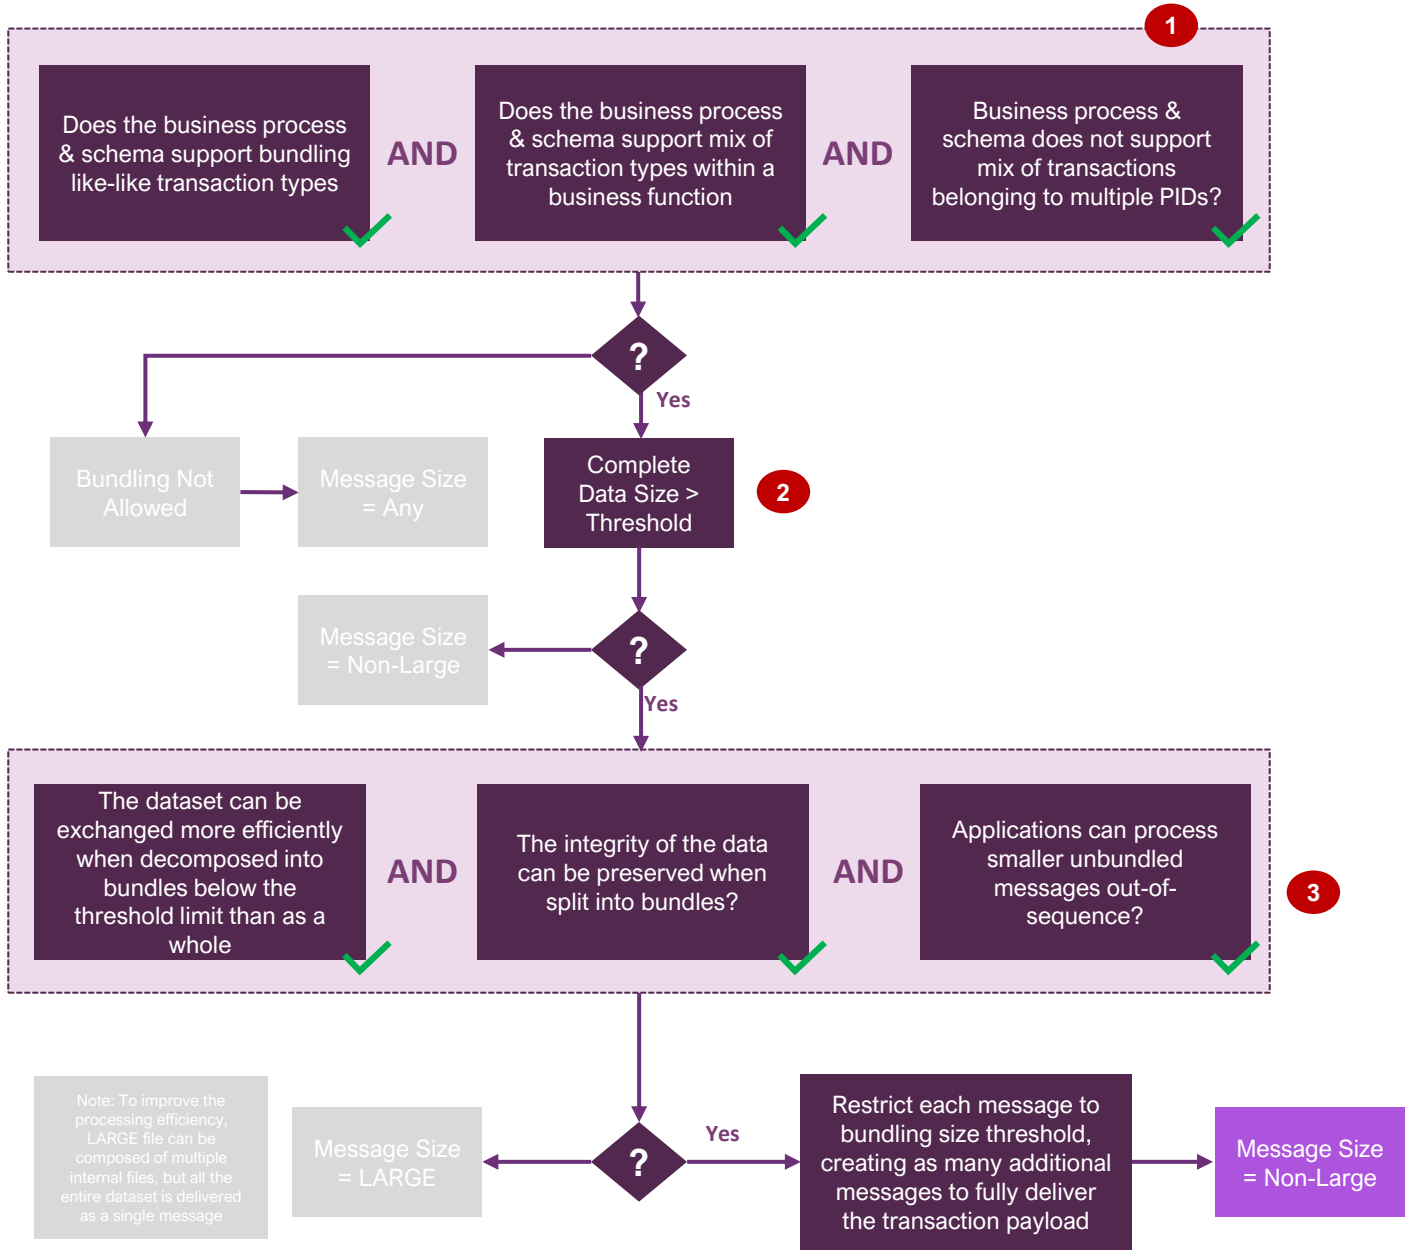

Use Case Description

Use Case: Delivery of Meter Read Data – MTRD or MDMT

For illustration: Message size threshold = 1MB

Transaction Name: Meter Data Notifications
Interval Meter Read Data: NEM12
Basic Meter Read Data: NEM13

Interval meter reads (AMI) are delivered to AEMO for all the NMIs that are owned by the MDP

Assume MDP sends interval meter reads for ~800,000 NMIs daily

Applying Decision Tree for the Use Case

1a. Meter Read Data support bundling of like-like transactions
1b. Meter Read Data allow mix of transaction types (e.g. interval and basic reads) within a business function
1c. Meter read dataset are to be delivered to a participant and/or AEMO

2. Complete Data Size (may be Gigs) > Threshold (1MB for this example)

3a. Dataset can be exchanged more efficiently when decomposed into bundles than as a whole
3b. Integrity of the data can be preserved when split into smaller bundles
3c. Applications can process smaller unbundled messages out-of-sync

Message Bundling Worked Example: Delivery of CATS Messages

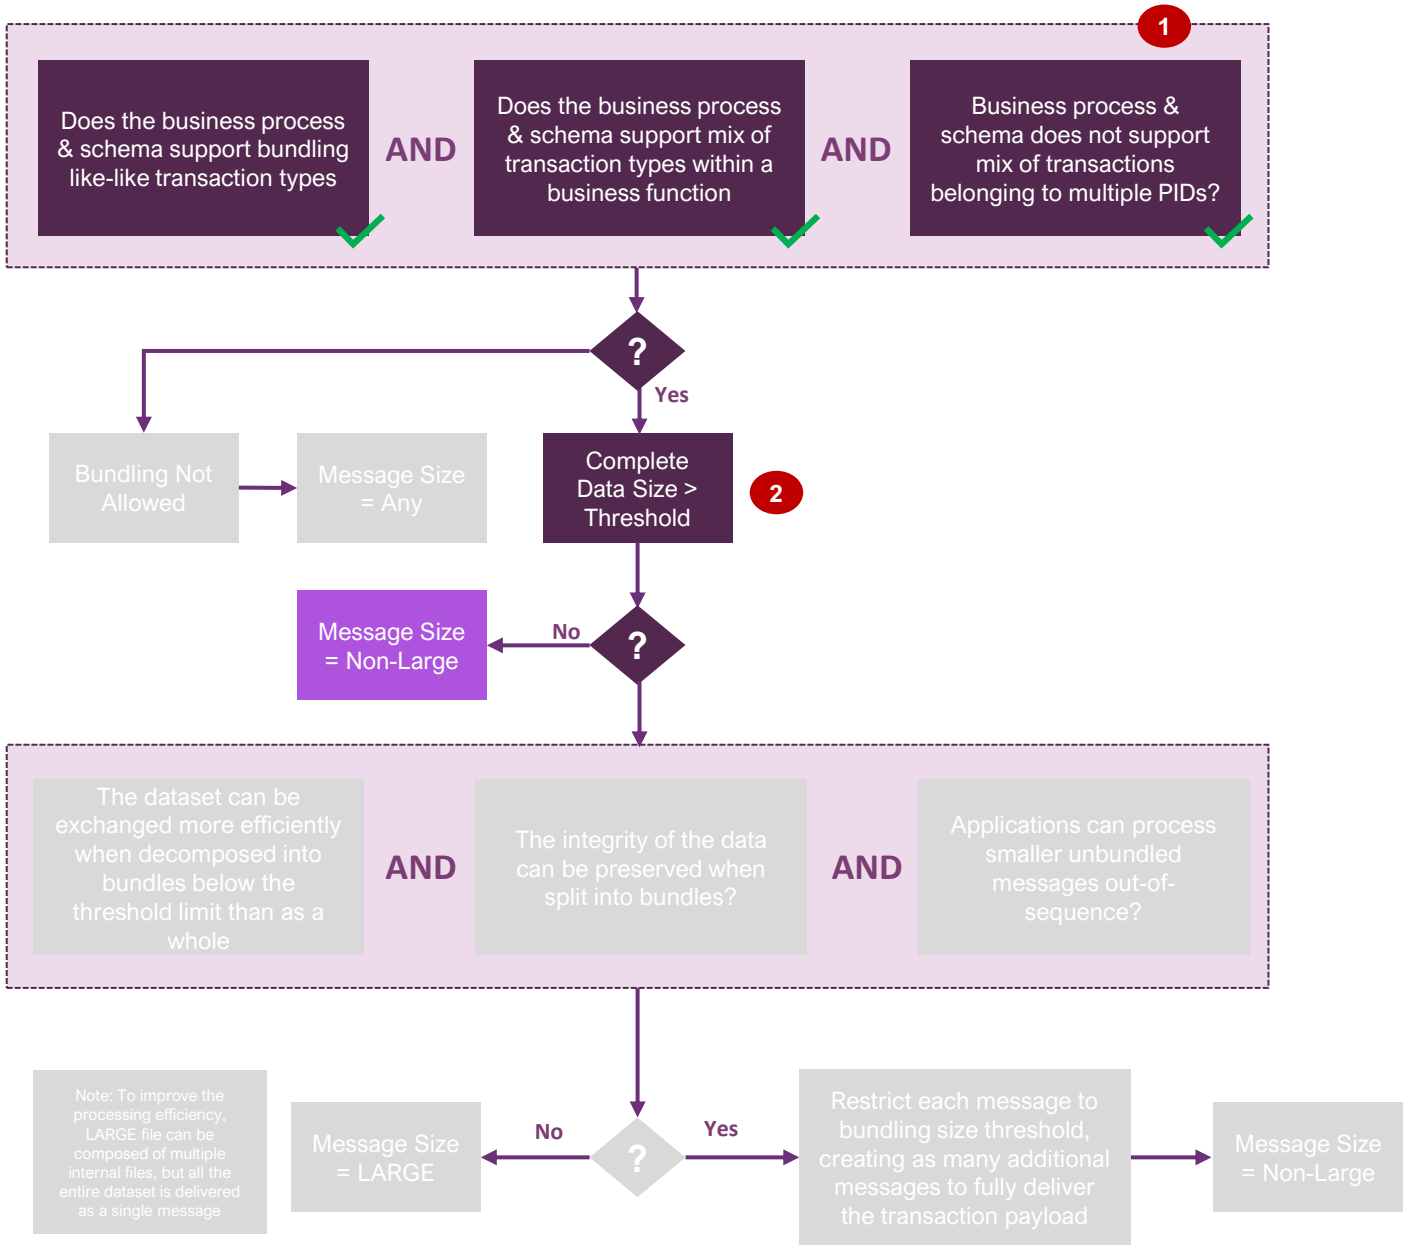

Use Case Description

Use Case: Delivery of CATS Messages

For illustration: Message size threshold = 1MB

Transaction Name: Change Request

System waits for 5 min to bundle all <ChangeRequest> transactions. After bundling <ChangeRequest> transactions in that 5 min duration, the message size do not exceed Threshold

Applying Decision Tree for the Use Case

1a. CATS data support bundling of like-like transactions
1b. CATS Data allow mix of transaction types (e.g. change responses, change notifications) within a business function
1c. CATS dataset are to be delivered to a participantID

2. Complete Data Size < Threshold (1MB for this example). Accumulation/bundling of transactions are time bound (e.g. 5 min), the chance of message size exceeding 1MB is not high

Message Size Threshold

Message Size Classifications for Decision Tree

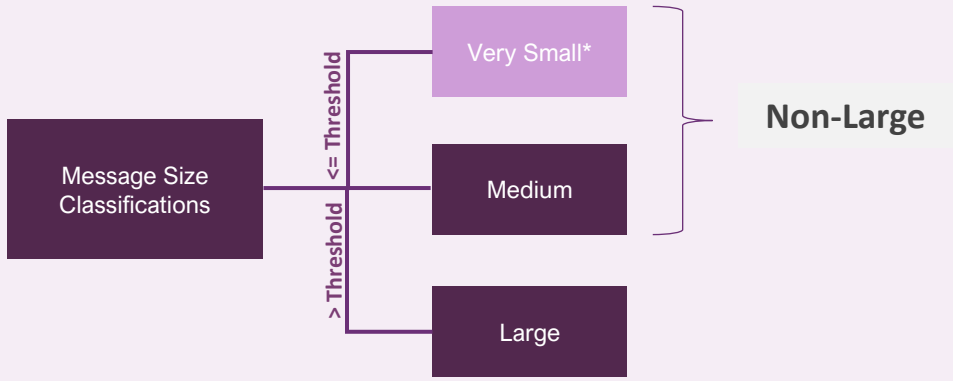

** Very small message sizes will be applied to fire & forget data exchange pattern; requiring high message volume (frequency) and low response times*

Factors that will limit the message size threshold value

1) Limitations imposed by the integration technology components (e.g. API Gateway)
2) Limitations imposed by the applications; time taken to process (performance) the message payloads (e.g. B2B application managing routing & transforms)
3) Infrastructure capacity of the Participants' integration components and applications. Agreed threshold values must meet the infrastructure requirements of all Participant types (e.g. retailers, MDPs, generators etc) and sizes (large and small Participants); not imposing overheads to Participants
4) Industry best practices

Options to determine the value of Message Size Threshold

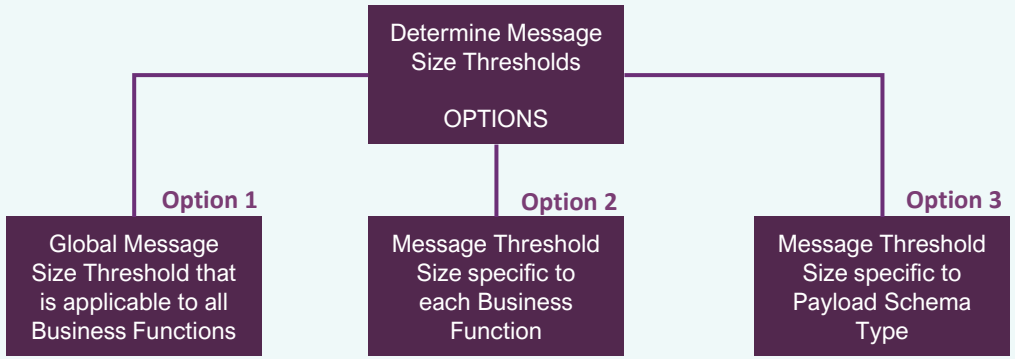

1) Message size greater than the threshold will be classified as 'Large' and message size less than or equal to the threshold will be classified as 'Non-Large'
2) 'Message Size' is a core attribute in the decision tree that will determine the selection of a channel for the data exchange
3) To determine an optimal value of the message size threshold, following options are being assessed
   - Option 1: Global threshold value that is a common value across all business functions
   - Option 2: Threshold value that is specific to each business function
   - Option 3: Threshold value that is specific to payload schema type

Message Size Threshold Option 3 – Recommended option

Option 3: Threshold Value @ Schema Type Level

1) Threshold value will be set based on the schema type that supports the business function(s)
2) Message exchanges that use JSON schema types will all share a common threshold value and message exchanges that use AEMOCSV schema types will share a common threshold value
3) API Gateways operate efficiently when the message size is optimal. Considering the industry best practices and the limitations of the API Gateway technology, threshold value of 10MB is the maximum that can be set for RESTful channel.
4) Worked example shows a variety of threshold values such as 10MB, 2MB; values are for illustration purposes. The optimal value of the threshold will be agreed in consultation with the participants post locking the message size threshold options

Note: The thresholds mentioned in this option refer to the thresholds for the 'Non-Large' category

Worked Example

Use Case	Target State Schema Type	Threshold Value
BF: energyFCASBids Resource: Submit Bids	JSON	2MB
BF: energyFCASBids Resource: /reports/public	AEMOCSV	10MB
BF: Service Order Resource: submissions	JSON (or XML)	2MB
BF: Meter Reads Resource: Send meter reads	AEMOCSV (or MDFF)	10MB
BF: Meter Reads Resource: PMD/VMD	JSON (or XML)	2MB

Pros

1) Participants are required to augment/uplift their infrastructure only if they consume the required business function and the associated payload type
2) Not all of the Participants' applications are required to manage large message sizes
3) Minimises operational and governance overheads in managing the threshold value for each business function

Cons

1) There may be instances where certain resources of a business function (that uses AEMOCSV format) may not require a higher threshold value such as 10MB. However, with this option a higher threshold value will be assigned to such resources that may lead to poor bundling practices

* Details on the message size threshold Options 1 and 2 can be found in Appendix C

Message Size Threshold – Options Comparison

AEMO requested Focus group members to share their feedback on the option that should be used to assign the threshold value. AEMO asked the Focus group to consider their organisation needs and provide inputs on how these options will improve their pain points and/or minimise the capex-opex costs

Criteria	Option 1: Global Value	Option 2: @ Business Function Level	Option 3: @ Schema Type Level
Minimises operational & governance overhead in managing the threshold value	●	◐	●
Minimises the impacts to performance of the message exchange e.g. (not limited to) schema validations, transformation, applying business & technical validations	◐	◔	◔
Minimises the impacts to Participants' infrastructure capacity uplifts. Ability to localise the uplifts only when consuming the business function / service	◔	◔	◐
Minimises the overhead to establish new business services	●	◐	●

The Focus group discussed these options in detail and the recommendation was that **Option 3 – Schema type level** provided the right balance factoring in flexibility and complexity

* Further details and examples available in appendix B

Approach to determining message size thresholds

Two potential approaches to determining message size thresholds were put forward in the Decision Tree Focus Group for consideration by the broader Working Group:

1. **Academic** – reference best practice / technical standards where available to determine thresholds

2. **Test & tune** – use best practice / technical standards to establish a target range, test and tune by measuring performance to confirm an optimal threshold
   - Performance testing could be performed internally by AEMO or in concert with industry as part of the pilot phase

POLL: We are seeking input from the MITEWG on the preferred approach

Are there other approaches we should consider?

Message Volume

Message Volume : Total volumes of messages for a data exchange between all participants and AEMO Hub for a particular Business Function use case

Message Volume Classifications for Decision Tree

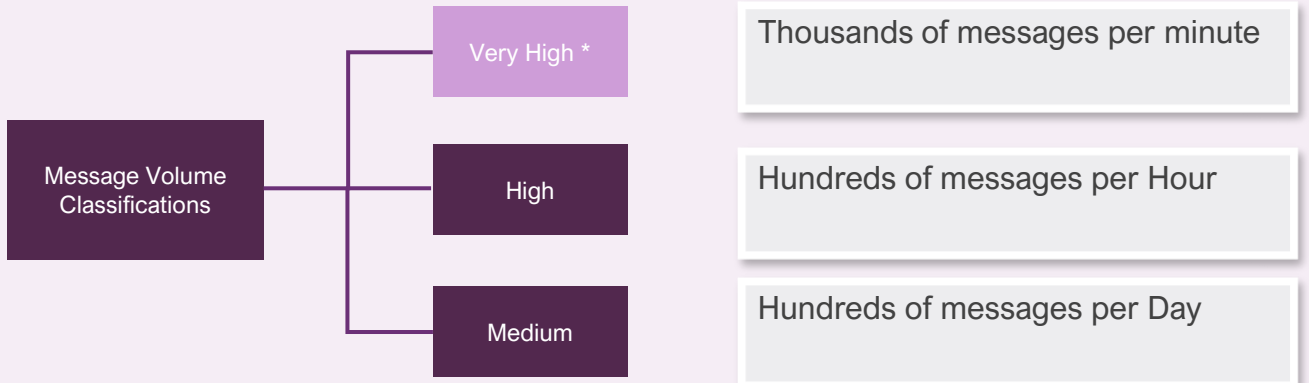

** Very High message volume will be applied to fire & forget data exchange pattern; For very small message size and requiring low response times*

Decision tree –Channel Definition

Channel	Definition
Restful API	A web-based communication channel that conforms to REST architectural principles, utilising standard HTTP methods (GET, POST, PUT, DELETE) to interact with resources. The representation of resources, available endpoints, and methods is defined in a RESTful API schema.
File Transfer	A file-based transfer channel designed to securely and reliably transfer large files or payloads between organisations that are too large to be sent via API (HTTP)
Graph QL	An API-based channel that enables clients to request specific data structures using a query language.
Streaming	A streaming pattern channel designed to securely and reliably transfer near real-time continuous collection and movement of data, handling high volumes, at scale, with high throughput and low latency. ***IDX Foundation is proposing not to build this capability; rather we are considering the definition of this channel to cater for future use cases**

NOTE 1: We will discuss protocols for each channel in later MITE sessions
NOTE 2: MarketNet vs public internet options are being reviewed by AEMO
Are there other channels we should be considering?

Decision trees – Synchronous & Asynchronous

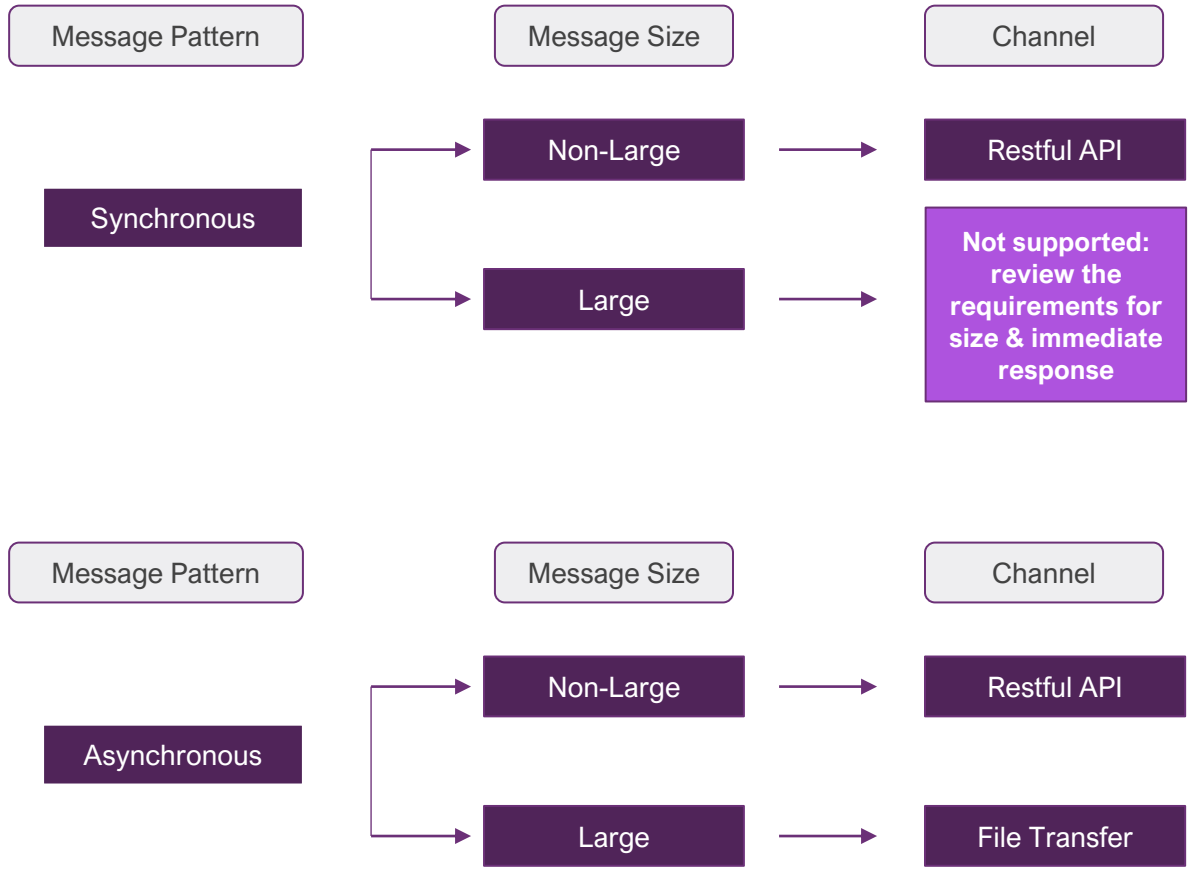

Decision tree – Fire & Forget

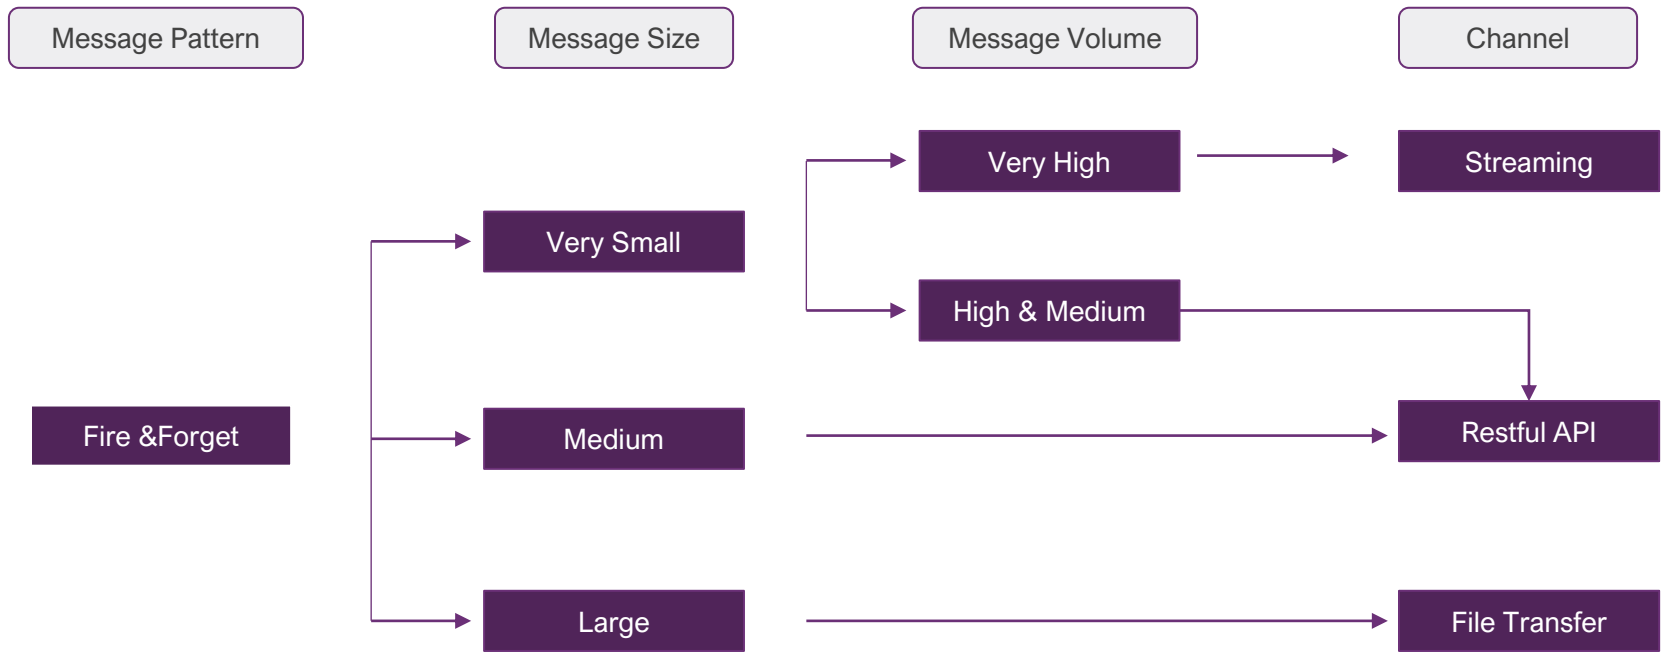

Decision trees – Inquiry

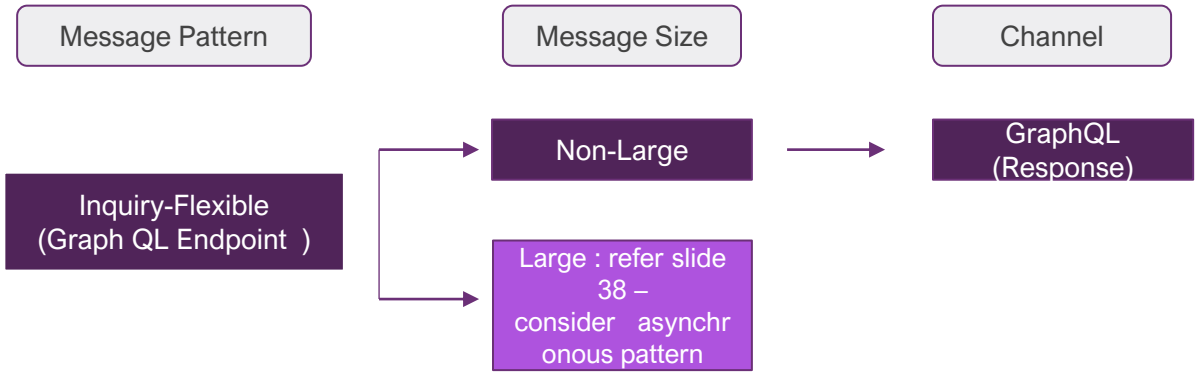

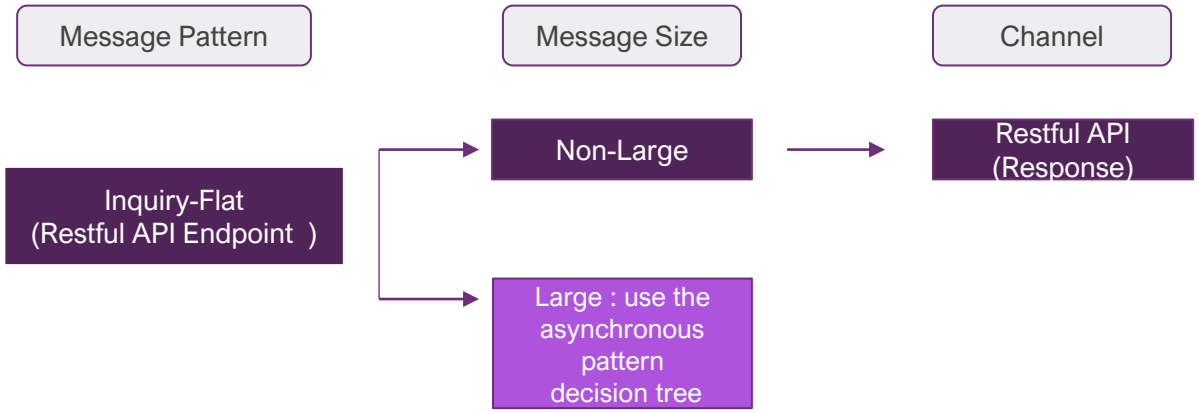

Decision tree – worked example: Bid submission

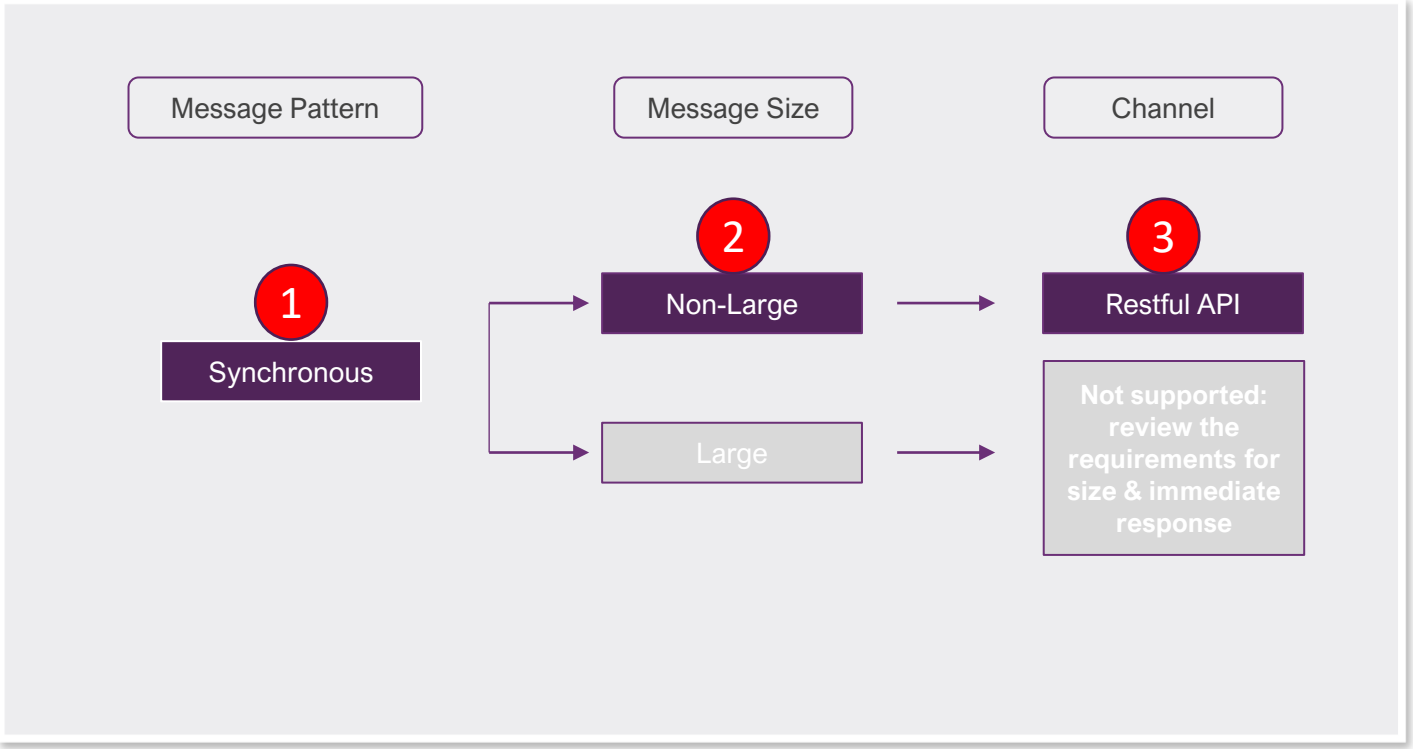

Use Case: Bids Submission

1. When Bids submission request is received, AEMO validates, stores and sends the business response - fits the Synchronous pattern
2. Using 'Message Size Decision Tree' in slide#17; bidding data can be bundled into payload chunks of 1MB i.e. message size = 'Non-Large'
3. Channel to be used is Restful API

Decision tree – worked example: Service Orders

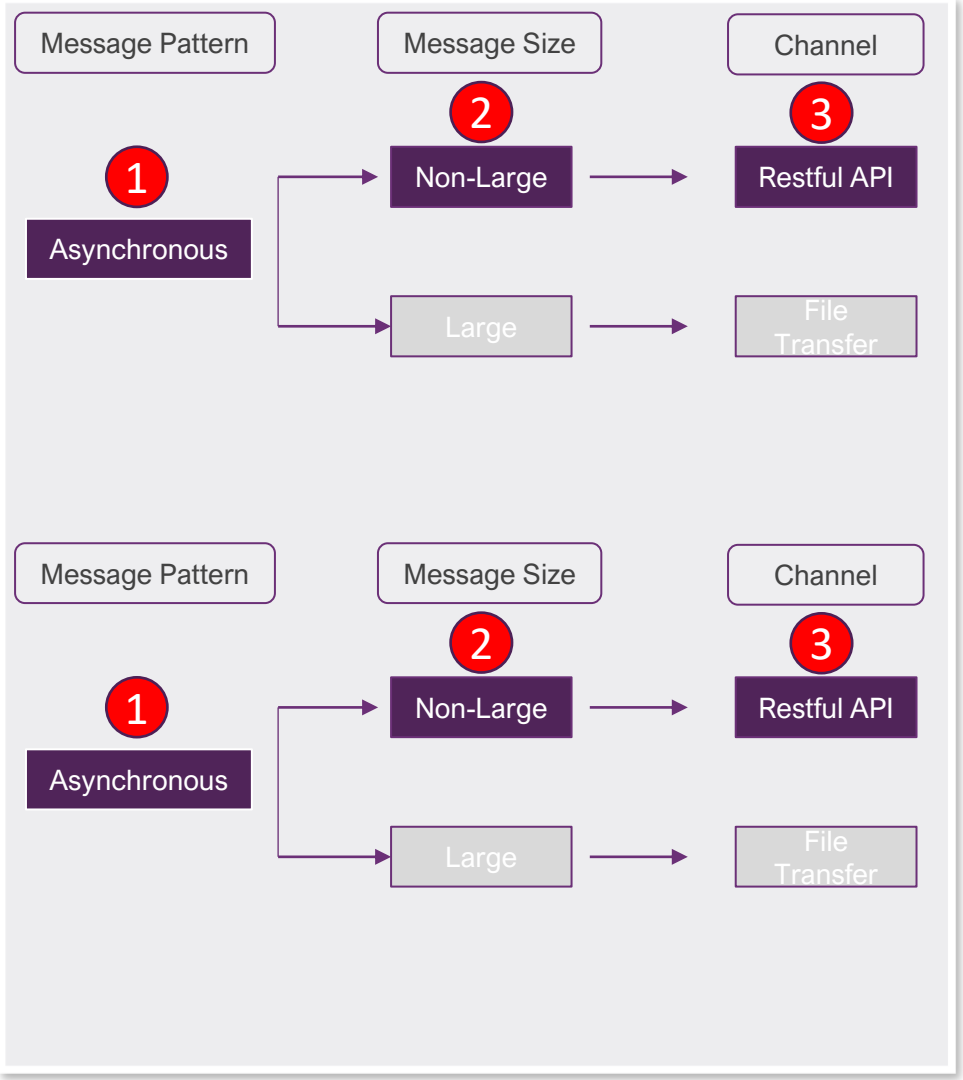

Use Case: Delivery of Service Order Request

Leg 1

1. Service Order business function fits the Asynchronous pattern
2. ServiceOrderRequest transactions can be bundled and limited to 1MB file size. Integrity of the data can be maintained by chunking the payloads into sizes of 1MB. Using the 'Message Size Decision Tree' in slide#17, message size = 'Not Large'
3. Channel to be used is Restful

Leg 2

1. It is another leg in an overall Asynchronous pattern
2. Message size = 'Not Large' as identified in Leg 1
3. Channel to be used is Restful

Likewise, channels for the rest of the Legs are to be assessed

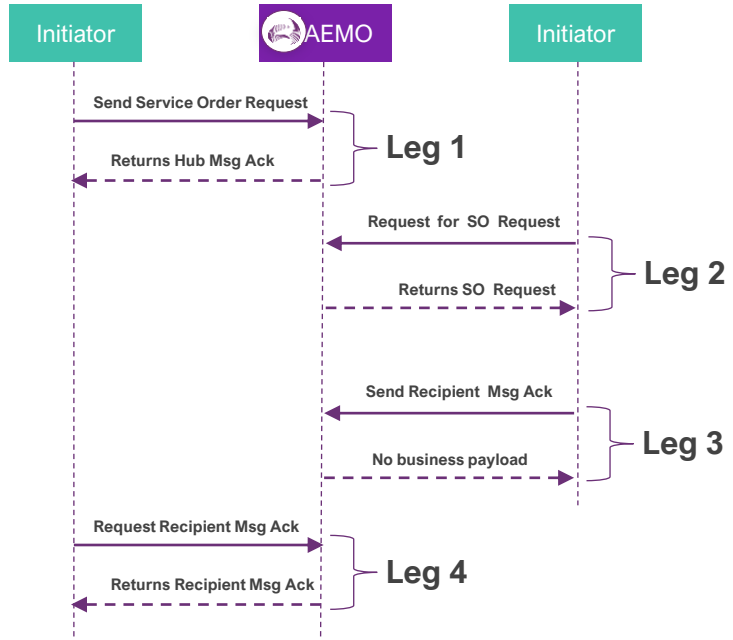

Decision tree – worked example: MIRN Listing

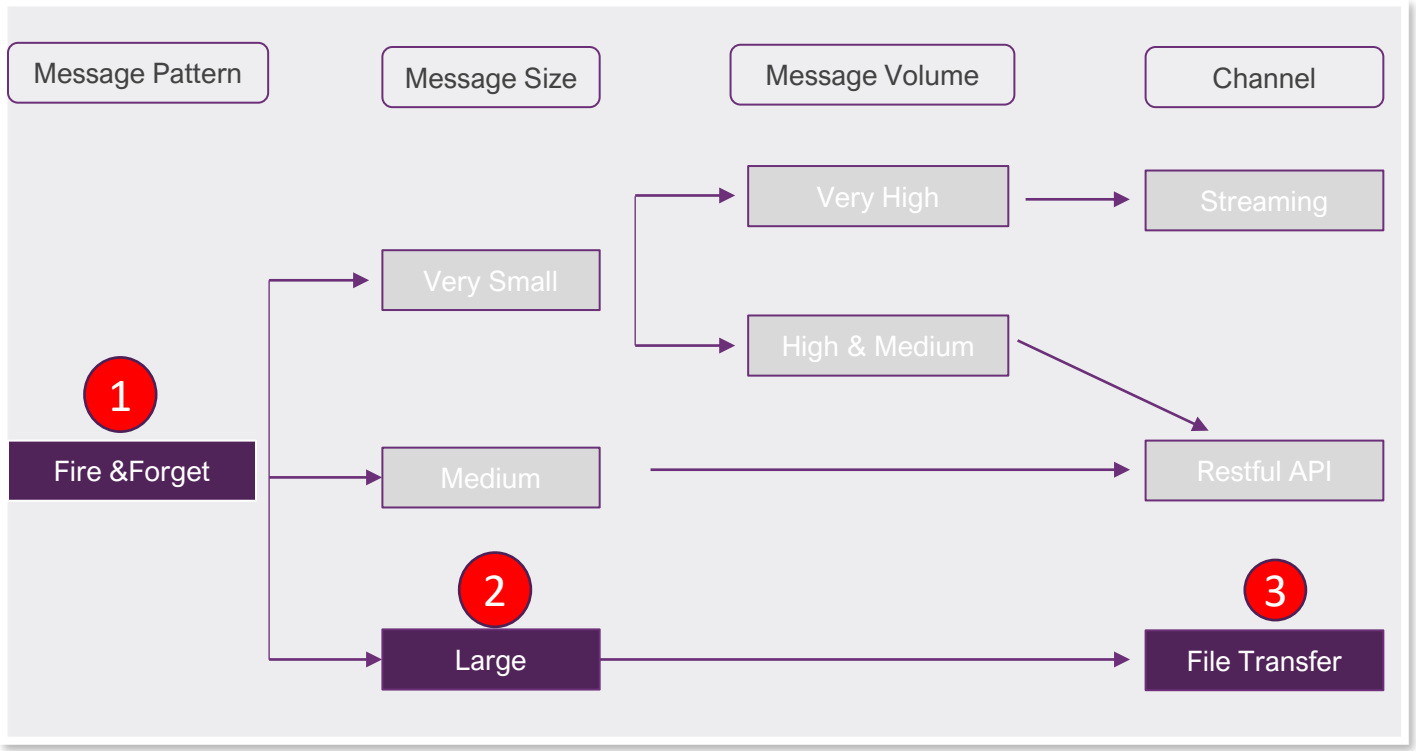

Use Case: Delivery of MIRN Listing / INT reports

1. When Participant submits MIRN Listing to AEMO (i.e. B2B) there is no expectation of a business response, so it is Fire & Forget pattern
2. Message payload size is 'Large'
3. Channel to be used is File Transfer

Notes

Sri Gundu spoke to the IDX Decision tree playback with lots of great questions and feedback.

The following questions were raised:

- What is the rationale behind the Inquiry flexible and Inquiry flat? **AEMO Response**: the rationale was to bring in the flexibility of GraphQL which will add additional capabilities to the IDX Platform.
- Will Publish-subscribe and Streaming be discussed later? **AEMO Response**: Publish-subscribe was discussed in the focus group with the reasons why this was discounted now published on slide 129 of this pack. Streaming will be assessed after IDX Foundation on when a suitable use case can be identified.
- Will participants need to host an API to received acknowledgements for an Asynchronous process? **AEMO Response:** AEMO will host the endpoints and participants will submit the acknowledgements. Further discussions on this will be discussed in the Asynchronous focus group, with a Notification Channels being a topic to discuss.
- What are the benefits of bundling? **AEMO Response**: Where we've got a large volume of small transactions that could be bundled. There is an overhead making multiple API calls. There's an overhead both in the connection process as well as in the transport layer, it can be more efficient end to end to bundle those transactions together and allow them to be passed as a bulk load.
- Is there a timeline when we're going to discuss and finalize transactions based on these decision trees? **AEMO Response:** This is ultimately a post decision Point 2 discussion.
- Polling of endpoints to retrieve messages, is that the expectation? **AEMO Response**: The Asynchronous and Event notification FG that was deferred from this month will help to answer these questions, an event notification channel will be discussed in this focus group.

Other Key Notes:

- Message Threshold – Options comparison: Option 3 was recommended from the Decision tree focus group. No comments or concerns were raised in the WG so this will be the recommendation moving forward.
- Determining Message Size Thresholds approach Poll was conducted: 68% of organisations on the call voted for *OPTION 2: Test & tune – use best practice / technical standards to establish a target range, test and tune by measuring performance to confirm an optimal threshold.* AEMO will continue with this as the recommendation.

8. IDX Business Function Endpoints

Sri Gundu

Business Function API Endpoint – Options Comparison

In the IDX MITE WG on the 4th of September AEMO requested each organization provide a preference for the Business function endpoint target state based on the options below. A poll was conducted, and the result compiled by AEMO.

Criteria	Current State	Target State Option A	Target State Option B
Minimises mandatory schema changes for Participants who do not consume the impacted business function	○	●	●
Localises schema change to a specific sub-function of a business function when changes are introduced to a specific transaction within a business function	○	◔	●
Minimises the overhead of maintaining multiple schemas and versions within a business function; thereby improves operational efficiency	◐	●	◐
Efficient management of schemas for a business function e.g. leveraging a base schema that can inherited into sub-functions under a business function	●	◔	◐
AEMO & Participants' system changes (integration layer and downstream/upstream applications) to effectively apply schema validations in run time @ the business function level	◐	●	◔
MITEWG poll results	**N/A**	**9**	**9**

A total of 18 MITE WG organisations voted for their preference either directly in the poll or via email resulting in no clear preference from industry.

As requested by participants, AEMO has developed additional scenarios to explore the operation of these options in the following slides to help resolve which option to proceed.

NOTE: AEMO have consolidated the responses received from different organisations, considering one vote per organisation where we had multiple representatives voting from the same organisation

Business Function API Endpoint – Resources – Option A
Worked Example: B2B Meter Reads (MTRD Transaction Group)

Guiding Principle: A business function must only support a distinct schema for each payload type. Multiple transactions belonging to a business function and sharing the same payload type are bundled under one schema.

Assumption: For illustration purposes it is assumed that Business Function in the Retail space is "Transaction Group". However, this assumption will be assessed and refined during DP2 (Decision Point 2). Also, RemoteServiceRequest & RemoteServiceResponse are not shown in the list below; they belong to Transaction group of MRSR.

Market: NEM Retail
Business Function: Meter Reads (MTRD Transaction Group)
Business Function API: https://.../NEMRetail/v1/B2BMeterReads/<resource group>/<resources>
Supported functionalities required:

Use Case	API Method	API Definition	Proposed Payload Format
Send meter reads to the B2B Recipient (messages & TACKs)	POST	NEMRetail/**B2BMeterReads/**v1/transactions/meterDataNotification	AEMOCSV / MDFF
Retrieve meter reads from the B2B Sender (messages & TACKs)	GET	NEMRetail/**B2BMeterReads/**v1/transactions/meterDataNotification	AEMOCSV / MDFF
Send Provide Meter Data Request or Response (messages & TACKs)	POST	NEMRetail/**B2BMeterReads/**v1/transactions/provideMeterData	JSON
Retrieve Provide Meter Data Request or Response (messages & TACKs)	GET	NEMRetail/**B2BMeterReads/**v1/transactions/provideMeterData	JSON
Send Verify Meter Data Request or Response (messages & TACKs)	POST	NEMRetail/**B2BMeterReads/**v1/transactions/verifyMeterData	JSON
Retrieve Verify Meter Data Request or Response (messages & TACKs)	GET	NEMRetail/**B2BMeterReads/**v1/transactions/verifyMeterData	JSON

NEM_MTRD_AEMOCSV_r1

NEM_MTRD_JSON_Schema_r1

Note: MACKs are not represented here. During DP2, AEMO will consult with Participants to assess if AEMO must consider a generic MACK endpoint across all B2B business functions (or) MACKs must be part of each B2B business function

Business Function API Endpoint – Resources – Option B
Worked Example: B2B Meter Reads (MTRD Transaction Group)

Guiding Principle: Implement schemas specific to each of the sub-functions of a Business Function.

Assumption: For illustration purposes it is assumed that Business Function in the Retail space is "Transaction Group". However, this assumption will be assessed and refined during DP2 (Decision Point 2). Also, RemoteServiceRequest & RemoteServiceResponse are not shown in the list below; they belong to Transaction group of MRSR.

Market: NEM Retail
Business Function: Meter Reads (MTRD Transaction Group)
Business Function API: https://.../NEMRetail/v1/B2BMeterReads/<resource group>/<resources>
Supported functionalities required:

Use Case	API Method	API Definition	Proposed Payload Format
Send meter reads to the B2B Recipient (messages & TACKs)	POST	NEMRetail/**B2BMeterReads/**v1/transactions/meterDataNotification	AEMOCSV / MDFF
Retrieve meter reads from the B2B Sender (messages & TACKs)	GET	NEMRetail/**B2BMeterReads/**v1/transactions/meterDataNotification	AEMOCSV / MDFF
Send Provide Meter Data Request or Response (messages & TACKs)	POST	NEMRetail/**B2BMeterReads/**v1/transactions/provideMeterData	JSON
Retrieve Provide Meter Data Request or Response (messages & TACKs)	GET	NEMRetail/**B2BMeterReads/**v1/transactions/provideMeterData	JSON
Send Verify Meter Data Request or Response (messages & TACKs)	POST	NEMRetail/**B2BMeterReads/**v1/transactions/verifyMeterData	JSON
Retrieve Verify Meter Data Request or Response (messages & TACKs)	GET	NEMRetail/**B2BMeterReads/**v1/transactions/verifyMeterData	JSON

NEM_MTRD_AEMOCSV_r1

NEM_MTRD_PMD_JSON_Schema_r1

NEM_MTRD_VMD_JSON_Schema_r1

Note: MACKs are not represented here. During DP2, AEMO will consult with Participants to assess if AEMO must consider a generic MACK endpoint across all B2B business functions (or) MACKs must be part of each B2B business function

Business Function API Endpoint – Worked Example – Schema Changes Comparing Current State vs Option A vs Option B

For illustration purposes; the following B2B business functions are considered:

a. *Business Function 'B2B Meter Reads' with sub-functions of 'Provide Meter Data' & 'Verify Meter Data'*
b. *Business Function 'Site Notification' with sub-function of 'Site Access Notification' & 'Site Address Notification'*

Worked Example Steps	Current State	Option A	Option B
1) Representation of 'Business Function – Schema Association' for the worked example	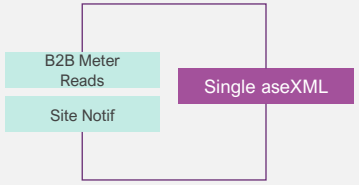	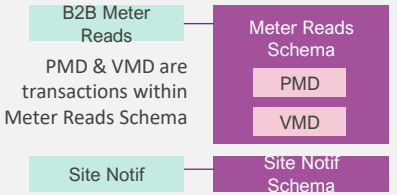	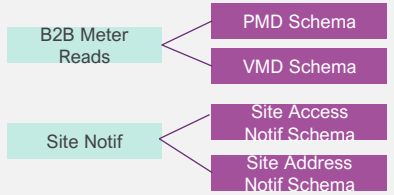
2) Representation of the Schema Versions of 'B2B Meter Reads' & 'Site Notification' Business Function for the worked example	Common_aseXML_V1	MTRD_JSON_V1 SITE_JSON_V1	MTRD_PMD_JSON_V1 MTRD_VMD_JSON_V1 SITE_ACC_JSON_V1 SITE_ADD_JSON_V1
3) Changes to Site Notifications Procedure - introduces one new field to existing <SiteAccessRequest> transaction. **Impacts to 'B2B Meter Reads' business function are:**	Common_aseXML_V2 **(impacts all B2B Business Functions)**	MTRD_JSON_V1 **(No change)** SITE_JSON_V2	MTRD_PMD_JSON_V1 **(No Change)** MTRD_VMD_JSON_V1 **(No Change)** SITE_ACC_JSON_V2 SITE_ADD_JSON_V1
4) Changes to Meter Data Process Procedure - introduces one new field to existing <MeterDataVerifyRequest> transaction **Impacts to 'B2B Meter Reads' business function are:**	Common_aseXML_V3 **(impacts all B2B Business Functions)**	MTRD_JSON_V2 **(Changes to the Business Function Schema. Does not impact schemas of other Business Functions)**	MTRD_PMD_JSON_V1 MTRD_VMD_JSON_V2 **(Change is limited to schema of VMD sub-function only)**
5) Changes to Meter Data Process Procedure - introduces a new transaction type **Impacts to 'B2B Meter Reads' business function are:**	Common_aseXML_V4 **(impacts all B2B Business Functions)**	MTRD_JSON_V3 **(Changes to the Business Function Schema. Does not impact schemas of other Business Functions)**	MTRD_PMD_JSON_V1 MTRD_VMD_JSON_V2 MTRD_NEW_JSON_V1 **(A new schema is introduced for B2B Meter Reads Bus Func)**
6) Common attribute (e.g. RequestID) that is applicable to all three sub-functions of 'B2B Meter Reads' requires a change **Impacts to 'B2B Meter Reads' business function are:**	Common_aseXML_V5 **(impacts all B2B Business Functions)**	MTRD_JSON_V4 **(Changes to the Business Function Schema. Does not impact schemas of other Business Functions)**	MTRD_PMD_JSON_V2 MTRD_VMD_JSON_V3 MTRD_NEW_JSON_V2 **(Change affects common element, all schema are incremented)**
7) Operational overheads of maintaining multiple schemas and versions	Monolithic schema and complex transformations that may be required for each business function	One schema at business function level	Multiple schemas with different versions for a given business function

Business Function API Endpoint – Worked Example – Inheritance of Base Schema
Comparing Current State vs Option A vs Option B

Steps	Current State	Option A	Option B
Worked Example of how base schemas are inherited	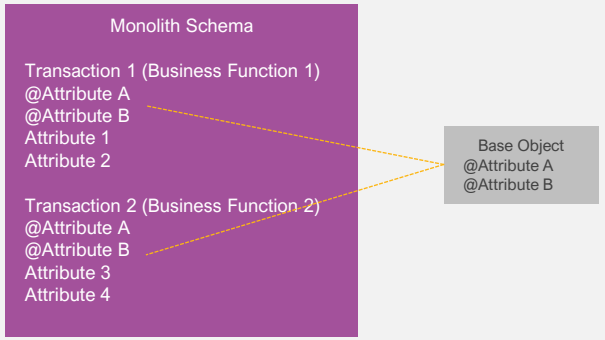	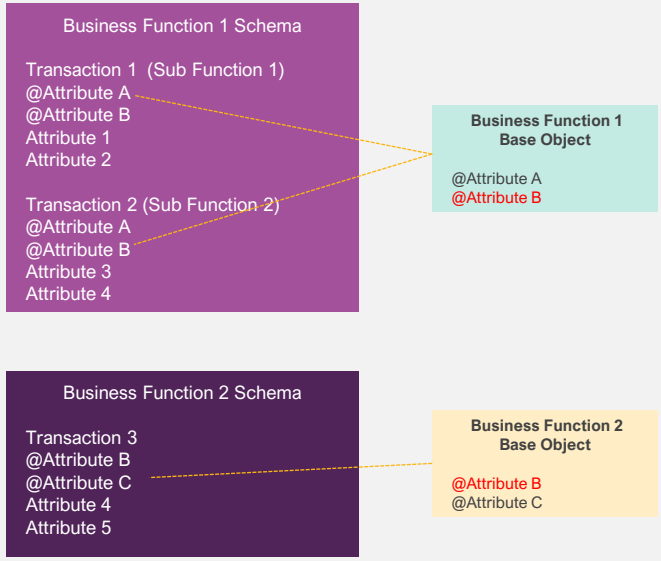	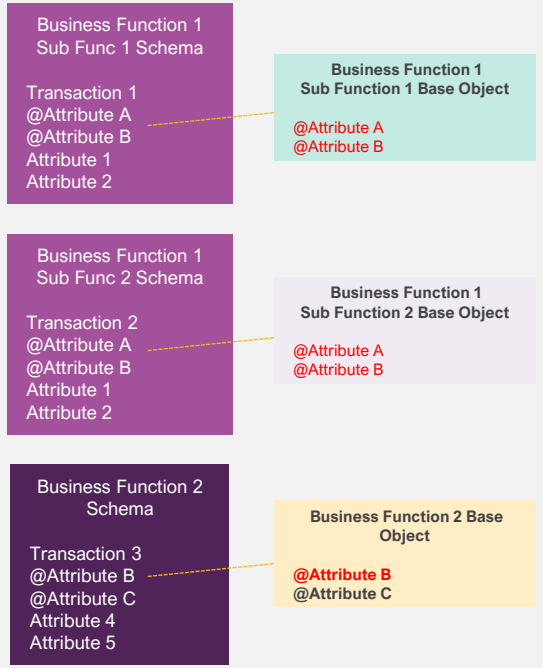
How efficiently are the base schemas inherited into sub-functions under a business function / across business functions?	Attributes of the base object are inherited into multiple business functions; avoids duplication	Attributes of the base object are not duplicated within the business function schema; however it could be duplicated across the business functions (e.g. Attribute B marked in red)	Attributes of the base object are duplicated within a business function and across business function schemas (e.g. Attribute A & B marked in red)
How efficiently are the changes to the base schema (versions) managed?	Efficiently managed; changes are inherited into multiple business functions	Efficiently managed @ the business function level. Additional operational overheads if the changed object/attribute is duplicated in other business functions.	Operational overheads in managing the changes to the base schema within the business function and across business functions

Business Function API Endpoint – Recommendation

The following Business Function Endpoint options were considered with an initial recommendation of Option A:

Criteria	Current State	Target State Option A	Target State Option B
Minimises mandatory schema changes for Participants who do not consume the impacted business function	○	●	●
Localises schema change to a specific sub-function of a business function when changes are introduced to a specific transaction within a business function	○	◔	●
Minimises the overhead of maintaining multiple schemas and versions within a business function; thereby improves operational efficiency	◐	●	◐
Efficient management of schemas for a business function e.g. leveraging a base schema that can inherited into sub-functions under a business function	●	◔	◐
AEMO & Participants' system changes (integration layer and downstream/upstream applications) to effectively apply schema validations in run time @ the business function level	◐	●	◔

Noting industry was unable to identify a preference between Option A and B at the last MITE WG, AEMO has provided these additional worked scenarios to better explore the impacts of these options. Noting AEMO's recommendation still aligns to option A, we would like to validate if this additional information allows a preference to be identified.

Notes

Sri Gundu spoke to the IDX Business Function playback with lots of great questions and feedback.

The following questions were raised:

- Would Option B have a lesser regression impact than option A? **AEMO Response:** For some scenarios is may, but for others it would have a much larger impact as there are many more schemas implemented as part of option B.
- Where do most changes sit historically 1-4 or 5-6? **AEMO Response:** It does vary across business functions, but AEMO view is that it is more likely to be in the no.6 worked example than in the no.4 worked example.
- Are there no linkage to access controls, regardless of which option is chosen? **AEMO Response:** Correct, for machine to machine you will define your service accounts, and you can filter down further with machine-to-machine API calls.
- Is there a plan to separate between you said gas and electricity is what if there is a change in the service order in gas? If there is any costing over so, then the electricity is not going to change. **AEMO Response:** In today's world between NEM and WEM there are differences in practically all transactions. We won't preclude a combined approach if there is a good rational for doing so.

Other Key Notes:

- Business Function API Endpoint – Recommendation: Based on some additional feedback and change in preferences (via email) and no major objections raised in the working group, AEMO will continue with the recommendation of Option A in relation to Business Function Endpoints.

9. IDX Archiving

Selwyn Sequeira

Archiving

Presently, there are several archiving patterns in use across the energy industry. During this session, we will define what archiving means, review the capabilities of existing archiving patterns and consider new capabilities as an to input to developing a single archiving solution as part of the IDX framework.

This topic will cover:

- Archiving definition and purpose
- Assessment of current state capabilities
- Discovery - new capability requirements
- Decision tree – when to archive, retention periods

The outcomes we are looking to achieve are:

- Confirming our understanding of the archiving use-cases currently employed and the capabilities they provide
- Provide further context to how archiving is used by participants that AEMO may not be aware of.
- Participate in reviewing and defining new capability requirements for archiving

Archiving – Definition and purpose

Definition

Archiving refers to the systematic process of *storing market transaction data exchanged between AEMO and market participants for a defined period*.

Purpose

Archiving allows for previously sent exchanges to be inspected and retrieved to facilitate:

- Participants to validate that their market systems are receiving and processing market messages
- BCP scenarios where participant systems are unable to recover all transacted data
- Reconciliation where participants wish to ensure they have processed all market messages

Archiving – inconsistent across markets

Background

- Across Energy Markets there are different solutions provided for NEM Retail, NEM Wholesale, Gas and WEM participants.

Current state*

- NEM Retail: Accessed via Open protocols (FTP) , LVI & Day Zip
- NEM Wholesale: Accessed via Managed interfaces - LVI and Machine to Machine
- GAS: GSH and GBB follows NEM Wholesale, MIBB & MIS archives available over FTP/LVI for a limited duration
- WEM: No access to a formal archive solution (participants are responsible for implementing their own archiving solution)

* Further details and examples available in appendix B

Archiving Capabilities

Assessment of current state capabilities and discovery of new capability requirements

Current State Capabilities

Capability	Description	Target state	Archive Scope
Message retention	The ability to store messages and the length of time those messages can be stored.	IDX Archiving	✔
Accessibility - LVI	The ability to access archive over Low Volume Interface (web browser).	IDX Archiving	✔
Accessibility – M2M	The ability to access archive via a Machine-to-machine mechanism (for example, FTP access to archive folders in Retail, archive re-request in Wholesale)	IDX Archiving	✔
Reconciliation	A mechanism to reconcile multiple messages (for example, Day Zip in Retail and Transaction Manifest in Wholesale)	IDX Archiving - subject to assessment on unification vs market separation	✔
Message search	The ability to search for a specific or set of messages by a unique identifier or other query parameters.	IDX Transaction logging	✘
Transaction search	The ability to search for a transaction by a unique identifier or other query parameters.	IDX Transaction logging	✘
Non-repudiation	Providing immutable proof of who-sent, and who-received a message or transaction.	Enabled by data exchange patterns & transaction logging	✘
Data sharing	Some participants leverage archive to enable sharing of messages (as opposed to more formal data sharing arrangements)	Not facilitated by archiving, leverage data sharing arrangements	✘

Have we missed any current state capabilities participants see as required for archiving?
Are there any existing capabilities we should not replicate?

Reconciliation & unification

Reconciliation is the process whereby participants ensure data exchanged within the industry is consistent and accurate (for example; reconciliation of NEM retail messages)

Wholesale	Retail
AEMO's PDR Software provides functionality to provide guaranteed delivery of outbound market data into a participant Data Model and automated recovery features fetching any missing transactions from archive. This functionality is retained in the target state.	DayZip functionality provides the ability to download a days' worth of archive material which can then be used to reconcile with participants systems. Participants can optionally build this reconciliation capability themselves

- For Participants using DayZip, what advantage does this provide over and above direct access to archive?
- What would the impact be if IDX only offered the Retail model for future state?
- What would the impact be if IDX only offered the Wholesale (Manifest) model for future state (noting this was intended for data model based exchanges)?
- Does your Organisation have a view as to whether AEMO should maintain the multiple models of reconciliation or whether a single model could be considered across wholesale/retail ?

New capabilities?

AEMO has reviewed the pain points identified as part of the prior business case phase and not identified any further capability requirements relevant to archiving.

- Do participants have any additional requirements or capability statements AEMO should evaluate as part of defining the Archiving solution?

Archiving Decision Tree

How and when to apply archiving

Archiving decision tree

For discussion, some data is more suitable to archiving than others. IDX will set parameters around what needs to be archived in a consistent manner.

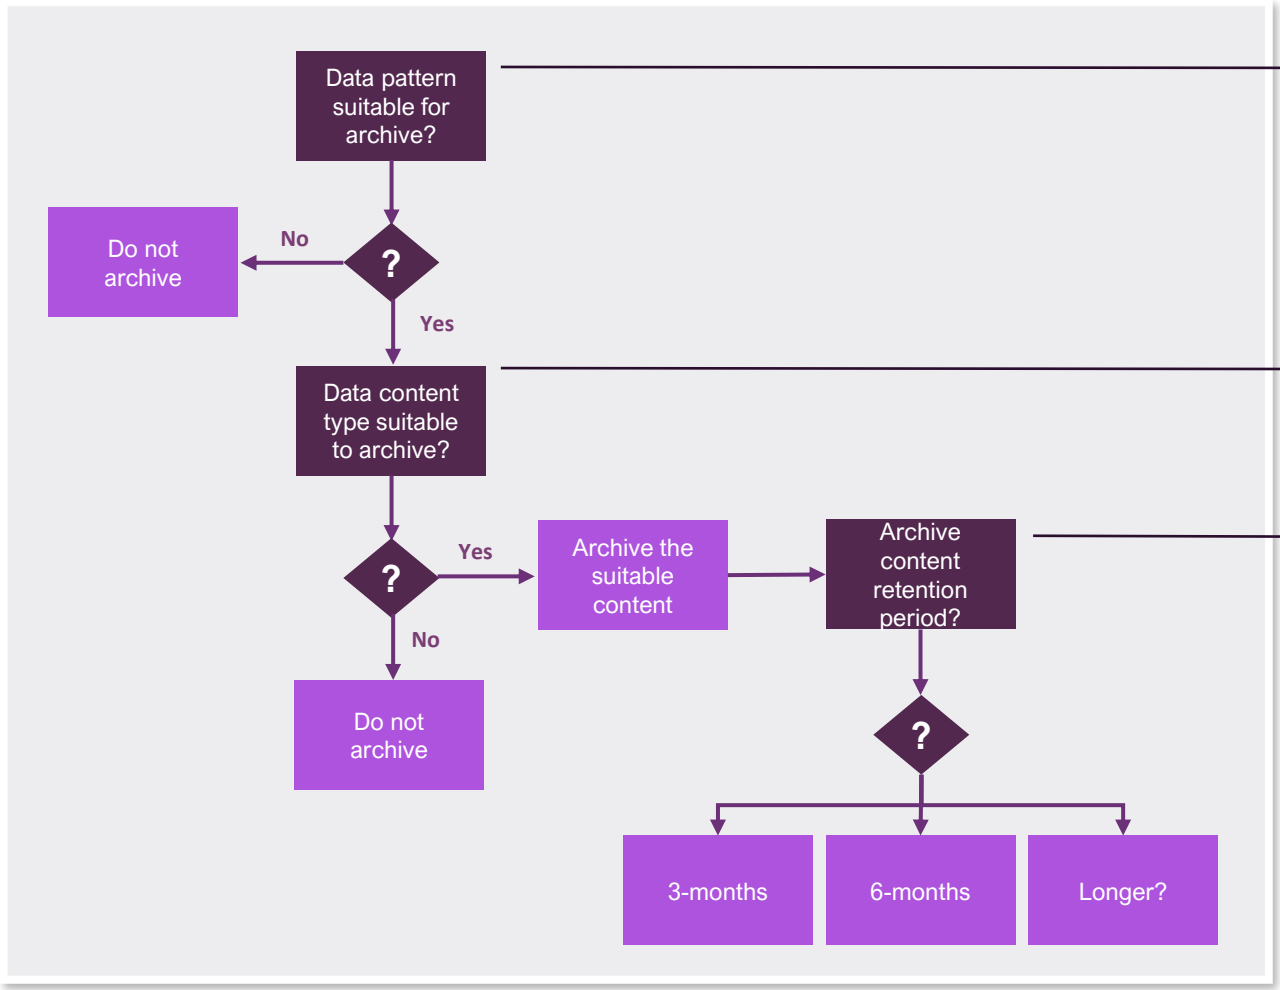

Some data patterns, such as Inquiry and Fire & Forget are not a fit for archiving, where-as Synchronous and Asynchronous patterns may be suitable

Some data content types may not need to be archived; for example, data which can be re-requested on demand

What is a reasonable length of time that the data should be archived, keeping in mind that longer retention times increase the market cost of archiving.

What data patterns are suitable for archiving?
What content should be archived?
What is a reasonable retention period for archived data?

Notes

This topic was deferred until the next Working Group.

10. IDX Future Topics

IDX WG-Session “Sync, Fire & Forget patterns ”

The objective of this focus group is to review an End-to-End Synchronous pattern , Fire & Forget Pattern and discuss technical topics related to this type of data exchange.

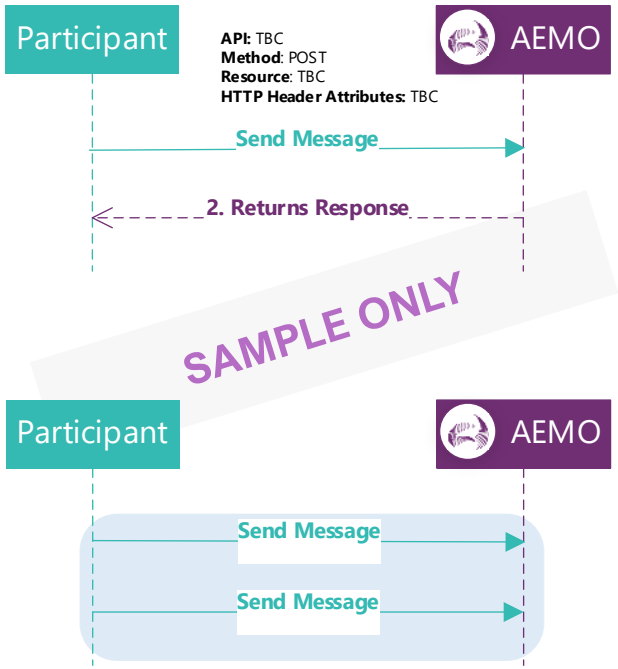

Synchronous data exchange patterns are used for business processes that require a multi-legged approach to facilitate a data exchange.

Synchronous data exchange discussion points:

- API, Method and Resource
- HTTP Header attributes
- Acknowledgements

Fire & Forget data exchange patterns are used for business processes that require a multi-legged approach to facilitate a data exchange.

Fire & Forget Inbound discussion points:

- Push data to IDX
- API, Method and Resource
- HTTP Header attributes

Fire & Forget Outbound discussion points:

- Pull data from IDX
- API, Method and Resource
- HTTP Header attributes

Audience Skill Set for Focus Group Discussion

- Technical Leads
- Integration Architecture Teams (Market Interface Specific)

Topics for Discussion

- A draft end to end business use case for Synchronous ,Fire & Forget flow
- Sequence diagram demonstrating the Synchronous pattern, Fire & Forget pattern (including events) to support the end to end business case.
- Interface requirements for Synchronous pattern, Fire & Forget pattern(including events).

This Focus group discussion will be relevant to all stakeholders who participate in exchanging data between AEMO and energy stakeholders

IDX MITE Focus Group Session "Payloads"

The objective of this working group session is to discuss the high-level overview of target state payloads.

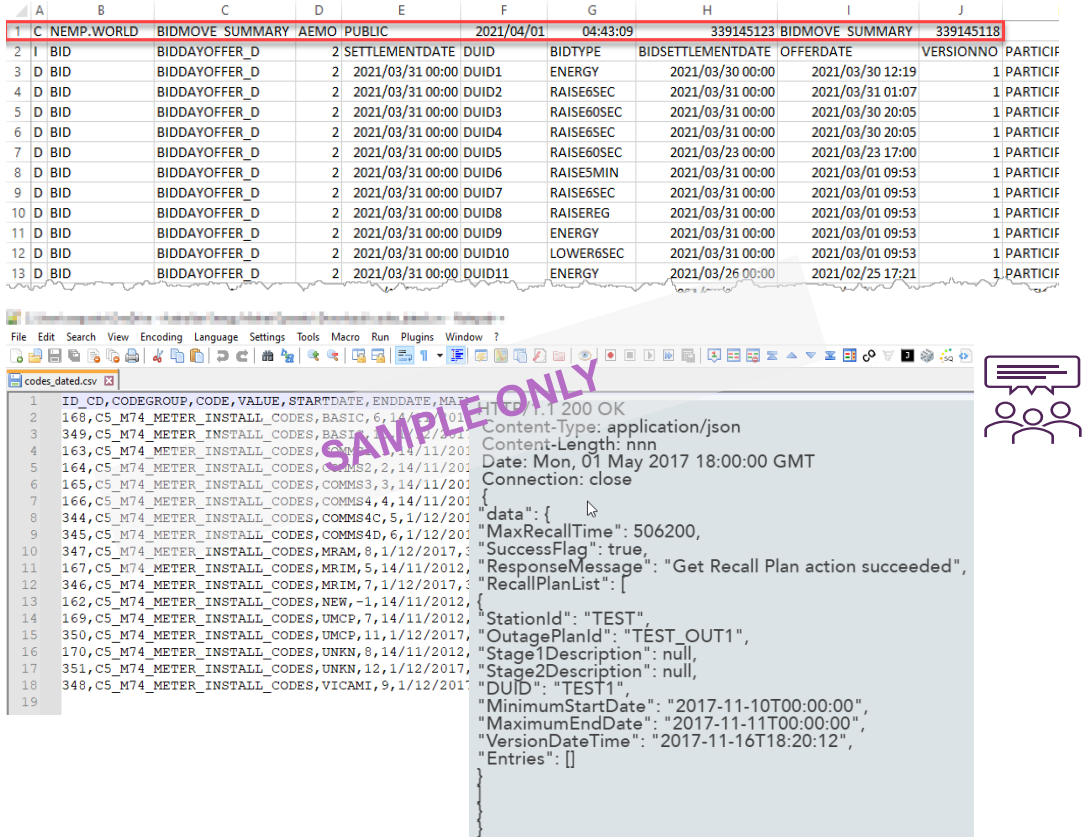

	A	B	C	D	E	F	G	H	I	J	
1	C	NEMP.WORLD	BIDMOVE_SUMMARY	AEMO	PUBLIC	2021/04/01	04:43:09	339145123	BIDMOVE_SUMMARY	339145118	
2	I	BID	BIDDAYOFFER_D	2	SETTLEMENTDATE	DUID	BIDTYPE	BIDSETTLEMENTDATE	OFFERDATE	VERSIONNO	PARTICIF
3	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID1	ENERGY	2021/03/30 00:00	2021/03/30 12:19	1	PARTICIF
4	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID2	RAISE6SEC	2021/03/31 00:00	2021/03/31 01:07	1	PARTICIF
5	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID3	RAISE60SEC	2021/03/31 00:00	2021/03/30 20:05	1	PARTICIF
6	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID4	RAISE6SEC	2021/03/31 00:00	2021/03/30 20:05	1	PARTICIF
7	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID5	RAISE60SEC	2021/03/23 00:00	2021/03/23 17:00	1	PARTICIF
8	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID6	RAISE5MIN	2021/03/31 00:00	2021/03/01 09:53	1	PARTICIF
9	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID7	RAISE6SEC	2021/03/31 00:00	2021/03/01 09:53	1	PARTICIF
10	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID8	RAISEREG	2021/03/31 00:00	2021/03/01 09:53	1	PARTICIF
11	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID9	ENERGY	2021/03/31 00:00	2021/03/01 09:53	1	PARTICIF
12	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID10	LOWER6SEC	2021/03/31 00:00	2021/03/01 09:53	1	PARTICIF
13	D	BID	BIDDAYOFFER_D	2	2021/03/31 00:00	DUID11	ENERGY	2021/03/26 00:00	2021/02/25 17:21	1	PARTICIF

SAMPLE ONLY

Payloads discussion points:

- Proposed Payloads
- Schema Validation
- Transition Implications
- Other considerations

Audience Skill Set

- Technical Leads / Gateway / Support teams
- Integration Architecture Teams (Market Interface Specific)

Topics for discussion

- High level overview of target state for Payloads
- Current Payload formats across NEM Retail , NEM Wholesale, Gas and WEM.
- Schema validation
- Drafting requirements/use cases for target state Payload formats for NEM Retail , NEM Wholesale, Gas and WEM.

This working group discussion will be relevant to all stakeholders who participate in exchanging data between AEMO and energy stakeholders

IDX Focus Group Session: "LVI"

The objective of this focus group session is to discuss capabilities to be delivered in the IDX Low Volume Interface (LVI).

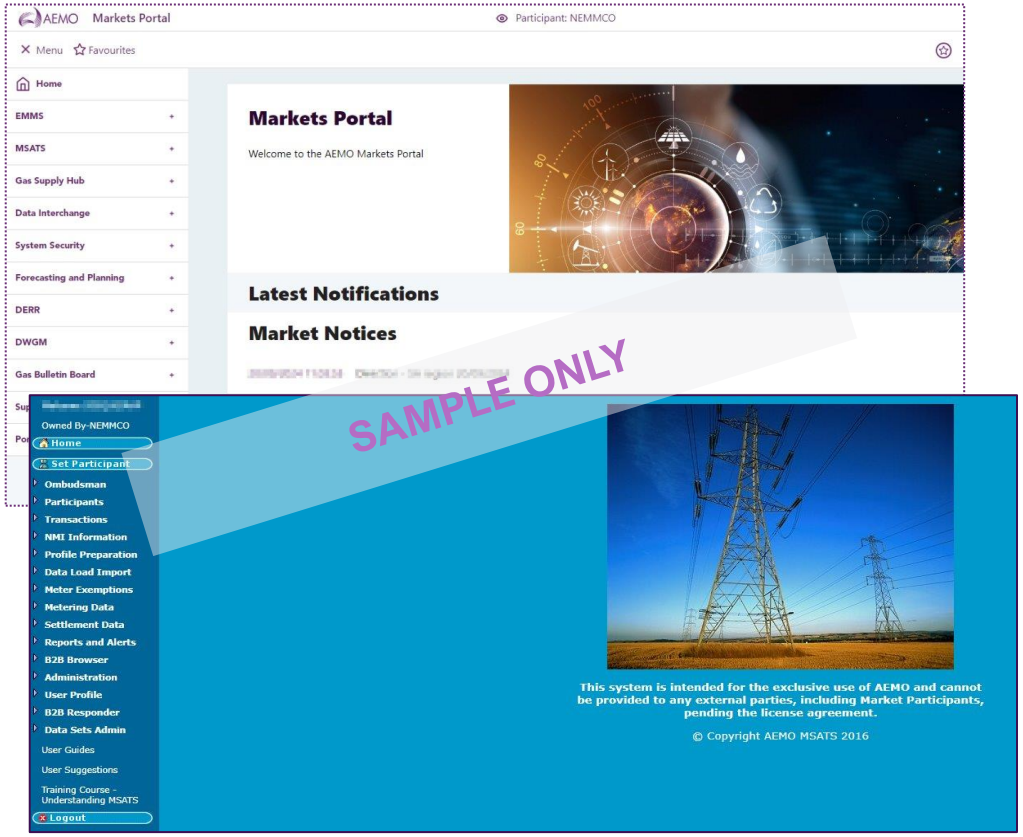

This focus group discussion will be relevant to all stakeholders who participate in exchanging data between AEMO and energy stakeholders via user interfaces

The **Low Volume Interface** is a user interface mechanism to allow participant to manage aspects of data exchange without the requirement of having a system-to-system interface.

Sample business drivers for consideration in the focus group are:

- Contingency (e.g. when participant systems fail)
- Operating small business without the need for system-to-system integration
- Accreditation

Audience Skill Set

- Users of existing market interfaces to input and receive market data.
- Technical/business leads who understand the use cases for existing User Interfaces

Topics for discussion

- Scope of the LVI
- Current user interface pain-points (e.g. as per the MSATS and Market Portal interfaces today)
- Capabilities required by the LVI e.g.:
  - Dashboards
  - Create, view and acknowledge transactions
  - Reports, analytics and insights.

IDX Focus Group: Large File Transfer

The objective of this focus group is to discuss the high-level overview and technical topics for the proposed Large File Transfer solution

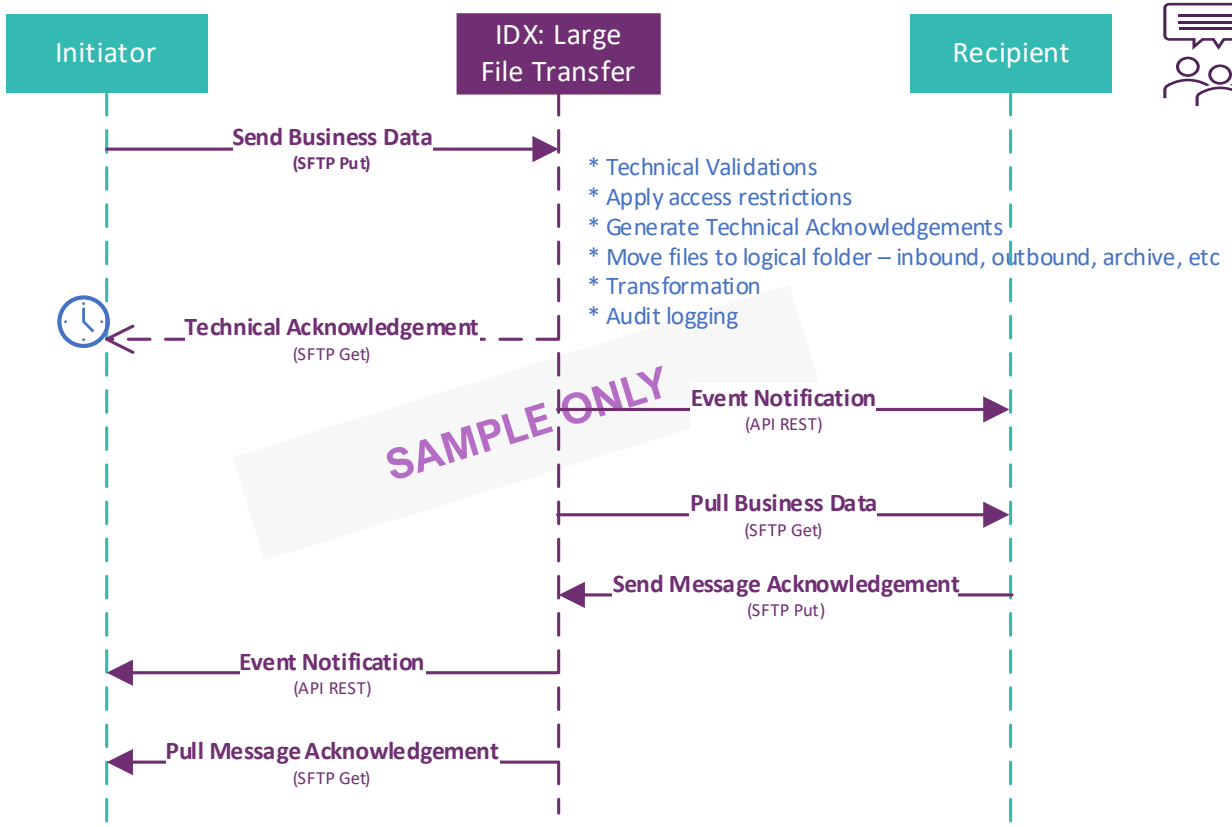

Audience Skill Set for Focus Group Discussion

- Technical Leads
- Integration Architecture Teams (Market Interface Specific)

Topics for discussion

- Principles used to develop Large File Share solution
- Overview of target state Large File Share solution
- Interface requirements for Large File Transfer
- High-level sequence diagram of the end-to-end business use case
- Roles and responsibilities
- Proposed folder structures and file naming conventions

This Focus group discussion will be relevant to all stakeholders who participate in exchanging data between AEMO and energy stakeholders

Notes

Andrew Bell spoke to the IDX Future topics

AEMO asked for participants to consider these future topics and committed to sending out a request for nominations.

11. General Business and Next Steps

NEMReform@aemo.com.au

General Business and Next Steps

MITE Working Group Forward Plan

Stream	Content	Timing
IDX and IDAM	**IDX** • Pilot-Lite scope • Decision Point 2 criteria **IDAM** • Future Topic Overview	**30 October**
IDX	**IDX** • Focus Group Playbacks: • Large File Share, Payloads, Low Volume Interface* • Sync and Fire & Forget Pattern	**27 November**

IDAM and IDX	• Focus Group Playbacks:
	  • IDX - AEMO Gateway Software, Inquiry Platform*
	  • IDAM - Organisation hierarchy, Security compliance obligations
	• Approach for 2025
		12 December

Focus Group Forward Plan

Stream	Topic	Timing
IDX	**IDX** Asynchronous Pattern Focus Group	~~**9 October**~~ **TBD**
IDAM	**IDAM** Organisation Hierarchy Use cases Focus Group	**18 October**
IDX	**IDX** Large File Share Focus Group	**WC 4 November**
IDX	**IDX** Payloads Focus Group	**WC 11 November**
IDAM	**IDAM** Security compliance obligations Focus Group	**18 November**
IDX	**IDX** Low Volume Interface Focus Group	**WC 18 November**

** Potential to be deferred dependent upon Focus Group content and feedback*

Focus Groups:
- **To be limited to 8 – 10 volunteers to ensure effectiveness, appropriate technical skillsets required (Refer to slides 85-88)**
- **AEMO to confirm attendees once nominations have close.**

AEMO welcomes Focus Group nominations to NEMReform@aemo.com.au **by 18 October**

AEMO | Market Interface Technology Enhancements

Notes

Blaine Miner spoke to General Business and Next Steps

Graeme Windley provided an update in regard CER data exchange industry co-design project. The IDX project is looking to provide some details around the proposed IDX Platform in the next CER Focus group. Graeme requested participants to consider engaging their internal teams who are participating in the CER Data Exchange Industry Co-design Project to align.

For more information visit

aemo.com.au

Appendix A

AEMO Competition Law - Meeting Protocol

AEMO Competition Law Meeting Protocol

AEMO is committed to complying with all applicable laws, including the Competition and Consumer Act 2010 (CCA). In any dealings with AEMO, all participants agree to adhere to the CCA at all times and to comply with appropriate protocols where required to do so.

AEMO has developed meeting protocols to support compliance with the CCA in working groups and other forums with energy stakeholders. Before attending, participants should confirm the application of the appropriate meeting protocol.

To access the full protocol at AEMO's website, visit: https://aemo.com.au/en/consultations/industry-forums-and-working-groups

Appendix B – IDAM End to End flow Example

Example: Bids Submission

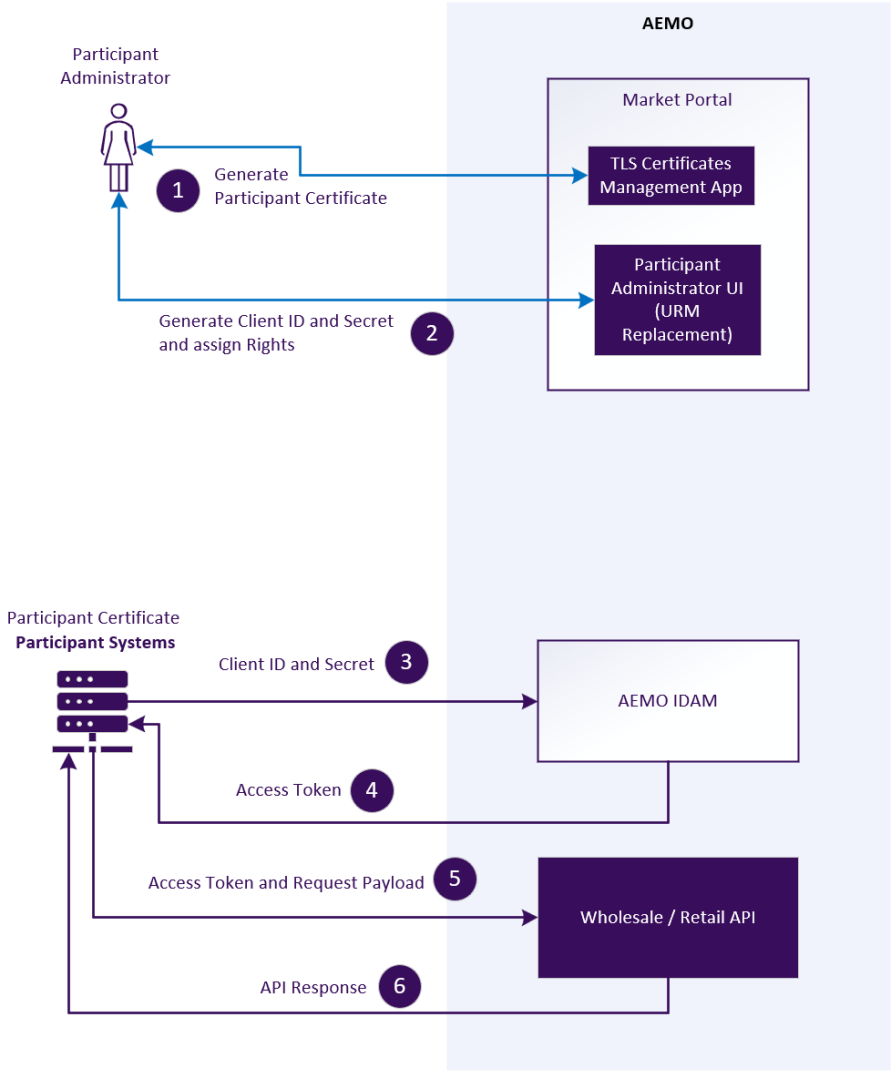

Steps	Pre-requisite step description
1	Participant administrator generates the Participant Certificate
2a	A client credential / service account is created using the Participant Administrator UI (replacement for URM). The credentials that get created will have a Client ID and a Secret.
2b	Participant administrator grants Rights to the newly created Client Credential using the Entitlements Management System, in this case to the **NEM Bidding and Dispatch** related Entities.

Steps	Run time step description
3	A participant system initiates a new session for the bid submission process by requesting a new Access Token for the NEM bid submission scopes (TBD) presenting the Participant Certificate and the Client Id and Secret from the steps above.
4	An Access Token is issued to the Participant System that is scoped to the Application Scopes (TBD) requested.
5	Participant System makes an API call to the NEM Bids Service on IDX presenting the Participant Certificate and Access Token to submit the bids.
6	Participant System successfully completes the submission and receives an Http Success response.

Note: Client Secrets based flow shown here is only shown for illustrative purposes, AEMO is looking into the feasibility of implementing RFC 8705 based Client Authentication

Failure Scenario: Requesting scopes that the Service Account do not have access to

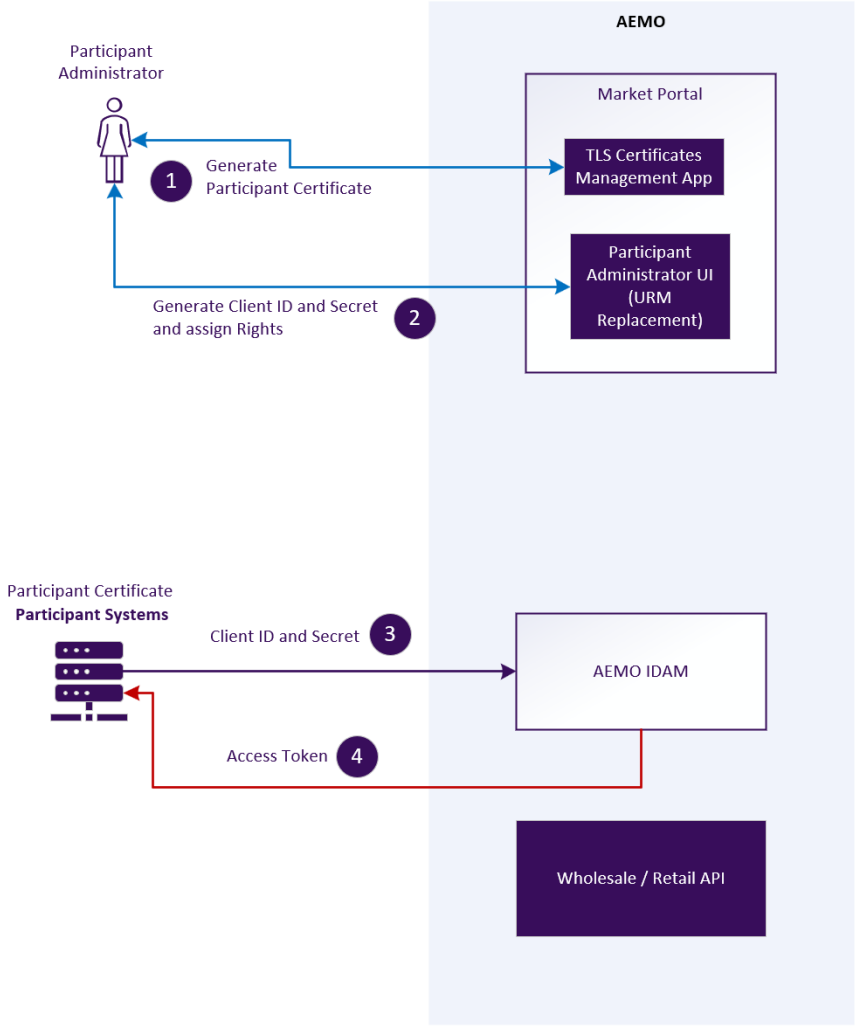

Steps	Pre-requisite step description
1	Participant administrator generates the Participant Certificate as a one time activity
2a	A client credential / service account is created using the Participant Administrator UI (replacement for URM). The credentials that get created will have a Client ID and a Secret.
2b	Participant administrator grants Rights to the newly created Client Credential using the Entitlements Management System, in this case to the **NEM Settlement** related Entities.

Steps	Run time step description
3	A participant system initiates a new session for a **bid submission** process by requesting a new Access Token for **NEM Bidding and Dispatch** scopes (TBD) presenting the Participant Certificate and the Client Id and Secret from the steps above.
4	Participant System Receives an HTTP 403 Forbidden indicating that, the while the Client Credential is valid, the NEM Bidding and Dispatch scopes requested are forbidden.

Note: Client Secrets based flow shown here is only shown for illustrative purposes, AEMO is looking into the feasibility of implementing RFC 8705 based Client Authentication

Failure Scenario: Invoking an API with an Incorrect Access Token

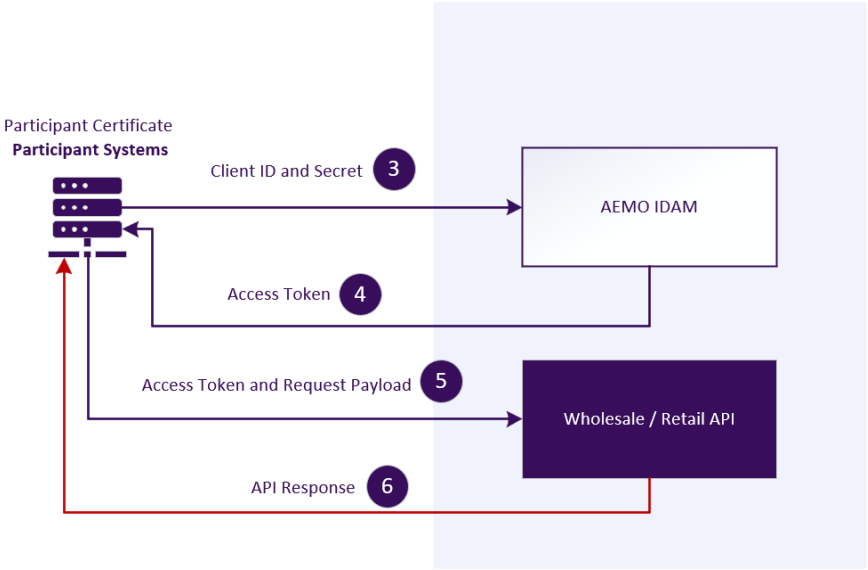

Steps	Pre-requisite step description
3	A participant system initiates a new session by requesting a new Access Token for the NEM bid submission scopes (TBD) presenting the Participant Certificate and the Client Id and Secret from the steps above.
4	An Access Token is issued to the Participant System that is scoped to the Application Scopes requested.

Steps	Run time step description
5	Participant System makes an API call to **NEM Settlement Service** on IDX presenting the Participant Certificate and incorrect Access Token for this service.
6	Participant System receives an HTTP 401 error indicating that the Access Token does not have access to the Settlement related functionality.

Note: Client Secrets based flow shown here is only shown for illustrative purposes, AEMO is looking into the feasibility of implementing RFC 8705 based Client Authentication

Appendix C – IDAM API Authentication and respone

Step 1: Token Request

Token Request Format

POST /oauth2/token HTTP/1.1

Headers
Host: identity.aemo.com.au (TBD)
Content-Type: application/x-www-form-urlencoded
Authorisation: Basic base64(client_id:client_secret) # Basic Auth for client credentials.

Body (Form Parameters):
grant_type=client_credentials
scope=nem-bids.crud nem-fcas-bids.crud (TBD)

If the service account / client credential has access to multiple Participant IDs, the access token requested will have access to all those PIDs for the requested scopes**.

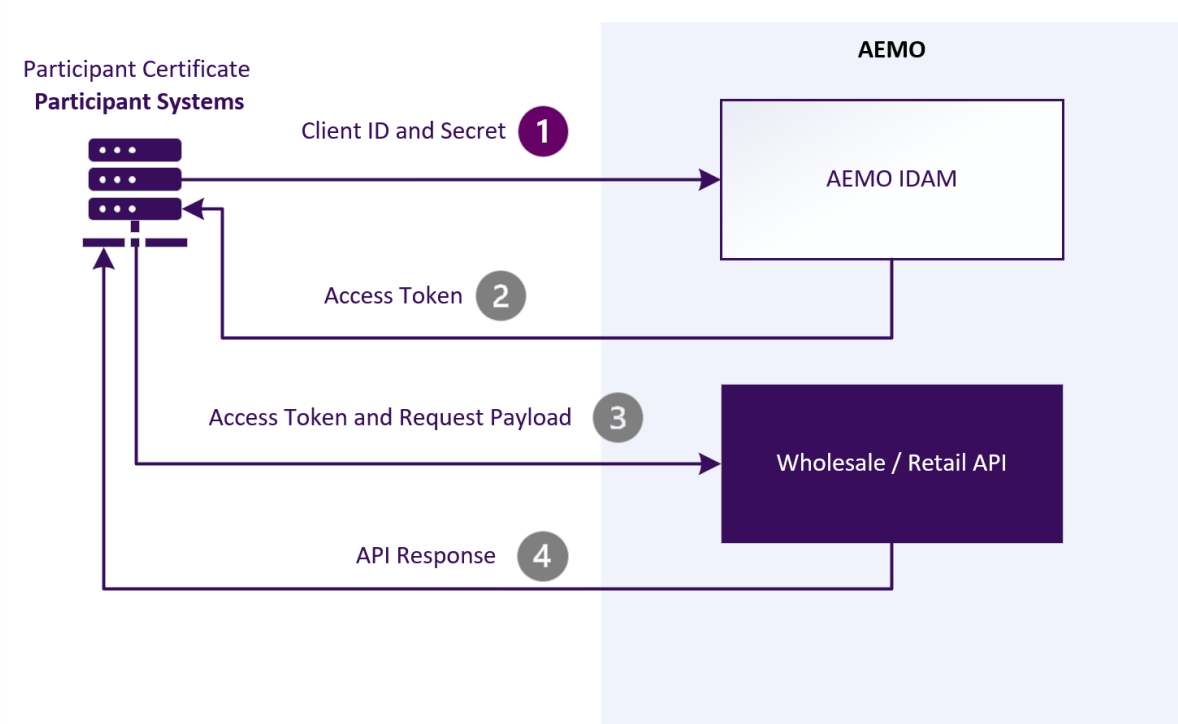

** ***Application scopes*** *are a fundamental concept in OAuth 2.x. Scopes define the level of access or permissions that an application (client) is requesting when it interacts with a resource server (API). Scopes allow the API provider to limit or grant access to specific actions or resources based on what the client is authorised to do. All scopes would be market specific.*

Discussion Points
- AEMO is seeking feedback from participants on the approach to requesting access to Multiple PIDs within the same Access token

Step 2: Token Issuance

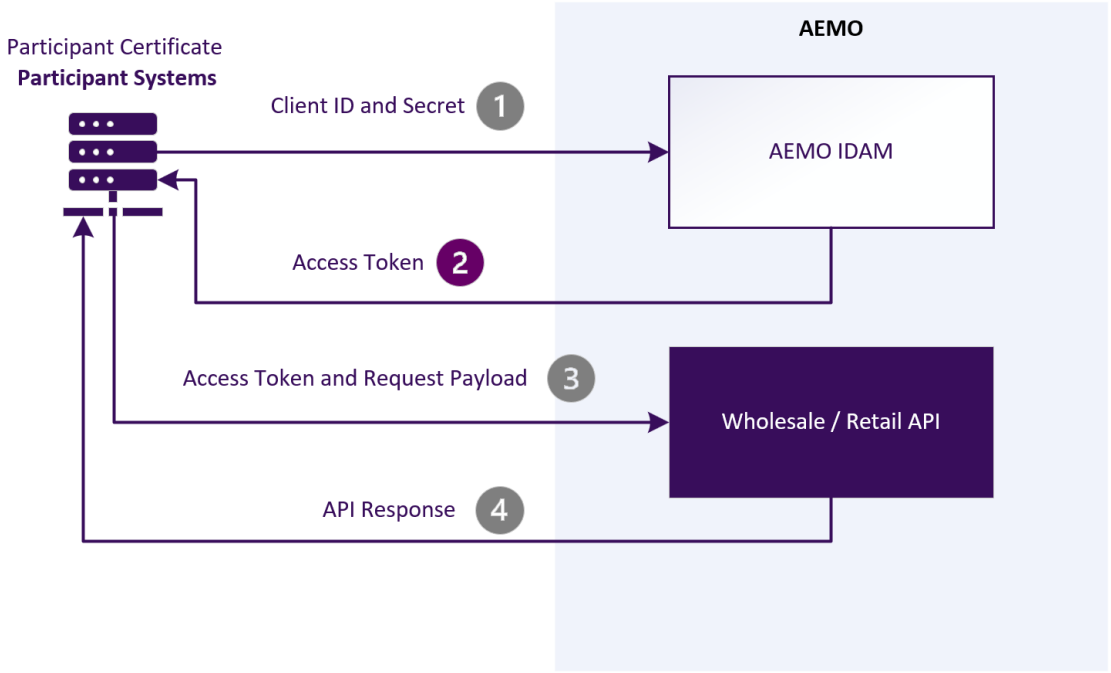

Token Issuance:

If the client credentials are valid, the authorisation server issues an **access token** to the client based on the requested scope in step 1.

The Access Tokens will have an agreed Time To Live, after which the token expires. When the token expires, a new Access Token can be requested by repeating Step 1 of the Token Request process.

Note: Client Secrets based flow shown here is only shown for illustrative purposes, AEMO is looking into the feasibility of implementing RFC 8705 based Client Authentication

Step 3: API Request (With Access Token)

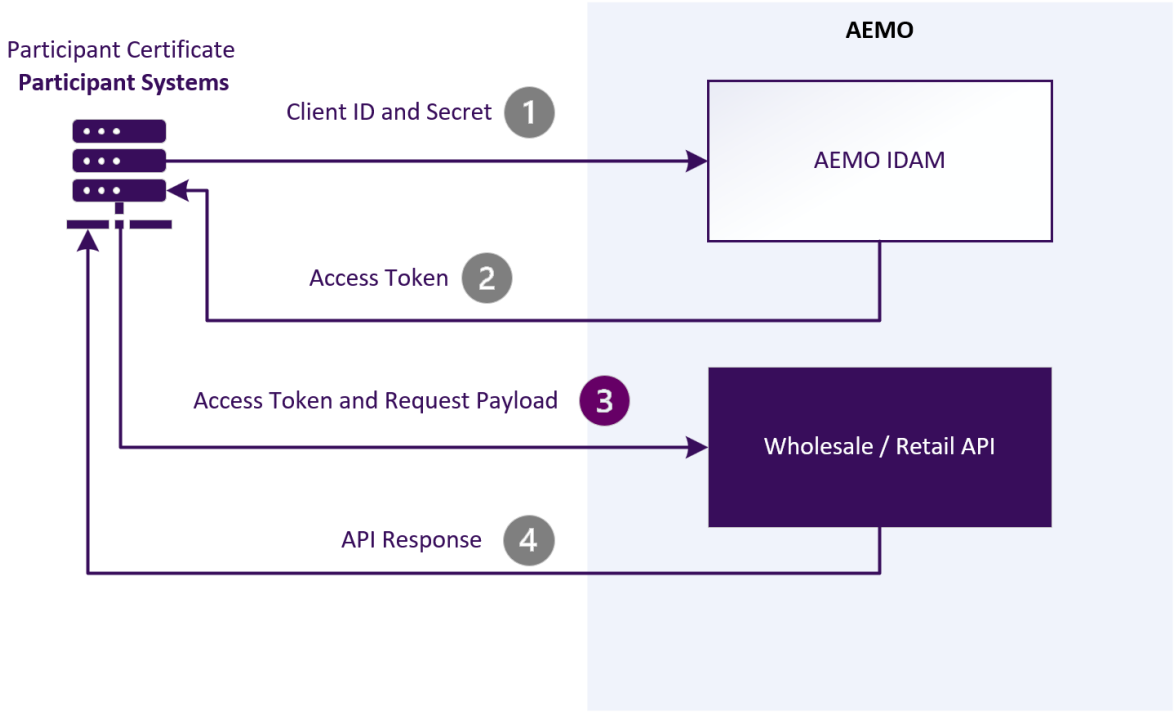

Note: Client Secrets based flow shown here is only shown for illustrative purposes, AEMO is looking into the feasibility of implementing RFC 8705 based Client Authentication

API Request (With Access Token):

- The client includes the access token obtained from the authorisation server in the **Authorisation: Bearer <access_token>** header when making API requests.
- The client makes these API requests over **mutual TLS** to the API server.

API Server Token validation:

- The API server validates the **access token** to ensure
  - The token is valid and not expired.
  - The token contains the correct scopes and permissions.

Step 4: Access to API Resources

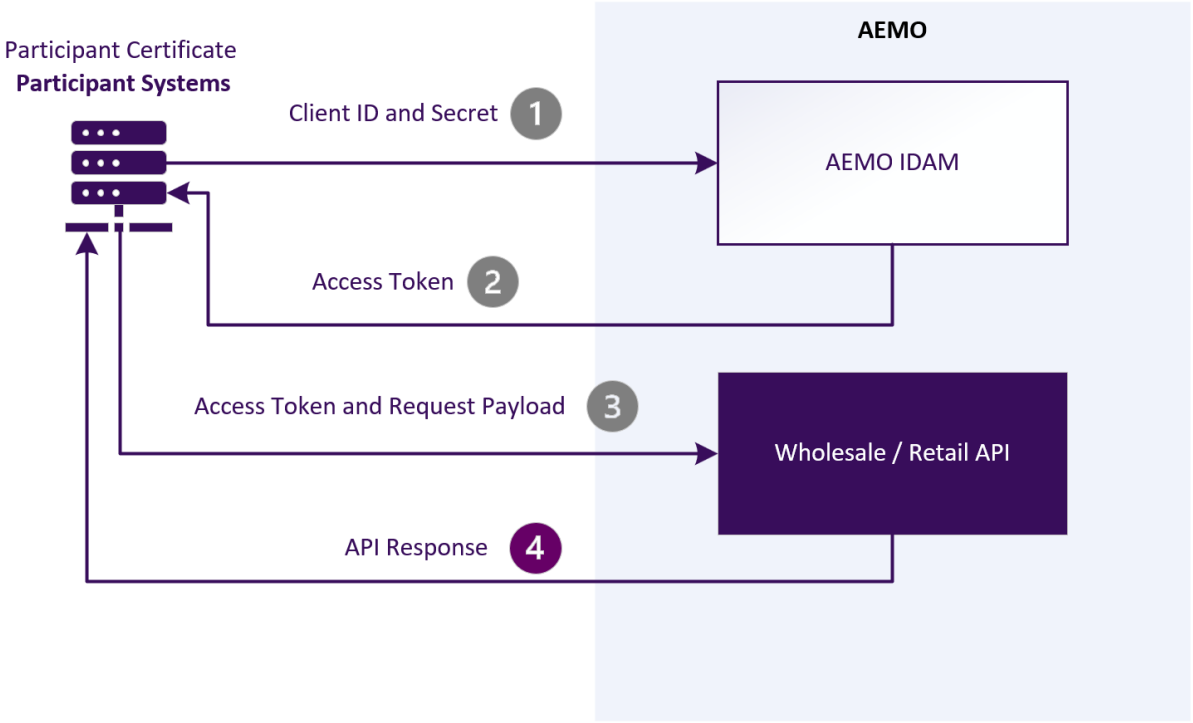

Note: Client Secrets based flow shown here is only shown for illustrative purposes, AEMO is looking into the feasibility of implementing RFC 8705 based Client Authentication

Access to APIs

Once the API server verifies the access token, it grants access to the requested resource.

If the token is invalid or mismatched, the server rejects the request, returning appropriate error responses (e.g., 401 Unauthorized or 403 Forbidden).

The client application will be able to initiate multiple API calls, in parallel or sequentially as required, using the same Access Token until the token expires.

Discussion Points

- AEMO is seeking feedback from participants on the approach for renewing access tokens when the Access Tokens expire.

Appendix D – Message threshold options

Message Threshold options in detail

Message Size Threshold – Option 1

Option 1: Global Threshold Value for all Business Functions

1) Global threshold value that is common for all business functions
2) Threshold value is a global value and does not vary based on other factors such as
    - Business function & use cases that support the business function
    - Schema type to support the business function
    - Message payload carrying transactional data or report data or event data
3) API Gateways operate efficiently when the message size is optimal. Considering the industry best practices and the limitations of the API Gateway technology, threshold value of 10MB is being recommended (for this option).
4) However, if Option 1 is preferred by AEMO & the Industry, AEMO will seek Participants' feedback on setting the threshold value to 10MB or a lesser value. For illustration purposes, a value of 10MB is considered.

Worked Example

Use Case	Target State Schema Type	Threshold Value
Submit Bids	JSON	10MB
Retrieve Bidding Public Reports	AEMOCSV	10MB
Retail B2B - Service Orders	JSON (or XML)	10MB
Retail B2B – Meter Reads	AEMOCSV (or MDFF)	10MB
Retail B2M – CATS	JSON (or XML)	10MB
Gas B2M & B2B data flows	JSON (or XML)	10MB

Pros

1) Minimises operational and governance overheads by managing a global threshold value for all business functions

Cons

1) Impact on performance when transforming large transactional messages in JSON (or XML) formats for a message size of 10MB
2) Participants need to augment/uplift the capacity of their infrastructure (@ the integration and/or application layer) to manage the message size of 10MB across the board; impacting their capital and operational costs

Message Size Threshold – Option 2

Option 2: Threshold Value @ Business Function Level

1) Threshold value will be set @ Business Function Level
2) Threshold set at the business function level will apply to each of the resources of the business function
3) API Gateways operate efficiently when the message size is optimal. Considering the industry best practices and the limitations of the API Gateway technology, threshold value of 10MB is the maximum that can be set for RESTful channel.
4) Worked example shows a variety of threshold values such as 10MB, 2MB; values are for illustration purposes only. The optimal value of the threshold will be agreed in consultation with the participants; post locking the message size threshold option

Worked Example

Use Case	Target State Schema Type	Threshold Value
BF: energyFCASBids Resource: Submit Bids	JSON	10MB
BF: energyFCASBids Resource: /reports/public	AEMOCSV	10MB
BF: Service Order Resource: submissions	JSON (or XML)	2MB
BF: Meter Reads Resource: Send meter reads	AEMOCSV (or MDFF)	10MB
BF: Meter Reads Resource: PMD/VMD	JSON (or XML)	10MB

Pros

1) Not all business functions are assigned a high message size threshold. Where the requirements demand a need, a high value of threshold is assigned
2) Participants who only consume the specific business functions/services are required to augment/uplift their infrastructure capacity to manage larger message sizes
3) Not all Participants' applications are required to manage large message sizes; the uplifts can be localised to the application that has to manage large message payloads

Cons

1) Operational and governance overheads in managing the threshold value for each business function
2) If the bundling of smaller transactions such as PMD/VMD are not time bound, chance of bundling too many smaller transactions into a 10MB message payload would lead to poor bundling practice

Appendix E - Current Archiving Use Cases

Archiving – NEM Retail

Current State	Retrieval	Retention
• For Messages sent by AEMO (outbound) - Messages are archived after an acknowledgement has been received from the Participant • For Messages sent by the Participant (inbound) – Messages are archived after an acknowledgement has been provided by AEMO • Eg CATS, B2B, MDMT and MTRD file types	• FTP, LVI (web interface) • DayZip functionality • Note: Not available for SMP (API's) • B2B: Message log and transaction log search txn and retrieve	• Data is retained in Archive folders for 13 months

Archiving – NEM Retail

NEM Retail provides archive folders where messages are archived. Participants access archived messages via the MSATS browser interface to navigate the archive and download files along with DayZip functionality

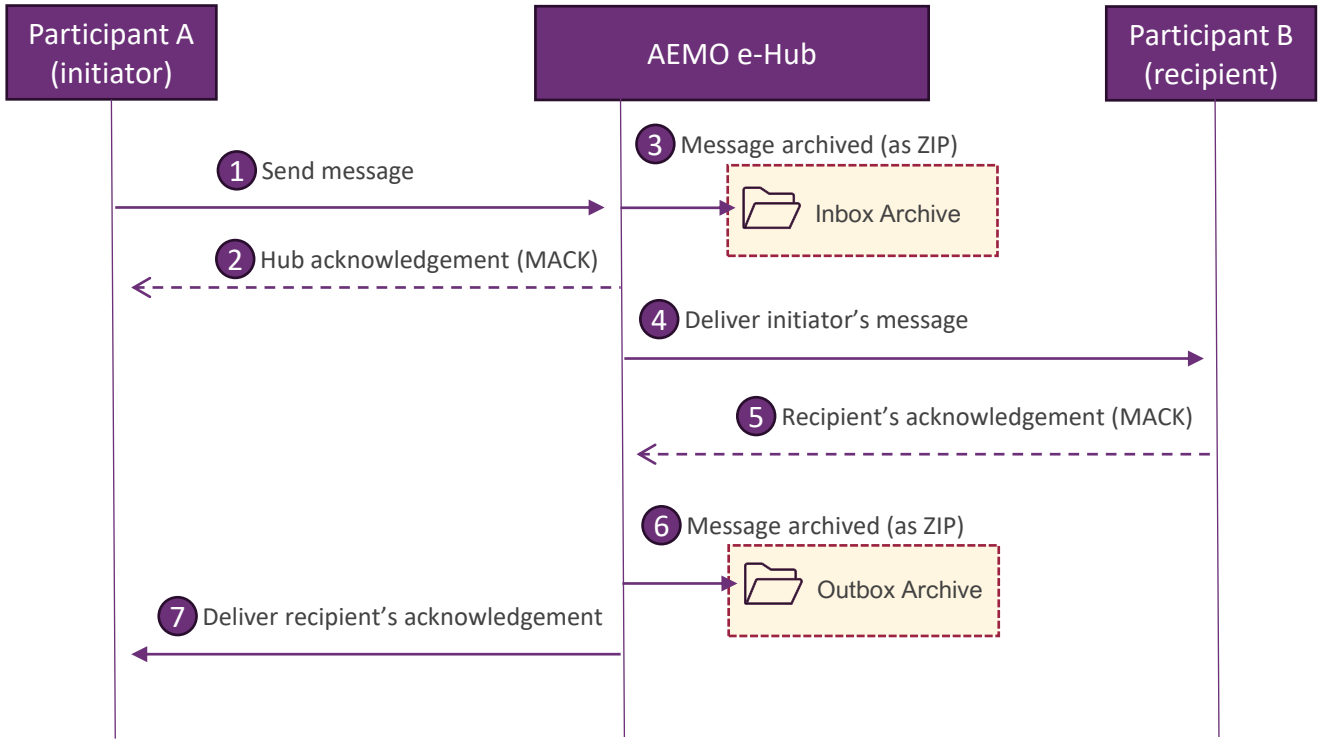

Sequence diagram example for B2BMessagingAsync API Archiving in SMP

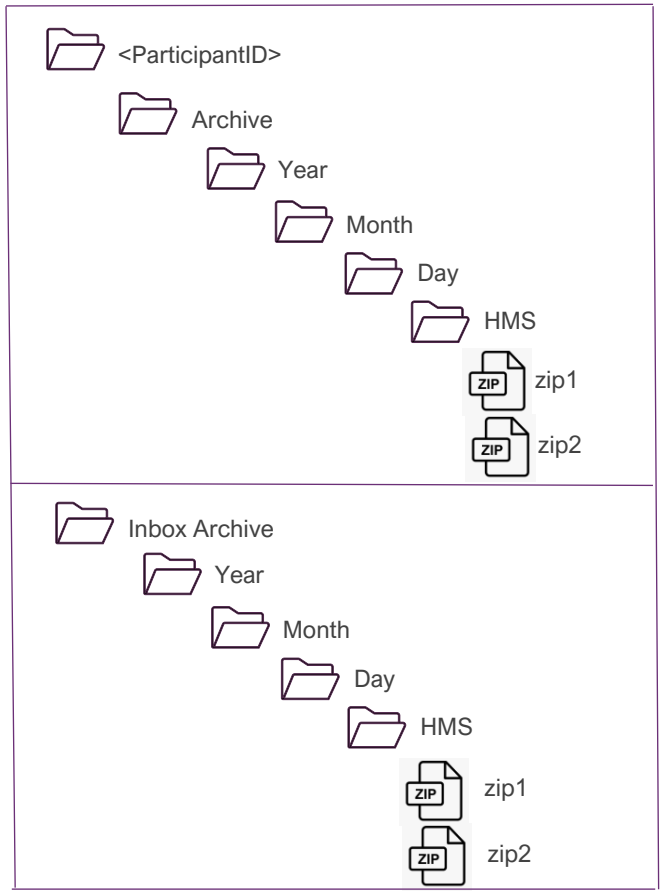

Archive Folder Structures visible to participants

Archiving – NEM Retail

The MSATS interface (Low Volume Interface, or LVI) is used by participants to navigate and retrieve archived messages.

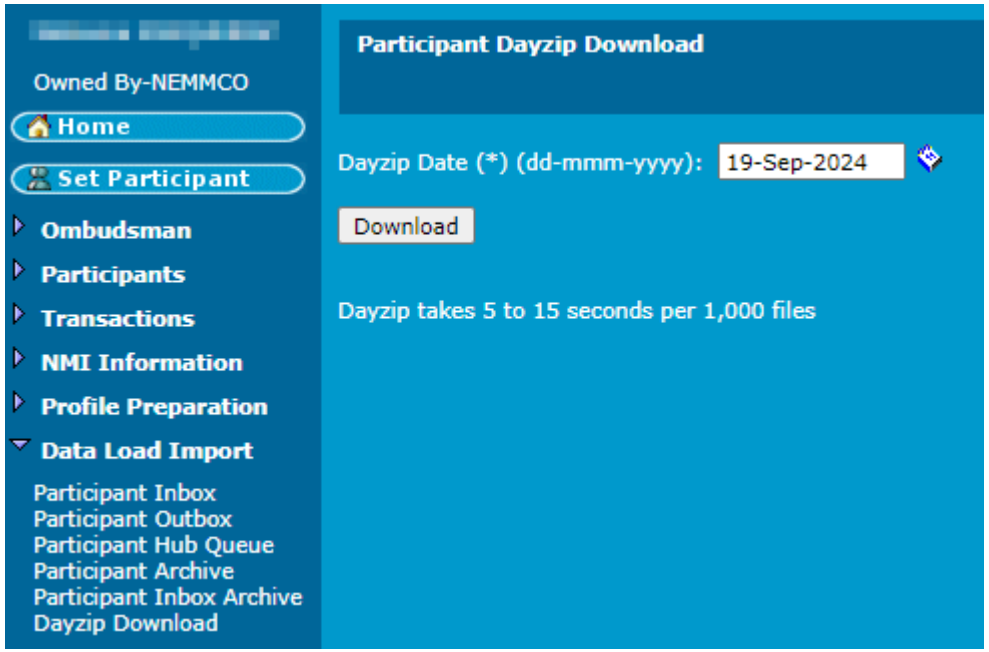

DayZip Retrieval using the LVI

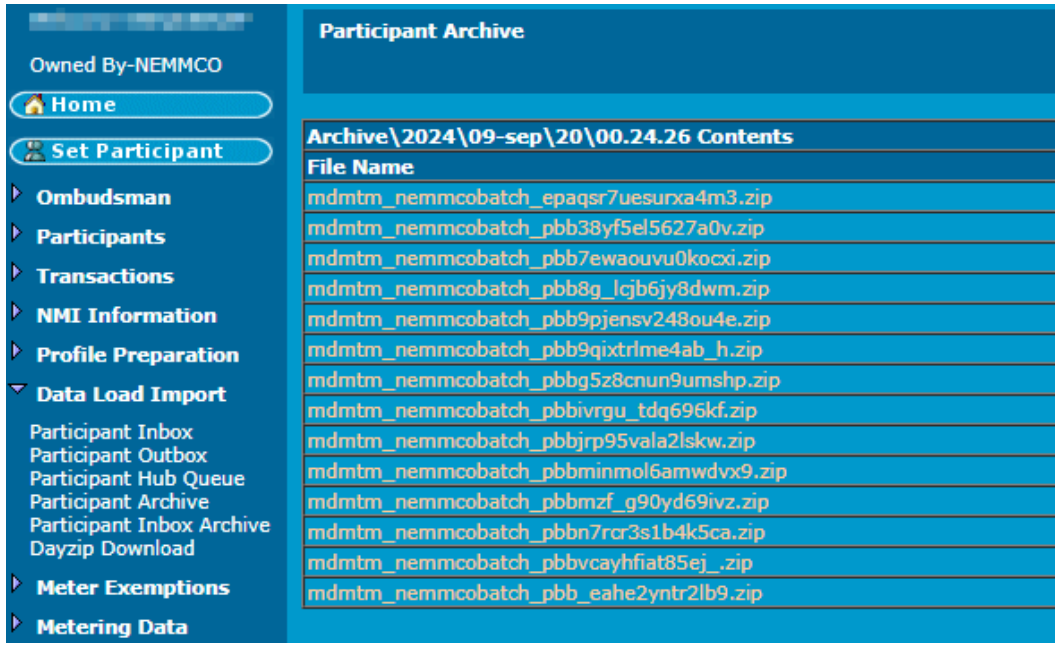

Retrieval using the LVI

Archiving – NEM Wholesale

Current State	Retrieval	Retention
• For Messages sent by AEMO (outbound) - Messages are archived and then distributed to the Participant • For Messages sent by the Participant (inbound) – Messages are archived after an acknowledgement has been provided by AEMO	• Files can be requested using PDR software and the manifest process (outbound only)	• Different report types have different retention policies defined. Minimum retention is 6 months

Archiving - NEM Wholesale
Outbound archiving

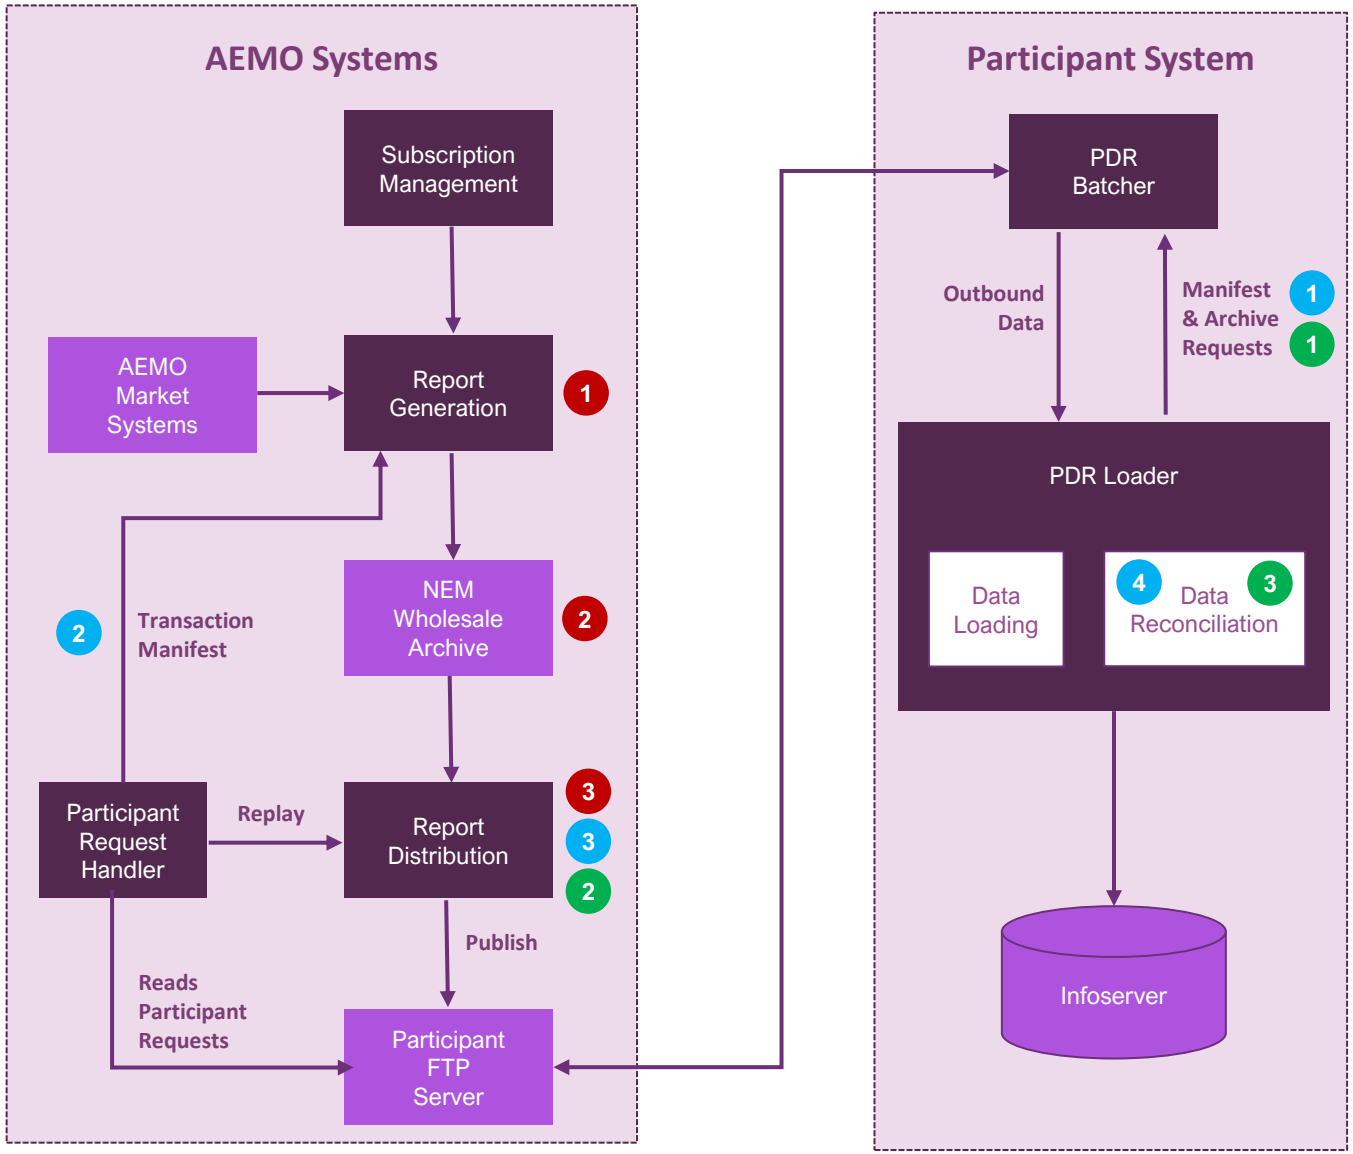

Primary Delivery Sequence

1. Reports are generated based on market events and subscriptions
2. Reports are deposited in the wholesale archive
3. Reports are distributed to participant facing channels

Transaction Reconciliation

1. Loader initiates a periodic transaction manifest request
2. Transaction manifest processed by AEMO
3. Transaction manifest delivered via standard outbound delivery process
4. Data reconciliation function detects any missing transactions

Archive Request

1. Loader initiates an archive request for missing transactions
2. Archive request processed by AEMO, outbound report delivered as replay event through outbound channel
3. Missing transaction marked as resolved

Archiving – Gas & WEM

Current State	Retrieval	Retention
• Files are copied to Participants base folder and archive • Only the most recent version of a file is retained in the base folder • Similar archiving methods are followed in Participant's Upload and Public folders. • Markets where this is used – DWGM, Gas/Vic (frc), QLD(frc) STTM – QLD hubs, Adelaide and Sydney Hubs • The WEM does not have a formal archiving capability	• Portals like the MIBB and MIS are used to browse and retrieve data • FTP can be used to download the data • For GSH and GBB files can be requested using PDR software and the manifest process (outbound only)	MIBB and MIS • Files in archive remain for 2 weeks • Zips of Archives are created and stored in the /Archive/Zip directory. These generally have 4 weeks of data GSH and GBB follows Nem Wholesale (Minimum retention is 6 months)

Archiving – Gas

Files for the DWGM and STTM markets are archived in Participant and Public folders. Only the most recent version of the files are retained.

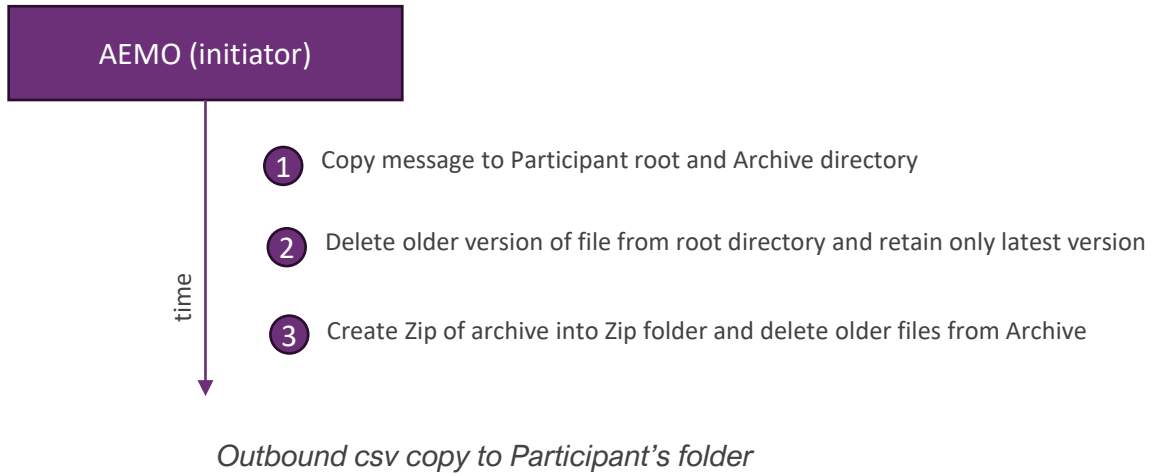

Outbound csv copy to Participant's folder

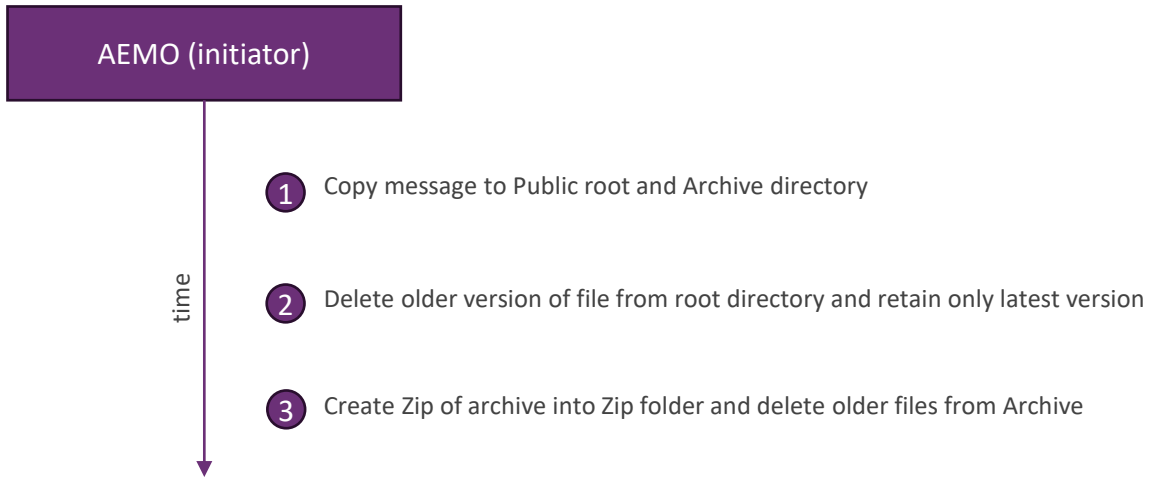

Outbound csv copy to Public folder

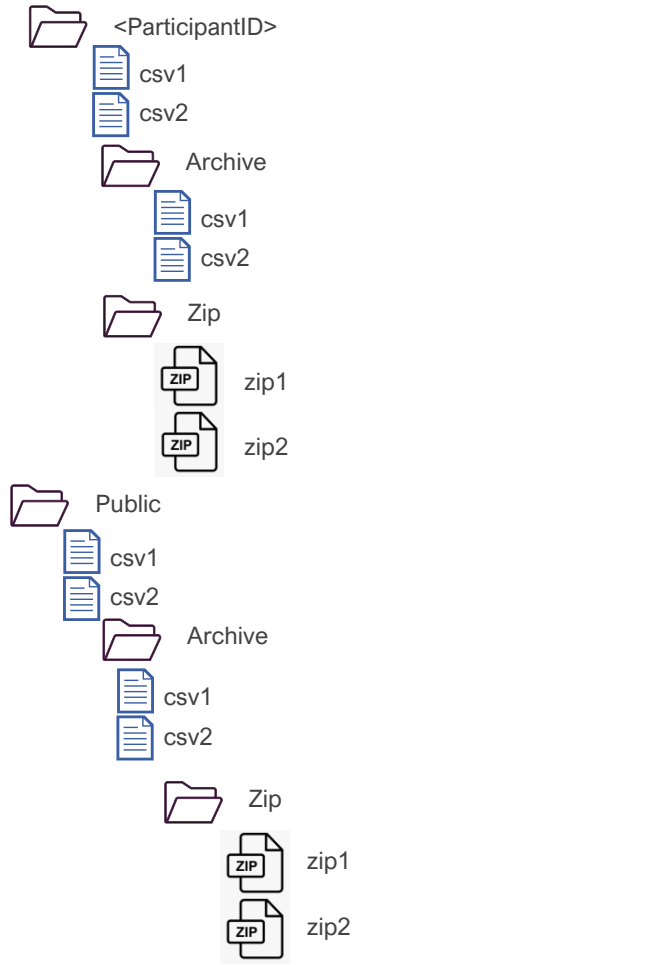

Gas DWGM and STTM use the hierarchical structure

Archiving –Gas

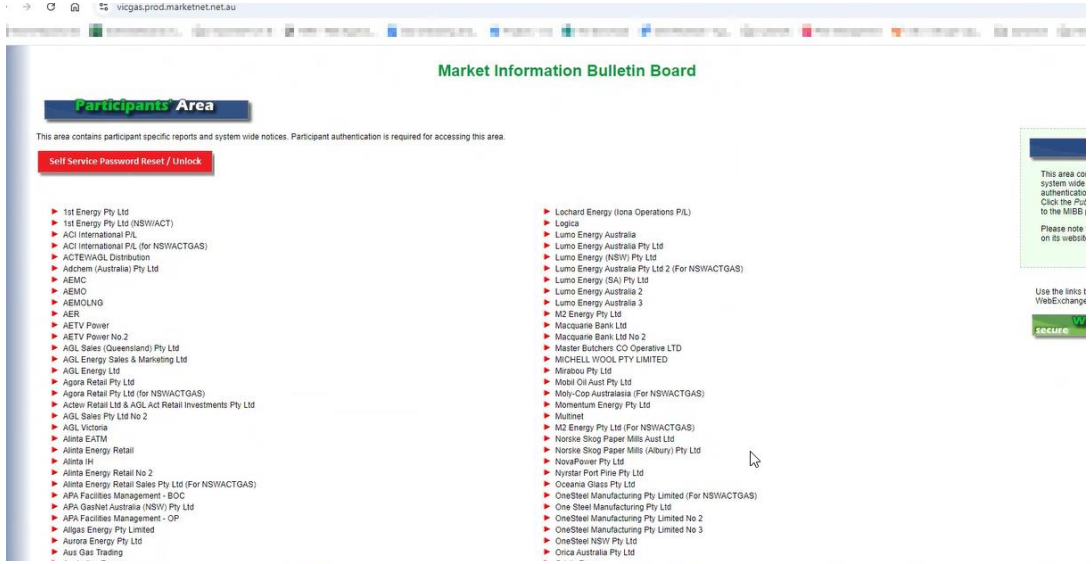

vicgas.prod.marketnet.net.au - /Public_Dir/Archive/Zip/

[To Parent Directory]

2/13/2023	1:30 PM	464062	Public_20230130_7.tmp
7/15/2024	1:41 PM	85635302	Public_20240701_7.zip
7/22/2024	1:30 PM	86215647	Public_20240708_7.zip
7/29/2024	1:36 PM	85132698	Public_20240715_7.zip
8/5/2024	1:39 PM	85876259	Public_20240722_7.zip
8/12/2024	1:40 PM	85917734	Public_20240729_7.zip
8/19/2024	1:05 PM	85978426	Public_20240805_7.zip
8/26/2024	1:39 PM	86437850	Public_20240812_7.zip
9/2/2024	1:32 PM	85693883	Public_20240819_7.zip
9/9/2024	1:41 PM	86332734	Public_20240826_7.zip

Current State Capabilities

When shown in a matrix view, it demonstrates the various approaches across markets for accessing archiving.

Mechanism to access Archive

Markets	FTP	Webpage (LVI)	Webservices (API)	Data Interchange Software
NEM Retail	✔	✔		
NEM Wholesale				✔
Gas	✔	✔	✔	✔
WEM				

Appendix F – Additional Slides

AMQP publish/subscribe Pattern assessment

- Does not offer the capability of transforming the outbound data on-request. AEMO is required to store the Participants' preference of the schema version; not resolving the pain points such as 'parkbox' process
- Does not offer flexibility for the Participants to configure the processing of outbound data based on parameters such as priority of the message
- Requires replicated architecture within the AEMO technology stack with associated performance impacts and increased implementation and run costs
- Cost overhead to introduce new Business Services.
- High Operational overheads for Participants & AEMO (to update client libraries) to manage the Data Delivery .

AEMO has consulted advisors including Gartner with regards to implementing the pub-sub pattern for external participant communication, their feedback :

- No known use cases of this pattern used for external communication.
- AMQP protocol is not ideal across network zones and quite unreliable over the internet.