DIGIT
Unit B4

Incubator report

Final Report

Date: 19/12/2014
Doc. Version: 2.0

Commission européenne, B-1049 Bruxelles / Europese Commissie, B-1049 Brussel - Belgium. Telephone: (32-2) 299 11 11.
Office: 05/45. Telephone: direct line (32-2) 2999659.

Commission européenne, L-2920 Luxembourg. Telephone: (352) 43 01-1.

Document Control Information

Settings	Value
Document Title:	Incubator Report
Project Title:	CISE
Document Author:	Bram De Schouwer, Gaëtan Vernaeve
Project Owner:	Vincent Dijkstra
Project Manager:	Vincent Dijkstra
Doc. Version:	2.0
Sensitivity:	Limited DG
Date:	19/12/2014

Document Approver(s) and Reviewer(s):

NOTE: All Approvers are required. Records of each approver must be maintained. All Reviewers in the list are considered required unless explicitly listed as Optional.

Name	Role	Action	Date
Joachim Beckh	MARSUR Coordination Group Chair, MARSUR Technical Working Group Vice Chair, DEU Delegate	*Page 8, change to read: "*such as a national MARSUR Node*", MARSUR is decentraliced and a MARSUR Node is always a national/organisational Node, never just MARSUR. Even when CISE interfaces the MEXS-to-MEXS side, it will see a participant, never the whole cooperation.*	12.12.2014

Document history:

The Document Author is authorized to make the following types of changes to the document without requiring that the document be re-approved:

- Editorial, formatting, and spelling
- Clarification

To request a change to this document, contact the Document Author or Owner.
Changes to this document are summarized in the following table in reverse chronological order (latest version first).

Revision	Date	Created by	Short Description of Changes
1.2	03/12/2014	DIGIT	Draft document
2.0	19/12/2014	DIGIT	Final document

TABLE OF CONTENTS

TABLE OF FIGURES

1 INTRODUCTION

1.1 Context

This document provides a list of recommendations and considerations based on the results of the CISE incubator. These recommendations and considerations intend to provide a valuable input for future work on the CISE specifications (e.g. the CISE Pre-Operational Validation project) and the implementation of those specifications. By listing a set of architectural recommendations, this document presents the challenges faced during the CISE incubator project.

This incubator started from some of the main outputs of the Cooperation Project, namely Work Package 2 (WP2) and Work Package 5 (WP5). The deliverables of these work packages define the CISE data model and specifications of the services to be offered by CISE in its initial release. Once translated into implementations, these services will enable different authorities to exchange information in an interoperable way. To verify that the specifications resulting from the Cooperation Project are suitable for implementation, these have been validated during the CISE Incubator project based on predefined usage scenarios, similar to real-life events. In doing so, it has been possible to validate whether or not the specifications can be realistically implemented into existing systems of different nature in the different public administrations of the Member States. Where the specifications are insufficient to implement the services, the required changes have been identified.

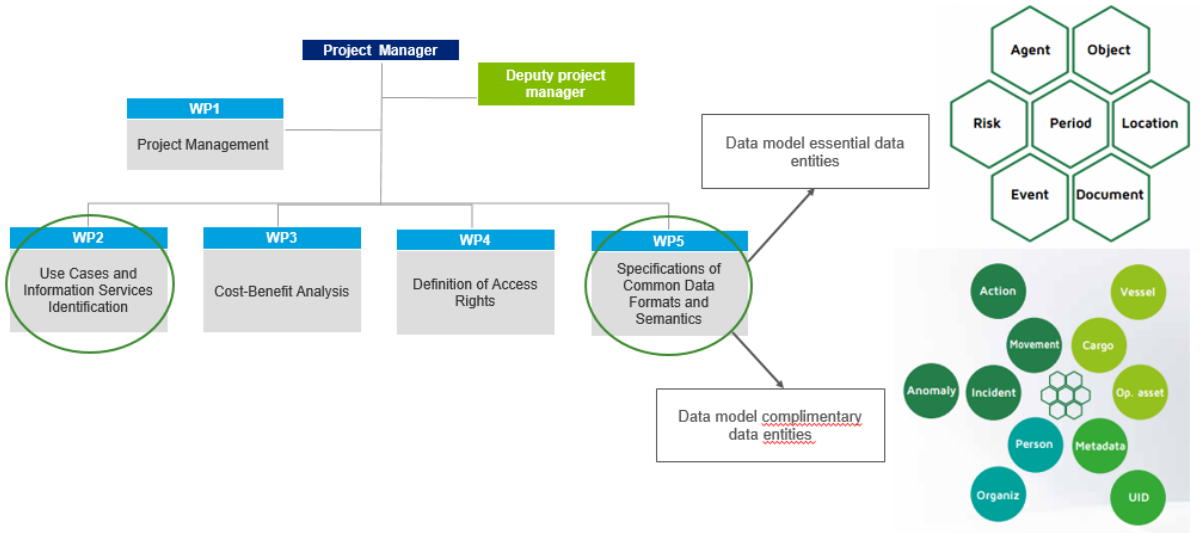

1.2 Purpose

The purpose of this document is to provide a valuable input for future work on the CISE specifications (e.g. the CISE Pre-Operational Validation project) in the form of recommendations and considerations learned during the CISE Incubator.

Chapter 2 provides an overview of the initially proposed architecture vision of CISE. This vision has been relevant to the CISE Incubator as it identifies the constraints and required flexibility of the operational CISE environment.

Chapter 3 provides more background information on the CISE Incubator, including its objectives, the different activities, approach and its constraints.

Chapter 4 describes the different recommendations and considerations which resulted from the CISE Incubator. These recommendations and considerations relate to both the validated specifications of the Cooperation Project and the Incubator itself.

Finally, Chapter 5 formulates a set of / lists overall conclusions.

1.3 Intended Audience

Due to the nature of the discussed subjects, this document is intended for people with a technical background having an interest in the results of the CISE Incubator. Before reading this document, it is recommended to read the Architecture Vision document of CISE (DIGIT, DG MARE, JRC, 2013) and the deliverables of the Cooperation Project.

2 ARCHITECTURE VISION

2.1 Description

The CISE Architecture Visions Document (DIGIT, DG MARE, JRC, 2013) analyses different visions to integrate User Communities. As a result of this study, a *hybrid* vision was selected as preferred solution by the Member States.

The hybrid vision allows Member States to decide whether to nominate a single or multiple providers of CISE services at national/sectorial level. This means that a provider of CISE services at national level may be nominated to deliver CISE services of interest for one or more User Communities. The hybrid vision applies a holistic and flexible governance model that takes into consideration both the national and the User Community perspectives.

2.2 Vision

The figure below shows how all elements will fit together in CISE's hybrid vision.

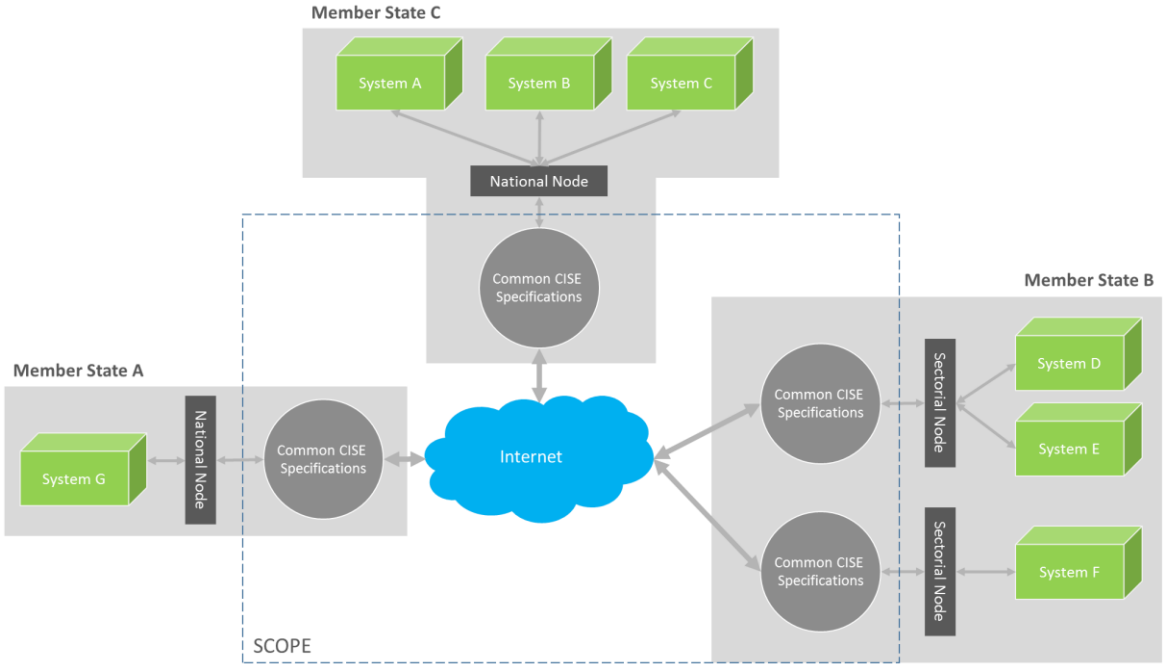

Figure 1 – CISE Hybrid Vision

The previous figure illustrates the hybrid vision, which will be used as a baseline for defining architectural recommendations. As a result of the hybrid vision, Member States are allowed to choose how they organise and connect their existing environment to the CISE network. Two main scenarios can be distinguished. In the first scenario, Member States connect a single national node aggregating one to multiple existing systems (e.g. Members State A and C). In the second scenario, the connection is made by connecting one or multiple sectorial nodes to the CISE network (e.g. Member State B). Whether a Member State connects to CISE through a sectorial or a national node should be transparent to the *CISE* network (see Figure 1 – CISE Hybrid Vision). In both scenarios, it is the full responsibility of the sectorial/national node to aggregate data from the underlying systems.

The goal of this document is to formulate architectural recommendations on the CISE architecture while supporting the previously described hybrid vision. The architectural considerations for the sectorial and national nodes are out of scope of this document since these are considered a Member State responsibility and will not be based on the CISE specifications. Finally, the scope of this

document is limited to the CISE environment and does not consider integration with other environments such as i.e. a national MARSUR Node.

2.3 Principles

The following list describes the main **architectural principles** of CISE. These are based on best practices and principles defined in section 4 of the Architecture Vision which apply to the CISE architecture. These principles should be validated and – if necessary – fine-tuned by all participating CISE stakeholders.

- CISE must allow interlinking any public authority in the EU and in the EEA involved in maritime surveillance.
- CISE must increase maritime awareness based on need-to-know and responsibility-to-share principles.
- CISE must privilege a decentralised approach at EU-level.
- CISE must allow interoperability among civilian and military information systems.
- CISE must allow interoperability among information systems at the European, national, sectorial and regional level.
- CISE must privilege reuse of existing tools, technologies and systems.
- CISE must allow seamless and secure exchanges of any type of information relevant for maritime surveillance.
- CISE must be system neutral.
- CISE must make it possible for information providers to change their service offering.

3 CISE INCUBATOR

3.1 Objectives

The concept paper of the CISE Incubator describes three main objectives:

- **Increase stakeholder engagement**
  By launching the Incubator, the Member States will be able to work in close collaboration with the Commission and will be able to provide valuable input for the finalisation of the specifications of CISE's services. The purpose is to increase collaboration between member states since the major part of the work will be performed between member states rather than with the European Commission.
- **Validate the feasibility of implementing a CISE service**
  One of the key inputs for the Incubator are the service specifications developed by the Cooperation Project. The Incubator will test the feasibility of implementing these services, as specified by the CP, so that potential issues are identified and discussed before the start of the POV.
- **Draw lessons for future implementation and pave the way for the POV**
  By coordinating the entire implementation process, collaborating with the participating Member States and by gathering feedback during the different stages of the process, the participating Member States and the Commission will be able to draw important lessons for the next steps. These lessons should serve as an input to the upcoming Pre-Operational Validation (POV) project and for the refinement of CISE's requirements.

3.2 Activities

To answer the objectives above, an approach has been defined for the CISE Incubator, which takes into consideration:

- current and future activities for establishing CISE;
- results of the Cooperation Project and;
- preparation of the POV.

The picture below depicts the relationship between the CISE Incubator, the Cooperation Project and the POV and their respective relation to the development of the CISE specifications (Identify, Verify and Detail and Implement).

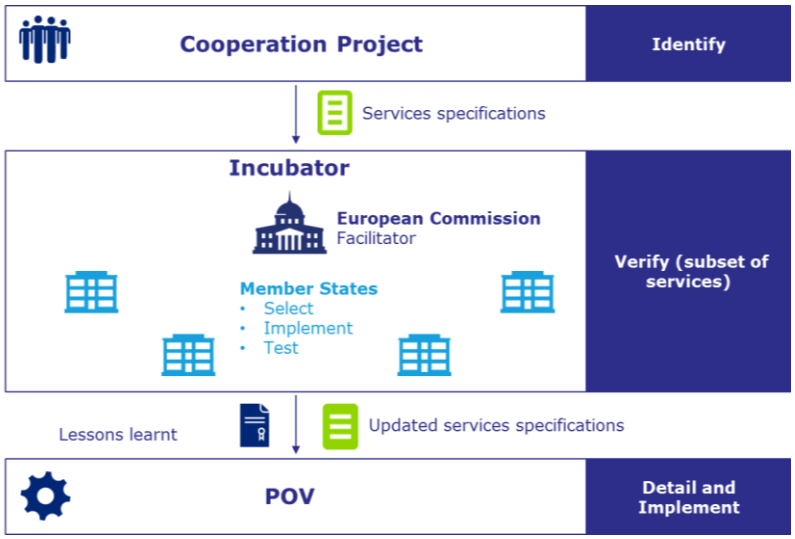

Figure 2 – Incubator positioning in the CISE planning

The Incubator tested the feasibility of implementing a subset of the CISE services on different systems of authorities in different Member States and exchanging information amongst these Member States. The different steps of the Incubator are visualised in the picture below:

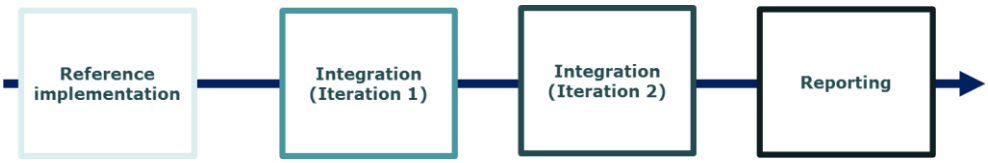

Figure 3 – Incubator Approach

1. **Reference implementation**
   From the deliverables of the Cooperation Project, a **selection of services** has been made to be implemented as part of the incubator. The selection of the services was based on the availability of required data by any of the participants. The service specifications have been refined into sufficient detail to allow their implementation. A **reference implementation** has been developed containing a common part for communication and message exchange. The reference implementation has been made available as a virtual environment to facilitate its deployment.
2. **Integration (Iteration 1 and 2)**
   During two iterations, the incubator participants developed different **services to be either provided or consumed**. A workshop was organised with the different participants at the end of each iteration to demonstrate the actual exchange of information. To mimic the value of CISE in a real environment, a **scenario** has been developed to simulate a real world situation in the context of maritime surveillance. By combining the different services of the Incubator participants, **cross-sectorial information exchange** has been made possible (see 7 Annex
3. Incubator sequence diagram).
4. **Reporting**
   During each of the activities throughout the CISE Incubator, **lessons learned** have been identified and documented. These will be further explained in chapter 4 of this document and have been presented to the CISE community.

3.3 Constraints

- CISE aims at interconnecting and facilitating the interoperability of both existing and new systems. These systems are characterised by the fact that they are all based on different technologies and protocols.
- The participation in CISE is voluntary and is not enforced by a directive for the Member States. This implies that the complexity of CISE participation should be minimised. Some scenarios might not be selected due to their complexity or high cost of implementation. Indeed, participations requiring complex integration and administrative burden often come with high costs which are not conducive to the purpose of CISE.
- According to the hybrid vision, different levels of granularity are possible when exposing Member States' systems to CISE. Some Member States will be using one node at national level whereas other Member States will rely on multiple sectorial nodes. It is therefore assumed that the number of connected nodes is unpredictable.
- Each of the CISE participants should be in control of their information. Consolidating all available information at a European central level is considered not to be an option.

4 RECOMMENDATIONS

4.1 Introduction

The CISE incubator, as described in the previous chapter, resulted in a variety of lessons learned. Many of these are related to the deliverables of the cooperation project, namely the CISE services, CISE data model and CISE message exchange patterns. In addition to these, the CISE incubator also resulted in lessons learned regarding the CISE infrastructure and the CISE topology. Finally, a number of lessons learned relate to the concept of the CISE Incubator itself.

The various lessons learned have been classified into different domains. The picture below positions these domains into a framework used for the categorisation of the different recommendations in the next sections.

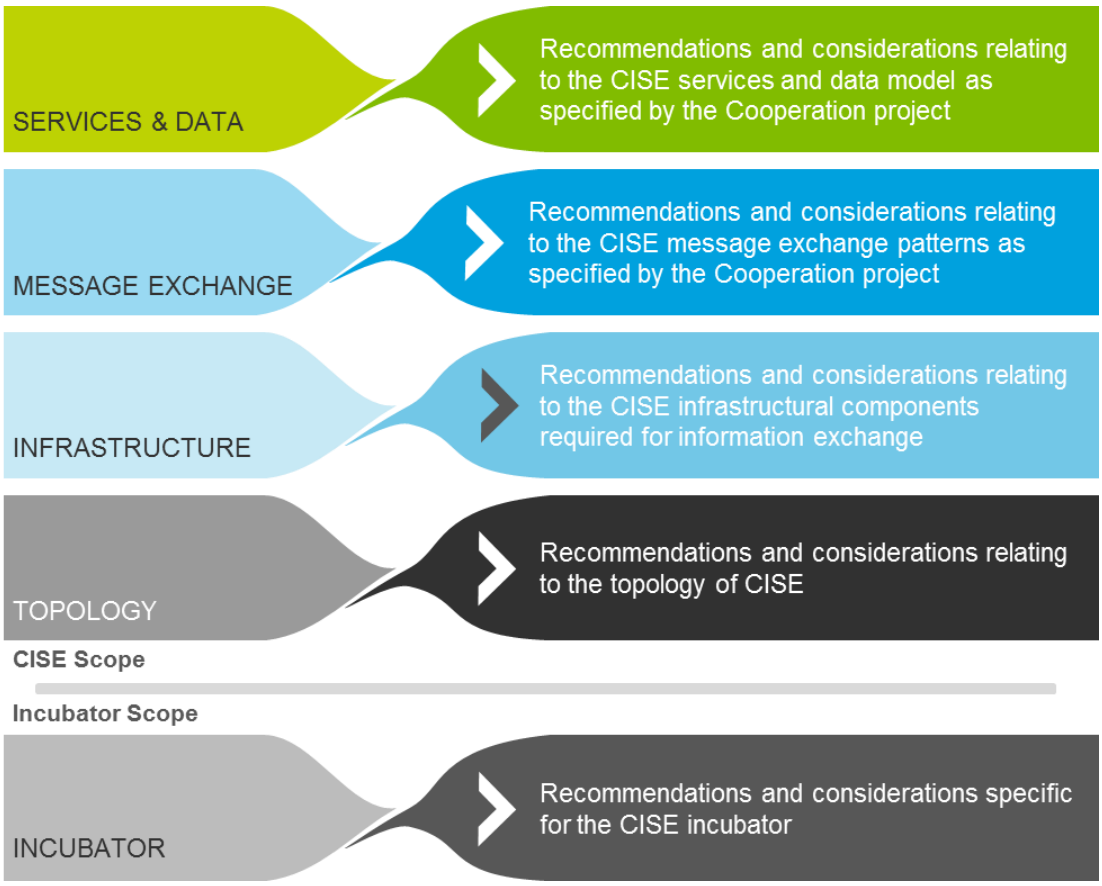

Figure 4 – Framework

4.2 Services & Data

This section provides recommendations and considerations relating to the CISE services and data model as specified by the Cooperation project.

4.2.1 RD-001 – Architectural Design Pattern

Issue Statement: How to integrate existing and new CISE systems using different technologies?

The concept of Service Oriented Architecture (SOA) allows for the organisation and distribution of capabilities across the different CISE participants. It allows for CISE participants to hide their

technology stack behind a standardised interface. Moreover, SOA enforces the use of standardized service contracts, which are decoupled from the service implementation of the different CISE participants. As a result, the desired level of interoperability is achieved between the CISE nodes.

Recommendation	Incubator	
Use of Service Oriented Architecture (SOA) as the preferred architectural design pattern.	☑	Implemented & Tested
	☐	Implemented
	☐	Discussed

4.2.2 RD-002 – Service Registry

Issue Statement: How can CISE participants find CISE services provided by other CISE partners?

A service registry, which acts as a repository for managing the life cycle of different services exposed by a CISE service provider can be used to communicate its offering towards other participants. This service registry contains the list of services that CISE participants are willing to provide to the CISE community as well as its service levels and other specifications. Whenever a CISE participant wants to consume services, it can query the service registry to retrieve the list of service providers that can fulfil its request.

A more detailed investigation on how and which information should be available in the service registry should be performed. ebXML is an example of a standard that defines and elaborates on the concept of a service registry but with a broader scope. Its applicability should be assessed and other market standards should be investigated as a next step.

Some of the questions which should be answered by this investigation are:

- What information should be published in the service registry? (e.g. a description of the service, the service profile, information about the service provider, …)
- How can information be retrieved from the service registry? (e.g. as a list of service descriptions, through a query interface, …)
- If a query mechanism is used, how flexible should this mechanism be? (e.g. query on function, region, service profile attributes, …)
- What is the best way to deploy service registries? (e.g. should every participant have its own registry or should a central and commonly shared one be deployed?)

Recommendation	Incubator	
The use of a service registry is recommended to allow CISE participants to expose different services and service levels and provide the ability to manage their own service offering on CISE. It allows CISE participants to have – at any time – a consolidated overview of all services available on CISE.	☐	Implemented & Tested
	☐	Implemented
	☑	Discussed

4.2.3 RD-003 – Service Profile

Issue Statement: How can service providers and consumers communicate about the service levels that they offer/expect when exchanging information?

As a result of WP5 of the Cooperation Project, the concept of a Service Profile has been introduced. The purpose of the Service Profile is to define a standardised structure to describe the service levels which can be offered and expected by respectively service providers and service consumers.

Typically, service profiles include attributes such as page size, quality levels, expected response time or maximum file size. For example, a service consumer having a bandwidth constraint due to the use of satellite connections, might want to use a service profile that limits the page size and maximum file size to optimize the data transfer.

It is recommended to use service profiles both on the service provider side and the service consumer side. At the service provider side, service profiles could be stored in the service registry to provide additional details about the service entry. The concept of service profiles is particularly useful in an environment like CISE where different systems are being interconnected, each with their own service levels.

To optimally use service profiles, it is recommended to make the different elements of the profile as quantifiable as possible. This will allow service consumers to match their expectations with the available service profile of the service provider. To guarantee successful message exchange, the requested service profile of the service consumer should be a subset of the service profile of the service provider.

It is recommended that the service registries foresee query functionalities to search for services based on the attributes of the service profiles as this will allow service consumers to only access services that are able to meet their requirements.

Recommendation	Incubator	
The use of service profiles is recommended to allow service consumers and providers to expose the service levels which they are able to offer. This is particularly useful in an environment like CISE where different systems are being interconnected, each with their own service levels.	☐	Implemented & Tested
	☐	Implemented
	☑	Discussed

4.2.4 RD-004 – Data Model to Services

Issue Statement: How to translate the data model and the relationships between the different data entities to a set of services?

Many different data entities, attributes and relationships could be made available by the different CISE participants. It is important that the services for exposing and exchanging this information are defined with the objective to provide maximum added value for the consumers of the service.

The data model as described in the deliverables of WP5 of the Cooperation project also describes the relationships between the different data entities. In order to translate this model into query services, different approaches are possible:

- **One service**: One query service is created that allows for querying the entire data model. By doing this, it will be possible to query all attributes of the different data entities and their relationships. The main challenge when defining a single query service based on a single data model is the fact that most CISE participants will only have a small part of the data model available. As a consequence of this approach, specific rules will have to be defined for queries on attributes which CISE participants do not have in their own data model. A possible solution could be to implement a “best match” algorithm. However, this solution involves complex implementations and will lower the value of search results.
- **Multiple services based on data entities**: By offering multiple more specific services, only parts of the data model can be queried. The main challenge when defining these finer grained services is to define the service granularity. It would for example be possible to define separate services for each data entities and for each relationship. The benefit of a more granular service offering is that service providers can only offer services which correspond to the data entities they have available;

- **Multiple services based on business requirements**: In this case, services are defined based on the business needs of the different CISE participants. Query services could be built based on data entities and relationships when there is a clear business case for this service to be available. This situation could also stimulate service providers to offer additional data to increase the business value of their services.

Recommendation	Incubator	
Based on the options above, it is recommended to define services based on business requirements. However, as this is a process which requires in-depth study and analysis of the different CISE participants, the other options could be a good starting point. By already offering services which facilitate information exchange in general, new business needs could be discovered in the future.	☐	Implemented & Tested
	☐	Implemented
	☑	Discussed

4.2.5 RD-005 – Optional Data Model Attributes

Issue Statement: Because the current CISE data model is based on data available in different maritime standards and thus different contexts, a solution needs to be found to identify what information is available from the different CISE service providers.

In the current data model, most attributes are optional to allow for the different standards to integrate their data with CISE. However, optional elements make the operational part of CISE more complicated than needed. It makes it hard for a service consumers to know which attributes will be returned when requesting information about data entities. Therefore it is important to find a good balance between mandatory and optional data since more optional attributes make it difficult for systems to interpret data and make valuable decisions if they cannot rely on specific information to be available. The solution to this problem is constrained by the fact that information will be coming from different sectors and contexts. The purpose of CISE should not be to adapt these systems to comply with mandatory requirements defined in the data model.

As an alternative solution, a context-aware data model could be developed. This means that for each sector/context, an investigation is performed to identify typical information which can be supplied by service providers in that sector or context. Based on this information, specific data models can be defined as a subset of the current CISE data model. These specific data models can then include a subset of mandatory attributes. As a consequence, context dependent services can be defined based on these context dependent data models.

This approach fits our previous recommendations regarding the translation of the CISE data model into query services as these services are a first step towards business oriented services. Through smart combination of these services, more value adding services can then be defined in the future.

Recommendation	Incubator	
The use of a context-aware data model is recommended to provide a better business alignment as well as improved transparency towards the service consumers.	☐	Implemented & Tested
	☐	Implemented
	☑	Discussed

4.2.6 RD-006 – Data Entity Relationships

Issue Statement: What is the best way to represent data entity relationships in a structured representation language?

In the current CISE data model, relationships have been defined between different data entities. In some cases, these relationships are also described by a set of attributes (e.g. the type of relationship). Different options exist to model the relationships between data entities:

- As a set of attributes of the entity: This is currently the approach as defined by WP5 of the Cooperation project. Related entities and the properties of the relationship are defined as part of the description of a data entity. The downside of this approach is that the same relationship is defined twice in case of a bi-directional relationship.
- A separate entity: It is recommended to adopt an approach where the relationship and its attributes are considered a separate data entity, referring to the two entities it is connecting. This approach is recommended as this aligns with database modelling best practices (a separate table in be defined to describe a relationship with attributes and one-to-many or many-to-many relationships). As a result, entities will be loosely coupled from their relationship which will facilitate the retrieval of data through services. For instance, a *getObject()* service will only return the attributes of the *Object* entity without any information on the *event* which is a sub-entity of the *Object* entity.

Recommendation	Incubator	
Modelling the data entity relationships is recommended to ease the mapping of the data model to the underlying data structures. As CISE will rely on many existing systems, the mapping of the CISE data model to an existing data model, should be made as easy as possible.	☑	Implemented & Tested
	☐	Implemented
	☐	Discussed

4.3 Messaging Patterns

This section provides recommendations and considerations relating to the CISE messages exchange patterns as specified by the Cooperation project.

4.3.1 RD-007 – Direct Communication

Issue Statement: How can direct communications be implemented efficiently?

For direct communication, CISE participants know the destination of their messages. Though the destination of a message is known, the underlying environment is not. Gateways of different CISE participants are connected to local environments, outside the scope of CISE and each with their own service levels in terms of response times, performance, etc. This aligns with the hybrid model described in section 2.2 and – as a consequence – enforces the use of an asynchronous solution. Indeed, sectorial or national nodes may "forward" requests to the appropriate systems, using any existing communication channel. It cannot be expected that any of those communications are of synchronous nature. Even in the case they would be, it would tightly couple the request with the processing in each node involved with getting the reply. CISE strives for a loosely coupled integration optimizing resource allocation in the nodes that are part of its network.

Asynchronous message exchange is well-suited for processes during which the requesting CISE participant is blocked, waiting for results from a CISE service provider. They allow the CISE requester to decouple the request from the response by using call-back mechanisms and allow the provider to balance the load. When multiple responses are expected, this approach suits well for gradual retrieval of results without having to wait for all responses.

Whenever an asynchronous request is sent, an acknowledgment containing a unique identifier (correlation ID) should be returned immediately by the receiving system. Finally, in an asynchronous approach each reply message should contain a correlation ID that indicates which request message a reply is for.

The following picture further describes this process:

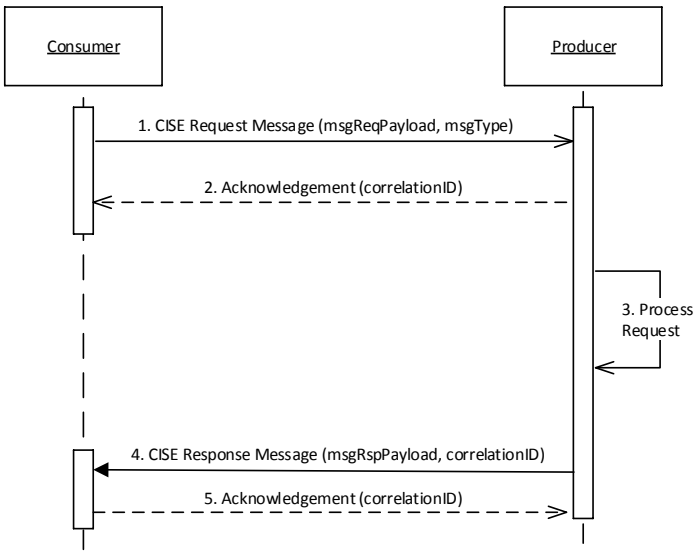

Figure 5 – Asynchronous processing

The following list describes the steps shown in the previous figure:

1. A CISE request message is sent to a provider. The message contains the message request payload and has a specific message type (e.g. Vessel query message)
2. Upon receipt of a request, the provider generates a correlation ID which is sent back in an acknowledgement message to the consumer.
3. The provider processes the received request while the consumer can continue processing something else.
4. The provider generates a response message containing the results of the request and attaches the correlation ID that has been generated in step 1. It is sent back to the consumer who will match the response to the request by using the correlation ID.
5. The consumer acknowledges receipt of the response message.

Recommendation	Incubator	
It is recommended to implement an asynchronous request/reply messaging pattern for all direct communications. This allows for a flexible and scalable solution where requests are decoupled from the reply messages. As a result, CISE gateways are not blocking transactions and can handle large amounts of requests by means of message queues.	☑	Implemented & Tested
	☐	Implemented
	☐	Discussed

4.3.2 RD-008 – Indirect Communication

Issue Statement: How can indirect communication be implemented efficiently?

Besides the request/reply messaging patterns described in WP5 (direct communications), CISE must support the exchange of notifications. Notifications are issued by one or more CISE participants and may relate to one or multiple topics. As CISE may have more than 300 connected authorities, the volume of notifications sent over the CISE network may increase exponentially and become a performance bottleneck when more authorities are involved. Moreover, CISE participants might not be interested in all notifications.

To implement a notification system, two use cases can be distinguished: receive notification and send notification. In the send notification use case, a CISE participant pushes information to other CISE participants, unaware of who is interested in the information. In the receive notification use case, a CISE participant knows what information it is interested in but is unaware of who has the information. The following sections further describe each use case.

Send Notification

When a CISE participant wants to issue a notification, topics describing the content of the message could be linked to this notification. This will allow receiving CISE participants to filter notifications based on topics in which they have expressed an interest. The process for the identification of the addressees is further elaborated in the next section.

Receive Notification

In this use case, three scenarios have been identified for the receipt of notifications:

1. Subscribe to all notifications that contain a specific topic.
2. Follow all notifications from a specific CISE participant.
3. Subscribe to all notifications that contain a specific topic and that were issued by a specific CISE participant.

Using this approach, CISE participants are responsible for subscribing themselves to specific notifications in which they have an interest. If they do not subscribe or follow any topics or participants they will not receive notifications. The analogy can be made with the Twitter example where a user can follow another user or subscribe to specific hashtags.

To build such a system, a subscription list containing the CISE participant's subscriptions and followers should be maintained. Depending on the chosen topology for the CISE architecture, this subscription list will be centralised or decentralised. Finally, further investigation is required to evaluate whether this subscription functionality should be embedded in the service registry.

Recommendation	Incubator	
It is recommended to use a publish / subscribe messaging pattern for notifications. Tags or topics describing notifications could be included as part of the published CISE notifications messages. This approach allows every CISE participant to define itself in which notifications it is interested by subscribing to specific topics or tags. Compared to a broadcast solution, this lowers data traffic and does not overload all CISE participants with messages in which they have no interest.	☐	Implemented & Tested
	☑	Implemented
	☐	Discussed

4.4 Infrastructure

This section provides recommendations and considerations relating to the CISE infrastructural components required for information exchange.

4.4.1 RD-009 – Gateway

Issue Statement: How can CISE provide efficient interoperability for CISE participants while having a clear separation between their local environments and the CISE specifications?

The CISE network should make use of gateways for the interconnection of the CISE participants. The CISE gateway should consists of a set of building blocks that implement common CISE specifications and will act as a bridge between the sectorial/national nodes and the CISE network.

As described by the hybrid architecture vision for CISE, each CISE participant will have to deploy one or more gateways. It is up to the CISE participant – in accordance with its own governance – to

define whether one or more gateways are used. In some cases, CISE participants may want to deploy a gateway per sectorial node whereas others may want to use only a single gateway for a national node.

The purpose of the gateway is to ensure reliable and secure communication between CISE participants. The gateway may consist of a number of architectural building blocks, such as a delivery manager, offering different functions. It is recommended that the gateway is kept message payload agnostic and thus does not implement any business logic. The three main reasons for this business agnosticism of the gateway are the following:

- It makes the gateway a set of reusable building blocks which could potentially be reused in environments similar to the CISE environment;
- It facilitates the deployment and maintenance of the gateway since no business/participant specific changes need to be taken into account when new versions of the gateway are deployed.
- The gateway is message payload agnostic, it can transport messages from one gateway to the other regardless of the message content.

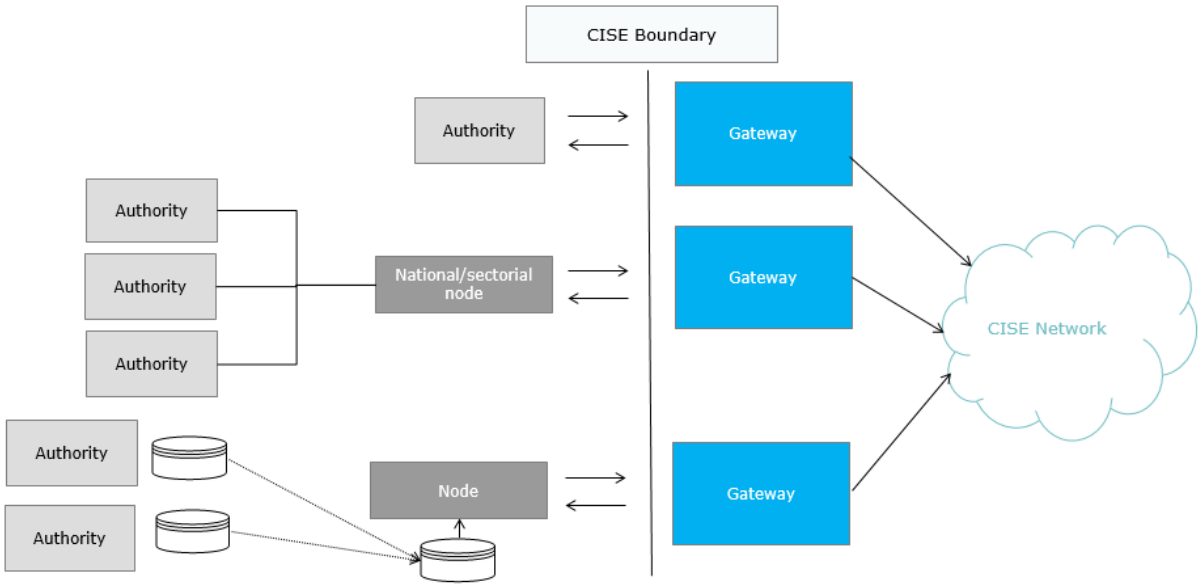

Figure 6 – CISE Gateway

As illustrated in the previous picture, the CISE gateway acts a single point of connection to the CISE network for one or multiple underlying national/sectorial nodes. The gateway should expose a common set of interfaces which support different protocols (i.e. JMS, SOAP, etc.) on the *Private* side, thus facilitating the connection of back office nodes to the gateway. The same is applicable on the *Public* side for gateway-to-gateway communications where multiple protocols such as AS1, AS2 or AS4 could be supported. By using a gateway, the implementation effort to connect to the CISE network does not have to be implemented in every node. This allows for transparently connecting secondary systems to the CISE network.
A gateway should be composed of two interfaces:

- The **incoming interface** connects the national or sectorial node to the gateway. A set of connectors should be included – such as queues or web services – to support the different technologies used by the existing the national or sectorial nodes. The preferred way to connect to the gateway should be configurable in the gateway's configuration files. It is not recommended to customize the gateway since it would imply that different flavours of the gateway exist in CISE which complicates its maintenance and evolution.

- The **outgoing interface** connects the gateway to the CISE network. It should be designed to ensure that a common behaviour is adopted when exchanging messages on the CISE network. The outgoing interface should not be customized and must implement the same specifications across the different CISE participants. It should make use of standardized, reliable and secure messaging protocols. Finally, by using common specifications for the outgoing interface the complexity to connect to the CISE network is embedded in a reusable gateway and thus facilitates the deployment of national or sectorial nodes to the CISE network.

The following recommendations will further describe the role of the gateway. Depending on the chosen architecture, the role of the gateway can be different. Finally, by using a standardized gateway, the interconnection to other environments such as the MARSUR community will be facilitated.

Recommendation	Incubator	
It is recommended to use gateways that will enforce the use of standardised interfaces between CISE participants. Moreover, this approach supports the business case of providing a reference implementation of the gateway, which could be reused by the different Member States.	☑	Implemented & Tested
	☐	Implemented
	☐	Discussed

4.4.2 RD-010 – Message Addressing

Issue Statement: How can CISE participants expose the same CISE service multiple times? This could be the case for a participant who has multiple service providers behind its CISE gateway that each offer the same or similar services.

Through CISE, service providers are able to expose a number of CISE services. As was recommended in section 4.4.1, gateways should be put in place, hiding the complexity of the CISE participant's systems. Combining this recommendation with the CISE hybrid architecture vision (which allows multiple systems to connect to a single gateway), results in the fact that CISE participants could expose services from multiple sources through one CISE gateway. The consequence of this situation, is that in some cases, the same CISE service could be exposed multiple times though the same gateway.

This approach raises the following question: how to address any of these services as they are all "hidden" behind the same gateway? An addressing mechanism is needed to distinguish the endpoint from the actual gateway. To answer this question, it is recommended to implement an addressing mechanism into the Service Registry described in RD-00. By adding addressing information to the service description in the CISE Service Registry, CISE participant nodes will be able to identify to which internal service to map the incoming service request.

Using this approach, there are different options for implementing the addressing mechanism. Since the mapping to the correct internal services will be the responsibility of the CISE service participant, it is sufficient to make each of the service uniquely identifiable in the Service Registry. It is recommended that the Service Registry allows for the specification of sufficient contextual information to distinct similar services, deployed in different endpoints, each hidden behind the same gateway.

Recommendation	Incubator	
Implementing an addressing mechanism in the CISE Service Registry still enables a CISE service provider to decouple the complexity of its own systems from the services offered through CISE. The information in the CISE Service Registry will allow a service consumer identify the required service and allows	☐	Implemented & Tested
	☐	Implemented
	☑	Discussed

the service provider to map the exposed CISE services to internal services.		

4.4.3 RD-011 – Message Translator

Issue Statement: Existing systems are not using standardised CISE messages.

To exchange messages between the different CISE participants, a common message definition has been developed (see CISE service model specifications document). To translate this common message definition to/from the CISE participant's local format, a message translator should be implemented. The CISE message translator should be implemented outside the CISE gateway to keep the gateway message agnostic.

The participants of the CISE incubator clearly indicated that the translation of CISE messages to the specifications of the participant should be done by a separate component which then is connected to the gateway. Whether or not this message translator is a component between the participant environment and the CISE gateway or is integrated in the participant environment is up to the participant.

If an existing system is already using the CISE message format, it will simply bypass the message translator.

Recommendation	Incubator	
By implementing the message translator in a component connected to the CISE gateway all messages coming from or sent out to a CISE participant are translated at a single location.	☑	Implemented & Tested
	☐	Implemented
	☐	Discussed

4.5 Topology

This section provides recommendations and considerations relating to the topology of CISE. The topology choices made in CISE will be particularly important in the next phases of the development of CISE. Indeed, the choice of the topology will have a large impact on the complexity, effectiveness and cost for CISE participants as it will be explained through several examples in this section. Advantages and disadvantages of both topologies are described from the point of view of the service registry and notification scenario.

To support the CISE architecture vision as described in chapter 2, this study identifies two main topologies (arrangement of various elements – nodes – in a computer network) a decentralised and a centralised topology. The CISE incubator learned that most recommendations can be applied to both topologies. A decision on which CISE components to centralise or decentralise should be taken on a case by case basis. Please note that a centralised topology does not imply that business data will necessarily be stored on a central component.

The following sections describe both topology options and introduce the concept of a hybrid topology which might be the result of the recommendations but will not be explicitly covered in this document.

4.5.1 Centralised topology

In the centralised topology CISE nodes will have a single connection to a central component (composed of several building blocks) and will not have direct connections to other CISE nodes. The main role of the central component is to orchestrate traffic between the CISE nodes and store common information. By using a central component, all information such as an address book or a

service registry can be stored in a single location. Moreover, it eases the process of connecting new CISE participants since they only need to connect to the central component.

However, using such topology requires a central authority that takes the responsibility to manage this central component and thus be responsible for delivering, routing or storing messages. Moreover, introducing a central component creates a single point of failure. Indeed, if the central component fails, communication between the different authorities will not be possible anymore. Therefore, redundant systems should be setup if this architecture is selected.

The following picture depicts the high-level overview of the centralised topology.

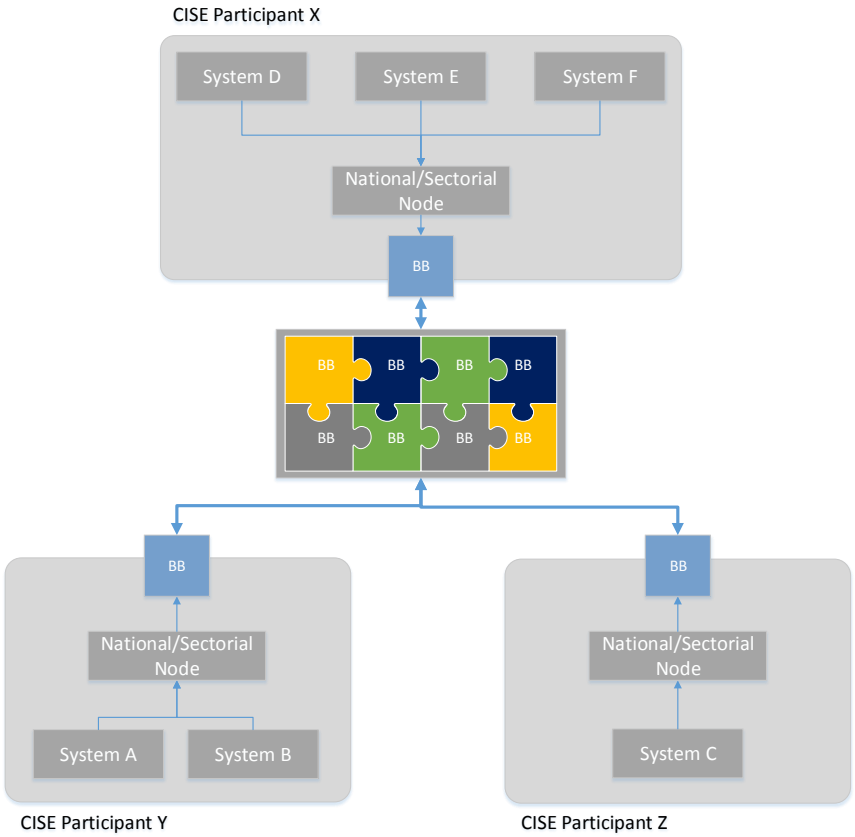

Figure 7 – Centralised Topology

As illustrated on the figure above, most of the building blocks are delegated to a central component. This results in light CISE nodes but introduces a single point of failure for the CISE network. If a central building block fails, the whole CISE community will be affected.

The impact of the centralised topology is further detailed through two examples describing the impact on the service registry (sections 4.2.2) and notifications system (section 4.3.2).

Notifications – Publish/Subscribe messaging pattern.

When using a centralised topology a single subscription list is maintained by the central component. This central component is responsible for dispatching notifications to the CISE participants based on this subscription list. CISE participants can subscribe to tags/topics and/or follow CISE participants via this central component. A subscription service should be created on the central component to allow participants to manage their subscriptions. In this approach, CISE participants do not have to know to whom the message has to be published.

Service Registry

In a centralised topology, the service registry is deployed as a building block in the central component. This is the most efficient solution, since there is only one copy of the service registry. It is the responsibility of the CISE participant to maintain its service registry in the central component. When new participants join the CISE platform all information can directly be published and made available from the central component. Unavailable or leaving CISE nodes have no impact on the availability of the service registry in this topology.

4.5.2 Decentralised topology

In a fully decentralised topology, each CISE participant is acting as a standalone CISE node encapsulating all functionality required to operate and orchestrate the CISE network. Therefore, no centralised building blocks are required to support requests or dispatching messages.

However, implementing such topology might introduce the need to duplicate common information – such as an "address book" – across the CISE participants. Common information in a decentralised environment can be simulated using synchronisation mechanisms but introduces a higher risk for inconsistencies and results in lower network efficiency than in a centralised topology.

The following picture depicts the high-level overview of the fully decentralised topology.

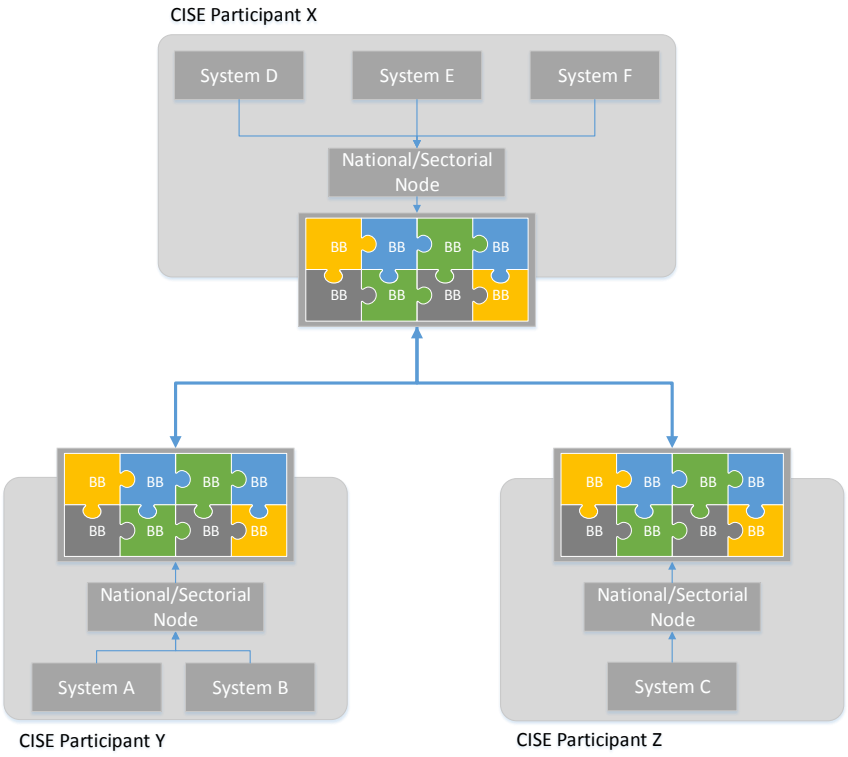

Figure 8 – Decentralised Topology

As depicted on the previous figure, all building blocks required to build the CISE topology are delegated to the CISE participants.

The impact of the decentralised topology is further detailed through two examples describing the impact on the service registry (sections 4.2.2) and notifications system (section 4.3.2).

Notifications – Publish/Subscribe messaging pattern.

Two potential scenarios for the decentralised topology are considered feasible in the context of CISE, a consolidated subscription list and a gateway-specific subscription list.

1. **Gateway-specific subscription list**: Each CISE participant maintains a list that contains subscriptions of other CISE participants to the gateway's notification messages. When a notification message has to be sent out, the CISE gateway will have to determine the relevant tags/topics and lookup the subscribers to these tags/topics. Additionally, the CISE gateway should retrieve the list of followers to determine the destination of the notification messages. In order to subscribe to tags/topics or follow a participant, the CISE participant will have to send a subscription message to the CISE participant's gateway to which it wishes to subscribe. If a CISE participant wants to subscribe to a tag/topic regardless of the source, it will have to send a subscription request to every CISE participant. In this scenario, each CISE participant should expose a subscription service to other participants. Each gateway only stores the subscriptions of other CISE participants to the gateway itself. In this scenario there is no synchronisation or consolidation of the subscription lists.
2. **Consolidated subscription list:** This approach is more complex since it is based on a distributed subscription list, which is synced and consolidated on a regular basis between the different CISE nodes. This approach is subject to a higher level of inconsistency due to potential synchronisation delays. As opposed to the gateway-specific subscription list, the process to subscribe to a tag or follow a participant requires less effort in this scenario. Indeed, instead of sending out subscription messages to every CISE gateway, a CISE participant will only have to update its own subscription list, which will then be synced and consolidated.

Service Registry

In a decentralised topology, CISE participants have to maintain distributed service registries in their own gateway. Therefore two options have been identified:

1. **Synchronised Service Registry**: When a synchronised service registry is used, a synchronisation mechanism is performed on a regular basis (e.g. daily sync). The use of this mechanism is rather complex to implement since it requires intelligent systems to merge and consolidate service registries. Moreover, this option results in a higher network traffic depending on the synchronisation frequency. Synchronisation mechanisms should support the events where CISE nodes are unavailable or leaving the network. When a new participant joins the CISE network, its service registry is synchronised automatically. Note, that choosing this option is a trade-off between synchronisation frequency and inconsistencies between the distributed service registries.
2. **Unsynchronised Service Registry**: In this option, broadcast messages are published by a gateway on the CISE network to inform other CISE participants on updates about their service registry. Unlike for the synchronised service registry, service registry update messages do only contain information about one single service registry. It is up to the CISE participants to ensure that their service registry is consolidated and updated correctly upon receipt of such a messages. One of the main risks identified is the introduction of a higher level of inconsistency across service registries. Finally, by using this option new joiners will have to gather service registries of all CISE members to build up their service registry. Leavers of the CISE network should make sure that they send an update to all service registries to remove their services.

4.5.3 RD-012 – Hybrid topology

Based on this topology study, it is recommended to opt for a hybrid topology which is a mix between the centralised and decentralised topologies.

For standard CISE message flows, the use of a decentralized approach with point-to-point communication is recommended. Indeed, this will reduce network traffic without having to rely on a central component. As a result, no single point of failure exists for this type of communication.

However, for the subscriptions and service registries it is recommended to use a central component that will store the subscription list and service registry. As described in the recommendations, using a centralised component for these two building blocks will be more efficient and will require less effort for CISE participants to join or leave the CISE network.

The following picture depicts a high-level overview of the suggested CISE topology.

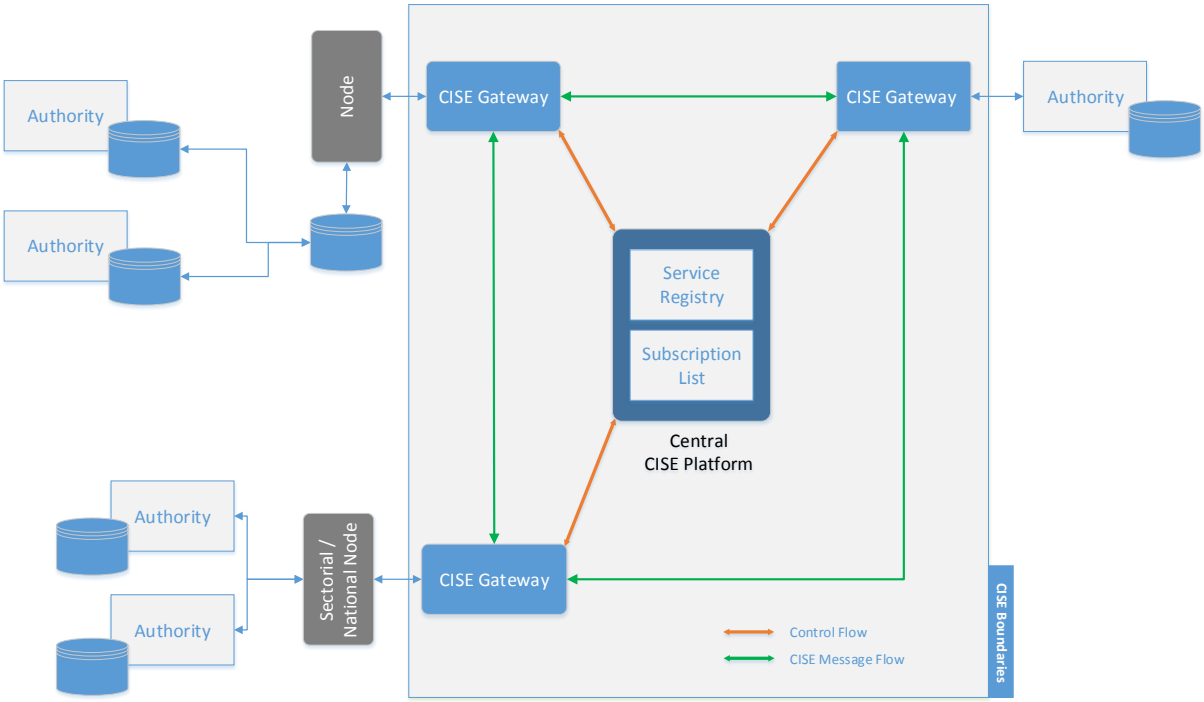

Figure 9 – Hybrid Topology

As a result of the workshops organised in the context of the incubator, an architectural blueprint has been developed by Portugal. This example architectural blueprint of Portugal connection's to CISE can be found in Annex.

Recommendation	Incubator	
It is recommended to implement a hybrid topology for CISE which will benefit from the advantages of a centralised as well as a decentralised topologies.	☐	Implemented & Tested
	☐	Implemented
	☑	Discussed

5 INCUBATOR CONCLUSION

The following results have been successfully achieved by the incubator:

- **Close collaboration and high participant motivation**

  By launching the incubator, the participating Member States were able to work in close collaboration with each other and the Commission and provided valuable input for the finalisation of the specifications of CISE's services. Several workshops have been organised to bring all participants together which proved to be a valuable tool to increase the progress made on the incubator and to keep the momentum and motivation of the participants at a high level.

- **Successful cross-sectorial information exchange**

  One of the key outcomes of the incubator was the successful exchange of information between systems from different sectors. A virtual development and collaboration environment has been setup, including online documentation to kick-start the integration work between the participants. The use of such an environment combined with practical workshops, boosted the speed of the integration work and allowed to tackle many of the challenges of the integration by working closely together.

- **Elaborate a list of recommendations for the POV**

  By developing an incubator, test it with Member States and aligning it with industry best practices, recommendations and lessons learned have been documented areas such as architecture, security or governance of CISE. Since different practical challenges of the implementation of the CISE specifications were tested during the incubator, this automatically resulted in fruitful discussions about the best ways to overcome these challenges. Given the limited scope of the incubator, not every challenge could be tested through a practical implementation. These untested challenges however, still resulted in lessons learned and food for thought for the future steps of the CISE project.

- **Evaluate required effort to connect to CISE**

  The workshops, during which Member States had the opportunity to setup a gateway and connect their environment, provided a good understanding of the required effort to connect to CISE. Indeed, this approach allowed Member States to identify issues they might face when integrating their environment. This contributed to a well-defined strategy for the participating Member States to join CISE based on practical experience.

- **Ease deployment of the gateway and its underlying infrastructure**

  One of the main accelerators of the incubator was the use of virtualisation technology to deploy gateways. This was a key element since it allowed Member States to quickly deploy a fully preconfigured environment allowing Member States to focus on the interconnection of their existing environment.

Finally, one of the key take away of the incubator project is that **trust establishment will be the key success factor to the success of CISE**. The incubator clearly showed that mutual trust is the basis for good collaboration and fast progress.

Please note that although the main building blocks have been successfully tested and implemented, this report is not exhaustive and does not cover all areas required to successfully implement and operate CISE. Other areas such as deployment, change management or capacity management should be further investigated within this context.

6 REFERENCES

CISE Architecture Visions document

The CISE architecture visions are explained in great detail in the CISE Architecture Visions document available for download on the Maritime Forum:

https://webgate.ec.europa.eu/maritimeforum/content/3099

CISE Architecture Visions - Hybrid Vision document

The CISE hybrid architecture vision is explained CISE Architecture Visions – Hybrid Vision document available for download on the Maritime Forum:

https://webgate.ec.europa.eu/maritimeforum/content/3347

OASIS SOA Reference Architecture Foundation for Service Oriented Architecture

The analysis model used in the impact assessment is SOA-centric. The model of OASIS is used to describe the 'service' concept. More information about OASIS' reference architecture is available online:

http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/soa-ra.html

Commission Impact Assessment Guidelines

The Commission Guidelines give general guidance to the Commission services for assessing potential impacts of different policy options.

http://ec.europa.eu/governance/impact/commission_guidelines/docs/iag_2009_en.pdf

White Paper on Integrating Maritime Surveillance: the Implementation of the Common Information Sharing Environment

This document describes an indicative roadmap for the implementation of CISE.

http://ec.europa.eu/governance/impact/planned_ia/docs/2012_mare_002_cise_en.pdf

JoinUp collaborative environment for CISE

https://joinup.ec.europa.eu/software/digit_cise/home

7 ANNEX

7.1 Incubator sequence diagram

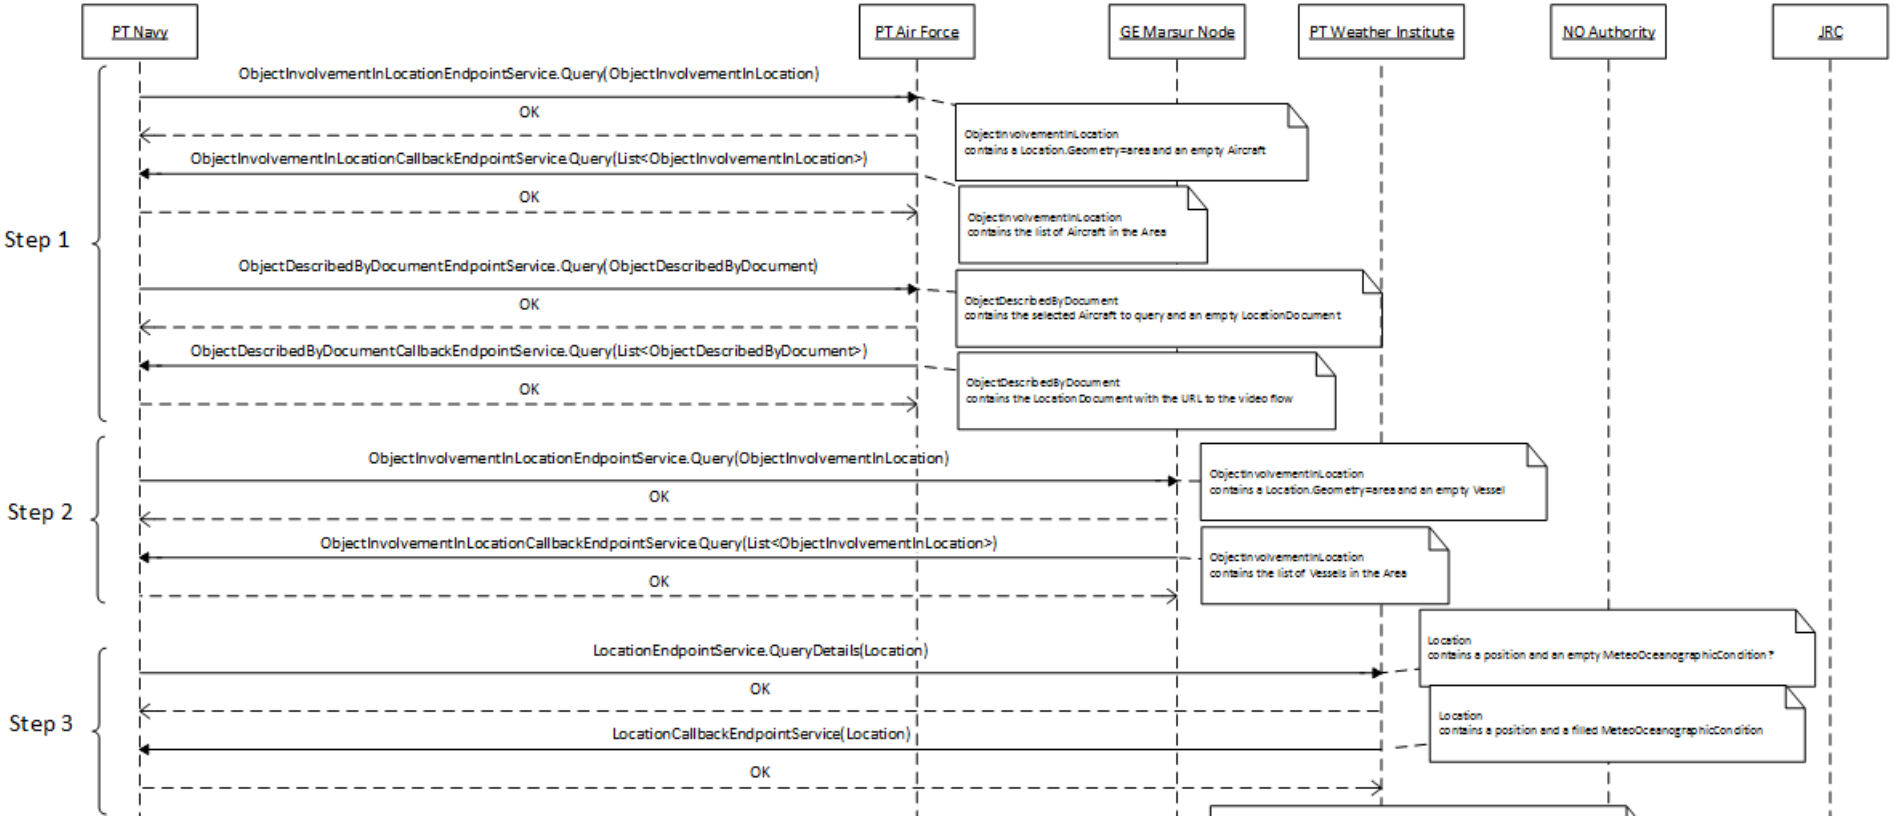

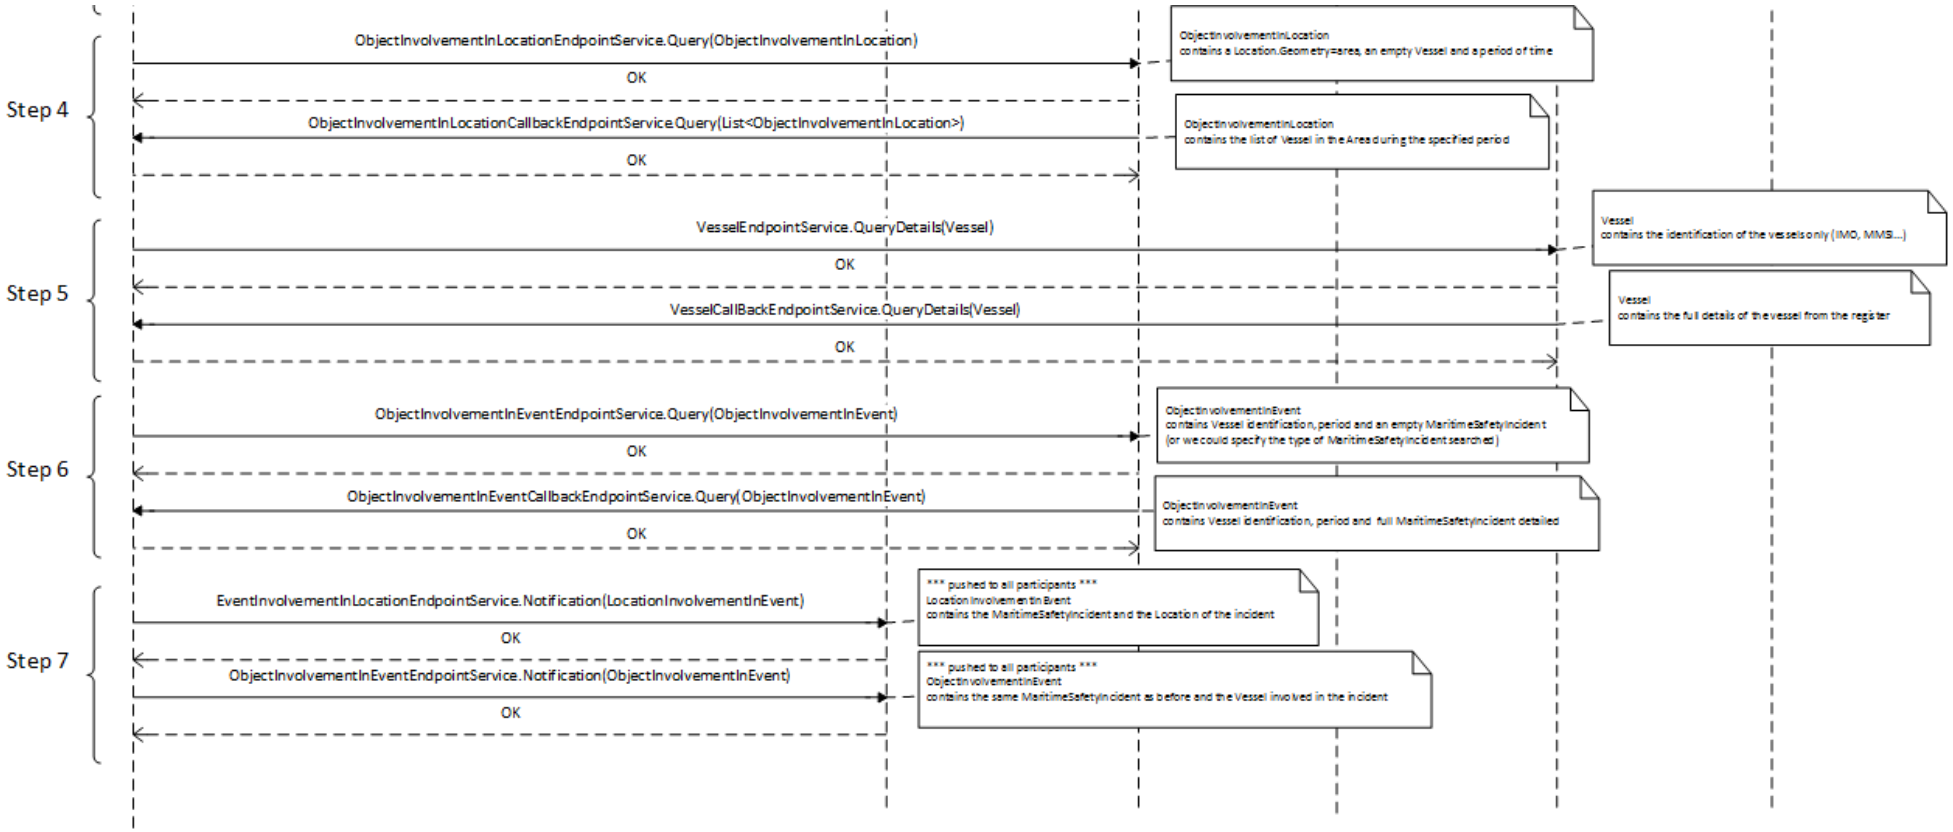
Step 4
ObjectInvolvementInLocationEndpointService.Query(ObjectInvolvementInLocation)
ObjectInvolvementInLocation
contains a Location.Geometry=area, an empty Vessel and a period of time
OK
ObjectInvolvementInLocationCallbackEndpointService.Query(List<ObjectInvolvementInLocation>)
ObjectInvolvementInLocation
contains the list of Vessel in the Area during the specified period
OK
Step 5
VesselEndpointService.QueryDetails(Vessel)
Vessel
contains the identification of the vessels only (IMO, MMSI...)
OK
VesselCallBackEndpointService.QueryDetails(Vessel)
Vessel
contains the full details of the vessel from the register
OK
Step 6
ObjectInvolvementInEventEndpointService.Query(ObjectInvolvementInEvent)
ObjectInvolvementInEvent
contains Vessel identification, period and an empty MaritimeSafetyIncident
(or we could specify the type of MaritimeSafetyIncident searched)
OK
ObjectInvolvementInEventCallbackEndpointService.Query(ObjectInvolvementInEvent)
ObjectInvolvementInEvent
contains Vessel identification, period and full MaritimeSafetyIncident detailed
OK
Step 7
EventInvolvementInLocationEndpointService.Notification(LocationInvolvementInEvent)
*** pushed to all participants ***
LocationInvolvementInEvent
contains the MaritimeSafetyIncident and the Location of the incident
OK
ObjectInvolvementInEventEndpointService.Notification(ObjectInvolvementInEvent)
*** pushed to all participants ***
ObjectInvolvementInEvent
contains the same MaritimeSafetyIncident as before and the Vessel involved in the incident
OK

7.2 Architecture Blueprint - Portugal