

**REQUEST FOR INFORMATION**

for

**THE PROVISION OF CLOUD HOSTING SERVICES**

FOR

**ECMWF, SHINFIELD PARK, READING, UK**

**REF: COP/C3S\_23B**

# CONTENTS

---

## CLAUSE

1.	BACKGROUND.....	1
1.1	<b>Introduction to the Request for Information</b> .....	1
1.2	<b>Introduction to ECMWF</b> .....	1
1.3	<b>Background and Scope</b> .....	1
	(a) <b>Software Architecture</b> .....	2
	(b) <b>Motivations</b> .....	5
	(c) <b>Assumptions</b> .....	6
	(d) <b>Estimated Hardware Requirements</b> .....	7
1.4	<b>Anticipated Timetable for Tenders</b> .....	7
1.5	<b>Enquiries and contact procedure</b> .....	8
1.6	<b>Submission of Responses</b> .....	8
2.	REQUIRED INFORMATION .....	9
2.1	<b>Company and Contact details</b> .....	9
2.2	<b>Company Background</b> .....	9
2.3	<b>Approach to Delivery</b> .....	9
2.4	<b>Requirements</b> .....	9
2.5	<b>Additional matters</b> .....	9
	ANNEX 1 - REQUIRED INFORMATION.....	10
	<b>General Questions</b> .....	10
	(a) Transition to/from the Cloud .....	10
	(b) Customer Service/Support.....	11
	(c) Availability & Performance.....	11
	(d) Data Ownership: .....	11
	(e) Security.....	12
	(f) Cloud Interoperability.....	12
	(g) Pricing .....	12
	(h) Schedule .....	12

## 1. **BACKGROUND**

### 1.1 **Introduction to the Request for Information**

This Request for Information (RFI) has been prepared by the European Centre for Medium-Range Weather Forecasts, (governed by its Convention and associated Protocol on Privileges and Immunities which came into force on 1 November 1975, and was amended on 6 June 2010) ("ECMWF") for the purpose of gathering information about the technical and financial feasibility of different scenarios and the solutions provided by cloud hosting services, together with the capabilities of the cloud hosting providers, and the potential benefits that each of these might bring to support the Copernicus Climate Change Service ("C3S") infrastructure at ECMWF

ECMWF is not looking for detailed proposals at this stage, and there will be no commitment on the part of respondents. Responses will however assist ECMWF in understanding the market interest and available options, and hence in influencing their decision to tender for this service and its approach to such.

As part of the assessment of the responses to this RFI, ECMWF may look to meet with respondents in order to clarify and expand on their responses. Finally, and based on all the inputs received, ECMWF intends to issue a tender in the final quarter of 2016 for providers of cloud compute services.

### 1.2 **Introduction to ECMWF**

ECMWF is an independent intergovernmental organisation supported by 34 States. Information on ECMWF's activities can be found at <http://www.ecmwf.int/en/about> with further detail on the Copernicus project under the 'what we do' / 'environmental services' page.

ECMWF has two sites in Reading, UK.

### 1.3 **Background and Scope**

Copernicus is a European system for monitoring the Earth. The European Commission (EU) has entrusted ECMWF with the implementation of the Copernicus Climate Change Service (C3S) for the period of 2014 to 2020. The Copernicus Climate Change Service's principal aim is to provide information that will lead to an increase in the knowledge base to support climate change adaptation and mitigation policies.

The Climate Data Store (CDS) constitutes the core building block of the C3S and will be at the heart of the C3S infrastructure. The system is designed to provide a consolidated view and point of access to climate data that is distributed over multiple data suppliers. All C3S data will be catalogued in the CDS, and augmented with quality information that allows the data users to discover, appraise and select appropriate data. The CDS will also provide a set of tools (toolbox), workflows and applications that will allow users to perform processing, computation, transformation and visualisation of the catalogued data. Users will be presented with options to download data without modification but with options to build workflows based on concatenation of one or more data sources and parameterised tools, submit these to the system, directly visualise data or even create ad-hoc applications through available functionalities. The web portal and backend systems manage all interfacing and brokering between internal components and provider institutions. The user experiences a single portal environment with consistent style, interface and controls.

As a principle, computation must take place close to the sources in order to minimize data transfers, and in case of non-availability, data might be pushed into the CDS compute area for further processing. In the first stage of the CDS life-cycle it is not envisaged to offer a permanent repository of data but just a temporal staging area for requested and computed data, including cache functionality. All this, nevertheless, could be reconsidered in the coming future as C3S is envisaged as a dynamic and evolving system.

In order to assure the highest quality of service, the Copernicus Climate Change Service aims to benefit from the latest state-of-the-art technologies, hardware and software, setting up a reference framework infrastructure able to accommodate evolving needs and requirements coming from a wide diversity of users and sectors, from experts and scientists to policy makers, while keeping the integrity and robustness of the system.

In this context, the purpose of this RFI is to gather information about the technical and financial feasibility of different scenarios, the solutions provided by cloud hosting services, the capabilities of the cloud hosting providers, and the potential benefits that each of these might bring to support the Copernicus Climate Change Service infrastructure on the development of its aims and expectations.

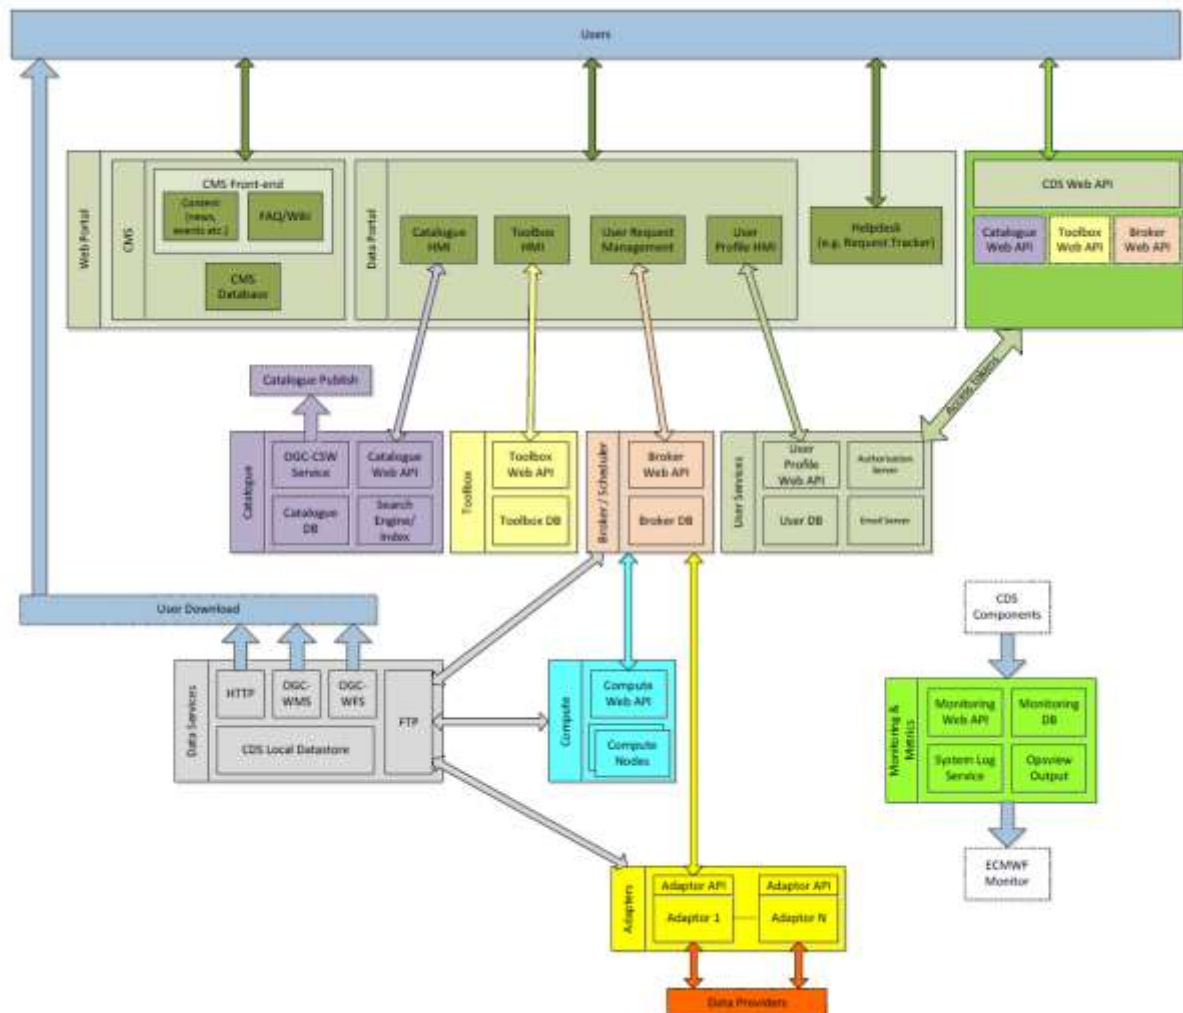
(a) **Software Architecture**

The C3S Infrastructure is built following the bases of a Service Oriented Architecture in which each functional component is engineered as a self-contained entity that exposes a well-defined interface through which its services may be invoked.

This architecture represents a flexible solution in which each service can be deployed on its own dedicated platform, or platforms can play host to multiple services. The CDS

will use CentOS Linux, but in practice it should be compatible with many modern Linux distributions.

Additionally, this approach lends itself to the use of Virtual Machines for the deployment of discrete services in response to the challenges regarding platform independence; scalability and extensibility.



**Fig.1: Envisaged Services Architecture (not including toolbox)**

The CDS Infrastructure services are briefly summarized here:

- Web Portal:** The CDS-Toolbox user interface is a single consolidated front-end bringing together the Drupal CMS and the data portal user interface and presented as an enterprise web application. A fully integrated data portal user interface is implemented by utilizing the Web APIs provided by each service – supplying information for display and functional interfaces for requests.

- **Catalogue:** The Catalogue provides the inventory of all datasets available across all data providers. The Catalogue functions are supported by a dedicated search engine and index that provide fast faceted search capabilities.
- **Toolbox:** The Toolbox provides an inventory of all tools available in the CDS Toolbox. For each tool the inventory describes the parameterization of the inputs, the format of the outputs and the means of invoking.
- **Broker/Scheduler:** The Broker/Scheduler is a single component that sits at the heart of the system workflow. All user requests for data retrieval and processing are mediated and managed by the Broker/Scheduler. It is responsible for directing the user requests to the appropriate endpoints, either to data providers via the corresponding adaptor, or execution of tools in the CDS Compute area. The Broker handles the outputs of these operations and drives the workflow by feeding the outputs to the next step in the flow. Finally the Broker returns the URL of the resultant dataset to the user. The Scheduler maintains queues that implement the Quality of Service rules and user priority rules that control the order and simultaneity of requests within the CDS.
- **Compute:** Provides a cluster of nodes for execution of processing tasks using tools from the toolbox. The compute nodes are 'fronted' by a compute service that abstracts the details and management of the individual nodes. Compute functions may also be supported by some data providers, in which case the Broker directs the request via the appropriate adaptor.
- **Adaptors:** The interface between the Broker/Scheduler and the External Data Providers. They provide an interface translation that allows the CDS to request data retrievals, data sub-setting and processing requests. An adaptor is implemented as a service that conforms to a common interface.
- **Data Services:** provides a so called “Local Data-store” within the CDS for the temporary storage of data, intermediate products and processing results. CDS components push, pull and manage data in the local store. Data retrieved from external providers is pushed into the Local Data-store by adaptors. This data may be input to further processing steps, or can be provided to users via the appropriate download service.
- **User Services** maintains the User Profile Database that captures the details of each user, including their privileges and priorities regarding access to data/processing resources.

- **CDS Web API.** An Application Programming Interface that provides a language agnostic programmatic interface to the CDS functions, including Catalogue, Toolbox and Broker/Scheduler.
- **Monitoring & Metrics:** Each component in the CDS is responsible for maintaining relevant status and metrics regarding its operations to be reported on demand. Each component makes this information available via its service Web API. The Monitoring & Metrics service routinely interrogates each component to gather information. In addition to the monitoring of system status, each component logs messages regarding their activities, and C3S staff have the capability to access and interrogate these logs.

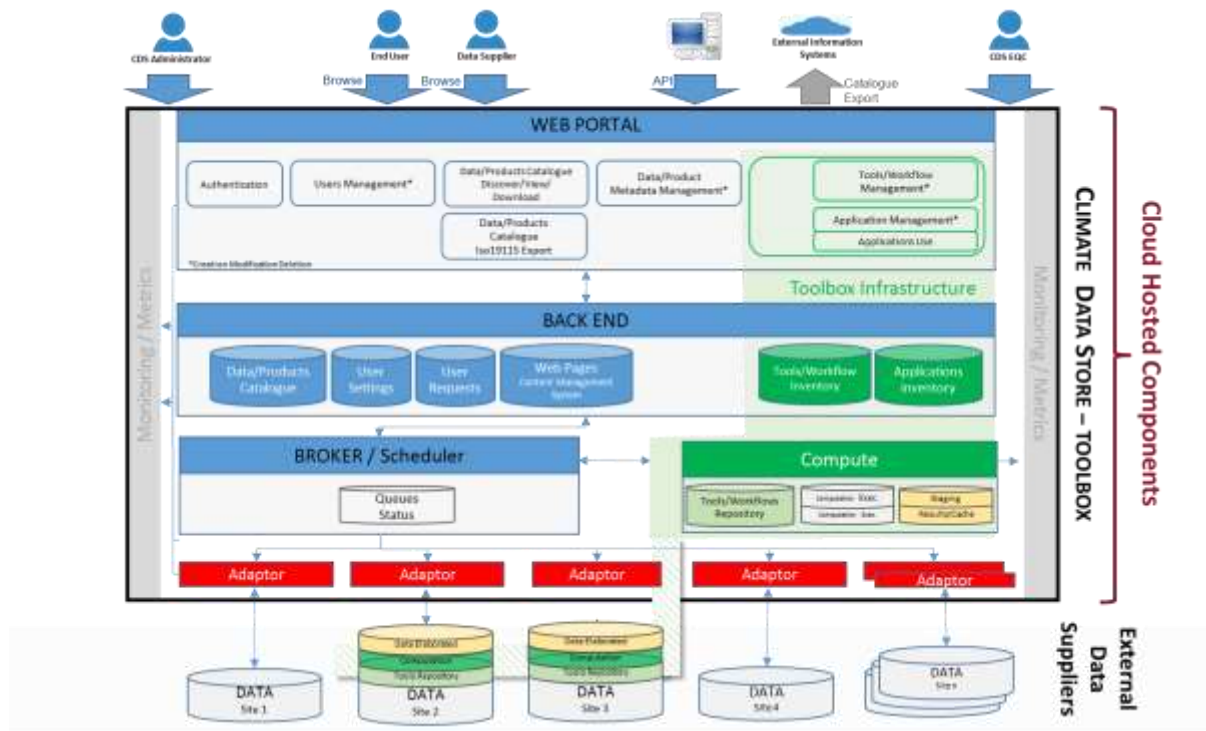


Fig.2: Scope of Hosted Components

(b) **Motivations**

The following motivations have been initially considered for hosting the C3S Infrastructure in the Cloud:

- Major scalability, flexibility and adaptability, leading the service towards increasingly demanding requirements and functionalities.

- Detach the future evolution of the system from the constraints and limitations imposed by proprietary hardware or software assets.
- Optimize time and resources by the automation of routine processes such as upgrades or backups.
- Elasticity of resources and flexibility of the billing basis to cut down operational costs, adapting these in a proportional way to the C3S changing requirements.
- Achieve enhanced and simplified IT management and maintenance capabilities through central administration of resources, provider managed infrastructure and SLA backed agreements.
- Assure reliability and consistency of the infrastructure by establishing Service Level Agreements which covers 24/7/365 and 99.99% availability.
- Reach the highest stability / cost ratio performance while ensuring a scalable, stable, fast and secure environment.
- Benefit from a massive pool of redundant IT resources, as well as quick failover mechanism manageability.
- Enjoy a simple web-based interface for accessing assigned cloud environments, deployed software, applications and services.
- Delegation of routine IT administration tasks considered out of the C3S operational business core.
- Assure the compliance of the system with quality standards such as ISO27001 security levels.

(c) **Assumptions**

- At this stage in the system life-cycle is difficult to estimate potential data volumes, compute workloads, user access and future requirements. Until the system operation reaches a mature status, the hosting solution should be as elastic as possible. Nevertheless it is envisaged to refine this information jointly with external contractors as long as the development project evolves.
- The system is designed to cope with a number of active users of around 1000, with up to 10,000 registered accounts.
- Handled data volume might range from Megabytes to Terabytes. Data requests and Workflows are created ad-hoc by end users which makes it difficult to estimate the potential dimension and concurrency. Quality of Service (QoS) is managed by rules defined at system level aiming to protect the systems from overloads, concurrences, and failures.
- CDS will be a one-stop shop, 24/7/365 available.
- User perception of the system performance is a cornerstone issue to be considered.
- Data transfers must be minimized. Computation should take place close to the data sources whenever feasible.



- There will be at least three different environments: Development, Integration and Production.
- CDS is not envisaged to host a permanent data repository but a temporal staging area. Nevertheless, this could be later reconsidered. The CDS Local Data Store will provide temporary storage for:
  - Data/results retrieved from external data providers.
  - Processing results output from CDS Compute.
  - Intermediate data used within a workflow.

(d) **Estimated Hardware Requirements**

The C3S Infrastructure will consist of three different environments: Development, Integration and Production. The following hardware assumptions would be global covering the three of them. The splitting and allocation of resources for each particular environment should be adaptable.

- Hardware:
  - Computing resource equivalent to 45 x 2.0 GHz Cores.
  - 192GB RAM.
  - Total Initial Storage: 50TB.
  - 100 simultaneous users
  - Virtualised deployment (flexibility, shared resources, over-provisioning)
- Use of SAN facilities VM failover for availability
- Industry Standard Hardware – i.e. no special-to-type HW.

1.4 **Anticipated Timetable for Tenders**

Whilst this is currently only a request for information, C3S envisages the following timetable for any subsequent tender that may come as a result:

- Tender published: Last quarter 2016
- Evaluation of tenders and negotiations with preferred tenderer(s): First quarter 2017
- Award of contract: Early second quarter 2017

## 1.5 Enquiries and contact procedure

Any enquiries or requests for clarification of any matters arising from this RFI should be sought from the Procurement Section at ECMWF and should be made in writing by e-mail as follows:

E-mail: [procurement@lists.ecmwf.int](mailto:procurement@lists.ecmwf.int)

Emails should be headed '**RFI - C3S\_23b**'.

## 1.6 Submission of Responses

Before submitting your responses to this RFI, please be aware of the following considerations:

- Responses to this RFI are voluntary.
- This RFI is for the purpose of collecting information, and should not be construed as a solicitation or as an obligation on the part of ECMWF in any way.
- C3S will use the feedback received to ascertain the period of service required, however, it is likely that this will be between 1 and 3 years.
- C3S does not intend to make any type of award based on responses to this RFI or to pay for either the preparation of information submitted or the use of such information.
- C3S will use the information submitted in response to the RFI at its discretion within standard confidentiality rules.
- Respondents are advised that C3S is under no obligation to acknowledge receipt of the information provided and may not provide feedback to respondents.
- The information submitted will be analysed and may be shared internally or incorporated into future documentation, as appropriate and at the discretion of C3S.
- Proprietary, classified, confidential, or sensitive information should be highlighted in the response.

**Responses to this RFI should be submitted by email as per the instructions in Clause 1.5 by 12:00 UK local time on 7 October 2016.**

## 2. **REQUIRED INFORMATION**

Interested parties to this RFI are requested to respond to the following. Responses should not be overly detailed (in the event that we progress to tender, full proposals will be sought at that stage), but should be sufficient to enable C3S to understand the options available.

### 2.1 **Company and Contact details**

Please confirm the full name and address of your company, together with details of the person at your organisation who can be the contact point in relation to your response, and any clarification questions we may have. Please give their name, title, address and location, telephone number and e-mail address.

### 2.2 **Company Background**

Please describe in brief terms your company's history with regards to the provision of cloud hosting services.

### 2.3 **Approach to Delivery**

Please describe in brief terms the type of package(s) you could offer in line with the Motivations, Assumptions and Estimated Hardware Requirements stated above.

This should include an indication of an annual budget.

### 2.4 **Requirements**

Your reply should take the form of a *brief, but sufficiently detailed, response* to each of the questions set out in **Annex 1** to enable C3S to fully understand what services are available, can be delivered and how.

### 2.5 **Additional matters**

Please set out any additional information or other relevant matters which you think have not been adequately addressed in the RFI and/or merit further consideration in your response.

## **ANNEX 1 - REQUIRED INFORMATION**

Computational infrastructure for the Copernicus Service is envisioned as employing cloud computing services. A “conformant provider” of cloud services for the C3S should offer an array of options, such as Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS), meeting the ECMWF - Copernicus defined requirements for business relationships, interfaces, capacity, networking and connectivity, information assurance and authentication and authorization.

**C3S is now requesting information on an initial set of proposed questions, in order to assess the range and availability of such services with a view to tendering for such during the last quarter of 2016.** Providers are therefore requested to respond to the following questions using the numbering system indicated.

Please do not include any marketing or similar brochures in your response unless they are specific to your response to one of the questions below.

### **General Questions**

1. What are the overall cloud hosting solutions and scenarios that, as a service provider, you envisage and recommend for the C3S?
2. What are the financial, legal, and operational advantages and disadvantages of the proposed cloud solutions and scenarios? Are there specific issues about which C3S should be aware when taking a decision?
3. What scale of operation would be required for the C3S Infrastructure to be attractive to your company to participate as potential bidders on any future tender? Is there a critical mass below which you might not be interested in bidding on a proposal for this approach? In case so, what metrics would you apply for that decision?
4. What issues, if any, would limit C3S ability to make use of presented cloud services?
5. What is your proposal for Standard Service Level Agreements (SLA), or equivalent?
6. Characterize your offerings as SaaS, PaaS or IaaS cloud services and expose the pros and cons, strength and weakness, requirements and constraints for each of them.
7. Please highlight any unique strategies or capabilities that you would provide to make our initiative successful.

#### **(a) Transition to/from the Cloud**

1. Describe the process for deploying applications and hosting C3S infrastructure in your cloud environment (“on boarding”), including testing, acceptance and cutover.
2. What guidelines can you provide to assist C3S in project planning (level of effort, timeline, decommissioning legacy services, etc.)?
3. Describe your method for protecting and returning infrastructure components either on demand or in case of contract termination (“off boarding”).
4. Please explain any application and data portability considerations (i.e. exit strategy for applications running in your cloud).

(b) Customer Service/Support

1. Please discuss your model for providing customer support, including charges for support contacts.
2. Describe your incident/problem reporting and tracking systems, and the ability for authorized C3S staff to access those systems directly.
3. What are the technical options to interoperate with your systems in order to access and integrate your information within ECMWF – C3S reporting and tracking systems?
4. What types of access to your customer support would be available (website, email, chat, telephone).
5. What level of automatic alerting can you provide to C3S staff/systems in the event of failure, degraded service, or exceeded planned utilization?

(c) Availability & Performance

1. Describe your mitigation strategies for potential availability and performance issues such as network outages, bandwidth shortages, or spikes in service demand?
2. Describe the redundancy features of your cloud services that ensure availability and performance.
3. Discuss your roles and responsibilities for system maintenance as the service provider and the roles and responsibilities that C3S staff are expected to assume.
4. Please include information about your procedures for operating system and other core software upgrades, patches, and service pack application.
5. Please provide past availability statistics for your cloud services that you feel may be representative for the C3S infrastructure
6. Discuss features of your cloud services that provide for scalability of C3S applications and data hosted in your environment.

(d) Data Ownership:

1. Describe your policies, roles and responsibilities regarding data ownership and control.
2. How do you handle data remanence (\*) once customer data is removed from your storage media?

*(\*) Data remanence refers here to the residual data remaining in magnetic storage devices used in electronic equipment such as hard drives, tape recorders and VCRs.*

3. Who owns the Intellectual Property for data hosted in your cloud and artefacts developed in or hosted in your cloud?
4. Are there any considerations about hosting jurisdictions on offer around the globe, as this pertains to data sovereignty, which might apply to the service?

(e) Security

1. Describe your approach to addressing IT security challenges in cloud computing, in particular dealing with hacker attacks, the potential for unauthorized access, and inappropriate use of proprietary data and IT applications.
2. What are your processes and solutions for preventing these challenges from occurring?
3. What controls are in place for administrative access, both internal to your company and for administrative access from government clients?
4. Please include discussion of administrator controls over provisioning.

(f) Cloud Interoperability

1. C3S have, or may in the future, establish internal cloud environments or external cloud services with other providers, i.e. private, public, hybrid or community cloud environments. Describe how your cloud services integrate with other cloud services to provide seamless interoperability for the end user.

(g) Pricing

1. Describe your pricing models and other relevant pricing factors such as CPU, memory, storage, bandwidth, and data transfers. Include break points for price changes and prices for transition services. Please ensure that bandwidth charges for uploading and downloading data are clearly described.
2. Describe your pricing model for loading data on a recurring basis and for transaction processing.
3. Describe your pricing model for on demand extraction of data by third parties and any approaches that would control cost for the C3S.
4. Describe the pricing differences for non-volatile data that requires only restore rather than recovery capability.
5. Please add any additional information relating to pricing or costs that would be useful in evaluating the overall cost to C3S in utilising the cloud services detailed in your response.

(h) Schedule

1. In terms of the initial ramp-up and potential scalability, are there any considerations that need to be taken into account by C3S in terms of lead or implementation times when ordering such services.