

Testbed-15 Call for Participation (CFP)

Version 1.4.5 - 1 March, 2019

Table of Contents

Main Body - Introduction	7
1. Benefits of Participation	10
2. Policies and Procedures	11
2.1. Initiative Roles	. 11
3. General Proposal Submission Guidelines	13
3.1. Questions and Clarifications	. 14
4. Proposal Evaluation Criteria	15
4.1. Management Criteria	. 15
4.2. Technical Criteria	. 15
5. Master Schedule	16
6. Sequence of Events, Phases, and Milestones	17
7. Summary of Initiative Deliverables	21
7.1. CFP Deliverables Summary Table	. 22
Appendix A: Management Requirements	25
A.1. Proposal Submission Procedures	. 25
A.2. Conditions for Participation	. 29
A.3. Proposal Evaluation and Invitations to Selected Bidders	. 30
A.4. Kickoff Workshop	. 32
A.5. Participant Communication and Reporting	. 33
A.6. Requirements for Proposing Deliverables	. 35
A.7. Project Monitoring and Control	. 38
A.8. Tips for New Bidders	. 39
Appendix B: Technical Architecture	43
B.1. Introduction	. 43
B.2. Testbed Baseline	. 43
B.3. Testbed Threads	. 47
B.4. Tasks	. 49
B.5. Machine Learning	. 49
B.6. Delta Updates	61
B.7. Data Centric Security	. 65
B.8. Federated Cloud Analytics	. 70
B.9. Portrayal	. 80
B.10. Earth Observation Process and Application Discovery	. 92

ESA ITT Released

IMPORTANTESA has currently released the Open Invitation To Tender "OGC Testbed
15 - ESA Sponsored Threads - Exploitation Platform"
[http://emits.sso.esa.int/emits/owa/emits_external.showsum?typ1=8317&
user=Anonymous&ctr=TELES_VEGA], which provides additional participation
opportunities! For further details, please see below.

Corrigenda

The following table identifies all corrections that have been applied to this CFP compared to the original release. Minor editorial changes (spelling, grammar, etc.) are not included.

Date	Section	Description
2019-01-09	Appendix A, Managemen t Requirement s	Term "memorialized" replaced with "developed"
2019-01-09	B5.4	"CGDI list of Web" replaced by "list of Web services provided by the CGDI"
2019-01-09	B6.1	Paragraph "OGC Third Generation API (OpenAPI) provides the ability for OGC standards, to incorporate delta updates utilising Feature ID and Timestamp elements of OpenAPI. This allows for data such as raster or vector to be updated reflecting the changes in the original version." replaced by "Emerging OGC interface specifications that use OpenAPI provide the ability for OGC standards to incorporate delta updates utilizing the Feature ID and Timestamp elements of OpenAPI. This allows for raster or vector data to be updated while still reflecting the changes in the original version."
2019-01-09	B7.2	"Implement a Data Centric Security approach using existing OGC standards and OGC third generation API (OpenAPI) implementation work as being documented in Testbed 14." has been replaced by "Implement a Data Centric Security approach using existing OGC standards and emerging OGC standards such as WFS 3.0 or WPS 3.0 that make use of OpenAPI) as being documented in Testbed 14."
2019-01-09	B7.4	"Testbed-15 shall develop an OGC Third Generation API (OpenAPI) implementation based on work in OGC Testbed 14." is replaced by "Testbed-15 shall develop an implementation of WFS 3.0 with support for data centric security.".
2019-01-09	B8.1	Added paragraph: "Of particular interest in this context are three aspects. Firstly, the handling of security in federations. Second, how the Testbed-13 and Testbed-14 research results of "bringing applications to the data" relate to SCALE [https://github.com/ngageoint/scale] and SEED [https://github.com/ ngageoint/seed]. SCALE is an open source system that provides management and scheduling of automated processing on a cluster of machines. It uses the SEED specification to aid in the discovery and consumption of processes packaged in a Docker containers. Third, the role of blockchain and distributed ledger technologies in the context of handling provenance in federations."
2019-01-11	B5.4	Link added to "list "list of Web services provided by the CGDI"

Date	Section	Description
2019-01-14	B9.1 & B9.4	"Vector Tiles" replaced by "tiled vector data"
2019-01-15	Beginning of CFP	Clarifications table added
2019-01-15	Main Body, Introduction	Link to ITT [http://emits.sso.esa.int/emits/owa/ emits_external.showsum?typ1=8317&user=Anonymous&ctr=TELES_VEGA] on EMITS added
2019-01-30	Machine Learning/Ai m	Paragraph added
2019-01-30	Clarification s table	Entry added
2019-02-04	Detailed Submission Instructions	OK to provide in-kind cost breakdowns, but avoid including any sensitive information in the "Proposed Contribution" text box.
2019-02-04	Main Body - Introduction	Add tips for Bidders wishing to access EMITS for the <i>Part 2 ITT</i> .
2019-02-08	<i>Tips for New Bidders</i>	Add an explanation of how to include an in-kind contribution from a subcontractor.
2019-03-01	Portrayal	"direct uptake" replaced by "consideration"

Clarifications

The following table identifies all clarifications that have been provided in response to questions received from organizations interested in this CFP.

Question	Clarification
Is the focus of the machine learning section on the development and optimization of models or more on their integration with OGC Web services?	The latter. To emphasize this aspect, an additional paragraph has been added to section Machine Learning/Aim.
What is the focus of the Data Centric Security Task? How can a WFS interact with a data store when all data is encrypted?	The task is envisioned to be the first part of a multi-initiative effort to develop interoperable solutions for Data Centric Security. The issue with WFS is one of the challenges that shall be explored. Goal of Testbed-15 is to experiment with Data Centric Security in the context of emerging OGC standards such as WFS 3.0 or WPS 3.0 and to do the foundational work necessary to build on in future initiatives.
Should a detailed breakdown of the "Inkind Contribution" amount be provided for each deliverable?	A breakdown (by cost category) of the "Inkind Contribution" may be included in the <i>Proposed Contribution</i> text box for each deliverable. However, please note that the content of this text box will be accessible to all Stakeholders and should contain no confidential information such as labor rates. If a Bidder wishes to provide confidential information that will only be accessible to OGC staff, it can be included in the <i>Attached Documentation</i> of the submission. Only one <i>Attached Document</i> may be submitted for the entire proposal (there is no way to submit a separate attachment for each individual deliverable).
<i>How can a (qualified) Bidder gain access to EMITS for the Part 2 ITT?</i>	See the <i>Tip</i> added under the Main Body - Introduction.

Question	Clarification
We would like to include an in-kind contribution from a subcontractor who will not be submitting their own separate bid. Should this contribution be included as part of our bid?	Since each Participation Agreement (PA) will be a bilateral contract between OGC and the bidding organization, in general OGC will not scrutinize the details of how the Bidder chooses to acquire any downstream subcontracting services so long as it promises to make timely delivery. So a Bidder's in-kind proposal could, for example, include an estimated market- value of a downstream subcontractor's contribution to the Bidder's own total proposed in-kind contribution. However, since cost-share funding may only be used for labor cost, any downstream subcontractor contribution amount may only be listed under the Bidder's in-kind contribution (not cost-share).

Quick Links

Here are several links for Bidders interested in quickly responding to the CFP:

- As indicated in the Master Schedule, the proposal submission deadline is **11 February, 2019**.
 - A Bidders Q&A Webinar is scheduled for Monday, 4 February at noon U.S. Eastern Time. The webinar is open to the public, but anyone wishing to attend must register beforehand at https://register.gotowebinar.com/register/6662012816865820675.
- The process for completing a Proposal will essentially be embodied in an online Proposal Submission website [https://proposal.opengeospatial.org/]. Once this site is fully prepared to receive submissions (likely within several days of the CFP release), it will include a series of web forms, one for each deliverable of interest. See Proposal Submission Procedures for a more detailed process description.
- The Technical Deliverables open for bidding are summarized under the CFP Deliverables Summary.
- Deliverable details can be found in the Appendix B Technical Architecture.
- New to OGC? Consider reviewing the Tips for New Bidders first.
- Please visit the Testbed-15 project page [http://www.opengeospatial.org/projects/initiatives/ testbed15] for ongoing news.

Abbreviations

The following table lists all abbreviations used in this CFP.

3DPS	3D Portrayal Service
ABI	Activity Based Intelligence
AOI	Area of Interest
AMQP	Advanced Message Queuing Protocol
AR	Augmented Reality
AtomPub	Atom Publishing Protocol
AVI	Aviation
BBOX	Bounding Box
BPMN	Business Process Model and Notation
CDR	Content Discovery and Retrieval
CIS	Coverage Implementation Schema
CITE	Compliance Interoperability and Testing
CF	Climate and Forecasting
CFP	Call for Participation
CMD	Command Center
CSMW	Community Sensor Model Working Group
CSW	Catalog Service Web
CTL	Compliance Testing Language
DAAC	Distributed Active Archive Centers
DAP	Data Access Protocol
DASS	NASA's Data Analytics and Storage System
DCAT	Data Catalog Vocabulary
DDIL	Denied, Degraded, Intermittent, or Limited Bandwidth
DGIWG	Defense Geospatial Information Working Group
DISA	Defense Information System Agency
DWG	Domain Working Group
EMITS	EMITS is the system used by ESA for publishing Invitations To Tender and information about the ESA procurement process [http://emits.sso.esa.int/emits/owa/emits.main]
EO	Earth Observation
EOWCS	Earth Observation Profile Web Coverage Service
ER	Engineering Report
ESA	European Space Agency
ESGF	(US Department of Energy) Earth System Grid Federation
EXI	Efficient XML Interchange format
FGDC	Federal Geographic Data Committee
FIXM	Flight Information Exchange Model
FO	Field Operations
GDAL	Geospatial Data Abstraction Library

GEOINT	Geospatial intelligence
GeoXACML	Geospatial XACML
GIBS	Global Imagery Browse Services
GML	Geography Markup Language
GUI	Graphical User Interface
HDF	Hierarchical Data Format
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext transfer protocol secure
ISO	International Organization for Standardization
ITT	Invitation to Tender
JSON	JavaScript Object Notation
JSON-LD	JSON Linked Data
KML	Keyhole Markup Language
LiDAR	Light detection and ranging
MEP	Mission Exploitation Platform
МТОМ	Message Transmission Optimization Mechanism
NASA	National Aeronautics and Space Administration
NetCDF	Network Common Data Form
NetCDF-CF	NetCDF Climate Forecasting
NSG	National System for Geospatial Intelligence
OAuth	Open Authorization
OBP	Object Based Production
OGC	Open Geospatial Consortium
OPeNDAP	Open-source Project for a Network Data Access Protocol
PKI	Public Key Infrastructure
PMT	(ShapeChange) Profile Management Tool
POI	Points-of-interest
PubSub	Publication Subscription
RDF	Resource Description Framework
SAML	Security Assertion Markup Language
SOS	Sensor Observation Service
SPARQL	SPARQL Protocol and RDF Query Language
SSO	Single Sign On [https://en.wikipedia.org/wiki/Single_sign-on]
STAC	SpatioTemporal Asset Catalog specification
SWAP	Size, Weight, and Power
SWE	Sensor Web Enablement
SWG	Standards Working Group
T-nn	Testbed-nn
TEAM	Test, Evaluation, And Measurement Engine
TEP	Thematic Exploitation Platform [https://tep.eo.esa.int]
TSPI	Time-Space-Position-Information Standard
TWMS	Tiled Web Mapping Service

US	United States
UML	Unified Modeling Language
USGS	U.S. Geological Survey
W3C	World Wide Web Consortium
WCPS	Web Coverage Processing Service
WCS	Web Coverage Service
WFS	Web Feature Service
WIS	Web Integration Service
WKT	Well Known Text
WMS	Web Mapping Service
WMTS	Web Mapping Tile Service
WPS	Web Processing Service
WS	Web Service
WSDL	Web Services Description Language
XACML	eXtensible Access Control Markup Language
ХОР	XML-binary Optimized Packaging
XXE	XML External Entity Injection

Main Body - Introduction

The Open Geospatial Consortium (OGC[®]) is releasing this *Call for Participation* ("CFP") to solicit proposals for the OGC Testbed-15 initiative ("Testbed" or "initiative") under the OGC Innovation Program [http://www.opengeospatial.org/ogc/programs/ip] (OGC-IP).

The OGC-IP provides global, hands-on, collaborative prototyping for rapid development and delivery of proven candidate specifications to the OGC Standards Program, where these candidates can then be considered for further action. Participants in OGC-IP initiatives collaborate to examine specific geo-processing interoperability questions posed by the initiative's sponsors. These initiatives include testbeds, experiments, pilots, and plugfests – all designed to foster the rapid development and adoption of open, consensus-based standards.

The OGC recently reached out to potential initiative sponsors to review the OGC standards baseline [http://www.opengeospatial.org/standards] (also see Standards Baseline), discuss results of prior initiatives, and identify this initiative's particular requirements.

This CFP is soliciting proposals from "Bidders" to address these requirements in the initiative. Received proposals ("bids") will be evaluated, and selected Bidders will become cost-share "Participants" to receive funding to help offset the cost of performing initiative work.

Please note the Requirement to Propose an In-Kind Contribution for making a cost-share request for a funded deliverable. In-kind contributions are an important part of the proposal evaluation criteria.

It is expected that most deliverable proposals will include a combination of a cost-share request and an in-kind contribution. However, any proposed deliverable offering a *purely in-kind* contribution (i.e., requesting no cost-sharing at all for that deliverable) will certainly be welcomed.

The solicitation is being issued in two parts due to specialized sponsor procurement requirements:

- Part 1 this CFP document ("Part 1 CFP" or just "CFP")
- Part 2 an Invitation To Tender pack ("Part 2 ITT" or just "ITT") under the European Space Agency (ESA) Electronic Mailing Invitation to Tender System [http://emits.sso.esa.int/emits/owa/ emits_external.showsum?typ1=8317&user=Anonymous&ctr=TELES_VEGA] (EMITS).

Please note the following tips relating to the *Part 2 ITT*:

	• Only entities from the listed countries can bid for ESA cost-sharing funds. Any bid that will be <i>purely in-kind</i> would be better-suited for submission under the Part 1 CFP.
	 Only registered entities may bid. See http://www.esa.int/About_Us/ Business_with_ESA/How_to_do, especially http://www.esa.int/About_Us/ Business_with_ESA/How_to_do/New_Tenderers.
	• A Bidder who doesn't have an EMITS / ESA bidding account must first register the entity and obtain an account before proceeding. A user would then visit the EMITS front page [http://emits.esa.int] and follow these steps to access the documents:
	\circ Login (login button at the top of the page).
ΓΙΡ	 Select Entities in the top Menu.
	 In the Entities page scroll down and select "Telespazio Vega UK Limited".
	 Press OK in the pop-up warning.
	 On the left menu, select "Open Invitations to Tender (ESA Programmes)".
	 Select AO700000 on the List of Open Invitations to Tender.
	 The documents are the bottom of the "Open Invitation to Tender AO700000" page.
	• Although <i>Part 2 ITT</i> Bidders will make a bid to an ESA activity, work execution will (unless indicated otherwise) follow the same OGC testbed policies and procedures as <i>Part 1 CFP</i> .

Part 2 ITT is described separately from this document and has distinct response requirements. So any Bidder wishing to respond to both Part 2 ITT (described externally) and the Part 1 CFP (described herein) should deliver two separate proposals, one for each set of response requirements. A Bidder wishing to respond to one or the other (but not both) would submit only one proposal.

-

To provide a complete description of the Testbed scope, this CFP combines all deliverables from both parts (*Part 1 CFP* and *Part 2 ITT*). All work items marked as belonging to *Part 2 ITT* have distinct response requirements that are provided as part of the ITT. Bidding on these ITT requirements and work items requires an independent proposal! ITT requirements and work items are only copied into the *Part 1 CFP* solicitation documents for completeness.

The *response requirements* described here in the CFP body and in Appendix A Management Requirements apply to *Part 1 CFP* requirements and deliverables only.

Under this *Part 1 CFP*, the OGC will provide cost-sharing funds on behalf of sponsoring organizations ("Sponsors") to partially offset expenses uniquely associated with this initiative. This CFP requests proposals from organizations ("Bidders") wishing to participate in delivery and, in some cases, to receive cost-sharing funds. Any Bidder interested in initiative participation should respond by submitting a proposal per the instructions provided in the Proposal Submission Procedures.

In general, Bidders should propose specifically against the list of deliverables described under the Summary of Deliverables section below. Only content that *directly* addresses CFP requirements should be included in a Proposal. These are the only items that will be eligible for cost-sharing funds.

Any additional in-kind scope should be offered outside the formal bidding process, where an independent determination can be made as to whether it should be included in initiative scope or not. Items deemed out-of-initiative-scope might be more appropriate for inclusion in another OGC-IP initiative.

IMPORTANT

IMPORTANT

Discussions are ongoing for potential sponsorship of additional requirements that are not yet firm enough for inclusion in this CFP. Should funding for these additional requirements be provided later, a follow-on CFP with a compressed response timeline might also be issued if the work can be included without interfering with the original Master Schedule or Appendix B Technical Architecture.

Chapter 1. Benefits of Participation

This initiative provides a business opportunity for stakeholders to mutually define, refine, and evolve service interfaces and protocols in the context of hands-on experience and feedback. The outcomes are expected to shape the future of geospatial software development and data publication. The Sponsors are supporting this vision with cost-sharing funds to partially offset the costs associated with development, engineering, and demonstration of these outcomes. This offers selected Participants a unique opportunity to recoup a portion of their initiative expenses.

Chapter 2. Policies and Procedures

The will **Bylaws** initiative be conducted within the framework of OGC's Intellectual [http://www.opengeospatial.org/ogc/policies] and Property Rights Policy [http://www.opengeospatial.org/about/ipr], as agreed to in each selected Participant's Membership with the OGC-IP **Policies** and Procedures Agreement, and in accordance [http://www.opengeospatial.org/ogc/policies/ippp] and the OGC Principles of Conduct [http://www.opengeospatial.org/ogc/policies/conduct], the latter governing all related personal and public interactions.

Two key requirements are summarized below for ready reference:

- Each selected Participant will be required to agree to notify OGC staff if it is aware of any claims under any issued patents (or patent applications) which would likely impact an implementation of the specification or other work product which is the subject of the initiative. Participant need not be the inventor of such patent (or patent application) in order to provide notice, nor will Participant be held responsible for expressing a belief which turns out to be inaccurate. Specific requirements are described under the "Necessary Claims" clause of the Intellectual Property Rights Policy [http://www.opengeospatial.org/about/ipr].
- Each selected Participant will be required to refrain from making any public representations that draft Engineering Report (ER) content has been endorsed by OGC before the ER has been approved in an OGC Technical Committee (TC) vote.

OGC will create a Participation Agreement (PA) contract with each selected Participant. With regard to any overlapping subject matter, the PA Statement of Work (SOW) will take precedence over CFP Clarifications, which will take precedence over this CFP. Each Participant will also be required to provide more-detailed description of the requirements for its assigned deliverables (and to coordinate with other initiative Participants) at the Kickoff Workshop.

One initiative objective is to support the OGC Standards Program [http://www.opengeospatial.org/ ogc/programs/spec] in the development and publication of open standards. Each Participant will be required to allow OGC to copyright and publish documents based in whole or in part upon intellectual property contributed by the Participant during initiative execution. Specific requirements are described under the "Copyrights" clauses of the OGC Intellectual Property Rights Policy [http://www.opengeospatial.org/about/ipr].

2.1. Initiative Roles

The roles generally played in any OCG-IP initiative are defined in the OGC-IP Policies and Procedures [http://www.opengeospatial.org/ogc/policies/ippp], including Sponsors, Bidders, Participants, Observers, and the Innovation Program Team ("IP Team"). Additional explanations are provided in the Tips for New Bidders.

The IP Team for this initiative will include an Initiative Director (also sometimes called "initiative manager"), an Initiative Architect, and multiple Thread Architects. Unless otherwise stated, the

Initiative Director will serve as the primary point of contact (POC) for the OGC.

Thread Architects will work with Participants, Sponsors, and each other to ensure that initiative activities and deliverables are properly assigned and performed. They are responsible for scope and schedule control, as well as for within-thread and cross-thread communications. They will also provide timely escalation to the Initiative Director regarding any serious issues or risks that arise.

Chapter 3. General Proposal Submission Guidelines

This section presents general guidelines for submitting a CFP proposal. Detailed instructions for submitting a response proposal using the *Bid Submission Form* web page can be found in Appendix A Management Requirements.

Proposals must be submitted before the appropriate response due date indicated in the Master Schedule.

Bidders responding to this CFP must be OGC members and must be familiar with the OGC mission, organization, and process. Proposals from non-members will be considered *provided that* a completed application for OGC membership (or a letter of intent to become a member) is submitted prior to (or with) the proposal.

Information submitted in response to this CFP will be accessible to OGC and Sponsor staff members. This information will remain in the control of these stakeholders and will not be used for other purposes without prior written consent of the Bidder. Once a Bidder has agreed to become a Participant, they will be required to release proposal content (excluding financial information) to all initiative stakeholders. Commercial confidential information should not be submitted in any proposal and should generally not be disclosed during initiative execution.

Bidders will be selected to receive cost sharing funds on the basis of adherence to the requirements stipulated in this CFP and the overall quality of their proposal. The general initiative objective is for the work to inform future OGC standards development with findings and recommendations surrounding potential new specifications. Bidders are asked to formulate a path for producing executable interoperable prototype implementations that meet the stated CFP requirements, and for documenting the findings and recommendations arising from those implementations. Bidders not selected for cost sharing funds may still be able to participate in addressing the stated CFP requirements on a purely in-kind basis.

However, to help maintain a manageable process, Bidders are advised to avoid attempts to use the initiative as a platform for introducing new requirements not included in Appendix B Technical Architecture. Any additional in-kind scope should be offered outside the formal bidding process, where an independent determination can be made as to whether it should be included in initiative scope or not. Items deemed out-of-initiative-scope might be more appropriate for inclusion in another OGC-IP initiative.

Each selected Participant (including pure in-kind Participants) will be required to enter into a Participation Agreement contract ("PA") with the OGC. The reason this requirement applies to pure in-kind Participants is that other Participants will be relying upon their delivery to show component interoperability across organizational boundaries.

Each PA will include a Statement of Work ("SOW") identifying Participant roles and responsibilities. The purpose of the PAs is to encourage and enable Participants to work together to realize initiative goals for the benefit of the broader OGC community.

3.1. Questions and Clarifications

Once the original CFP has been published, ongoing authoritative updates and answers to questions can be tracked by monitoring the CFP Clarifications [https://proposal.opengeospatial.org/ clarification.php] page. Instructions for accessing this page are included under Proposal Submission Procedures.

Bidders may submit questions via timely submission of email(s) to the OGC Technology Desk (techdesk@opengeospatial.org [mailto:techdesk@opengeospatial.org]). Question submitters will remain anonymous, and answers will be regularly compiled and published in the clarifications.

A Bidders Q&A Webinar is scheduled for Monday, 4 February at noon U.S. Eastern Time. The webinar is open to the public, but anyone wishing to attend must register beforehand at https://register.gotowebinar.com/register/6662012816865820675.

The webinar will review any clarifications received to date and will provide an opportunity for follow-on questions.

Chapter 4. Proposal Evaluation Criteria

Proposals will be evaluated according to criteria that can be divided into two areas: *Management* and *Technical*.

4.1. Management Criteria

- Bidder willingness and ability to comply with Appendix A Management Requirements,
- Feasibility of proposed solution utilizing proposed resources,
- Proposed in-kind contribution in relation to proposed cost-share funding request.

4.2. Technical Criteria

- Understanding of and compliance with requirements as stated in Appendix B Technical Architecture,
- Quality and suitability of proposed design, and
- Where applicable, proposed solutions are OGC-compliant.

Chapter 5. Master Schedule

The following table details the major initiative milestones and events:

Milestone	Date (2019)	Event
M01	7 January	CFP Release (ITT release expected 7 Jan.)
M02	31 January	Questions for CFP Bidders Q&A Webinar Due (ITT details forthcoming)
М03	4 February	CFP Bidders Q&A Webinar (ITT details forthcoming). Register at https://register.gotowebinar.com/ register/6662012816865820675.
M04	11 February	CFP Proposal Submission Deadline (11:59pm U.S. Eastern time)
M05	31 March	All CFP Participation Agreements Signed
M06	2-4 April	Kickoff Workshop
M07	31 May	Initial Engineering Reports (IERs)
M08	June (precise dates TBD)	Interim Workshop (new for Testbed-15)
M09	30 June	Component Interface Designs
M10	31 July	IER-to-DER status check; TIE Connectivity Test; Early implementations of any component on which another component depends
M11	31 August	TIE Readiness Review
M12	30 September	TIE-Tested Component Implementations completed; Preliminary DERs complete & clean, ready for internal reviews
M13	31 October	Ad hoc TIE demonstrations (as requested during the month) & Demo Assets posted to Portal; Near-Final DERs posted to Pending & WG review requested
M14	November (specific date TBD)	Final DERs (incorporating WG feedback) posted to Pending to support WG & TC vote
M15	30 November	Participant Final Summary Reports
M16	December (specific dates TBD)	Focused Demonstration Events

Table 1. Master schedule

Chapter 6. Sequence of Events, Phases, and Milestones

The following diagram provides a notional schedule of major initiative events and milestones, and their approximate sequence of occurrence. The initiative will use rolling-wave project management whereby more detailed scheduling will take place as each milestone draws near.



Figure 1. Overview of Events and Milestones

Participant Selection and Agreements: Once the original CFP has been published, ongoing authoritative updates and answers to questions can be tracked by monitoring the CFP Clarifications [https://proposal.opengeospatial.org/clarification.php] page. Instructions for accessing this page are included under Proposal Submission Procedures.

Bidders may submit questions via timely submission of email(s) to the OGC Technology Desk (techdesk@opengeospatial.org [mailto:techdesk@opengeospatial.org]). Question submitters will remain anonymous, and answers will be regularly compiled and published in the clarifications.

OGC may also choose to conduct a Bidder's question-and-answer webinar to review these clarifications and invite follow-on questions.

Following the closing date for submission of proposals, OGC will evaluate received proposals, review recommendations with Sponsors, and negotiate Participation Agreement (PA) contracts, including statements of work (SOWs), with selected Bidders. Participant selection will be complete once PA contracts have been signed with all Participants.

Kickoff Workshop: A Kickoff Workshop ("Kickoff") is a face-to-face meeting where Participants, guided by thread architects, will refine the initiative architecture and settle upon specific interface models to be used as a baseline for prototype component interoperability. Participants will be *required* to attend the Kickoff, including the breakout sessions of each thread for which they were selected. Participants will be expected to use these breakouts to collaborate with other Participants and confirm the intended *Component Interface Designs* (broadly defined to include other types of design such as user-interfaces for portrayal).

After the face-to-face Kickoff, many initiative activities will be conducted remotely via teleconferences except the (new) Interim Meeting and possibly also SWG/DWG presentations at the TC Meetings [http://www.opengeospatial.org/event].

IMPORTANT The Interim Workshop is new for Testbed-15.

Interim Workshop: One innovation for Testbed-15 will be to conduct an Interim Workshop for Participants and other stakeholders. The purpose will be to conduct face-to-face discussions to refine our shared understanding of requirements, and to use this information to improve the quality and value of the deliverables.

The workshop is optional but highly recommended for all stakeholders able to attend, particularly those responsible for designing, agreeing, and documenting the interoperability interfaces.

The Interim Workshop for this initiative is likely to be held at a European location in the June timeframe.

Development of Engineering Reports, Change Requests, and Other Document Deliverables: Development of Engineering Reports (ERs), Change Requests (CRs) and other document deliverables will commence during or immediately after Kickoff. Participants will deliver an Initial Engineering Report (IER) plus several iterations of a Draft Engineering Report (DER). Full process details can be found in the OGC Testbed ER Development Process [https://portal.opengeospatial.org/ files/?artifact_id=73351].

Under the Participation Agreement (PA) contracts to be formed with selected Bidders, *ALL* Participants will be responsible for contributing content to the ERs. But the ER Editor role will assume the duty of being the primary ER author.

IMPORTANTER Editors should adopt a proactive approach to consuming and
documenting the knowledge being generated during component
implementation. In other words, ER Editors will serve as primary
authors in addition to their editor role. All Participants are required to
support the editors by making necessary documentation material
available.

Component Design, Development, and Preliminary Testing: Participants will continue documenting detailed and final *Component Interface Designs* (broadly defined to include other types of design such as user-interfaces for portrayal) using the selected collaboration tool (e.g., a GitHub wiki). This documentation will allow task Participants to confirm a mutual understanding of each other's interfaces so that subsequent interoperability testing can take place. A preliminary *Technology Integration Experiment* ("TIE", sometimes also referred to as a "Technology Interoperability Experiment") *Connectivity Test* milestone will be used to ensure that initiative service endpoints can be reached by initiative clients.

IMPORTANT

Each proposed *service* component implementation deliverable should be accompanied by identification of one or more particular datasets that would be suitable for the proposed service. Details are provided under Data Requirements for Proposing Service Component Implementation Deliverables.

The label "Component **Interface** Designs" is preferred over "Component Implementation Designs" because it's the interfaces between components that will have to be synchronized for the TIE experiments. The specifics of what is going on internally within any particular component is often irrelevant to the TIEs (i.e., an encapsulation approach). So, technically speaking, the full design of any particular component often does not need to be shared. Instead, what will be desirable in the TIEs is agreed interfaces that will enable seamless plug-and-play for all combinations of client-service instances. Ideally each interface design could be handed to two different developers (e.g., for one client + one service) who could then go off and start writing code independently, and then come back in a month or two with interoperating code. As a practical matter, these designs will still have to undergo some "agile" adjustment before being finalized.

TIE Readiness Review: A TIE Readiness Review will be conducted with a Thread Architect to confirm that each TIE Participant is prepared to start conducting TIE testing with counterpart Participants.

Component Interoperability Testing, and Acceptance: Participants should deliver *completed and fully TIE-tested* component implementations no later than the 30 September milestone (unless an earlier milestone is identified in Appendix B Technical Architecture). The primary acceptance criterion for a component implementation deliverable is the conduct and recording of the TIE test. This test can also help in the development of Demo Assets such as video recordings.

Draft Engineering Reports: Participants should also deliver *complete and clean* Draft Engineering Report (DERs) by the 30 September milestone. A *complete* DER is one for which all major clauses have been populated with meaningful content. A *clean* is one where all known editorial defects have been repaired. This milestone will impact *ALL* initiative Participants, even component implementers, who will be responsible for making necessary documentation material available to the ER Editor for use in authoring the ER. Initiative Sponsors and Thread Architects will review these DERs in the weeks following delivery.

DER Rework and SWG/DWG Reviews: Participants will be required to perform rework based on the reviews from Sponsors and Thread Architects. The ER editor will then make a request to the selected OGC SWG or DWG to perform its review and to consider making a request that the DER be voted on by the OGC Technical Committee (TC) for publication. The OGC 3-week rule must be followed if the DER is to receive a face-to-face vote at a TC Meeting. The DER must be posted to the members-only *OGC Pending Documents* directory to enable this TC vote. Participants will likely have to perform one final round of rework based on SWG/DWG feedback. The further along the DER is when submitted to the WG, the less rework will be required. **Final Summary Reports and Demonstration Events**: Participants Final Summary Reports will constitute the close of funded activity. Further development work might take place to prepare and refine Demo Assets (see the Demo Asset Requirements for details), which will be used to support various demonstration events such as presentations at OGC Technical Committee Meetings [http://www.opengeospatial.org/event].

Assurance of Service Availability: Participants selected to implement *service* components must maintain availability for a period of no less than one year after the Final Delivery Milestone. OGC might be willing to entertain exceptions to this requirement on a case-by-case basis.

Detailed requirements for meeting all these milestones are provided in Appendix A Management Requirements.

Chapter 7. Summary of Initiative Deliverables

The following tables show the full set of initiative deliverables, including ID, deliverable name, task, and funding status.

A deliverable's funding status can be funded ("F"), unfunded ("U"), or under negotiation ("Un-Neg"), depending on the current state of sponsor funding.

- For a deliverable with a funding status of "F", sponsor funding has already been confirmed.
- A deliverable with a funding status of "U" is within CFP scope, but has a lower priority and does not have any sponsor funding. It's possible that a Sponsor might later agree to provide funding. But no Sponsor funding is currently available, and any bids for these deliverables would necessarily be purely in-kind.
- A deliverable with a funding status of "Un-Neg" is one for which a sponsor intends to provide funding, but a final commitment of this funding is still pending.

Please also note the Requirement to Propose an In-Kind Contribution for bidding on any funded deliverable.

NOTE

A deliverable with a funding status of "*ITT*" can only be proposed under the *Part 2 ITT* solicitation. See the Main Body Introduction for details regarding how to bid on these deliverables.

Please note that each deliverable indicated as "F" or "Un-Neg" would be funded at most once. No deliverable should be interpreted as offering multiple instances. For any deliverable still under negotiation ("Un-Neg"), if funding for that deliverable ends up *not* being committed, any bid for cost-sharing on that deliverable will be dismissed. The deliverable would change to an unfunded "U" status, and a Bidder could potentially choose to contribute it purely in-kind.

IMPORTANT

Discussions are ongoing for potential sponsorship of additional requirements that are not yet firm enough for inclusion in this CFP. Should funding for these additional requirements be provided later, a follow-on CFP with a compressed response timeline might also be issued if the work can be included without interfering with the original Master Schedule or Appendix B Technical Architecture.

Initiative deliverables are organized into *threads*, which are described in detail in Appendix B Technical Architecture.

The task abbreviations are as follows:

• ML: Machine Learning

- DCS: Data-Centric Security
- DU: Delta Updates
- EOPAD: Earth Observation Process and Application Discovery
- FCA: Federated Cloud Analytics
- POR: Portrayal

7.1. CFP Deliverables Summary Table

The following table provides a summary of CFP deliverables and their funding status. Additional technical details can be found in the Technical Architecture. The following

ID	Document / Component	т	Funding Status
D001	Semantic ER	ML	F
D002	Machine Learning ER	ML	F
D010	Catalog ER (same deliverable as D010 listed below)	ML	ITT
D100	Petawawa cloud mosaicing ML model	ML	F
D101	Petawawa land cover classification ML model	ML	F
D102	New Brunswick forest ML model	ML	F
D103	Semantic Web link builder and triple generator	ML	F
D104	Quebec river-lake vectorization ML mode	ML	F
D105	Quebec model MapML WFS3.0	ML	F
D106	Quebec model MapML WFS3.0 Client	ML	F
D107	Arctic Discovery Catalog	ML	F
—	-	—	—
D004	Data Centric Security ER	DCS	F
D114	Data Centric Security Client	DCS	F
D115	Data Centric Security WFS	DCS	F
—	-	—	—
D005	Delta Updates ER	DU	F
D111	Delta Updates Client	DU	F
D112	Delta Updates WPS	DU	F
D113	Delta Updates WFS	DU	F
_	-	_	-
D121	Catalogue & Discovery Service SERVER (1 of 2)	EOPAD	F
D122	Catalogue & Discovery Service WEB CLIENT	EOPAD	F
D123	Catalogue & Discovery Service CONSOLE/DESKTOP CLIENT	EOPAD	F

ID	Document / Component	т	Funding Status
D124	Web Processing Service (1 of 2)	EOPAD	F
D117	Catalogue & Discovery Service SERVER (2 of 2)	EOPAD-ITT	ITT
D118	Web Processing Service (2 of 2)	EOPAD-ITT	ITT
D010 (same deliverable as D010 listed above)	Engineering Report	EOPAD-ITT	ΙΤΤ
-	-	—	—
D019	Federated Clouds Security ER	FCA	F
D020	Federated Clouds Analytics ER	FCA	F
D021	EOC, SCALE, and SEED ER	FCA	F
D022	Federated Clouds Provenance ER	FCA	F
D143	Federation Manager #1	FCA	F
D144	Federation Manager #2	FCA	F
D145	Scale Datacenter Environment	FCA	F
D146	WPS Analytics	FCA	F
D147	Federated Cloud Data Services	FCA	F
_	_	_	_
D011	Concept Model for Style Encoding & Metadata Model ER	POR	F
D012	Style API for WFS, WMTS, GeoPackage ER	POR	F
D013	WMTS ER	POR	F
D014	WMTS draft specification	POR	F
D015	WMTS-T ER	POR	F
D016	WMTS draft specification	POR	F
D017	Portrayal Summary Report	POR	F
D126	Visual Style Editor #1	POR	F
D127	Visual Style Editor #2	POR	F
D128	GeoPackage with Styles Encodings #1	POR	F
D129	GeoPackage with Styles Encodings #2	POR	F
D130	GeoPackage with Styles Encodings client #1	POR	F
D131	GeoPackage with Styles Encodings client #2 (mobile)	POR	F
D132	WMTS-T w Styles API #1	POR	F
D133	WMTS-T w Styles API #2	POR	F
D134	WMTS-T w Styles Client #1	POR	F
D135	WMTS-T w Styles Client #2	POR	F
D136	WFS 3.0 w Styles API #1	POR	F
D137	WFS 3.0 w Styles API #2	POR	F

ID	Document / Component	т	Funding Status
D138	WFS 3.0 w Styles Client #1	POR	F
D139	WFS 3.0 w Styles Client #2 (mobile)	POR	F

Table 2. CFP Deliverables Summary and Funding Status

< end of main body >

Appendix A: Management Requirements

This appendix presents detailed CFP management requirements for submitting a bid. It also covers procedures for participation during initiative execution. The Appendix B Technical Architecture presents the detailed CFP technical requirements.

All potential Bidders, *even experienced initiative Participants*, should read this appendix carefully from end-to-end (and notify OGC immediately if any defects are discovered). Bidders who are new to OGC initiatives are also encouraged to review the Tips for New Bidders.

The following sections describe the processes for a Bidder to submit a proposal and for a Participant (i.e., a selected Bidder) to perform against a Participation Agreement (PA) contract. The order of topics roughly parallels the sequence described in the Master Schedule. In general, the term "activity" is used as a verb describing work to be performed and the term "deliverable" is used as a noun describing artifacts to be developed, documented and delivered for inspection and use.

A.1. Proposal Submission Procedures

The process for a Bidder to complete a proposal is essentially embodied in the online Bid Submission Form [https://proposal.opengeospatial.org/]. Once this site is fully prepared to receive submissions (likely within several days of the CFP release), it will include a series of web forms, one for each deliverable of interest. A summary is provided here for the reader's convenience.

Clicking on the *Clarifications* link will navigate to the CFP Clarifications [https://proposal.opengeospatial.org/clarification.php] page, which is available to the public. Other functions will require that an account be created.

Once an online account has been created, the user will be taken to a home page indicating the "Status of Your Proposal." If any defects in the form are discovered, this page includes a link for notifying OGC. The user can return to this page at any time by clicking the *OGC logo* in the upper left corner.

Because the *Bid Submission Form* is still relatively new, it might contain some areas that are still brittle or in need of repair. Please notify OGC of any discovered defects. Periodic version updates will be provided as needed.

IMPORTANTPlease consider making backup local copies of all inputs in case any of
them need to be re-entered.

Please also note that once this form goes "live", any submitted bids will be treated as earnest submissions, even those submitted well before the response deadline. Be certain that you intend to submit your proposal before you click the *Submit* button on the *Review* page.

A.1.1. High-Level Overview

Clicking on the *Propose* link will navigate to the Bid Submission Form. The first time through, the user should complete all the organizational information fields and click *Update and Continue*.

This will navigate to an "Add Deliverable" page that will resemble the following:

OGC Proposal Manager × +									
1 C C L https://proposal.opengeospatial.org/deliverables.php		v C	+ 9		≡				
<u>OGC</u> Clarificat	tions Propose Review	Welc	ome gpercivall	Logou	ıt				
	Add Deliverable:								
	Deliverable*: Projected Labor Hours (for cost-share request only)*: Total Dollar Funding Request (may cover labor hours only)*: Inkind Labor Hours (in <u>addition to</u> Projected Labor Hours)*: Inkind Contribution (across all cost categories, not just labor)*: Proposed Contribution (Please include an NOTE - Images cannot be pasted into edit Buehler (gbuehler@opengeospatial.org): Image: Edit < View ≤ Insert ≤ Formati M → Formats ⊂ B I E E E	NISTO01 - Building Data USD) \$ USD) \$ USD) \$ (USD) \$ y proposed datasets) for. You can link to publicly available URLs if you need to include an image. If you have problems, contact Greg ✓							
	Developed and Maintained by OGC								

Figure 2. Sample "Add Deliverables" Page

The user should complete this form for *each* proposed deliverable.

On the far right, the *Review* link navigates to a page summarizing all the deliverables the Bidder is proposing. Note that this Review tab won't appear until the user has actually submitted at least one deliverable under the Propose tab first.

TIP

Consider regularly creating printed output copies of this *Review* page at various checkpoints during proposal creation in case an input error is made later.

Once a submission has been received, the system will send an email to the bidder and to OGC staff.

ТΙР

In general, up until the time that the user clicks this *Submit* button, the proposal may be edited as many times as the user wishes. However, this initial version of the form contains no "undo" capability, so please use caution in over-writing existing information.

The user is afforded an opportunity under *Done Adding Deliverables* at the bottom of this page to attach an optional *Attached Document* of explanation. If this attachment is provided, it is limited to one per proposal and there is a **5Mb file size limitation**.

This document could conceivably contain any specialized information that wasn't suitable for entry into a *Proposed Contribution* field under an individual deliverable. It should be noted, however, that this additional documentation will only be read on a **best-effort basis**. There is no guarantee it will be used during evaluation to make selection decisions; rather, it could optionally be examined if the evaluation team feels that it might help in understanding any specialized (and particularly promising) contributions.

A.1.2. Detailed Instructions

The *Propose* link takes the user to the first page of the proposal entry form. This form contains fields to be completed once per proposal such as names and contact information.

It also contains an optional *Organizational Background* field where Bidders (particularly those with no experience participating in an OGC initiative) may provide a description of their organization. It also contains a click-through check box where each Bidder will be required (before entering any data for individual deliverables) to acknowledge its understanding and acceptance of the requirements described in this appendix.

Clicking the *Update and Continue* button then navigates to the form for submitting deliverableby-deliverable bids. On this page, existing deliverable bids can be modified or deleted by clicking the appropriate icon next to the deliverable name. Any attempt to delete a proposed deliverable will require scrolling down to click a *Confirm Deletion* button.

To add a new deliverable, the user would scroll down to the *Add Deliverable* section and click the *Deliverable* drop-down list to select the particular deliverable of interest.

The user would then enter the required information for each of the following fields (*for this deliverable only*). Required fields are indicated by an asterisk ("*"):

- Estimated Projected Labor Hours* for this deliverable
- *Funding Request**: total U.S. dollar cost-share amount being requested for this deliverable (to cover burdened labor only)
- Estimated In-kind Labor Hours* to be contributed for this deliverable
- Estimated *In-Kind Contribution*: total U.S. dollar estimate of the in-kind amount to be contributed for this deliverable (including all cost categories)

Cost-sharing funds may **only** be used for the purpose of offsetting burdened labor costs of development, engineering, and demonstration of initiative outcomes related to the Participant's assigned deliverables. By contrast, the costs used to formulate the Bidder's in-kind contribution may be much broader, including supporting labor, travel, software licenses, data, IT infrastructure, etc.

Theoretically there is no limit on the size of the *Proposed Contribution* for each deliverable (beyond the raw capacity of the underlying hardware and software). But bidders are encouraged to incorporate content by reference where possible (rather than inline copying and pasting) to avoid overloading the amount of material to be read in each proposal. There is also a textbox on a separate page of the submission form for inclusion of *Organizational Background* information, so there is no need to provide this information for each deliverable.

A breakdown (by cost category) of the "Inkind Contribution" may be included in the *Proposed Contribution* text box for each deliverable. However, please note that the content of this text box will be accessible to all Stakeholders and should contain no confidential information such as labor rates. If a Bidder wishes to provide confidential information that will only be accessible to OGC staff, it can be included in the *Attached Documentation* of the submission. Only one *Attached Document* may be submitted for the entire proposal (there is no way to submit a separate attachment for each individual deliverable).

IMPORTANT

A Bidder proposing to deliver a *Service Component Implementation* should use the *Proposed Contribution (Please include any proposed datasets)* field to identify what suitable *datasets* would be contributed (or what data should be acquired from another identified source) to support the proposed service. Additional guidelines for this requirement can be found under Data Requirements.

This field *Proposed Contribution (Please include any proposed datasets)* could also be used to provide a succinct description of what the Bidder intends to deliver for this work item to meet the requirements expressed in Appendix B Technical Architecture. This language could potentially include a brief elaboration on how the proposed deliverable will contribute to advancing the OGC standards baseline [http://www.opengeospatial.org/standards] (also see Standards Baseline), or how implementations enabled by the specification embodied in this deliverable could add specific value to end-user experiences.

However, to help maintain a manageable process, Bidders are advised to avoid attempts to use the initiative as a platform for introducing new requirements not included in the Technical Architecture. Any additional in-kind scope should be offered outside the formal bidding process, where an independent determination can be made as to whether it should be included in initiative scope or not. Items deemed out-of-scope might be more appropriate for inclusion in another OGC-IP initiative.

The statements of work (SOWs) in the *Participation Agreement* (PA) contracts will ultimately require performance during initiative *execution* against the deliverables as described in the

Technical Architecture plus any subsequent Corrigenda.

A single bid may propose deliverables arising from any number of threads. To ensure that the full set of sponsored deliverables are made, OGC might negotiate with individual Bidders to drop and/or add certain deliverables from their proposals.

A.2. Conditions for Participation

A selected Bidder will become an initiative Participant by executing a *Participation Agreement* (PA) contract with OGC. But in order to be selected for cost-share funding for a particular deliverable in the first place, a bid must also meet the requirement to propose an in-kind contribution.

A.2.1. Requirement to Propose an In-Kind Contribution

Each cost-share request (i.e., a bid on a funded deliverable) is **required** to provide at least some level of **in-kind contribution** (i.e., activities requesting no cost-share compensation). As a rough guideline, a proposed deliverable should include at least one dollar of in-kind contribution for every dollar of cost-sharing compensation requested. All else being equal, **higher levels of in-kind contributions will be considered more favorably** during evaluation.

Some participation may be fully in-kind. However, to help maintain a manageable process, Bidders are advised to avoid attempts to use the initiative as a platform for introducing new requirements not included in the Technical Architecture. Any additional in-kind scope should be offered outside of the formal bidding process, where an independent determination can be made as to whether it should be included in initiative scope or not. Items deemed out-of-initiativescope might be deemed more appropriate for inclusion in another OGC-IP initiative.

Any item proposed as a *purely* in-kind contribution to meet a requirement already included in the Technical Architecture will likely be accepted if it meets all the other evaluation criteria.

Assuming that this in-kind requirement has been met for every proposed deliverable, a Bidder must also meet the following conditions in order to enter a PA contract:

- Each Bidder should use the Bid Submission Form to make their proposal.
- Proposals from non-OGC-members will be considered *provided that* a completed application for OGC membership (or a letter of intent to become a member) is submitted prior to (or with) the proposal. A non-member Bidder could conceivably defer payment of its membership application fee until after it has learned whether it has been selected for cost-share funds or not.
- No work facilities will be provided by OGC. Each Participant will be required to perform its PA obligations at its own provided facilities and to interact remotely with other initiative stakeholders.
- In general, a proposed component deliverable based on a product that has earned OGC Certification will be evaluated more favorably than one which has not.

- Bidders proposing to build interoperable components should be prepared to test and demonstrate interoperability with components supplied by other Participants. In particular, server-endpoints must be accessible over the public Internet during TIE experiments.
- Participants selected to implement component deliverables will be expected to participate in the full course of interface and component development, TIEs, and demonstration support activities throughout initiative execution.
- Participants selected as document Editors will also be expected to participate in the full course of activities throughout the initiative, recording implementation findings & recommendations and ensuring document delivery.
- Participants should remain aware of the fact that the components will be developed across many organizations. To maintain interoperability, each Participant should diligently adhere to the latest technical specifications so that other Participants may rely on the anticipated interfaces during the TIEs.
- All Selected Participants (both cost-share and purely in-kind) must send *at least one technical representative per assigned thread* to the Kickoff Workshop.
 - Participant representatives are encouraged to attend the Interim Workshop.
 - Participant representatives may also have opportunities to attend focused demonstration events that are likely to take place near the end of initiative execution.

Please note that although the bid submission form requires labor-hour estimates for each deliverable, the Participation Agreement (PA) will almost never be a time-and-materials contract to reimburse labor-hours. Instead, each PA will be a fixed-price contract with a particular Statement of Work (SOW) listing the required deliverable(s). A payment schedule will be established so that Participants may submit invoices at various milestones.

A.3. Proposal Evaluation and Invitations to Selected Bidders

Specific proposal evaluation criteria were presented earlier. Several process steps to be conducted by the IP Team are presented below to aid readers in understanding the overall process.

Soon after the proposal submission deadline, the IP Team and Sponsors will begin reviewing proposals, examining suitability of the proposed deliverables in light of the CFP requirements. During this analysis, the IP Team might need to contact Bidders to obtain clarifications and better understand what is being proposed.

At the *Decision Technical Evaluation Meeting I* (TEM I), the IP Team will present Sponsors with draft recommendations regarding which parts of which proposals should be offered cost-share funding. Sponsors will use this opportunity to suggest modifications.

The IP Team will notify each Bidder of their selection, including a description of the particular deliverables being offered, and OGC's intent to enter into a *Participation Agreement* (PA) with

them. Selected Bidders *must be available for these contacts to be made* to enable confirmation of continued interest. Bidder acceptance of the offer essentially *locks in the high-level scope* that will be used to form the SOW to be attached to the forthcoming PA. A Bidder may reject the offer (for example, if the offer is to deliver only a small subset of its proposed deliverables).

A *Decision Technical Evaluation Meeting II* (TEM II) meeting will be conducted to review the outcomes of initial offers to selected Bidders, and any modifications that were needed due to Bidder rejections.

Following TEM II, the IP Team will develop the formal SOW and attach it to the full PA for each selected Bidder. Selected Bidders *must be available* to enable finalization of their PA contract.

Each PA will include a final identification of all assigned deliverables. The specific requirements that these deliverables must meet will be those documented in Appendix B Technical Architecture (plus any subsequent corrigenda).

Because testbeds tend to deal with the lowest maturity level of OGC standards technology, they operate under the principle that "interoperable running code wins" (based loosely on an Internet Engineering Task Force founding belief [https://www.ietf.org/tao.html#rfc.section.2]).

In contrast with ordinary commercial contracts, which often focus on delivering a singlecustomer solution, the *Participation Agreement* contracts will encourage Participants to work in the **broader interests of the OGC community**. These interests are generally represented in the OGC Standards Working Groups (SWGs) and Domain Working Groups (DWGs). Initiative sponsorship enables a selected subset of these interests to be brought to bear on the production of *interoperable running code* and the reporting back of findings and recommendations as Change Requests (CRs) and document deliverables under the OGC Testbed ER Development Process [https://portal.opengeospatial.org/files/?artifact_id=73351].

In general, at least one Participant will be assigned for each client-server pair as an *Engineering Report* (ER) Editor, who will be responsible for authoring the ER, though there might be exceptions. See **Requirements** for Proposing Deliverables for details. An *ER Template* will be provided to assist in the breakdown of content into multiple chapters describing specific findings and recommendations, a summary chapter, and supporting chapters (e.g., containing references and revision history).

IMPORTANT

ER Editors should adopt a *proactive approach to consuming and documenting the knowledge* being generated during component implementation. In other words, ER Editors will serve as primary authors in addition to their editor role. All Participants are required to support the editors by making necessary documentation material available.

OGC members who who are not selected to make any deliverables may still observe all initiative activities by visiting the (members-only) OGC Observer Agreements [https://portal.opengeospatial.org/?m=public&orderby=default&tab=6] page and registering to become an

initiative Observer. The Observer Agreement will require that Observers follow the same intellectual property protection rules as apply to other initiative stakeholders (e.g., to avoid sharing other stakeholder's confidential information with individuals outside the initiative).

Stakeholders should also avoid making any representations that any of the intermediate technical decisions made internally during initiative execution reflect any sort of endorsed position of the OGC. Formal initiative recommendations will become "official" once the containing ER has received approval to be made public.

A.4. Kickoff Workshop

Initiative execution will commence with a Kickoff Workshop event ("Kickoff"). In general, the Kickoff ensures a common Participant understanding of the overall Technical Architecture and of all interfaces relevant to their assigned deliverables. Under the initiative's rapid pace, and guided by the Thread Architect, periods of development will be followed by synchronization among component developers, so that the likelihood of interoperation is maximized at each iteration. To maintain interoperability, each Participant should diligently adhere to the latest agreed specifications so that other Participants may rely on the anticipated interfaces in subsequent TIE testing.

Since Participants will be relying on one another to deliver interoperable components, all Participation Agreement (PA) contracts should be settled by the start of Kickoff. Any Participant that has not yet executed a PA by start of Kickoff will be required to **attest to its commitment** to a PA Statement of Work (SOW). The full PA must then be executed with OGC **no later than Kickoff completion** (two days after Kickoff start).

The Kickoff will include both plenary and thread-based breakout sessions. Thread breakouts will be used to ensure a mutual understanding of the detailed interfaces that will support component interoperability. The Kickoff will also present an opportunity to finalize regularly scheduled thread and subthread teleconferences ("telecons").

All selected Participants must send **at least one technical representative per assigned thread to the Kickoff Workshop**, and a representative should be present at the breakout sessions for **every thread** for which the Participant has been assigned a deliverable.

IMPORTANT

In-person attendance at these breakout meetings has been found to be critical to later initiative success. Remote attendance will not be an option.

These requirements will also apply to Participants not requesting any cost-share funding (i.e., purely in-kind contribution) since their deliverables are still being *relied upon* by other Participants.

A.5. Participant Communication and Reporting

A.5.1. Points of Contact

Each selected Participant should designate a primary point of contact ("Primary POC") who must remain available throughout initiative execution for communications regarding status. The POC should identify at least one alternative point of contact to support the Primary POC as needed. The POCs should provide contact information including their e-mail addresses and phone numbers.

All proposals should include a statement attesting to the POCs' understanding and acceptance of the duties described herein. This statement will be enabled by a simple check-box in the *Bid Submission Form*.

A.5.2. Kickoff Status Report

Selected Participants should provide a one-time Kickoff Status Report that includes a list of personnel assigned to support the initiative and assurance that the Participants understand the schedule for all of its deliverables. This report should be submitted in electronic form to a designated email address no later than the last day of the Kickoff event.

A.5.3. Monthly Progress Reports

Participant Business and Technical POCs should provide, respectively, monthly *business* and *technical* progress and status reports due on the 3rd of each month (or the first working day thereafter) and covering the previous month.

Monthly Technical Reports will be accessible to all Stakeholders and should contain no confidential information (e.g., labor rates). *Monthly Business Reports* may contain confidential information and should be shared only with identified OGC POCs.

Any Participant who has a reliable forecast of what will take place in the remaining days of any particular reported month may submit its report early and subsequently report any urgent, last-minute updates via a follow-on email.

Detailed monthly reporting requirements and procedures will be provided during PA contract negotiation. It is likely that the Portal-based *Actions* feature will continue to be used for monthly technical reporting.

The IP Team Thread Architects will review action item status on a weekly basis with assigned Participants, who should be available for these contacts to be made. Additional details regarding project management activities during initiative execution can be found at Project Monitoring and Control.

As a preview, monthly technical reports will likely require the following information for each deliverable:

- Deliverable ID and Name
- Health: G, Y, R (green, yellow, red)
- %complete (0%-100%)
- Work accomplished in reporting period (last month)
- Work anticipated in reporting period+1 (current month)
- Any known risks or issues (see definitions under Project Monitoring and Control)
- Response to mitigate the risk or remedy the issue

A.5.4. Final Summary Report

Participants should provide a *Final Summary Report* near the end of initiative execution. Detailed requirements and procedures will be provided during PA contract negotiation. These reports will likely require the following information:

- Describe, in detail, the work completed to fulfill the PA SOW items,
- Summarize Participant's overall contribution to the project, and
- Present recommendations for future OGC-IP and Standards Program efforts.

Any timely Participant who is up-to-date with all the Monthly Technical Status Reports in the Portal will likely be permitted to optionally satisfy the first "in detail" requirement by merely referencing the detail already contained in the Monthly Technical Reports. Timely Participants may also be permitted choose to build their Final Summary Report in the body of an email if desired (or, alternatively, submitted as a document attachment).

A.5.5. Regular Teleconferences

At least one of Participant Technical POC should be available for regularly scheduled telecons for *each* thread in which it is participating. Weekly thread telecons will be conducted and recorded, (with occasional exceptions for events such as for OGC Technical Committee Meetings [http://www.opengeospatial.org/event]).

These telecons are intended to accelerate understanding and action regarding relevant initiative activities, particularly status of any risks or issues that might block or are blocking progress toward timely delivery.

Participants assigned as ER Editors may, with permission from the relevant Thread Architect, lead the substantive discussions during any ad hoc or regularly scheduled sub-thread telecons.

A.5.6. Correspondence and Collaboration

Participants should be able to utilize the following correspondence and collaboration tools:
- Send and receive emails to/from thread-based email broadcast lists,
- Contribute technical content using the selected collaboration tool (e.g., GitHub [https://github.com/] Wiki),
- Participate in telecons using the GoToMeeting [https://www.gotomeeting.com/] tool,
- Edit ER source files in the Asciidoc [http://asciidoc.org/] format using the OGC ER Template [https://github.com/opengeospatial/er_template],
- Upload ER source files to the initiative GitHub [https://github.com/] repository.

A.6. Requirements for Proposing Deliverables

In general, Participants will be required to make deliverables meeting the requirements stated in CFP Appendix B Technical Architecture.

Under the Participation Agreement (PA) contracts to be formed with selected Bidders, all Participants will be responsible for contributing content to the ERs.

But the *ER Editor* role will assume the duty of being the *primary* ER author. ER Editors should adopt a *proactive approach to consuming and documenting the knowledge* being generated during component implementation. In other words, ER Editors will serve as primary authors in addition to their editor role. All Participants are required to support the editors by making necessary documentation material available.

In addition, *component implementers* (Participants selected to deliver component implementations) will be responsible for contributing *completed* technical artifacts for inclusion as ER annexes. These include UML diagrams, XML schemas and instances, and abstract test suites, all of which have their own specialized clauses in the ER Template. Component implementers will also be responsible for contributing demo assets such as video recordings of component operation.

ER Editors will be responsible for observing the component implementation work and documenting (in the assigned ER) all findings and recommendations (except the technical artifacts identified above). They will also be responsible for ensuring the document's overall integrity, authoring summary material and ensuring that global sections (e.g., References, Terms, Revision History, Bibliography) have all been completed.

ER Editors will also be responsible for enabling all required ER reviews (and subsequent rework), including internal reviews by Thread Architects and Sponsors, as well as requests for reviews by an appropriate OGC Standard Working Group [http://www.opengeospatial.org/projects/groups/swg] (SWG) or Domain Working Group [http://www.opengeospatial.org/projects/groups/wg] (DWG).

ER Editors will also be responsible for creating (and referencing) any **Change Requests** (CRs) arising from the work described in the assigned ER. CRs can be created using the OGC CR system [http://ogc.standardstracker.org/].

A.6.1. Specific Requirements for Proposing Document Deliverables

The primary acceptance criterion for a document deliverable (e.g., an ER) will be approval from a Thread Architect to post the document to an OGC (members-only) folder called the *OGC Pending directory*. Permission to post will not be granted until the full OGC Testbed ER Development Process [https://portal.opengeospatial.org/files/?artifact_id=73351] has been followed, including the step to request a review from an OGC SWG or DWG. All appropriate tooling should also have been utilized (ER Template, Asciidoc, GitHub), and the document should have achieved a satisfactory level of consensus among all assigned Participants.

A.6.2. Specific Requirements for Proposing *Component Implementation* **Deliverables**

In general, prototype *component implementation* deliverables (including services, clients, datasets, and tools) will be provided by methods suitable to its type and stated requirements. A service (e.g. a WFS instance or even an intermediary component) will be delivered by deployment for use in the initiative via an accessible URL. A client will be used to exercise a service to test and demonstrate interoperability. Components will be developed and deployed in all threads for integration and interoperability testing in support of agreed-upon thread scenario(s) and technical architecture. The Kickoff will provide the opportunity to establish many of the agreed-upon details. These technical artifacts could also be utilized for cross-thread scenarios in demonstration events.

More specifically, the first acceptance criterion for any *component implementation* deliverable will be a successful demonstration of capability in a *Technology Integration Experiment* ("TIE") test.

The next acceptance criterion will be assurance that the work-package ER Editor has access to sufficient documentation material to record all findings and recommendations in the ER and associated CRs.

Another criterion will be evidence that the implementer has *directly authored* full ER content for any relevant Abstract Test Suite (ER Template Annex A), XML Schema Document (ER Template Annex B), or UML Model (ER Template Annex C).

A Participant assigned to deliver a *client component implementation* should, in addition, **record TIE test execution** (as evidence of delivery).

A Participant assigned to deliver a *service component implementation* should, in addition, provide assurance that the service endpoint will remain available to OGC and Sponsors for a period of no less than one year after initiative execution. OGC might be willing to entertain exceptions to this requirement on a case-by-case basis.

A Participant assigned to deliver a *service component implementation* should, in addition, identify of one or more particular datasets that would be suitable for the proposed service. Details are provided under Data Requirements for Proposing Service Component Implementation Deliverables. Some requirements might request delivery of a packaged prototype component (e.g., using a virtual machine or container). Consult the specific language in the Appendix B Technical Architecture for details regarding these requirements.

A.6.3. Demo Asset Requirements for Proposing *Component Implementation* **Deliverables**

IMPORTANT

Demo requirements might vary between the Part 1 CFP and the Part 2 ITT. The guidelines here refer just to the CFP.

Under the CFP, all selected component implementers will be required to contribute "Demo Assets". These Demo Assets will be used to help compose outreach and communication artifacts toward the end of the initiative. There may also be other demonstration opportunities such as during an ER Editor's presentation to a Working Group (e.g., the target SWG/DWG that agreed to review the ER). It might also appear in a focused event such as an ESIP Meeting [http://www.esipfed.org/meetings].

These Demo Assets are not required to be professional-grade videos (which are not easily produced and can consume a lot of time and effort). Here a representative example of a Demo Asset from Testbed-13 [https://www.youtube.com/watch?v=MrXMLBwWYO0].

Also, some component implementers may seek permission to simply point to another Participant's Demo Asset (e.g., a video recording of a client interacting with a back-end service that was delivered by that Participant).

Component implementers will also be permitted to create more intricate demo assets that will be considered for upload to the OGC YouTube channel [https://www.youtube.com/user/ogcvideo]. But these additional videos will be optional, not mandatory.

There will be no independent dedicated initiative "Demo Event" as has taken place in some prior initiatives.

A.6.4. Data Requirements **for Proposing** Service Component Implementation **Deliverables**

A Bidder proposing to deliver a prototype *service component implementation* should provide *detailed information* regarding what *data* would be contributed (or what data should be acquired from another identified source) to support the proposed service. Some initiative Sponsors might provide data, but it might also be necessary to complement these with additional datasets.

A Bidder may propose any option that would be suitable to meet the technical requirement stated in the Appendix B Technical Architecture. Here are some examples to illustrate what options might be suitable:

- The San Francisco Bay served as the primary area of interest (AOI) for service component implementation data in Testbed-12. Full test dataset details are described in the document Testbed-12 Test Dataset Implementation with Documentation (16-136) [http://docs.opengeospatial.org/guides/16-136.html]. One of the primary sources for Testbed-12 data was the Homeland Infrastructure Foundation-Level Data (HIFLD) [https://hifldgeoplatform.opendata.arcgis.com/].
- The HIFLD datasets [https://respond-irma-geoplatform.opendata.arcgis.com/] covering the area of south Florida impacted by Hurricane Irma in September, 2017.
- The NGA Data Depot [https://www.geoint.community/data-depot] from the GEOINT Solutions Marketplace [https://www.geoint.community/].
- Testbed-12 [http://docs.opengeospatial.org/guides/16-136.html] reused OpenStreetMap feature data (which had actually originated in Testbed-11).

Any datasets that are not intended to be made public can be protected by a *click-through license* which restricts access solely for use in support of the initiative and which proscribes any additional release, reproduction, or use.

A.7. Project Monitoring and Control

Monitoring and control of initiative execution will follow an encapsulation principle. As long as a Participant continues [1] supporting activities on which other initiative stakeholders are relying (e.g., telecons, TIEs, ER inputs) and [2] making timely delivery of their own deliverables, the question of what's "going on" inside the Participant organization will likely not be raised.

Otherwise, any issues (or risks) that block (or threaten to block) timely delivery will enter a *remediation* status requiring the following responses:

- Development, implementation, and reporting of a credible mitigation or recovery response to get back on track,
- Recording in the Monthly Participant Technical Status Report a "Red Flag" Action Status indicating that the deliverable is in remediation (plus either "Red" or "Yellow" health in the Action Description),
- Notification of the relevant Thread Architect(s), who will review the recovery plan and report status to the Initiative Director and Sponsors.

The following definitions will be used to communicate and track deliverable status:

- For purposes of this initiative, a *risk* can be defined as any uncertain future event that has a sufficiently high exposure to threaten Participants' ability to make on-time delivery of the promised deliverable scope.
- An *issue* is a risk that has actually occurred (no longer any uncertainty surrounding whether it will occur or not).

- A *mitigation* response is one that reduces a risk's exposure (e.g., finding an alternative source for meeting a data dependency).
- A *recovery* response is one that remedies the damage caused by the issue (e.g., getting the deliverable back on schedule or being excused from delivery of some particular requirement).
- A deliverable whose timely delivery is threatened should have an overall status of being *in remediation* and should reflect a *Red* or *Yellow* health.
- A *Red* deliverable health is appropriate when the response has not yet been created and agreed, or it has been created and agreed but is not believed to be sufficient to make timely delivery.
- A *Yellow* deliverable health is appropriate when the response has been created and agreed, and is in the process of execution (but not yet fully completed).
- Once the response has been fully completed, the deliverable health should return to *Green* and the deliverable would no longer be considered in remediation.

In general, risks and issues that are contained within a particular thread will be monitored and managed by the assigned Thread Architect, who will report ongoing status to other stakeholders. Any risks or issues that cross threads should be brought to the attention of all relevant Thread Architects and the Initiative Director.

A.8. Tips for New Bidders

Bidders who are new to OGC initiatives are encouraged to review the following tips, many of which have been drawn from the clarifications of prior initiatives.

- The roles generally played in any OCG-IP initiative are defined in the OGC-IP Policies and Procedures [http://www.opengeospatial.org/ogc/policies/ippp], from which the following definitions are derived and extended.
 - Sponsors are OGC member organizations that contribute financial resources to steer initiative requirements toward rapid development and delivery of proven candidate specifications to the OGC Standards Program. These requirements take the form of the deliverables described herein. Sponsors representatives help serve as "customers" during initiative execution, helping ensure that requirements are being addressed and broader OGC interests are being served.
 - Bidders are organizations who submit proposals in response to this CFP. A Bidder selected to participate will become a Participant through the execution of a Participation Agreement contract with OGC. Most Bidders are expected to propose a combination of cost-sharing request and in-kind contribution (though purely in-kind contributions are also welcomed).
 - Participants are OGC member organizations that generate empirical information through the definition of interfaces, implementation of prototype components, and documentation of all related findings and recommendations in Engineering Reports,

Change Requests and other artifacts. They also make substantial in-kind contributions to an initiative. Participants assign business and technical representatives to represent their interests throughout initiative execution.

- Observers are individuals from OGC member organizations that have agreed to OGC intellectual property requirements in exchange for the privilege to access initiative communications and intermediate work products. They may contribute recommendations and comments, but the IP Team has the authority to table any of these contributions if there's a risk of interfering with any primary initiative activities.
- The Innovation Program Team (IP Team) is the management team that will oversee and coordinate the initiative. This team is comprised of OGC staff, representatives from member organizations, and OGC consultants. The IP Team communicates with Participants and other stakeholders during initiative execution, provides initiative scope and schedule control, and assists stakeholders in understanding OGC policies and procedures.
- The term Stakeholders is a generic label that encompasses all initiative actors, including representatives of Sponsors, Participants, and Observers, as well as the IP Team. Initiative-wide email broadcasts will often be addressed to "Stakeholders".
- Suppliers are organizations (not necessarily OGC members) who have offered to supply specialized resources such as capital or cloud credits. OGCs role is to assist in identifying an initial alignment of interests and performing introductions of potential consumers to these suppliers. Subsequent discussions would then take place directly between the parties.
- Non-OGC member organizations must become members in order to be selected as Participants.
- Any individual wishing to gain access to initiative intermediate work products in the restricted area of the Portal (or attend the private initiative working meetings / telecons) must be a member-approved user of the OGC Portal system.
 - The reason for this restriction is that members are bound by the OGC bylaws, which offer protections to other members. Affording the same privileges to non-members could have a chilling effect on what are intended to be free and innovative initiative discussions.
- Individuals from any OGC member organization that does not become an initiative Sponsor or Participant may still (as a benefit of membership) quietly observe all initiative activities by registering as an *Observer*.
- Prior initiative participation is not a direct bid evaluation criterion.
 - However, prior participation could accelerate and deepen a Bidder's understanding of the information presented in the CFP.
- All else being equal, preference will be given to proposals that include a larger proportion of in-kind contribution.
- All else being equal, preference will be given to proposed components that are certified OGC compliant.

- All else being equal, a proposal addressing all of a deliverable's requirements will be favored over one addressing only a subset.
 - Each Bidder is at liberty to control its own proposal, of course. But if it does choose to propose only a subset for any particular deliverable, it might help if the Bidder prominently and unambiguously states precisely what subset of the deliverable requirements are being proposed.
- Sponsors will be given an opportunity to review selection results and offer advice, but ultimately the Participation Agreement (PA) contracts will be formed bilaterally between OGC and each Participant organization. No multilateral contracts will be formed.
 - Beyond this, there are no restrictions regarding how a Participant chooses to accomplish its deliverable obligations so long as the Participant's obligations are met in a timely manner (e.g., with or without contributions from third-party subcontractors).
 - Since each PA will be a bilateral contract between OGC and the bidding organization, in general OGC will not scrutinize the details of how the Bidder chooses to acquire any downstream subcontracting services so long as it promises to make timely delivery. So a Bidder's in-kind proposal could, for example, include an estimated market-value of a downstream subcontractor's contribution to the Bidder's own total proposed in-kind contribution. However, since cost-share funding may only be used for labor cost, any downstream subcontractor contribution amount may only be listed under the Bidder's in-kind contribution (not cost-share).
- In general, only one organization will be selected to receive cost-share funding per deliverable, and that organization will become the *Assigned Participant* upon which other Participants will rely for delivery.
 - Optional in-kind contributions may be made provided that they don't disrupt delivery of the required, reliable contributions from *Assigned Participants*.
- A Bidder may propose against any or all threads.
 - Participants in past initiatives have often been assigned to make only a single deliverable. At the other extreme, there is theoretically no upper limit on the number of deliverables a single organization could be assigned to make.
- The period of performance for PAs is expected to run through December.
- In general, the PAs will not require delivery any component source code to OGC.
 - What is delivered instead is the behavior of the component in the TIEs, and the corresponding documentation of findings, recommendations, and technical artifacts in the related ER(s).
 - In some instances, a Sponsor might expressly require a component to be developed under open-source licensing, in which case the source code would become publicly accessible outside the initiative as a by-product of implementation.
- Results of other recent OGC initiatives can be found in the OGC Public Engineering Report Repository [http://docs.opengeospatial.org/per/].

• A Bidders Q&A Webinar will likely be conducted shortly before the proposal submission deadline. The webinar will be open to the public, but prior registration will be required.

< end of appendix >

Appendix B: Technical Architecture

B.1. Introduction

This Annex B provides background information on the OGC baseline, describes the Testbed architecture and thread-based organization, and identifies all requirements and corresponding work items. For general information on the Testbed, including deadlines, funding requirements and opportunities, please be referred to the CFP Main Body.

Each thread aggregates a number of tasks. Each task specifies a number of requirements that need to be fulfilled by work items. The work items are funded by different sponsors.

An overview of all threads and assigned tasks is provided further below.

B.2. Testbed Baseline

B.2.1. Types of Deliverables

OGC Testbed threads require several types of deliverables. Deliverable indications "funded" or "unfunded" are provided in the CFP Main Body. Please make sure you understand your tasks and provision requirements according to Annex A.

Documents

Engineering Reports (**ER**) and **Change Requests** (**CR**) will be prepared in accordance with OGC published templates. Engineering Reports will be delivered by posting on the OGC Portal Pending Documents list when complete and the document has achieved a satisfactory level of consensus among interested participants, contributors and editors. Engineering Reports are the formal mechanism used to deliver results of the Innovation Program to sponsors and to the OGC Standards Program [http://www.opengeospatial.org/ogc/programs] and OGC Standard Working [http://www.opengeospatial.org/projects/groups/swg] Group or Domain Working Groups [http://www.opengeospatial.org/projects/groups/wg] for consideration.

IMPORTANT

It is emphasized that participants delivering Engineering Reports must also deliver Change Requests that arise from the documented work.

Implementations

Services, **Clients**, **Datasets**, **Models** and **Tools** will be provided by methods suitable to its type and stated requirements. For example, services and components (e.g. a WFS instance) are delivered by deployment of the service or component for use in the Testbed via an accessible URL. A Client software application or component may be used during the Testbed to exercise services and components to test and demonstrate interoperability; however, it is most often not delivered as a license for follow-on usage. Implementations of services, clients and data

instances will be developed and deployed in all threads for integration and interoperability testing in support of the agreed-up thread scenario(s) and technical architecture. The services, clients, and tools may be invoked for cross-thread scenarios in demonstration events.

B.2.2. OGC Reference Model

The OGC Reference Model (ORM) [http://www.opengeospatial.org/standards/orm] version 2.1, provides an architecture framework for the ongoing work of the OGC. Further, the ORM provides a framework for the OGC Standards Baseline. The OGC Standards Baseline consists of the member-approved Implementation/Abstract Specifications as well as for a number of candidate specifications that are currently in progress.

The structure of the ORM is based on the Reference Model for Open Distributed Processing (RM-ODP), also identified as ISO 10746. This is a multi-dimensional approach well suited to describing complex information systems.

The ORM is a living document that is revised on a regular basis to continually and accurately reflect the ongoing work of the Consortium. We encourage respondents to this CFP to learn and understand the concepts that are presented in the ORM.

This Annex B refers to the RM-ODP approach and will provide information on some of the viewpoints, in particular the Enterprise Viewpoint, which is used here to provide the general characterization of work items in the context of the OGC Standards portfolio and standardization process, i.e. the enterprise perspective from an OGC insider.

The Information Viewpoint considers the information models and encodings that will make up the content of the services and exchanges to be extended or developed to support this Testbed. Here, we mainly refer to the OGC Standards Baseline, see section Standards Baseline.

The Computational Viewpoint is concerned with the functional decomposition of the system into a set of objects that interact at interfaces – enabling system distribution. It captures component and interface details without regard to distribution and describes an interaction framework including application objects, service support objects and infrastructure objects. The development of the computational viewpoint models is one of the first tasks of the Testbed, usually addressed at the kick-off meeting.



Figure 3. Reference Model for Open Distributed Computing

The Engineering Viewpoint is concerned with the infrastructure required to support system distribution. It focuses on the mechanisms and functions required to:

- a. support distributed interaction between objects in the system, and
- b. hides the complexities of those interactions.

It exposes the distributed nature of the system, describing the infrastructure, mechanisms and functions for object distribution, distribution transparency and constraints, bindings and interactions. The engineering viewpoint will be developed during the Testbed, usually in the form of TIEs, where Testbed participants define the communication infrastructure and assign elements from the computational viewpoint to physical machines used for demonstrating the Testbed results.

B.2.3. OGC Standards Baseline

The OGC Standards Baseline [http://www.opengeospatial.org/standards] is the complete set of member approved Abstract Specifications [http://www.opengeospatial.org/docs/as], Standards [http://www.opengeospatial.org/docs/is] including Profiles [http://www.opengeospatial.org/docs/profile] and Extensions [http://www.opengeospatial.org/docs/is], and Community Standards [http://www.opengeospatial.org/standards/community].

OGC standards are technical documents that detail interfaces or encodings. Software developers use these documents to build open interfaces and encodings into their products and services. These standards are the main "products" of the Open Geospatial Consortium and have been developed by the membership to address specific interoperability challenges. Ideally, when OGC

standards are implemented in products or online services by two different software engineers working independently, the resulting components plug and play, that is, they work together without further debugging. OGC standards and supporting documents are available to the public at no cost. OGC Web Services (OWS) are OGC standards created for use in World Wide Web applications. For this Testbed, it is emphasized that all OGC members have access to the latest versions of all standards. If not otherwise agreed with the Testbed architects, these shall be used in conjunction with - in particular - Engineering Reports resulting from previous Testbeds.

Any Schemas (xsd, xslt, etc.) that support an approved OGC standard can be found in the official OGC Schema Repository [http://schemas.opengis.net].

The OGC Testing Facility [http://cite.opengeospatial.org/teamengine/] Web page provides online executable tests for some OGC standards. The facility helps organizations to better implement service interfaces, encodings and clients that adhere to OGC standards.

B.2.4. OGC Best Practices and Discussion Papers

OGC also maintains other documents relevant to Innovation Program initiatives, including Engineering Reports [http://www.opengeospatial.org/docs/er], Best Practice Documents [http://www.opengeospatial.org/docs/bp], Discussion Papers [http://www.opengeospatial.org/docs/ discussion-papers], and White Papers [http://www.opengeospatial.org/docs/whitepapers].

B.2.5. Data

All participants are encouraged to provide data that can used to implement the various scenarios that will be developed during the Testbed. A number of Testbed sponsors will provide data, but it might be necessary to complement these with additional data sets. Please provide detailed information if you plan to contribute data to this Testbed. Additional guidelines can be found under Data Requirements.

B.2.6. Services in the Cloud

Participants are encouraged to provide data or services hosted in the cloud. There is an overarching requirement to provide cloud-hosting capabilities to allow thread participants to move services and/or data to the cloud.

B.2.7. Dependencies between Threads

Dependencies between threads have been minimized. In some rare cases one thread uses components from another thread. In this case, the thread providing components used by another thread need to provide the relevant components in time. No cross-thread dependencies exist where component developers from disjunct threads need to agree on a common interface. The interface is always defined within a single thread exclusively.

B.3. Testbed Threads

The Testbed is organized in a number of threads. Each thread combines a number of tasks that are further defined in the following chapters. The threads integrate both an architectural and a thematic view, which allows to keep related work items closely together and to remove dependencies across threads.



Figure 4. Overview of all tasks and assignment to Testbed threads

- Thread 1: Secure Data and Federated Clouds (SFC)
 - Data Centric Security
 - Federated Cloud Analytics
- Thread 2: Cloud Processing and Portrayal (CPP)
 - Earth Observation Process and Application Discovery
 - Open Portrayal Framework
- Thread 3: Machine Learning and Delta Updates (MLD)
 - Machine Learning
 - Delta Updates

B.4. Tasks

The following sections describe each task in detail.

NOTE

Please note that some documents referenced below may not have been released to the public yet. These reports require a login to the OGC portal [https://portal.opengeospatial.org]. If you don't have a login, please contact OGC staff.

B.5. Machine Learning

B.5.1. Problem Statement & Research Questions

Recent advancement in Artificial Intelligence (AI) has demonstrated the value of using Machine Learning (ML) approaches for automated image and vector data processing. However, the availability of data to efficiently train ML models is still an issue. A large variety of geospatial data are available through standardized OGC interfaces that could facilitate the discovery and access to datasets used to feed ML tools.



The Testbed-15 Machine Learning (ML) task will help to explore the utility of existing OGC Web services (OWS) to support a large scope of ML tools including EO data processing, image classification, feature extraction and vector attribution. The key research question is how these

various ML models can be integrated best within standards-based infrastructures. These infrastructures include OGC Web services that interface any kind of data repository from rather stable image archives to Big data sensor data archives or real time systems.

It is expected that this task approaches the machine learning model development in a holistic way, focusing on the models itself as well as their integration into OGC Spatial Data Infrastructures equivalently.

All machine learning models shall provide WPS interfaces, ideally of the latest WPS generation that provide OpenAPI-type resource-oriented interfaces. Though the full implementation of Web services providing the access to the results of the ML models is not expected in this Testbed, the design of standardized ML interfaces shall explore how ML tools would access and publish their input and outputs using OWS. This research may also suggest a path for the development of future OWS.

B.5.2. Aim

Develop a set of machine learning models and explore their usage within OWS Web service based environments. Though the various machine learning models are in focus of this task, their integration within OWS Web service infrastructures, usage of catalogs and dynamic binding of data and services are equally important. It is emphasized that this task is expected to approach the machine learning model usage in a holistic way.

The interface component of machine learning models development should focus on inputs and outputs based on standards including OGC standards. This testbed should demonstrate and describe how OGC Web Services available in the list of CGDI Web services [https://www.nrcan.gc.ca/earth-sciences/geomatics/canadas-spatial-data-infrastructure/10783] and elsewhere can be utilized; it should also inform how OGC standards can be used and advise whether they should be amended to better support Machine Learning and Linked Data environments. Other existing and emerging standards and open platforms from ISO, IHO, W3C, WMO, Global Biodiversity Information Facility, Arctic SDI and others should also be leveraged.

B.5.3. Previous Work

The following reports serve as a baseline for this task:

- 18-038r2 OGC Testbed-14: Machine Learning Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81977&version=1]
- 18-094r1 OGC Testbed-14: Characterization of RDF Application Profiles for Simple Linked Data Application and Complex Analytic Applications Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81923&version=1]
- 18-097 OGC Environmental Linked Features Interoperability Experiment Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81777&version=1]
- 18-022r1 OGC Testbed-14: SWIM Information Registry Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81786&version=1]

- 18-090r1 OGC Testbed-14 Federated Clouds Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81863&version=1]
- 16-059 Testbed-12 Semantic Portrayal, Registry and Mediation Engineering Report [http://docs.opengeospatial.org/per/16-059.html]
- 15-054 Testbed-11 Implementing Linked Data and Semantically Enabling OGC Services Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=64405&version=2]
- 14-049 Testbed 10 Cross Community Interoperability (CCI) Ontology Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=58974&version=2]

B.5.4. Scenarios & Requirements

This task has five different scenarios. Each scenario focuses on a different machine learning challenge and needs to be implemented as an individual demonstration scenario. Nevertheless, all five scenarios shall cooperate on the integration challenge mentioned above. It is emphasized that the fourth model, the *Richelieu Hydro Linked Data Harvest*, focuses more on the Semantic Web link building and triplet generation aspects, whereas the fifths model, *Arctic web services discovery*, focuses on discovery and cataloguing aspects.

Each model shall be integrated with different types of OGC Web service instances for training and execution. Partly, these instances are provided by the sponsor, partly the participants shall access other resources available online. One option that shall be explored in Testbed-15 is feeding services provided by the CGDI [https://www.nrcan.gc.ca/earth-sciences/geomatics/canadas-spatial-data-infrastructure/10783] as potential data sources to all machine learning models to fill data gaps and to reduce the need for new data acquisition that will lead to more application development eventually. For sure, all participants are invited to setup additional Web service instances.

A. Petawawa Super Site research forest change prediction ML model

Cloud cover is an unavoidable presence in remotely sensed data. Estimates of the mean annual global cloud cover over land surfaces ranges from 35% (Ju and Roy, 2008) to 66%(Zhang et al, 2004). Clouds affect remotely sensed data dramatically.

As a first step towards an automated forest change prediction system, participants will demonstrate the use of Machine Learning to remove clouds and high altitude cloudets (popcorn clouds) from historical datasets from the Petawawa super site [http://www.forestresearch.ca/ index.php?option=com_content&view=article&id=272&Itemid=83].



Figure 5. Petawawa Research Forest; located along Highway 17, just east of Chalk River, Ontario, Canada

Existing products and algorithms have helped to identify and segment clouds on Landsat Level-1 data products. These algorithms utilize a multi-pass algorithm that use decision trees to label pixels in scenes, validate and discard and create masks for shadows. These algorithms work well on certain clouds, cloud shadows, snow and ice pixels. However the identification of high altitude cloudets over bright objects such as snow, ice or lakes are unsatisfactory. By introducing expert knowledge and supervised classification along with a decision tree algorithm into a ML methodology, we hope to see improvements in the cloud detection process.

Testbed-15 shall deliver a component of a forest change detection ML model for the Petawawa Super Site research forest through a live demo and document such model in an OGC Engineering Report.

The ML model work shall:

- 1. Discover datasets of the area (Petawawa) over a given time frame from OGC Catalogue (CSW)
- 2. Discover which images are usable (less than 70% cloud coverage from Landsat and possibly Sentinel 2 products)
- 3. Use machine learning and convolutional neural networks to detect features, segment cloud and shadow pixels based on training data provided (surface reflectance and expert knowledge); ideally using TensorFlow
- 4. Create of a cloud free mosaic within a given timeframe by assembling best non-cloud segments
- 5. Produce a Land cover classification from cloud-free mosaic based on ML using provided land classification training data plus DTM/DSM and LiDAR data where applicable. Study how optical imagery (RGB or multi-spectral) can be combined with height above ground information (derived from LiDAR) to feed ML Land Cover classification
- 6. Continue the work of cloud computing from Testbed 13, 14 with work on WPS and possibly CWL (Common Workflow Language), and continue the LiDAR best practices work from Testbed 14.
- 7. Be compatible with NRCan's Boreal Cloud (OpenStack cloud environment). Access to the Boreal Cloud will be negotiated.
- 8. Ideally uses open source libraries such as TensorFlow, PyTorch or others for Machine

Learning

9. Should put special emphasis on best practices for distributing and managing big data (LIDAR and imagery), and processing datasets using machine learning algorithms for auto classification and with future use in change detection

The following data sets will be made available:

- The Canadian Forest Service has collected and curated a mass volume of data surrounding the Petawawa Research Forest. Data from the late 70's to present is being made available via OGC Web services. This includes historical Landsat (Collection 1 Level-1) scenes from Landsat 5, 7, 8 from 1983 to 2018, Sentinel 1, Sentinel 2, LiDAR, multispectral data and aerial photography missions.
- A set of classified cloud image segments based on surface reflectance CFMask products (Zhu and Woodcock algorithm) and expert knowledge for high altitude (popcorn) cloudets can be used for training data. This dataset will be published through a WFS and WMS service.
- A set of manual classified Land cover products with associated raw landsat 8 datasets for training data for land classification use. This dataset will be published as a WFS and WMS
- The Canadian Geospatial Data Infrastructure CGDI [https://www.nrcan.gc.ca/earth-sciences/ geomatics/canadas-spatial-data-infrastructure/10783] list of Web services may contain other data sources
- NRCan forestry data can be accessed via OGC services such as Catalogue services, Web Feature, Web Mapping and Web Coverage service: https://saforah2.nfis.org/index.html



The following figure illustrates the work items and deliverables of this scenario.

Figure 6. Petawawa scenario work items

The scenario includes two WPS instances. The first one interfaces the selection, mosaicing, and cloud detection model; the second one the forest classification model. Both services shall be complemented with a simple client application that allows executing the services. Each service needs to connect to various OWS instances at the backend.

All results will be captured in the Machine Learning Engineering Report.

B. New Brunswick forest supply management decision maker ML model

Testbed-15 shall deliver a forest supply management decision maker ML model for the province of New Brunswick forested areas through a live demo and document such model in an OGC Engineering Report.



Figure 7. New Brunswick (red), one of four Atlantic provinces on the east coast of Canada (source: wikipedia)

The ML model work shall:

- 1. Recommend the most efficient optimized path from forest to market -"wood flow model"; this model addresses optimal allocation of resources and flow of products to secondary and tertiary processing.
- 2. Recommend new road construction that will be the most efficient over time and safety being considered. Work may include life-cycle analysis for bridge/culvert infrastructure.
- 3. Recommend the best time for road conservation closure that will cause the least negative impact
- 4. Utilize spatial data such as secondary infrastructure (secondary roads, logging roads, bridges, wood lots, sort yards), primary infrastructure (highways, bridges, dams, power, mills, paper mills) and current market prices of lumber and fuel/energy
- 5. Explore the use of existing OGC, ISO and other Best practices and standards such as WPS, WCS, WFS, WMS, along with upcoming standards to handle Big Data dissemination
- 6. Continue the work of cloud computing from Testbed 13, 14 with work on WPS and possibly CWL (Common Workflow Language), and continue the LiDAR best practices work from Testbed 14.

- 7. Shall take results from OGC Testbed-14 Point Cloud Data Handling Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81608&version=1] into account
- 8. Be compatible with NRCan's Boreal Cloud (OpenStack cloud environment). Access to the Boreal Cloud will be negotiated.
- 9. Ideally uses open source libraries such as TensorFlow, PyTorch, Hadoop, or others for Machine Learning

The following data sets will be made available:

- Enhanced forest inventory (imagery, LIDAR, Forest Resources Information), terrain models (slope, elevation, visual quality) available from the province of New Brunswick
- NRCan and provincial data access via http://nb.nfis.org/ and other CGDI [https://www.nrcan.gc.ca/earth-sciences/geomatics/canadas-spatial-data-infrastructure/10783] sources
- The list of Web services provided by the CGDI [https://www.nrcan.gc.ca/earth-sciences/geomatics/ canadas-spatial-data-infrastructure/19359] may contain other data sources

The following figure illustrates the work items and deliverables of this scenario.



Figure 8. New Brunswick scenario work items

The scenario includes a single WPS instance that serves as an interface to the forest management model. The service/model needs to connect to various OWS instances at the backend. The service provider shall provide a simple client application that allows executing the service.

All results will be captured in the Machine Learning Engineering Report.

C. Quebec lake - river differentiation ML model

Image and LiDAR analysis techniques are often used to classify an image into two classes, water and land. Vectorization and correction of the resulting classification can produce a nice cartographic vector dataset but will not determine where the boundary exists that separates a lake from a large river entering or exiting it, or where the respective boundaries are at the confluence of two large rivers. An automated means of creating lake and river features from general water boundaries is required.



Figure 9. Manicouagan Reservoir (also Lake Manicouagan) is an annular lake in central Quebec, Canada, covering an area of 1,942 km2 (750 sq mi). Landsat-8 imagery courtesy of NASA and USGS, edited by Pierre Markuse

Testbed-15 shall deliver an ML model to delineate lake and river features from an undifferentiated waterbody vector dataset. The ML model will be applied against undifferentiated hydrographic network data from Québec (géobase du réseau hydrographique du Québec (GRHQ), with adjacent lakes and double-line rivers merged together to form general waterbody features), with the intent of generating an GRHQ equivalent with the lake and double-line rivers differentiated as distinct features. Supporting data includes DEM data and place names with locations. The work should proceed with the data just described, and then separately considering the National Hydro Network (NHN) [https://open.canada.ca/data/en/dataset/a4b190fe-e090-4e6d-881e-b87956c07977] for the same area.

The ML model work shall:

- 1. Recommend if a given waterbody needs to be split into lake/river features.
 - a. If a split does not need to be made, then determine whether the feature is a lake or a river.
 - b. If a split should be made, recommend where the division(s) between lake and river features should be placed based on terrain topography provided through DEM or LiDAR data.
- 2. Evaluate the confidence level of each recommendation.
- 3. Apply the recommendations, defining the feature type and the boundary of each waterbody feature.
- 4. The final geospatial results will be provided in vector format; thus, if raster processing is used, conversion to vector format must be carried out.
 - a. No gaps or overlaps between adjacent features are allowed. This applies to lakes and double-line rivers, as well as to these features and the adjacent land.
 - b. Not all lakes or rivers will have place names, and for those that do, the place names do not necessarily fall on the feature.

c. The data at its different stages should be displayed using MapML served by WFS 3.0.

The following data sets will be made available:

- GRHQ hydrographic network data for Québec, with the lakes and double-line rivers already differentiated (as the expected result)
- GRHQ hydrographic network data for Québec, with the lakes and double-line rivers not differentiated (as the dataset to be differentiated)
- National Hydrographic Network (NHN) for Québec, with the lakes and double-line rivers already differentiated
- Gridded elevation dataset (DEM GeoBase CDEM) for Québec
- Gridded elevation dataset (DEM from lidar) for Québec
- Place names with associated point locations for Québec
- The list of Web services provided by the CGDI [https://www.nrcan.gc.ca/earth-sciences/geomatics/ canadas-spatial-data-infrastructure/19359] or the Open Maps site may contain other data sources

All datasets will be made available for a given area within Québec, with the specific area to be determined. The size of this area will be determined by the participants in consultation with the sponsor at the kick-off meeting. The input vector dataset (except DEM) are vector products. They may or may not require rasterization for use in a ML environment; this is at the discretion of the participants.



The following figure illustrates the work items and deliverables of this scenario.

Figure 10. Quebec scenario work items

The scenario includes a single WPS instance that serves as an interface to the vectorization model. The service/model needs to connect to various OWS instances at the backend. The WFS instance shall serve the results of the vectorization process as MapML. A dedicated client application shall interact with both service instances and display the MapML data.

All results will be captured in the Machine Learning Engineering Report.

D. Richelieu River hydro linked data harvest model

To understand important aspects of the water cycle, it is often vital to relate many kinds of hydro features, from the subsurface to the surface and atmosphere. This involves hydro data that are

stored in distributed databases and are relatively isolated, making it difficult to capture relations between features and keep them current.

By creating Linked Open hydro data, we anticipate a series of business benefits, including improved discoverability, access, and connectivity, as well as enhanced interoperability and openness. Furthermore, to understand water related socio-economic issues, it is often necessary to relate various statistical data with various water feature geospatial data.



Figure 11. Richelieu River near Saint-Marc-sur-Richelieu (image by Jean-Philippe Boulet, map by Kmusser, both wikipedia)

Testbed-15 shall harvest hydrological relations from the web for the Richelieu River / Watershed area. The harvesting might shall use different types of data sources. Using artificial intelligence techniques (e.g. machine learning), **the model shall**:

- 1. Semantic Web: mine existing Semantic Web infrastructures for relevant relations, e.g. DBPedia, GeoNames, or others as determined.
- 2. Online geospatial data: search existing geospatial databases, geospatial web resources (e.g. MapML) or hydrological models for relevant information.
- 3. The Web: identify relevant relations on the web and/or in relevant publications.

The following data sets will be made available:

- Ontologies (schemas) for types of features and relations to be harvested.
- Specification of format for results (as triples).
- Specification of URL identifier pattern to be used in triples.
- Example file containing existing relations in the specified format.

The following figure illustrates the work items and deliverables of this scenario.



Figure 12. Richelieu scenario work items

The scenario includes a single WPS instance that serves as an interface to the triple generator and delivers triples in a format to be determined. The service/model needs to connect to various backends. The service provider shall provide a simple client application that allows executing the service.

The link building shall be restricted to the specified region (Richelieu) and types of features and relations as provided by the sponsor. The use of open source software where possible, especially for any code to be transferred, is highly encouraged. The work shall respect the Spatial Data on the Web Best Practices [https://w3c.github.io/sdw/bp/] in its latest version.

All results related to Machine Learning will be captured in the Machine Learning Engineering Report. All results related to Semantics will be captured in the Semantic Engineering Report.

E. Arctic web services discovery ML model

There are thousands of geospatial services available in .ca domain. How many of such services represent a specific type of forest and its health condition in a specific location? Which service provide data on Richelieu River water features? What is the relationship of these features with others? Which services provide lake and river maps over the province of Quebec? Do they differentiate lakes and rivers? What is best of approach to leverage Canada's investment in these Web services?



Figure 13. Arctic coast

Testbed-15 shall deliver a component capable to build an evergreen catalogue of relevant arctic circumpolar Web services. The goal is to develop a **machine learning model that will:**

1. Discover OGC and ESRI REST Web services that have some relevance to circumpolar science.

- 2. Evaluate the confidence level of each recommended service using both metadata and data parameters.
- 3. Include the review of multiple datasets within services in this evaluation.
- 4. Make results available at a OGC Catalogue Service implementation (Catalogue Service for the Web v2.0.2)
- 5. Ideally uses open source libraries such as TensorFlow, PyTorch, Hadoop, or others for Machine Learning
- 6. Be compatible with other OGC Testbed 15 cataloguing requirement and solutions (see Discovery of EO Processing Services and Applications).

The following data sets will be made available:

- A training list of services deemed relevant to circumpolar arctic science
- A boundary file that can serve as a definition of the Arctic

The following figure illustrates the work items and deliverables of this scenario.



Figure 14. Arctic scenario work items

The scenario includes a single CSW instance that serves as an interface to the data produced by the Quality of Service and Confidence model in the background. The service/model needs to connect to various backends. The service provider shall provide a simple client application that allows executing the service.

All results related to Machine Learning will be captured in the Machine Learning Engineering Report. All results related to cataloguing and discovery will be captured in the Catalog and Discovery Engineering Report.

B.5.5. Work Items & Deliverables

All work items have been illustrated above. The following list identifies all deliverables that are part of this task. Detailed requirements are stated above. All participants are required to participate in all technical discussions. Thread assignment and funding status are defined in section Summary of Testbed Deliverables.

Engineering Reports

- **D001 Semantic ER** Engineering Report capturing all results and experiences that address semantics from this task.
- **D002 Machine Learning ER** Engineering Report capturing all results and experiences that address machine learning from this task.
- **D010 Catalog ER** Engineering Report capturing all results and experiences that address catalogs and discovery from this task. The home of this work item is task Discovery of EO Processing Services and Applications.

Components

- D100 Petawawa cloud mosaicing ML model Machine Learning model with WPS interface
- **D101 Petawawa land cover classification ML model** Machine Learning model with WPS interface
- D102 New Brunswick forest ML model Machine Learning model with WPS interface
- **D103 Semantic Web link builder and triple generator** Machine Learning model with WPS interface
- **D104 Quebec river-lake vectorization ML mode** Machine Learning model with WPS interface
- D105 Quebec model MapML WFS3.0 WFS 3.0 instance serving MapML data
- D106 Quebec model MapML WFS3.0 Client Client instance interacting with D105
- D107 Arctic Discovery Catalog Machine Learning model with CSW interface

B.6. Delta Updates

B.6.1. Problem Statement & Research Questions

Dissemination of GeoINT data in a Denied, Degraded, Intermittent and Limited (DDIL) Bandwidth environment is a challenging problem. Providing prioritised delta updates can be a good approach to resolve this problem.

A delta update is defined as an update that only requires the system to download the new changes and not the whole database. Depending on DDIL conditions (or other constraints), the server might decide to further reduce the amount of update data. Depending on the applied classification schema, the server may only send top priority updates in bad DDIL conditions and further updates once conditions improve. This is important as it significantly saves time and bandwidth.



Emerging OGC interface specifications that use OpenAPI provide the ability for OGC standards to incorporate delta updates utilizing the Feature ID and Timestamp elements of OpenAPI. This allows for raster or vector data to be updated while still reflecting the changes in the original version.

The key research question is how to implement reliable and secure delta update mechanisms with OGC next generation Web Services such as WFS 3.0 and WPS 3.0. The research shall include different mechanisms that either require enhancements to existing WFS 3.0 instances, or use to be developed WPS 3.0 instances to realize similar functionality without touching existing data access services.

B.6.2. Aim

Describe and demonstrate how OGC standards can be used to provide delta updates to data using services in a Denied, Degraded, Intermitted or Limited Bandwidth (DDIL) environment.

B.6.3. Previous Work

The following reports serve as a baseline for this task:

- OGC Geo-Synchronization Service [https://portal.opengeospatial.org/files/?artifact_id=76559& version=1]
- OGC Testbed-13: Disconnected Networks Engineering Report [http://docs.opengeospatial.org/ per/17-026r1.html]
- Testbed-12 Web Feature Service Synchronization [http://docs.opengeospatial.org/per/16-044.html]
- OGC Testbed-13: Asynchronous Services ER [http://docs.opengeospatial.org/per/17-028.html]
- Testbed-12 Implementing Asynchronous Services Response Engineering Report [http://docs.opengeospatial.org/per/16-023r3.html]
- Testbed-11 WFS-T Information Exchange Architecture [https://portal.opengeospatial.org/files/? artifact_id=66905&version=2]

B.6.4. Scenario & Requirements

The diagram below shows different scenarios that shall be all considered in Testbed-15. In the first scenario (interactions (1) and (2)), the Client node requests an initial set of data directly at the Data Server. Two alternatives shall be considered: First, the server stores references to all initially synchronized features and provides a set identifier (SetID) that can be used by the client in future calls. This requires the server to handle some form of state to manage the sets. The client then requests updates to these sets: "Provide me all top priority updates for SetID-1 since 2018-11-26T12:00:00".

Alternatively, the data server does not provide any set identifiers. Then, client nodes need to use feature identifiers in future calls to identify interesting features, e.g. "Provide me all top priority updates for the following features since 2018-11-26T12:00:00; featureID-1, featureID-2,...".

To leave existing data servers untouched, an alternative scenario two exists that makes use of a WPS instance to handle the synchronization between the client and the server node (steps 3-5). Here, the client identifies the initial synchronization set and the WPS instance handles all delta updates.

Testbed-15 shall research the different options as illustrated in the two scenarios. Further on, it shall address interrupted communication situations where synchronization messages are interrupted or lost. Clients as well as servers shall be able to handle prioritization of delta updates. Based on some sort of classification system, delta updates will be prioritised by the client and or the server. The client may indicate bad communication ahead and requests only top priority data, or the server has knowledge about the client's context and reduces delta updates autonomously. Once DDIL conditions improve, delta updates shall ensure fully synchronized client and server nodes again.

Prioritization can depend on different classification schemas. Either data is already tagged as e.g. normal, priority, and top priority; or depends on the client's context. An example of context could be the client is an autonomous vehicle and it requires the latest transport updates relevant to its route.



Figure 15. Delta updates scenario

Testbed-15 shall implement the scenarios using a WFS 3.0 and a WPS 3.0. The implementation

- shall focus on both the server and client side;
- is ideally based on Geoserver and delivered as Open Source code
- shall be delivered as a Docker or VM container including installation scripts and instructions
- shall be documented in an Engineering Report that further captures all research results and lessons learned from the implementations.

B.6.5. Work Items & Deliverables

The following figure illustrates the work items and deliverables of this task.



Figure 16. Delta updates work items

The following list identifies all deliverables that are part of this task. Detailed requirements are stated above. All participants are required to participate in all technical discussions. Thread assignment and funding status are defined in section Summary of Testbed Deliverables.

• D005 Delta Updates ER - Engineering Report capturing all results and experiences from this

task.

- D111 Delta Updates Client Client implementation that supports delta updates
- **D112 Delta Updates WPS** WPS 3.0 implementation with support for delta updates
- D113 Delta Updates WFS WFS 3.0 implementation with support for delta updates

B.7. Data Centric Security

Data-centric security is an approach to security that emphasizes the security of the data itself rather than the security of networks, servers, or applications. Data-centric security is evolving rapidly as enterprises increasingly rely on digital information to run their business and big data projects become mainstream. (Wikipedia)



With more and more data stored in the cloud and made accessible and dynamically exchanged by cloud services, data centric security approaches become more relevant in the context of standards based geospatial applications as well.

Security of Open Geospatial Consortium (OGC) Standards could be defined in three levels:

- Service Level
- Layer Security (within a service context)
- Feature Level

OGC has successfully implemented Service Level and Layer Security through the OGC Web Services Security Standard [https://portal.opengeospatial.org/files/77693]. OGC Testbed-14 Security Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81774&version=1] has further developed security thinking in the OGC.

Security at the Feature Level can be explained by the concept of Data Centric Security as defined by the definition above. Data Centric Security implements a data classification scheme combined with a set of rules which are applied to individual data objects. In this instance, security and business requirements lead to an overall business data classification (BDC), which represents the labels or attributes of data that are used to determine the data classes. Data will also be classified by criteria such as ownership, origin time and location. The data classification together with policy rules are encoded into data control rules (DCRs). These DCRs represent the unified data handling policies expressed in terms of BDC. They are used to establish appropriate access policies and practices to support data handling policies.

In summary, it is a combination of capabilities that allow the data host to inherently know what data is appropriate for each user and enforce security controls. Data can be labeled with degree of sensitivity (or level). For example, in government and defense applications, data might be labeled unclassified, secret, or top secret. Among other aspects, access controls are enforced by comparing a data classification label with a user's access clearance. Access clearance can be thought of as an extension to standard database privileges and roles.

B.7.1. Problem Statement & Research Questions

In Testbed-15, it is of interest how security works at a Feature Level as well as its implications as a data burden on the network. With focus on the actual interactions and general workflows, Testbed-15 shall answer the question how data centric security can be applied to OGC standards based architectures:

- How does data centric security works with OGC standards and best practices?
- Which elements are already supported and how?
- Which modifications to existing standards or best practices are necessary to exploit full potential of data centric security?

The Testbed-15 work on data centric security is considered being a first step towards fully exploited data centric security in standardized geospatial data infrastructures. This work in Testbed-15 shall lay the foundation for those future activities by clearly describing the challenges and opportunities of data centric security and provide a critical inventory of existing OGC specifications and approaches.

Testbed-15 shall capture and discuss security aspects of the full data lifecycle as illustrated below.



Figure 17. Security lifecycle (slide by Rich Mogull [https://securosis.com/assets/library/presentations/ DataSecurityLifecycle-PDF.pdf])

One important aspect is how existing OGC solutions, such as GeoXACML and security framework elements such as context handlers, policy administration points, or policy information points can be mapped to data centric security terminology. For example, data centric security terms such as Identity and Access Management (IdAM), Key Management, or Object Release are used in the scenario description below.

B.7.2. Aim

Implement a Data Centric Security approach using existing OGC standards and emerging OGC standards (such as WFS 3.0 or WPS 3.0 that make use of OpenAPI) as being documented in Testbed-14. Demonstrate how data centric security can be implemented for a scenario described below. Provide a detailed discussion that can be considered as foundational work for future OGC IP activities that discusses the research questions posted above.

B.7.3. Previous Work

The following reports serve as a baseline for this task:

- OGC Web Services Security Standard [https://portal.opengeospatial.org/files/77693]
- OGC Testbed-14 Secure Client Test Engineering Report [https://portal.opengeospatial.org/files/? artifact_id=81797&version=1]
- OGC Testbed-14 Security Engineering Report [https://portal.opengeospatial.org/files/?

artifact_id=81774&version=1]

- 13-101 OGC Geospatial eXensible Access Control Markup Language (GeoXACML) 3.0 GML
 3.2.1 Encoding Extension, https://portal.opengeospatial.org/files/?artifact_id=55232& version=2
- 18-090r1 OGC Testbed-14 Federated Clouds Engineering Report, https://portal.opengeospatial.org/files/?artifact_id=81863&version=1
- 16-048r1 Testbed-12 OWS Common Security Extension ER, http://docs.opengeospatial.org/ per/16-048r1.html
- 15-022 Testbed-11 Engineering Report: Implementing Common Security Across the OGC Suite of Service Standards, https://portal.opengeospatial.org/files/?artifact_id=63312& version=2
- 15-010r4 Testbed-11 WFS-T Information Exchange Architecture, https://portal.opengeospatial.org/files/?artifact_id=66905&version=2
- 12-118 OWS-9 Security ER, https://portal.opengeospatial.org/files/?artifact_id=51806& version=1

B.7.4. Scenario & Requirements

Data Centric Security in the Testbed-15 scenario has separated the feature and the metadata as separate entities, both of which shall be separately encrypted with different encryption keys. The scenario in the diagram below shows a number of different clients. All are viewing the same map, but due to the sensitivity of the data, each client can only see the features on the map that the client is authorized to see. This is reducing complexity of the map, while protecting the sensitivity of the different data features. The general interaction between client and cloud based server is illustrated in the figure below.



Figure 18. Data centric security scenario.

Testbed-15 shall develop an implementation of WFS 3.0 with support for data centric security. This implementation is to learn from recent OGC Security work in OGC Testbed - 14: Security Engineering Report (18-026) and OGC Web Services Security Standard (17-007r1). Where necessary also reference previous Security work in other OGC Innovation Initiatives. This work will also reference STANAG 4774 – Metadata Confidentiality Label Syntax and STANAG 4778 – Metadata Binding Mechanism. Geospatial eXtensible Access Control Markup Language should be referenced.

Testbed-15 shall implement the scenario using a WFS 3.0 with tagged data and metadata, both encrypted with different encryption keys. The implementation

- shall focus on both the server and client side;
- shall consider implications of data centric security in feature and feature metadata implementations as an increased data burden on the network;
- is ideally based on Geoserver and delivered as Open Source code
- shall be delivered as a Docker or VM container including installation scripts and instructions
- shall support encryption of feature and metadata with different keys and shall allow to understand the implication of this approach as a data burden on the network. *shall be documented in an Engineering Report that further captures all research results and lessons learned from the implementations.

B.7.5. Work Items & Deliverables



The following figure illustrates the work items and deliverables of this task.

Figure 19. Data centric security work items

The following list identifies all deliverables that are part of this task. Detailed requirements are stated above. All participants are required to participate in all technical discussions. Thread assignment and funding status are defined in section Summary of Testbed Deliverables.

- **D004 Data Centric Security ER** Engineering Report capturing all results and experiences from this task.
- D114 Data Centric Security Client Client implementation that supports data centric security
- **D115 Data Centric Security WFS** WFS 3.0 implementation with support for data centric security.

B.8. Federated Cloud Analytics

As discussed in the OGC Testbed-14 Federated Clouds Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81863&version=1], the advent of the cloud computing era has fundamentally changed how people and organizations view computing — and more specifically how people and organizations interact with the resources that they care about, i.e., data and services. All computing resources, including clouds, exist in some type of administrative domain wherein access management can be done. As long as resources are all in the same administrative domain, managing access is straight-forward.

However, with the continued development of our interconnected world, it is becoming increasingly common that data and services desired by a user exist across different administrative domains. Organizations — each with their own administrative domain — that need to collaborate will also need a way to securely bridge those domains to selectively share data and services to achieve their joint organizational goals. Easily accessing resources distributed across different administrative domains is a challenge. The naive approach is for an individual to maintain n different accounts and credentials for n different organizations. A more effective
Simply put, a federation enables a set of participating organizations to selectively share data and resources for specific purposes. The goal is to make federated environments as seamless, transparent, and easy to use as a single centralized environment. More precisely, a federation is a security and collaboration context wherein participants can define, agree up on, and enforce joint resource discovery and access policies. A federation is not necessarily owned by any one organization, or located at any one site. In the current cloud computing era, many federations will be cloud-based. However, since federations can be managed at any level in the system stack, it is certainly possible to address federation at the Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS) levels. To manage SaaS-level federations means to manage arbitrary, application-level services. That is to say, federations can be used to manage collaborations for arbitrary business functions.

B.8.1. Problem Statement & Research Questions

Previous OGC Testbeds have addressed a number of issues related to supporting analytic workflows where the data and analytics are hosted or deployed in an ad-hoc manner on multiple heterogeneous clouds that belong to different administrative domains. It is time for the OGC to assess the sufficiency of that body of work and identify areas were additional work is needed. This assessment should be performed through a proof of concept executing a non-trivial analytic mission leveraging data and analytics hosted on two or more clouds.

Of particular interest in this context are three aspects. Firstly, the handling of security in federations. Second, how the Testbed-13 and Testbed-14 research results of "bringing applications to the data" relate to SCALE [https://github.com/ngageoint/scale] and SEED [https://github.com/ngageoint/seed]. SCALE is an open source system that provides management and scheduling of automated processing on a cluster of machines. It uses the SEED specification to aid in the discovery and consumption of processes packaged in a Docker containers. Third, the role of blockchain and distributed ledger technologies in the context of handling provenance in federations.



This task is organized in four separate sub-tasks. The following research questions shall be addressed:

- Federated Security: Can the NIST/IEEE Federated Cloud architecture be validated (or invalidated) in a typical federated clouds analytics scenario that includes separate cloud environments? How can the Mediation Server concept developed in Testbed-14 be further enhanced to a fully functional Federation Manager in the sense of NIST/IEEE? What are the advantages and disadvantages, and how does this extended functionality fit within the OGC family of standards?
- Federated Cloud Analytics: How to bring SCALE and SEED into the family of cloud architectures supported by OGC standards? What role does WPS play? What catalog solutions work best?
- **EOC, SCALE, and SEED**: How to handle the different approaches for cloud processing? Where are harmonization opportunities, what needs to remain separate?
- Federated Clouds Provenance: How can Blockchain and distributed ledger technologies be used to protect the integrity of different types of provenance data?

B.8.2. Aim

The aim of this task is to understand how emerging technologies and architectures featuring distributed cloud environments can be successfully integrated with OGC standards to support non-trivial analytic missions.

B.8.3. Previous Work

The following reports serve as a baseline for this task:

- 18-090r1- OGC Testbed-14 Federated Clouds Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81863&version=1]
- 18-026r1 OGC Testbed-14: Security Engineering Report [https://portal.opengeospatial.org/files/? artifact_id=81774&version=1]
- 18-030 OGC Testbed-14: Secure Client Test ER [https://portal.opengeospatial.org/files/? artifact_id=81834&version=1]
- 18-085 OGC Testbed-14: BPMN Workflow Engineering Report [https://portal.opengeospatial.org/ files/?artifact_id=81593&version=1]
- 18-050r1 OGC Testbed-14: ADES & EMS Results and Best Practices Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81794&version=1]
- 18-049r1 OGC Testbed-14: Application Package ER [https://portal.opengeospatial.org/files/? artifact_id=81790&version=1]

B.8.4. Scenario & Requirements

This task has four different sub-tasks. Each sub-task focuses on a different federated cloud challenge and needs to be implemented as an individual demonstration scenario. The list of sub-tasks includes:

- Federated Security,
- Federated Cloud Analytics,
- EOC, SCALE, and SEED, and
- Federated Clouds Provenance.

Federated Security

The U.S. National Institute of Standards and Technology (NIST) and IEEE are developing a Federated Cloud architecture [https://www.nist.gov/publications/nist-cloud-computing-referencearchitecture]. The Testbed-14 Federated Clouds Engineering Report [https://portal.opengeospatial.org/ files/?artifact_id=81863&version=1] illustrates and describes what federation actually entails as follows (see Figure 1): "In a typical, stand-alone administrative domain, an Identity Provider (IdP) issues identity credentials to a User. When that User requests service from a Service Provider (SP), that SP validates the User's credentials with the IdP. The User's request is either honored or declined based on the User's valid credentials."



Figure 20. The Essence of Federation (source: OGC18-090r1 [https://portal.opengeospatial.org/files/? artifact_id=81863&version=1])

A *federated environment* essentially crosses the boundary between administrative domains, e.g., A and B. Here a UserA must be able to discover (find) any useful services that SPB may wish to make available to a federation. When UserA invokes SPB, SPB must have some way of validating UserA's credentials and making a valid access decision. Supporting just these key functions has a large design space when one considers the issues of deployment models, governance models, and simply scale.

In Testbed-14 the OGC made progress in implementing parts of that architecture. In particular, they demonstrated the ability to enable access to Web Services where each service implements a different authentication technique. The infrastructure required to enable this capability is referred to as a Federation Manager in the Federated Cloud Architecture. OGC18-090r1 [https://portal.opengeospatial.org/files/?artifact_id=81863&version=1] defines a *Federation Manager* as the entity that manages different, multiple Federation Instances, or simply federations, among two or more Administrative Domains. Multiple Federation Managers may cooperate to form larger,

distributed federation infrastructures. The Testbed-14 implementation of the Mediation Service functioned somewhat like a centralized Federation Manager. However, it still had some aspects of a distributed implementation since there is a Mediation Service in every Security Environment. Further on, the aspect of resource discovery using the Federation Manager was not addressed in OGC Testbed-14 Federated Clouds anv detail. 18-090r1-Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81863&version=1] in conjunction with 18-026r1 - OGC Testbed-14: Security Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81774& version=1] provides a detailed discussion of achieved and ignored elements in Testbed-14 that need further research in future initiatives.

Testbed-15 will take the next step, federating Federation Managers. At least two Federation Managers will be deployed. All components in the Federated Cloud Analytics sub-task will participate in one or more Federations and will exercise the exchange of authentication data between the Federation Managers and validate (or invalidate) the NIST/IEEE architecture. Participants are invited to propose the most important NIST/IEEE architecture aspects that shall be addressed in Testbed-15.

Though services from other sub-tasks will be available, participants in this sub-task shall be prepared to setup simple test servers and clients to fully exercise the Federation and needed capabilities. At least one of the Federation Managers should support TLS Mutual Authentication using X.509 PKI certificates.

The following tasks shall be executed:

- 1. Deploy at least two Federation Managers
- 2. For each deployment model described in the NIST Federated Cloud Reference Architecture
 - a. Configure the Federation Managers in accordance with the deployment model
 - b. Clearly define and demonstrate how federated identity can be consistently managed and used.
 - c. Clearly define and demonstrate how the scope of attributes and authorizations can be used to consistently manage federated environments.
 - d. Clearly define and demonstrate how resource discovery and access can be consistently managed across all participating administrative domains

The following figure illustrates the work items and deliverables of this sub-task.



Figure 21. Work items and deliverables of the Federated Clouds Security sub-task (green and red). Grey items are part of the connected Federated Clouds Analytics sub-task.

The sub-task includes two *Federation Managers*. At least *Federation Manager #1* shall support TLS mutual authentication using X.509 PKI certificates. Participants are invited to define the appropriate interface. Federation Manager providers are required to provide additional test servers as necessary, as it cannot be assumed that the *Federated Clouds Analytics* sub-task provides all necessary functionality to the fully exercise the tests of this sub-task. Nevertheless, components from the *Federated Cloud Analytics* sub-task should be part of the exercise.

All results will be captured in the *Federated Clouds Security Engineering Report*.

Federated Cloud Analytics

The NGA Datacenter platform turns hundreds of thousands of processes. This platform is built on SCALE [https://github.com/ngageoint/scale], SEED [https://github.com/ngageoint/seed], and Docker. SCALE is an open source system that provides management and scheduling of automated processing on a cluster of machines. It allows users to define jobs, which can be any type of script or algorithm. These jobs run on ingested source data and produce product files. The produced products can be disseminated to appropriate users and/or used to evaluate the producing algorithm in terms of performance and accuracy. SCALE uses the SEED specification to aid in the discovery and consumption of processes packaged in Docker containers. SEED is a general standard to aid in the discovery and consumption of a discrete unit of work contained within a Docker image. The SEED specification is available online [https://ngageoint.github.io/seed/seed.html].

This task will bring SCALE into the family of cloud architectures supported by OGC standards. It is built around a proof of concept executing a non-trivial analytic mission with an emphasis on moving analytics to the data. The workflow will also explore the use of Analysis Ready Data principles and technologies including:

- Cloud Optimized GeoTIFF (COG)
- SpatioTemporal Asset Catalog (STAC)
- HTTP Range Requests (RFC 7233)

The following tasks shall be executed:

- 1. Deploy a Scale Datacenter Environment (installation guidelines [https://ngageoint.github.io/ scale/quickstart.html]), ideally version 6 or higher
 - a. Investigate how to use WPS 2.0 as a standard API for Cloud analytics (Scale provides a RESTful HTTP interface for its own web UI and for any external applications that would like to connect to Scale, see API documentation [https://ngageoint.github.io/scale/docs/rest/ index.html]).
 - b. Investigate and test the use of OGC services, such as WMS, WFS and WCS, to serve as data sources and sinks for a WPS analytic.
 - c. Evaluate STAC vs. OpenSearch, vs. CSW as a cloud-resident catalog. Provide recommendations on which should be used, where, and for what resource types.
 - d. Investigate and test WPS 2.0 as a workflow automation service for managing job execution involving multiple containers; and provide a mechanism to retrieve workflow results as an integral function in the Scale Datacenter environment
 - e. Provide further recommendations to advance the architecture, integration and implementation strategies for use of OGC standards.

Sponsor Supplied Resources

The purpose of the Federated Analytic Cloud thread is to mature the platform, not to develop analytics. Therefore, Government Furnished Information (GFI) will be made available to the Participants. The GFI will include:

- 1. One or more workflow definitions to be executed on the Scale Datacenter Environment
- 2. Analytic software for use in executing a workflow
- 3. Input data for use in exercising a workflow
- 4. Sample output products for evaluating the success of a workflow

GFI will be identified prior to the kick-off meeting. The workflows selected will serve to identify specific deliverables making up the Scale Datacenter Environment.

The following figure illustrates the work items and deliverables of this sub-task.



Figure 22. Work items and deliverables of the Federated Clouds Analytics sub-task (green and red). Grey items are part of the connected Federated Clouds Security sub-task.

This sub-task includes a WPS Analytics as interface to the SCALE datacenter environment. The participant developing the WPS shall provide two identical instances that are deployed in two different security environments. Both, WPS and the SCALE datacenter environment can connect to a number of OGC data services such as WMS, WCS, or WFS. These services are bundled as a single deliverable. These services serve as data sources for SCALE datacenter applications, and as data sinks for persistent storage and interface to access processing results. It is emphasized that the *Federated Cloud Data Services* shall be provided as two identical sets of services to be deployed in two different security domains. The *SCALE datacenter environment* is deployed only once and shall be hosted for a period of 6 months after the final demonstration of the Testbed. It shall be available for further experimentation by Sponsor authorized personnel.

All results will be captured in the Federated Clouds Analytics Engineering Report.

EOC, SCALE, and SEED

The Testbed-14 Earth Observation Cloud thread [https://portal.opengeospatial.org/files/? artifact_id=77327#ExploitationPlatform] has been working on requirements similar to those of the SCALE Datacenter Environment and SEED. Multiple incompatible solutions are not in the interest of any of the Sponsors. This thread will evaluate the two platforms and make recommendations on a way forward. That may include recommendations for harmonization as well as operational rationales for maintaining any differences between the platforms.

The following tasks shall be executed:

1. Evaluate the work produced by the Testbed-14 Earth Observation Cloud thread and the documentation for the Scale Datacenter Environment.

- a. Identify opportunities for harmonization or standardization between them.
- b. Identify any features which must remain separate and the rationale for why this should be so.
- c. Identify and describe any hard problems which will require additional work.
- d. Identify and describe any opportunities which should be examined in future initiatives.

The following figure illustrates the work items and deliverables of this sub-task.



Figure 23. Work items and deliverables of the EOC, SCALE and SEED sub-task

This subtask includes a single deliverable only, which documents the approach, results, and conclusions reached by this study.

Federated Clouds Provenance

Analytic systems work with data from a variety of sources including internal, government, commercial, and open sources. As the number of sources grows alongside the use of machine learning algorithms, engineers and analysts need a means of ensuring that their workflows take into account the pedigree of veracity.

Testbed-15 seeks experimental ideas on the use of blockchain and/or distributed ledger technologies to address this issue. How might sources and methods for data collection be recorded and reported along the supply chain of geospatial data in such a way that removes the need for trust between suppliers and consumers? What draft standards or recommendations for the geospatial community might advance the widespread adoption of technology or processes to ensure data provenance?

There are a number of standards for provenance. Selecting a provenance data standard is not the purpose of this task. The purpose is to explore Blockchain and distributed ledger technologies to protect the integrity of that provenance data

Some provenance metadata models:

- LI-Lineage from ISO 19115 and 19115-2: This metadata is designed to capture all of the process steps, their inputs, outputs, and governing parameters used to generate an image from remote sensing data.
- **OGC Context**: A 'context document' specifies a fully configured service set which can be exchanged (with a consistent interpretation) among clients supporting the standard.
- **W3C Prov**: A family of documents published by the W3C addressing the problem of provenance on the Web.

The following tasks shall be executed:

- 1. Contractor shall work with an NGA Subject Matter Expert to identify a notional workflow and the appropriate provenance data to represent that workflow. This workflow shall contain multiple steps with a full suite of provenance data captured at each step.
- 2. Contractor shall fabricate a test set of provenance data reflecting the conclusions developed through Task 1.
- 3. Contractor shall examine alternate techniques for using Blockchain technologies to:
 - a. Assure the integrity of the provenance data
 - b. Allow extraction of provenance data from any stage of the workflow
 - c. Provide a trusted assertion of when each change to the provenance data was made and the identification of the entity that made the change.
- 4. Contractor shall document the approach, results, and conclusions used in this effort.

The following figure illustrates the work items and deliverables of this sub-task.



Figure 24. Work items and deliverables of the Federated Clouds Provenance sub-task

This subtask includes a single deliverable only, which documents the approach, results, and conclusions reached by this study. It shall further describe and contain any data or software artifacts used in the Provenance study.

B.8.5. Work Items & Deliverables

All work items have been illustrated above. The following list identifies all deliverables that are part of this task. Detailed requirements are stated above. All participants are required to participate in all technical discussions. Thread assignment and funding status are defined in section Summary of Testbed Deliverables.

Engineering Reports

- **D019 Federated Clouds Security ER** Engineering Report capturing all results and experiences that address security from this task. This Engineering Report also documents all recommendations for new or modifications to existing OGC standards arising from this effort.
- **D020 Federated Clouds Analytics ER** Engineering Report capturing all results and experiences that address analytics from this task. This Engineering Report also documents all recommendations for new or modifications to existing OGC standards arising from this effort.

- **D021 EOC, SCALE, and SEED ER** Engineering Report capturing the approach, results, and conclusions reached by this study.
- **D022 Federated Clouds Provenance ER** Engineering Report capturing the approach, results, and conclusions reached by this study. It shall further describe and contain any data or software artifacts used in the Provenance study.

Components

- **D143 Federation Manager #1** Federation Manager instance that shall support TLS mutual authentication using X.509 PKI certificates.
- **D144 Federation Manager #2** Federation Manager instance.
- **D145 Scale Datacenter Environment** SCALE datacenter environment with support for vendor provided applications. The environment shall be hosted for a period of 6 months after the final demonstration of the Testbed. It shall be available for further experimentation by Sponsor authorized personnel.
- **D146 WPS Analytics** WPS instance interfacing the SCALE datacenter environment; component shall be deployed twice in two separate security environments
- **D147 Federated Cloud Data Services** OGC data services such as WMS, WCS, or WFS (details to be specified by participant). The services are bundled as a single deliverable. These services serve as data sources for SCALE datacenter applications, and as data sinks for persistent storage and interface to access processing results. Identical sets of services shall be deployed in two separate security domains.

B.9. Portrayal

Interoperable dynamic geospatial portrayal is still challenging across multiple computing, rendering, communications and display environments. Testbed-14 investigated the situation since the release of Style Layer Descriptors back in 2002 and analyzed alternative rendering concepts and solutions.



Section 5.1 of the Testbed-14 Symbology Engineering Report [https://portal.opengeospatial.org/files/? artifact_id=81833&version=1] provides a detailed summary of all activities since Testbed-10, when the formalization of portrayal ontologies started. These efforts have been continued during Testbed-11, 12, 13, and 14, and the (currently ongoing) extended Vector Tiles Pilot [https://portal.opengeospatial.org/index.php?m=projects&a=view&project_id=586&tab=0]. As styling aspects as well as WMTS and WFS API design are in focus of the pilot, participants interested in this task are strongly advised to sign up to the pilot's email reflector to become knowledgable of the latest discussions.

B.9.1. Problem Statement & Research Questions

Despite the previous efforts referenced above, what is still missing in OGC is a robust conceptual model and APIs capable of supporting multiple style encodings and the style encodings themselves. This task will develop an **Open Portrayal Framework** (**OPF**) to address this issue. The OPF will leverage recent OGC portrayal efforts including Testbed-14 and the Vector Tiles Pilot. It will also address methods to style tiled vector data, Elevation Data, Routes and Points of Interest (online and - equally important - offline; including GeoPackage) to increase real-world utility, i.e. it clearly extends beyond tiled vector data. The result of the Open Portrayal Framework initiative will be a candidate international and NSG standard for geospatial portrayal, aligned with modern Style encoding, tools and approaches. The OPF shall be exercised in various ways, featuring GeoPackage, WFS, and WMTS technologies.

This task also recognizes that WMTS is one of the most implemented OGC technologies within geospatial enterprises and essential to online/offline display environments - and future Open Portrayal Frameworks. Its tile-based operations are also used by major data providers including satellite imagery organizations. Therefore, the second goal of this task is the definition of an OpenAPI based WMTS, following principles set by WFS 3.0 [https://cdn.rawgit.com/opengeospatial/WFS_FES/3.0.0-draft.1/docs/17-069.html].

Finally, this task recognizes that Portrayal display environments must have the latest satellite and aerial imagery – which will be challenging as hundreds of new commercial satellites and other sources are deployed. To meet this challenge a Transactional WMTS (WMTS-T) will be defined.

B.9.2. Aim

The aim of this task is to create an Open Portrayal Framework, to advance WMTS to support OpenAPI, and to add transactional capabilities to WMTS. the Open Portrayal Framework shall be exercised in various ways to strengthen the robustness of the conceptual and metadata models and to ensure robust interoperability in real world situations.

B.9.3. Previous Work

The following reports serve as a baseline for this task:

- 18-029 OGC Testbed-14: Symbology Engineering Report [https://portal.opengeospatial.org/files/? artifact_id=81833&version=1]
- 18-083 OGC Vector Tiles Pilot: WMTS Vector Tiles Extension Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81734&version=1]
- 18-078 OGC Vector Tiles Pilot: WFS 3.0 Vector Tiles Extension Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81698&version=1]
- 18-074 GeoPackage Vector Tiles Extensions Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81680&version=1]
- 18-045 OGC Testbed-14: Next Generation Web APIs WFS 3.0 Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81793&version=1]
- 17-045 OGC Testbed-13: Portrayal Engineering Report [http://docs.opengeospatial.org/per/17-045.html]
- 17-094r1 OGC Portrayal Concept Development Study Engineering Report [http://docs.opengeospatial.org/per/17-094r1.html]
- 05-078r4 OpenGIS Styled Layer Descriptor Profile of the Web Map Service Implementation Specification [http://portal.opengeospatial.org/files/?artifact_id=22364]
- Mapbox Styles Specification [https://www.mapbox.com/mapbox-gl-js/style-spec/]

B.9.4. Scenario & Requirements

The Portrayal scenarios integrate the three topics introduced below, the Open Portrayal Framework, the WMTS 2.0 with OpenAPI support, and its transactional extension. It is emphasized that all groups participate in the development of the conceptual model and the metadata model as well as implementation tests. The list of tasks provided below provides additional requirements for the various teams to work together.

Open Portrayal Framework

OPF in Testbed 15 will deliver a single conceptual model and APIs capable of supporting multiple style encodings, and the style encodings themselves. OPF will be tested using NSG data and symbols in scenarios simulating multiple environments - including computing (web, mobile, desktop), rendering (server, client), communications (connected, disconnected) and display (imagery background, day and night views, variable bearing and pitch, see figure below). Methods to style tiled vector data, Elevation Data, Routes and Points of Interest offline/online will also be assessed and delivered to increase real-world utility.



Figure 25. Examples of topographic, satellite overlay, and night styles from the Vector Tiles Pilot

The OPF development starts with the definition of the single conceptual model for Style encodings and the metadata model to support Style resource discovery. Both models shall be developed by all participants of this task. Further on, methods shall be developed and tested to associate Style resources with sets of feature tiles, layers, elevation as well as features without tiles.



Figure 26. Development of the Open Portrayal Framework, including conceptual model, metadata model, and implementation tests for various types of data among all teams

Next, the WFS and WMTS server and client providers shall develop and prototype extensions for Vector Tiles WMTS 2.0 and WFS 3.0 to publish, share, discover and use Style resources via a Styles API implemented according to Next Generation API methods (GET, POST, PUT, DELETE) defined in Testbed-14 and the Vector Tiles Pilot.



Figure 27. Development of a Styles API for WFS 3.0 and WMTS 2.0

An example of the current discussion in the Vector Tiles Pilot is provided below.

Tiles Access to data (features), partitioned into a hierarchy of tiles.	>
Styles Access to styles.	\sim
GET /styles retrieve all available styles from the dataset	
FOST /styles add styles to the dataset	
GET /styles/{styleId} retrieve a style of the dataset by id	
PUT /styles/{styleId} replace a style from the dataset	
DELETE /styles/{styleId} delete a Style from the dataset	
GET /collections/aeronauticcrv/styles retrieve all available styles from the collection Aeronautical (Curves)	
POST /collections/aeronauticcrv/styles add styles to the collection Aeronautical (Curves)	
GET /collections/aeronauticcrv/styles/{styleId} retrieve a style of the collection Aeronautical (Curves) by id	
PUT /collections/aeronauticcrv/styles/{styleId} replace a style from the collection Aeronautical (Curves)	
DELETE /collections/aeronauticcrv/styles/{styleId} delete a Style from the collection Aeronautical (Curves)	

Figure 28. Style API is currently discussed in the Vector Tiles Pilot

The providers of Visual Style Editors shall encode NSG symbols according to the conceptual model. The symbols shall be hosted on the WFS and WMTS servers. The Visual Style Editor components shall be extended to support the Style APIs developed for WFS and WMTS



Figure 29. Development of NSG symbols according to the conceptual model

An exmple of an Visual Style Editor from the Vector Tile Pilot is provided below.



Figure 30. Visual Style Editor Integrated Into Vector Tiles Pilot Test Client for MB and SLD Styles

Once all models and APIs are designed and implemented, all components shall undergo detailed testing. For that purpose, a set of experiments shall be executed to exercise the conceptual model, metadata and APIs based on the following variables: a. Computing environment (web, mobile, desktop), b. Rendering location (server, client), c. Communications (connected and disconnected), d. Display (imagery background, day and night views, variable bearing and pitch) e. Style encodings (Styled Layer Descriptor and Mapbox Style Specification) f. Assessment of methods to handle multi-layer Styles, including style sheets

The following figures illustrate a possible flow of events and interactions. It is emphasized that the full list provided above shall be exercised. The figures serve illustrative purposes and are not exhaustive.

In the figure below, a user connects with the Visual Style Editor to a WFS/WMTS to discover a Style resource. The user makes use of the metadata model developed above. The discovered resource is downloaded, modified, and the updated Style resource is uploaded to the server again. Alternatively, the user develops a new Style resource and inserts the resource into a WFS/WMTS by means of the Style APIs.



Figure 31. User discovers, modifies, and makes Style resources available online

In the second scenario, different types of clients (mobile, browser based, desktop app based) discover data and corresponding styles on WFS or WMTS servers, exercise the styles, and make

use of the transactional WMTS interface to keep server data up-to-date.



Figure 32. Different client types exercising the new Styles APIs and Transactional Capability of WMTS

GeoPackage

The GeoPackage client and GeoPackage component providers play a critical role in the development of the Open Portrayal Framework. The OPF must support both online and offline operations. Most portrayal work done to-date has been focused on Web Services (online). This may result in an OPF with a strong on-line bias. GeoPackage participants should counter that bias by identifying the technical constraints which differentiate the on-line from the off-line environments. Those constraints shall be used throughout the OPF development effort to assure that the OPF maximizes common concepts and capabilities across on-line and off-line implementations. GeoPackage participants shall also develop and test methods to implement the OPF in GeoPackage for off-line use and use the results to further refine the OPF.



Figure 33. Development of a multiple Style resource encodings in GeoPackage for offline use

OpenAPI WMTS

This task enhances data supply chains for Raster Tiled Map/Imagery Data and future tiled vector data-based maps with modern, simple APIs. This capability would be beneficial to multiple NSG partners for provisioning of both raster data and vector data tiles. In addition, this effort will align WMTS will modern API approaches for handling Styles and enhance joint and coalition interoperability.

The general approach used by WFS 3.0 should apply easily to an OpenAPI WMTS, and the existing "RESTful" version of WMTS should be simple to port to an OpenAPI-based approach. This could be achieved by keeping the tile access URL and templates and replace the *Capabilities* document with an equivalent of landing page, API and tiles endpoints based on reusable OpenAPI components. The effort would leverage lessons learned in Testbed-14 and reuse the definitions of

Tiling Scheme resources as applicable – and prepare 'WMTS 2.0' to handle Styles APIs in an Open Portrayal Framework.

Transactional WMTS

Finally, this Work Item recognizes that Portrayal display environments must have the latest satellite and aerial imagery – which will be challenging as hundreds of new commercial satellites and other sources are deployed. To meet this challenge a Transactional WMTS (WMTS-T) will provide the ability to INSERT, UPDATE and DELETE individual images in a WMTS 2.0 tile server. Without the WMTS-T, every change to a tile set may require the regeneration of the full tile set. For some systems this can take a significant amount of time and resources. It also introduces increased risk that errors may be introduced into the tile set. Incremental changes greatly reduce the risk and expense of maintaining the tile set. Tactical users and First Responders may also wish to add their own mission specific information to the tile set. Yet they frequently must work over denied, degraded, intermittent, and low bandwidth (DDIL) networks. Due to the limitations of DDIL networks, updating the tile set is not practical if the entire tile set must be rebuilt. However, tile updates can be made using a WMTS-T as network connectivity allows.

The following list summarizes the tasks that shall be executed:

- 1. Continue development of the single conceptual model for Style encodings
- 2. Develop and test methods to associate Style resources with sets of feature tiles, layers, elevation data and features without tiles
- 3. Develop and prototype extensions for Vector Tiles WMTS and WFS to publish, share, discover and use Style resources via a Styles API implemented according to Next Generation API methods (GET, POST, PUT, DELETE) defined in Testbed-14 and the Vector Tiles Pilot
- 4. Develop and test methods to support multiple Style resource encodings in GeoPackage for offline use
- 5. Define and test metadata to support Style resource discovery
- 6. Encode NSG symbols according to the conceptual model. Host the symbols on the prototype Vector Tile WMTS and WFS. Support the metadata defined above.
- 7. Extend Visual Style Editor components to discover, use, publish and share styles from multiple sources implementing Style APIs including the prototype Vector Tiles WMTS and WFS
- 8. Develop a set of experiments to exercise the conceptual model, metadata and APIs based on the following variables:
 - a. Computing environment (web, mobile, desktop),
 - b. Rendering location (server, client),
 - c. Communications (connected and disconnected),
 - d. Display (imagery background, day and night views, variable bearing and pitch)
 - e. Style encodings (Styled Layer Descriptor and Mapbox Style Specification)

- f. Assess methods to handle multi-layer Styles, including style sheets
- 9. Develop a proposed version 2.0 of WMTS based on OpenAPI and the design patterns established by WFS 3.0.
- 10. Develop server implementations of WMTS 2.0.
- 11. Develop client implementations of WMTS 2.0
- 12. Develop a Transactional extension to the WMTS (WMTS-T) 2.0 specification.
- 13. Extend the WMTS 2.0 Servers for WMTS-T.
- 14. Extend the WMTS 2.0 clients to exercise the transaction extension.
- 15. Exercise the clients against both Servers to validate the extension in terms of functionality and implementability.
- 16. Execute the experiments and demonstrate the results.
- 17. Document the results.

B.9.5. Work Items & Deliverables

The following figure illustrates the work items and deliverables of this task.



Figure 34. Portrayal work items overview

The following list identifies all deliverables that are part of this task. Detailed requirements are stated above. All participants are required to participate in all technical discussions. Thread assignment and funding status are defined in section Summary of Testbed Deliverables.

Engineering Reports

- D011 Concept Model for Style Encoding & Metadata Model ER Engineering Report documenting the Conceptual Model for Style Encoding and the Metadata Model for Style resources.
- **D012 Style API for WFS, WMTS, GeoPackage ER** Style API for WFS 3.0, WMTS 2.0, and GeoPackage Engineering Report
- **D013 WMTS ER** WMTS 2.0 Engineering Report capturing all results and experiences that address the WMTS Styles API. This Engineering Report also documents all recommendations for new or modifications to existing OGC standards arising from this effort.
- **D014 WMTS draft specification** WMTS 2.0 draft specification for consideration by the OGC Standards Program. Includes Styles API and OpenAPI design elements.
- **D015 WMTS-T ER** WMTS 2.0 Engineering Report capturing all results and experiences that address the WMTS Transactions Extension. This Engineering Report also documents all recommendations for new or modifications to existing OGC standards arising from this effort.
- **D016 WMTS-T draft specification** WMTS 2.0 draft specification for consideration by the OGC OGC Standards Program.
- **D017 Portrayal Summary Report** Summary Report of all activities, conducted experiments, and results from this task.

Components

- **D126 Visual Style Editor #1** Visual Style Editor that connects to Style APIs and delivery of NSG Symbols according to conceptual model (1/2)
- **D127 Visual Style Editor #2** Visual Style Editor that connects to Style APIs and delivery of NSG Symbols according to conceptual model (2/2)
- **D128 GeoPackage with Styles Encodings #1** GeoPackages implementing the evaluated Style encodings(1/2)
- **D129 GeoPackage with Styles Encodings #2 (mobile)** GeoPackages implementing the evaluated Style encodings(2/2), mobile device app
- **D130 GeoPackage with Styles Encodings client #1** GeoPackage client supporting the evaluated Style encodings (1/2)
- **D131 GeoPackage with Styles Encodings client #2 (mobile)** GeoPackage client supporting the evaluated Style encodings (2/2), mobile device app.
- **D132 WMTS-T w Styles API #1** WMTS 2.0 Server with Style API Implementation and Transactions support (1/2)
- **D133 WMTS-T w Styles API #2** WMTS 2.0 Server with Style API Implementation and Transactions support (2/2)
- D134 WMTS-T w Styles Client #1 WMTS 2.0 Client with Style API Implementation and

Transactions support (1/2)

- **D135 WMTS-T w Styles Client #2** WMTS 2.0 Client with Style API Implementation and Transactions support (2/2)
- D136 WFS 3.0 w Styles API #1 WFS 3.0 Server with Style API Implementation (1/2)
- D137 WFS 3.0 w Styles API #2 -WFS 3.0 Server with Style API Implementation (2/2)
- D138 WFS 3.0 w Styles Client #1 WFS 3.0 Client Implementation (1/2)
- D139 WFS 3.0 w Styles Client #2 (mobile) WFS 3.0 Client Implementation (2/2), mobile device app

IMPORTANT

The following task Earth Observation Process and Application Discovery (EOPAD) only partly funded through this CFP and partly through an ESA tender. Special conditions apply for proposals for the work items highlighted as being part of the ESA tender! Please review the Main Body and the Work Items and Deliverables section towards the end of the chapter very carefully!

B.10. Earth Observation Process and Application Discovery

Several platforms emerge that provide access to Earth Observation data and processing capacities. These platforms host very large datasets, which makes a paradigm shift from data download and local processing towards application upload and processing close to the physical local of the data more and more important. To interpret peta- or exascale scientific data, capabilities of these platforms need to be combined in future.

Testbed 13 & 14 included threads related to packaging, deployment and execution of applications in cloud environments that expose standardized interfaces (Web Processing Service, WPS). Previous Testbeds used ESA's cloud platform environment, but the developed approach is agnostic to the target cloud platform as long as a dedicated standardized interface (e.g. a Web Processing Service, WPS), a container execution environment (e.g. Docker), and data access (ideally supporting dynamic mounting into and out of containers) are provided. In order to for end-users to exploit such applications and already deployed processing capacities, facilities must be provided for users to discover the applications and to be provided with detailed descriptions and invocation instructions. It is emphasized that this task always addresses both applications to be deployed as well as already deployed applications that are available behind WPS interfaces.

Hence, the focus of Testbed-15 work is to define the building blocks through which such applications and related services can be exposed through a Catalogue service. The following research questions shall be addressed:

• How can an application and application data catalog be established without inviting yet another catalog specification?

- What are the differences of the various catalog models, their limitations and key characteristics?
- How to link apps and data? Are apps by definition bound to a specific data server? If not, how to establish the links between apps and data?
- Catalogs pointing to data instances often point to the access interface, but not necessarily to the data itself. How to use catalogs so that an app can dynamically load data from the data source?
- How to annotate data sets? Imagine a Quality of Service process that add additional metadata, or a QoS process that develops a scoring or performance value for specific entries
- How to support the incremental search process, where users use a sequence of queries to identify their target application?
- How to check links between applications (for matching I/O settings) to ensure workflows being executable? How to ensure this aspect across platforms?
- How to facilitate access to the actual data using the applications?
- How to enable fusion of multiple data and visualization of resulting products?

B.10.1. Aim

Describe and demonstrate how OGC standards can be used or have to be extended to provide for discovery and use of EO data processing applications that can be deployed and executed by the user or are already deployed and available behind standardized OGC interfaces. Demonstrate how existing and emerging systems as deployed by e.g. NASA (e.g. NASA DAAC [https://earthdata.nasa.gov/about/daacs]s and NASA DASS [https://www.nccs.nasa.gov/sites/default/files/ NCCS%20DASS%20Overview.pdf]), ESA (ESA TEP [https://tep.eo.esa.int/]s) or systems that have already integrated various nodes such as the Earth System Grid Federation [https://esgf.llnl.gov/2018-F2F.html] (ESGF) can be federated to allow for cross-platform analysis and visualization of data.

B.10.2. Previous Work

The following reports serve as a baseline for this task (not all of these reports are publicly available yet, as the OGC voting and release process is currently ongoing. They can be made available upon request):

- 18-050r1 OGC Testbed-14: ADES & EMS Results and Best Practices Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81794&version=1]
- 18-049r1 OGC Testbed-14: Application Package ER [https://portal.opengeospatial.org/files/? artifact_id=81790&version=1]
- 18-057 OGC Testbed-14: Authorisation Authentication and Billing ER [https://portal.opengeospatial.org/files/?artifact_id=81795&version=1]
- 18-090r1 OGC Testbed-14 Federated Clouds Engineering Report [https://portal.opengeospatial.org/files/?artifact_id=81863&version=1]

- 17-035 OGC Testbed-13: Cloud ER [http://docs.opengeospatial.org/per/17-035.html]
- 17-023 OGC Testbed-13: EP Application Package ER [http://docs.opengeospatial.org/per/17-023.html]
- 17-024 OGC Testbed-13: Application Deployment and Execution Service ER [http://docs.opengeospatial.org/per/17-024.html]
- 17-018 OGC Testbed-13: Data Quality Specification Engineering Report 2018 [http://docs.opengeospatial.org/per/17-018.html]
- 17-040 OGC Testbed-13: DCAT/SRIM Engineering Report [http://docs.opengeospatial.org/per/17-040.html]
- 16-059 Testbed-12 Semantic Portrayal, Registry and Mediation Engineering Report [http://docs.opengeospatial.org/per/16-059.html]
- 16-062 Testbed-12 Catalogue and SPARQL Engineering Report [http://docs.opengeospatial.org/ per/16-062.html]
- OGC Catalog Service specification [http://www.opengeospatial.org/standards/cat]
- 07-045r2 OpenGIS Catalogue Services Specification 2.0.2 ISO Metadata Application Profile: Corrigendum [https://portal.opengeospatial.org/files/?artifact_id=80534&version=1]

B.10.3. Scenario & Requirements

The discovery solution in principle comprise the following building blocks:

- **Data Model**: Providing the data structure through which the application and/or service is described and presented as a resource through the catalogue.
- **Service Interface**: Providing the call interface through which users can discover applications and services through facetted and textual search, and then retrieve application/service metadata providing more detail.
- Service Management Interface: Providing the call interface through which users, operators or administrators can Create, Update, Delete the information about applications.

It is important to recognize that the Catalogue Service must serve the needs of browser-based users requiring a user interface for their catalogue interactions, as well as the machine-to-machine interface that can be exploited for scripted and automated interactions.

It should also be considered that Discovery is not simply limited to a single cloud platform, but should also consider the requirement to catalogue external services – for example, a catalogue that federates the offering of multiple platforms by crawling their contents or query them in real time. The Testbed-15 solution will be tested with platforms provided by NASA, ESA, and ideally the ESGF.

A **Data Model** should be defined that provides:

- A unique identifier for each application or service
- Descriptive information such as title/abstract etc.
- Classification information that could be used for facetted search / browse
- Descriptions of the data types / constraints of the processing inputs and outputs. *These* should be formally described such that machine-decisions can be made regarding suitability of input data and pipelining of tool workflows.
- Accuracy / error information regarding the algorithms
- Version information and associated application provenance
- Interactive or non-interactive
- Execution environment constraints (e.g. docker version, hardware requirements) or Execution end-points (e.g. URLs)
- Metrics for costs estimation (data sizing, processing costs)
- Any other element as mutually agreed upon at the kick-off meeting

The application packaging strategies devised in Testbed 13/14 go some way to describe aspects of this model. These should be analyzed as a starting point, potentially re-using any useful concepts and thus maintaining consistency.

The Service (Management) Interface should provide the following capabilities:

- Search interface for discovery by facet and free-text in combination
- Request interface for retrieval of facet dictionary to support browse
- Request interface for retrieval of Application detailed metadata, e.g. for display of an application landing page in catalogue web User Interface (UI)
- Management interface for the Creating, Update and Deletion of records in the catalogue

Many possible approaches to this problem are possible and the participants should be ready to analyze as many as possible, weight Pros and Cons and decide on an approach that reduces as much as possible the burden to future end-users and service operators. The following are existing standards, concepts and/or models that should be taken into consideration:

- OGC CSW [http://www.opengeospatial.org/standards/cat] (Catalog Service Web)
- OGC OpenSearch (including existing EO related profiles [http://docs.opengeospatial.org/is/13-026r8/13-026r8.html] and extensions [https://portal.opengeospatial.org/files/?artifact_id=56866])
- STAC [https://github.com/radiantearth/stac-spec] (SpatioTemporal Asset Catalog specification)
- WFS 3.0 [https://github.com/opengeospatial/WFS_FES]
- DCAT [https://www.w3.org/TR/vocab-dcat/] (specially GeoDCAT-AP [https://inspire.ec.europa.eu/ documents/geodcat-ap])

- INSPIRE [https://inspire.ec.europa.eu/] in general
- Semantic Web/RDF/SPARQL; including schema.org and similar taxonomies;
- SRIM (Semantic Resource Information Model; see Testbed-12 [http://docs.opengeospatial.org/ per/] and Testbed-13 [http://www.opengeospatial.org/pub/Testbed13/er.html] results)
- Application-based catalogs such as Google Play Store [https://play.google.com/store] or Apple's App Store [https://developer.apple.com/app-store/]
- Digital Object Interface Protocol [https://www.dona.net/sites/default/files/2018-11/ DOIPv2Spec_1.pdf]

Special emphasis is put on schema.org, which was recently discussed in several presentations at the AGU Fall meeting [https://fallmeeting.agu.org/2018/], December 2018, in Washington DC for science data discovery or as a baseline for JSON-Based Structured Data for Discovery, Access, and Usability of Earth observation data and products (see link [https://agu.confex.com/agu/fm18/ meetingapp.cgi/Search/0?sort=Relevance&size=10&page=1&searchterm=schema%2Eorg])

Key Use Cases

The use cases apply equally to processing service end-points (such as an existing WPS service), as to 'packaged' applications deployed by users (as described by Testbeds 13/14). Thus, the data model shall handle the invocation target generically, and hence the term 'application' can be taken to be any such processing service end-point.

It is emphasized that all use cases shall support both chaining of *to be deployed* as well as *already deployed applications* that are available at WPS interfaces.

A. Application/Service lifecycle management from user/operator

In this use case Alice, the owner of the application or the operator of the service, registers her application in the Catalogue Service. After some time, she updates the application and consequently the information in the service. Finally, at the end of life of the application, Alice removes it from the catalogue.

For example, this use case could be used to describe a set of processing functions provided by an existing WPS service – to catalogue them and make them discoverable by end users.

The interaction between Alice and the Catalogue Service is reported below.

- 1. Alice activates the Catalogue Service management interface to register her application (passing all information required by the data model).
- 2. The service checks that Alice is authorized to register the new application and then fulfils her request.
- 3. Alice performs a simple query to be presented with a list of only her applications and checks that the application is registered, asking for its details.

- 4. The service returns the record of the application.
- 5. Alice updates the metadata information of her application.
- 6. The service checks that Alice is authorized to update the information of the application and then modifies it.
- 7. Alice searches for the application she previously registered.
- 8. The service returns the record of the application.
- 9. Alice requests the deletion of the application in the Catalogue Service.
- 10. The service checks that Alice is authorized to delete the information of the application and then deletes it.

B. Application creation/deletion via M2M interaction

In this use case Alice, the owner of the application, deploys it on a cloud platform, using for instance an EMS as defined in Testbed-14 [https://portal.opengeospatial.org/files/?artifact_id=81794& version=1]. After the deployment, the cloud platform registers the application information in the Catalogue Service.

At last Alice un-deploys the application from the cloud platform which therefore deletes the record relative to the application from the Catalogue Service.

The interaction between Alice, the cloud platform and the Catalogue Service is reported below.

- 1. Alice prepares an application package and interfaces with a cloud platform to deploy it.
- 2. The cloud platform checks that Alice is authorized to deploy the application and then fulfills the request.
- 3. The Exploitation Platform, on behalf of Alice, requests to the Catalogue Service to register the application providing all the required information.
- 4. The service checks that Alice is authorized to register new applications and then fulfills her request.
- 5. Alice requires to the cloud platform to un-deploy the previously registered application
- 6. The cloud platform checks that Alice is authorized to un-deploy the application and then fulfills the request.
- 7. The Exploitation Platform, on behalf of Alice, sends a deletion request to remove the application from the Catalogue Service.
- 8. The service checks that Alice is authorized to remove the application and then fulfills the request.

C. Application discovery

In this use case Bob, a user of applications, uses the Catalogue Service to find an application to

perform a specific processing algorithm (e.g. an NDVI). He fills a form with the available search criteria or uses a free text form to specify some key words and then the service returns all the applications which matches the specified criteria. The interaction between Bob and the Catalogue Service is mediated by a client of the service which provides a GUI to ease the user experience of Bob.

The interaction between Bob, the Catalogue Service and its Client is reported below.

- 1. Bob accesses the client which is linked to the Catalogue Service.
- 2. The client requests to the Service which are the facets it supports.
- 3. The Service returns the facets to the client.
- 4. The Client presents Bob with a query interface based on the available facets.
- 5. Bob fills the query form or perform a textual search for key words Bob is able to incrementally narrow the search results by selection of facet filters in combination with textual search terms.
- 6. The client sends query requests to the service using the inputs provided by Bob.
- 7. The Service returns a minimal set of information for each of the application matching the search criteria (*).
- 8. The client presents the return information for Bob to see.
- 9. Bob asks for more details on one application.
- 10. The client sends a request to the service for the complete metadata of the select application.
- 11. The server returns the information.
- 12. The client presents the information to Bob.

(*) this step has the following variants:

- a. the service searches for matching applications internally.
- b. the service forwards the request to all federated Catalogue Services
- c. the service searches for matching applications internally but asynchronously has previously collected all available application information from all federated Catalogue Services.

D. Support to application and service chaining

In this use case Bob, a user of applications, uses the Catalogue Service to find first an application to perform a processing algorithm (e.g. classification) and then another application which can be fed with the output of the previous one (e.g. change detection) in order to create a complex application (a workflow or chain) to run systematically over a certain area of interest. The interaction between Bob on one side and the Catalogue Service or the cloud platform on the other is mediated by a client of the service which provides a GUI to ease the user experience of Bob. The client supports also Bob in verifying that the two applications are compatible (i.e. the input type of the second correspond to the output of the first) thanks to the information returned

by the Catalogue Service.

The interaction between Bob, the Catalogue Service, the EMS and the client is reported below. The interaction needs to consider applications and data being provided by different platforms.

- 1. Bob accesses the client which is linked to the Catalogue Service.
- 2. Bob performs two times the use case C to identify two applications A1 and A2.
- 3. Bob uses the client to create a complex application containing both applications in a chain $(A1 \rightarrow A2)$.
- 4. The client verifies that the output of the A1 is compatible with the input of A2 in order to validate the application chaining.
- 5. The complex application is prepared and can be deployed (as per Testbed 14 EOC).

The use-case would continue with the deployment of the application in the Cloud Platform, for this activity these steps are not required.

The "complex application" referenced above needs to be designed by participants in mutual agreement with sponsors at the kick-off meeting. It should be based on science application scenarios using various platforms, such as ESA Thematic Exploitation Platforms, NASA's Earth Science Data Systems, or Earth System Grid Federation climate data.

B.10.4. Work Items & Deliverables

The following figure illustrates the work items and deliverables of this task.

IMPORTANT	The description of the work items and deliverables uses different colors to differentiate work items that are part of the OGC CFP from work items that are part of the ESA ITT. All work items colored in green with red work item identifiers are part of the OGC CFP. All work items colored in orange with purple work item identifiers are part of the ESA ITT. Both parts require explicit responses as described in CFP Main Body at http://bit.ly/Testbed-15.
IMPORTANT	It is emphasized that all work items, independently of being part of the OGC CFP or ESA, require Technical Interoperability Experiments (TIEs) with other components as indicated by black connectors. That means that components of the OGC CFP will need to interact with components of the ESA ITT and vice versa!



Figure 35. Overview of all work items and deliverables of this task. Green work items require a response to the OGC CFP, orange work items require a response to the ESA ITT. Arrows indicate interaction between client and service, direct lines indicate usage relationships. Not all connections are shown in the diagram. Both demo apps as well as WPS instances need to be registered at both catalogs.

The following list identifies all deliverables that are part of this task. Detailed requirements are stated above. All participants are required to participate in all technical discussions. Thread assignment and funding status are defined in section Summary of Testbed Deliverables.

OGC CFP Work Items:

- D121 Catalogue & Discovery Service SERVER (1 of 2) Implementation of Catalogue & Discovery Service SERVER in accordance with the Data Model (D010). The catalog server shall make use of the data model discussed above and be called from both client applications. WPS instances as well as demo apps need to be registered at the catalog.
- D122 Catalogue & Discovery Service WEB CLIENT Implementation of Catalogue & Discovery Service WEB CLIENT in accordance with the Data Model (D010). The client shall interact with both catalog instances as well as both WPS instances. It shall support catalog as well as WPS application deployment and registration requests.
- **D123 Catalogue & Discovery Service CONSOLE/DESKTOP CLIENT** Implementation of Catalogue & Discovery Service CLIENT in accordance with the Data Model (D010).). The client shall interact with both catalog instances as well as both WPS instances. It shall support catalog as well as WPS application deployment and registration requests.
- **D124 Web Processing Service (1 of 2)** WPS instance with already deployed process. The WPS needs to be registered at both catalog services.

ESA ITT Work Items:

- D117 Catalogue & Discovery Service SERVER (2 of 2) Implementation of Catalogue & Discovery Service SERVER in accordance with the Data Model (D010). The catalog server shall make use of the data model discussed above and be called from both client applications. WPS instances as well as demo apps need to be registered at the catalog.
- **D118 Web Processing Service (2 of 2)** WPS instance and two demo applications that can be deployed dynamically and be compliant with Application Package specification. The WPS as well as both demo apps need to be registered at both catalog services.
- **D010 Engineering Report** Engineering Report covering all results of this task including the development and description of the data model.