OGC Mixed Reality to the Edge Concept Development Study

Table of Contents

1. Executive Summary	4
1.1. Mixed Reality to the Edge Concept Development Study overview	4
1.2. The problem in a nutshell	5
1.3. Key findings	5
1.3.1. Context, foundation data, and real time inputs at the edge	5
1.3.2. Speed	6
1.3.3. Standards	6
1.3.4. Training	7
1.3.5. Human factors	7
1.3.6. Bandwidth and updates.	8
1.3.7. Security	8
1.3.8. Authoritative versus non-authoritative sources (Crowdsourcing)	9
1.3.9. Power at the edge	9
1.3.10. Technology	9
1.3.11. AI/ML/DL	10
1.4. Recommendations	10
1.5. Prior-After Comparison	11
1.6. Recommendations for Future OGC Work	12
1.7. Document contributor contact points	13
1.8. Foreword	13
2. Acknowledgements	14
3. Terms and definitions	
3.1. Technical Tools, Applications, and Platforms	17
3.2. Standards and Specifications Mentioned in the Responses	
3.3. Abbreviated terms	20
4. Overview	
5. RFI Respondent Characteristics	23
5.1. Type of Organization	23
5.2. Provider or Consumer?	
5.3. Communities your organization Supports	
6. Summary of responses to the Architecture Questions	
6.1. Response to Question 4.2.1. Automated processes - general	
6.2. Response to Question 4.2.2. Automated Processes for content integration	27
6.3. Response to 4.2.3 - What are key technology components?	
6.4. Response to 4.2.4 - Controlled Vocabularies.	30
6.5. Responses to 4.2.5: Provision of interoperable, automated 3D workflows	31
6.6. Responses to 4.2.6 - Provenance and metadata	32
6.7. Responses to 4.2.7 - Provenance and the end user.	33

6.8. Responses to 4.2.8 - Does your organization have a ML/AI plan	. 34
6.9. Responses to 4.2.9 - Geospatial Standards	. 35
7. Summary of Responses to the Data Analytics Questions	. 37
7.1. Responses to 4.3.1: What 3D/4D data does your organization already use?	. 37
7.2. Responses to 4.3.2: 3D/4D geospatial content and ML/AI analytics	. 38
7.3. Responses to 4.3.3: What 3D formats/encodings do you use?	. 39
7.4. Responses to 4.3.4: What interoperability capabilities are missing?	. 39
7.5. Responses to 4.3.5: Fit for use?	. 40
7.6. Responses to 4.3.6: Are there adequate tools?	. 41
7.7. Responses to 4.3.7: Experience level of users	. 42
7.8. Responses to 4.3.8: Use of domain models.	. 42
8. Summary of Responses to the Scenarios and Use Cases Questions	. 44
8.1. Planit provided use case	. 44
8.2. Responses to 4.4.1: Key Workflows	. 44
8.3. Responses to 4.4.2: Key use cases for provision of 3D/4D content	. 45
9. Summary of Responses to the Operation and Organization Questions	. 47
9.1. Responses to 4.5.1: Policy, organizational, and administrative challenges	. 47
9.2. Responses to 4.5.2: Unique Needs	. 48
10. Summary of Responses to the Technology and Applications Questions	. 49
10.1. Responses to 4.6.1: Portals?	. 49
10.2. Responses to 4.6.2: What other capabilities need to be developed?	. 49
11. Summary of responses to the Other Factors Questions	. 51
11.1. Responses to 4.7.1: What sensors and sensor capabilities?	. 51
11.2. Responses to 4.7.2: Effects of a DDIL environment	. 52
11.3. Responses to 4.7.3: Other success factor?	. 52
Appendix A: Revision History	. 54
Appendix B: Bibliography	. 55

Publication Date: 2019-08-20

Approval Date: 2019-06-28

Submission Date: 2019-06-05

Reference number of this document: OGC 19-030r1

Reference URL for this document: http://www.opengis.net/doc/PER/mr-cds

Category: OGC Public Engineering Report

Editor: Carl Reed, PhD

Title: OGC Mixed Reality to the Edge Concept Development Study

OGC Public Engineering Report COPYRIGHT

Copyright © 2019 Open Geospatial Consortium. To obtain additional rights of use, visit http://www.opengeospatial.org/

WARNING

This document is not an OGC Standard. This document is an OGC Public Engineering Report created as a deliverable in an OGC Interoperability Initiative and is not an official position of the OGC membership. It is distributed for review and comment. It is subject to change without notice and may not be referred to as an OGC Standard. Further, any OGC Public Engineering Report should not be referenced as required or mandatory technology in procurements. However, the discussions in this document could very well lead to the definition of an OGC Standard.

LICENSE AGREEMENT

Permission is hereby granted by the Open Geospatial Consortium, ("Licensor"), free of charge and subject to the terms set forth below, to any person obtaining a copy of this Intellectual Property and any associated documentation, to deal in the Intellectual Property without restriction (except as set forth below), including without limitation the rights to implement, use, copy, modify, merge, publish, distribute, and/or sublicense copies of the Intellectual Property, and to permit persons to whom the Intellectual Property is furnished to do so, provided that all copyright notices on the intellectual property are retained intact and that each person to whom the Intellectual Property is furnished agrees to the terms of this Agreement.

If you modify the Intellectual Property, all copies of the modified Intellectual Property must include, in addition to the above copyright notice, a notice that the Intellectual Property includes modifications that have not been approved or adopted by LICENSOR.

THIS LICENSE IS A COPYRIGHT LICENSE ONLY, AND DOES NOT CONVEY ANY RIGHTS UNDER ANY PATENTS THAT MAY BE IN FORCE ANYWHERE IN THE WORLD. THE INTELLECTUAL PROPERTY IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE DO NOT WARRANT THAT THE FUNCTIONS CONTAINED IN THE INTELLECTUAL PROPERTY WILL MEET YOUR REQUIREMENTS OR THAT THE OPERATION OF THE INTELLECTUAL PROPERTY WILL BE UNINTERRUPTED OR ERROR FREE. ANY USE OF THE INTELLECTUAL PROPERTY SHALL BE MADE ENTIRELY AT THE USER'S OWN RISK. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR ANY CONTRIBUTOR OF INTELLECTUAL PROPERTY RIGHTS TO THE INTELLECTUAL PROPERTY BE LIABLE FOR ANY CLAIM, OR ANY DIRECT, SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM ANY ALLEGED INFRINGEMENT OR ANY LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR UNDER ANY OTHER LEGAL THEORY, ARISING OUT OF OR IN CONNECTION WITH THE IMPLEMENTATION, USE, COMMERCIALIZATION OR PERFORMANCE OF THIS INTELLECTUAL PROPERTY.

This license is effective until terminated. You may terminate it at any time by destroying the Intellectual Property together with all copies in any form. The license will also terminate if you fail to comply with any term or condition of this Agreement. Except as provided in the following sentence, no such termination of this license shall require the termination of any third party end-user sublicense to the Intellectual Property which is in force as of the date of notice of such termination. In addition, should the Intellectual Property, or the operation of the Intellectual Property, infringe, or in LICENSOR's sole opinion be likely to infringe, any patent, copyright, trademark or other right of a third party, you agree that LICENSOR, in its sole discretion, may terminate this license without any compensation or liability to you, your licensees or any other party. You agree upon termination of any kind to destroy or cause to be destroyed the Intellectual Property together with all copies in any form, whether held by you or by any third party.

Except as contained in this notice, the name of LICENSOR or of any other holder of a copyright in all or part of the Intellectual Property shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Intellectual Property without prior written authorization of LICENSOR or such copyright holder. LICENSOR is and shall at all times be the sole entity that may authorize you or any third party to use certification marks, trademarks or other special designations to indicate compliance with any LICENSOR standards or specifications.

This Agreement is governed by the laws of the Commonwealth of Massachusetts. The application to this Agreement of the United Nations Convention on Contracts for the International Sale of Goods is hereby expressly excluded. In the event any provision of this Agreement shall be deemed unenforceable, void or invalid, such provision shall be modified so as to make it valid and enforceable, and as so modified the entire Agreement shall remain in full force and effect. No decision, action or inaction by LICENSOR shall be construed to be a waiver of any rights or remedies available to it.

None of the Intellectual Property or underlying information or technology may be downloaded or otherwise exported or reexported in violation of U.S. export laws and regulations. In addition, you are responsible for complying with any local laws in your jurisdiction which may impact your right to import, export or use the Intellectual Property, and you represent that you have complied with any regulations or registration procedures required by applicable law to make this license enforceable.

Chapter 1. Executive Summary

"Mixed Reality (MR) [https://medium.com/@northof41/what-really-is-the-difference-between-ar-mr-vr-xr-35bed1da1a4e], also referred to as hybrid reality, is the merging of real and virtual worlds to produce new environments and visualizations where physical and digital objects co-exist and interact in real time". MR has great potential in enhancing situation awareness and otherwise augmenting the experiences and performance of humans on the go.

This OGC Engineering Report summarizes information and findings collected during the *Mixed Reality at the Edge* Concept Development Study (CDS). Specifically, this report presents the significant findings concerning the state-of-the-art and potential of employing MR in modern systems, with a focus on discussing the state of needed interoperability and standards.

The term *mixed reality* was originally introduced in a 1994 paper by Paul Milgram and Fumio Kishino, "A Taxonomy of Mixed Reality Visual Displays." [1] [2].

1.1. Mixed Reality to the Edge Concept Development Study overview

The primary objective of executing an OGC Concept Development Study (CDS) is to assess emerging technologies and architectures capable of supporting possible future OGC Interoperability Initiatives and Standards activities. A CDS may also examine alternative mechanisms that enable commercial technology to interoperate to meet the requirements of the organization(s) that sponsored the CDS activity.

This CDS managed under the auspices of the OGC Innovation Program was [https://www.opengeospatial.org/ogc/programs/ip]. The activity was sponsored by **SOFWERX** [https://www.sofwerx.org/]. The goal of this CDS was to demonstrate to stakeholders the diversity, richness and value of new and emerging technologies for 3D and 4D processing, curation, and analytics in support of users in potentially disconnected computing environments. Specifically, consider from an interoperability perspective what data sources, technologies, analytics and associated Information Technology (IT) services are required for addressing the needs for the convergence of geospatial 3D modelling, simulation, and gaming integrated with machine learning (ML) for automated 3D workflows for such activities as enhanced decision support, mission rehearsal, and/or situational understanding. This includes the role of Artificial Intelligence (AI) along with Augmented Reality (AR) and Virtual Reality (VR) for enhanced visualization and decision support.

The Mixed Reality at the Edge CDS information gathering process consisted of a Request for Information [https://portal.opengeospatial.org/files/?artifact_id=82438] (RFI) phase and a subsequent Workshop held at SOFWERX in Tampa Florida. The RFI focused on seeking information on technologies that enable, enhance, and/or support automation of workflows that curate, discover, process, integrate, and publish (disseminate) 3D geospatial content to users at the edge, now and in the future. Specifically, there is a critical requirement to provide curated, actionable 3D content potentially fused and/or integrated from multiple sources to the user at the "edge". Users at the edge may have denied, degraded, interrupted, or limited (DDIL) internet access. There is also the

requirement to provide the right data at the right time to the right user. The timely sharing of resources and collaboration are important needed capabilities.

Fifteen organizations responded to the RFI. These organizations represented the commercial, government, and academic sectors. There were also responses from entrepreneurial startups, small companies, and very large enterprises. Workshop participants and attendees represented a wide range of organizations. In all, dozens of domains and information communities were represented, providing a broad sampling of input for this CDS.

1.2. The problem in a nutshell

The following statement was made at the Mixed Reality workshop, "We are not talking about 'the edge' where everybody just sits at the campsite and doesn't move. That is nobody's 'user at the edge' scenario. The user is moving across a piece of terrain that's complex and continuously changing. Not only is the terrain changing, the adversary is moving and changing or the thing that's causing angst is changing." This might be a wildfire driven by high winds, changes in weather as the user climbs a rock face, or the enemy firing at the user. Then the user's own configuration, their own logistics, about the environment at a specific time and location becomes critical to the user's decision making. This is *Context*! Combine constant change with the fact that the user at the edge may be under considerable stress and then human factors becomes critical to how change - including alerts - are communicated to the user.

So, who - and what - technologies and other factors can be leveraged to address the issues identified in the above use case scenario? The balance of this Executive Summary provides key findings, suggestions for research, and recommendations.

1.3. Key findings

A well-defined grouping of requirements and related recommendations coalesced based on the responses to the RFI. This information was enhanced by the workshop discussions. There is a common requirement that permeates all the responses and discussions: context. In a broad sense, context is about providing the right data in the right form at the right time to the individual at the edge. "Context" also incorporates the concept of knowing characteristics of the individual at the edge combined with the current (and changing) situation at the edge. This includes location and temporal awareness. There was general agreement that AI (Machine Learning, Deep Learning, and Machine Reasoning) has a critical role to play in processing, then understanding and providing the right data at the right time to the user at the edge.

Keeping in mind the general context framework, the following sections define the groupings of requirements:

1.3.1. Context, foundation data, and real time inputs at the edge

There were considerable discussions and suggestions related to how to best provision data to the user at the edge. A recurring theme is to provision all users going to the edge with foundation data prior to a mission start. This has several advantages: 1.) All users have the same foundation data and 2.) This helps mitigate the need to push large volumes of data to and around devices at the edge. However, the foundation data needs to be as up to date as much as possible and at the fidelity

required for the mission. Context again!

Once in the field, additional data (information) assets can be either fused with the foundation data or injected into the visualization pipeline. These tend to be more dynamic, time-sensitive forms of information that are crucial to the mission (e.g., intelligence, location tracking data, etc.). Context again! Typically, these additional data assets will be based on real time sensor observations. These sensors could be in-situ (e.g. weather stations) or dynamic (e.g. drones with on-board sensor systems). The questions then become how to best manage and process these real time sensor observations at the edge (as close to the point of need as possible) - especially in a disconnected environment. These edge cases pose a number of challenges, including not just processing power but also human factors implications (see below).

1.3.2. Speed

The ability to collect, process, and provision geospatial data for the user at the edge as quickly as possible was a common thread in many of the responses but especially in the discussions at the workshop. The speed - and quality - of processing was suggested for both backend server and client workflows. Numerous times participants mentioned that data should be ready and useable during the event and not just for mission postmortems.

Speed also again highlights the need to accommodate situational and/or environmental changes at the edge. As long as the data has to be developed by traditional means and then packaged and moved to the user (this usually takes too long), we have not really solved the challenge of what happens with the last-minute change(s) to the data that occur after the mission starts. This situation is being exacerbated by the fact that we now have the ability to collect ever increasing volumes of geospatial at ever higher fidelity.

There is general agreement that this is one area that AI can really have an impact - filtering new and existing content and then processing and provisioning that content based on the context of the situation that the user is experiencing.

1.3.3. Standards

Along with *context*, perhaps the most often mentioned requirement for evolving the ability to bring a high-quality mixed reality experience (and information) to the edge is standards. The role of standards cannot be overstated regardless of where the user is in the information/application provision pipeline. Even with existing standards and models for sharing geospatial data, a concerted community standards effort is needed to bring all important aspects together in support of mixed reality to the edge. This was very clear based on the range of standards mentioned in the RFI responses.

RFI respondents and workshop participants all agreed that standards in the following areas are critical:

- 3D and 4D models
- Application Programming Interfaces (APIs)
- Streaming formats
- Content formats that are lightweight

• Symbology

Related to the importance of standards was an agreement that well-known controlled vocabularies are required. The use of controlled vocabularies is of paramount importance in automated workflows, especially in the use of machine learning, and for packaging, sharing and deployment of ML models in "moving the process to data at the edge" type scenarios.

1.3.4. Training

Another recurring theme in the RFI responses and workshop discussion was training. However, training discussions covered a broad spectrum from knowing how to use backend home office applications, to pre-mission rehearsal training, to training on applications and data deployed to the edge.

Participant feedback for backend support services ranged from minimal training to requiring experienced people with special skills. The range of responses is dependent on the type of processing, analytics, or visualizations required. For example, developing applications based on Deep Learning (DL) for specific types of feature recognition or patterns requires skilled technicians with domain knowledge.

There was agreement that tools and applications used at the edge needed to be easy to use with light user training requirements. One suggestion that was discussed numerous times is to look at the human interfaces and training requirements for gaming systems when designing client (field deployed) tools and applications.

There is general agreement that training is tightly coupled with human factors

1.3.5. Human factors

Many participants mentioned that the community is evolving toward a synergistic human/Artificial Intelligence (AI) (aka human/machine) relationship in support of users at the edge. In this relationship, human factors become highly relevant. Part of this evolution requires designers, developers, and analysts to focus on how the human at the edge really uses the geospatial data and applications. To then provide unambiguous man machine interfaces targeted to an individual's current context requires a deep understanding of human cognition, cognitive burden, and spatial perception in potentially stressful situations.

A great example from the firefighting domain is while there must be integration across technologies, there is also the need to understand human beings' ability to capture and process information. For example, as a first responder showing up at a building on fire, there is a certain level of detail they need to know. Give the user too much information too fast and it is of no value. As an incident progresses, there may be a need for deeper and deeper pieces of information. It is more important to give users the right actionable information, when they need it. Context again! This will require integration of other technologies and other sources of geospatial data. For example, turn an AI application loose on a three-dimensional model of this building that will tells the responder how the air flow is moving through the building, so that he can best determine how the fire will behave.

Many participants also stressed that the human will remain "in the loop" in all stages of any

workflow. Participants strongly believe that the human computer (field technician) is the best present-day technology to recognize and make decisions to teach algorithms to automatically learn what certain features, such as clusters in point clouds, represent in the real world.

1.3.6. Bandwidth and updates

A key consideration for supporting users at the edge with the best geospatial content is the restrictions imposed by communications bandwidth. Often, these users are operating in limited or denied communications environments. When there is connectivity, careful consideration needs to be given to what prioritized content and/or alerts need to be uploaded to devices at the edge is critical. Setting priorities as to what is uploaded has to be based on situational and user context. In general, respondents and participants tended to focus on three key technology areas: 1.) Geosynchronization (aka Delta Updates) with priorities 2.) Research approaches to data caching, such as symbol encodings and content streaming approaches that significantly reduce transmission impact, and 3.) Analytical services provisioned at the edge rather than from some backend server. The latter requirement is further discussed in the Power at the Edge section.

Geosynchronization [https://portal.opengeospatial.org/files/?artifact_id=39476] is the process of communicating and synchronizing geospatial content updates between mobile devices at the edge and backend servers at a remote back office or operations center. Such updates are bidirectional. For example, a user at the edge may report the status of a road or bridge. Or the operations center may have critical updates from airborne sensors on the current extent and movement of a wildfire. In all cases, these updates need to be shared so that all parties have the same information base, enabling better decision making supported by enhanced collaboration. A key aspect of these geosynchronization activities is setting information exchange priorities. This is especially important in Denied, Disconnected, Intermittent and Limited bandwidth (DDIL) situations. Given the potential complexity of setting these context-based priorities, there is agreement that this is an area in which AI can provide major benefits.

In OGC Testbed 15 [https://www.opengeospatial.org/projects/initiatives/testbed15], there is a
NOTE Delta Updates work activity that is building on previous OGC geosynchronization work.

1.3.7. Security

Interestingly, security was not a hot topic of discussion. There may be a basic assumption that the communications networks are secure. There may also be an assumption that foundation data loaded to devices prior to mission deployment can be trusted. Finally, there may an assumption that geospatial content communicated during an event may be encrypted or in a binary format that is not easily hacked.

That said, there is a general concern about location and/or identity spoofing, such as a bad actor inserting false locations into a data stream. Given the criticality of location awareness and context, such spoofing is a major concern. This is why the internet community as part of the Next Generation 911 work spent considerable time looking at such threats to the integrity of the network and emergency calling (e.g. https://tools.ietf.org/html/rfc7378)

Additional discussion pointed out that due to a variety of concerns and potential issues at the edge,

paper maps and the discipline of map reading may never disappear.

1.3.8. Authoritative versus non-authoritative sources (Crowdsourcing)

This was an interesting topic of animated discussion at the workshop. There were also several points made in responses to the RFI related to authoritative (curated, trusted) geospatial data sources versus non-authoritative source, such as OpenStreetMap. A key point in this discussion is that from a usability or trust perspective. There is a clear distinction between crowdsourced map data (aka volunteered geographic information) and social media. Much crowdsourced location-based content is "self-healing" due to the number of people "measuring" the same phenomena. This is the Wikipedia model. Social media on the other hand can, at best, be taken with a grain of salt. Again, AI is believed to have a place in separating fact from fiction in social media sources.

As an indication of the growing importance of non-traditional open geospatial data sources, there is the NSG Open Mapping Enclave (NOME). This is one example of an attempt to take an open data source and fuse that content with a curated geospatial data store. The community is going to see incremental versions of private-public partnerships to provide the same capability.

A key aspect to enabling the trusted use of open sources is provenance. Overwhelmingly, the consensus is that capturing, curating, and using provenance and metadata information in automated workflows is critical. The reasons for the consensus vary but in general focused on trust, using the right data at the right time, determining whether data is fit for purpose, calibrating machine learning models, and tracking uncertainty and error propagation.

1.3.9. Power at the edge

A number of RFI responses and workshop discussion points spoke to requirements of "power at the edge". In reality, there are two types of power discussed: One has to do with the battery power of the device and the other has to do with moving analytical processing power from the backend server farms to devices at the edge.

With regard to power consumption, the primary issue is short battery life. This issue occurs when applications such as AR draw heavily on the device's battery. Further, as many AR/VR applications make heavy use of Graphics Processing Units (GPUs), devices can overheat.

At the same time, many participants spoke to the need of moving more analytical processing capability to devices at the edge. This is the other type of "power" at the edge. Such capabilities could be for processing real time sensor feeds with AI based algorithms and then fusing the new content with the foundation data. Moving processing power to the edge removes some of the reliance on the back office (e.g. command and control center) and helps reduce issues related to limited or no bandwidth.

The participants mentioned technologies that could be used to move processing power to the edge. These include mesh networks, mini-clouds based on a FOG architecture, and new hardware components such as Amazon's Snowflake and Snowball technologies.

1.3.10. Technology

Obviously, technology discussions permeated the entire CDS process. However, discussions and

responses to the technology question tended to focus on three key functional areas: 1.) The ability to provide delta updates (geosynch) in a DDIL communications or network environment 2.) Enhanced Machine Learning/Deep Learning (ML/DL) workflows to facilitate the processing, delivery, and visualization of enhanced 3D/4D content to users at the edge and 3.) New and enhanced tools for field data collection and processing of 3D content, especially indoors. These focus areas included recommendations on having "lightweight" encodings and highly efficient - and secure transmission compression. There was also discussion on being able to train and use ML/DL based workflows in the field.

However, there was broad agreement that while technology has evolved to the point where strong benefits to the user at the edge can be measured, we still have a long way to go.

1.3.11. AI/ML/DL

There was general agreement that the use of AI/ML/DL is critical to all aspects of supporting the user at the edge. However, there was also general agreement that the industry has a long way to go before AI applications are advanced enough to enable complete synergistic human/AI workflows with "trust" and information assurance. There needs to be a larger volume of curated data for algorithm developers to test with. Then, a "sandbox" of training data that allows developers to plugin to the data in an easier way and run on new data could be used.

There was a fair amount of discussion about where AI capabilities occur in workflows. Are they staged in the backend servers in the back office or are they available on devices (or meshes of devices) at the edge. Generally, the participants felt that research is needed to determine what AI capabilities are best staged in the back office and which AI capabilities are available at the edge.

For example, possible ML/DL analytics in the back office would enhance the value of 3D and 4D geospatial content in many areas, such as classification, ranking, object detection and semantic segmentation. ML can be used to process imagery to determine and classify the features within. Given a set of parameters ML can also be used to find 3D content relevant to a mission. Conversely, ML/AI capabilities at the edge could focus on analytics of real time sensor streams available to the user. For example, point clouds collected using SLAM should be coupled with ML software hosted on a Fog network to generate near real time indoor navigation maps.

1.4. Recommendations

The following recommendations were derived from RFI responses and interactions with attendees at the workshop.

- Recommendation 1: Engage in geospatial and related standards activities. Especially focus on the interoperability gaps and pain points, including geospatial standards for content models, controlled vocabularies, metadata, provenance, and visualization. Also focus on standard geospatial APIs that enable not just interoperability but also interchangeability of elements of workflows.
- Recommendation 2: In general, there is a sense that more research and technical capabilities are required to support more automated conversion and integration workflows with a strong trust and confidence in the provenance and quality of the final product. The use of AI to determine which content is fit for purpose based on context is an area for research.

- Recommendation 3: The present environment is plagued with many disjoint models and tools, with users manually passing outdated file schemas around. To make progress we will need to abstract away from legacy approaches to get to the next level of capability. A fresh new "3D/4D capability layer of data and applications" is needed so that we can incrementally and progressively move towards more "automated workflows" and improved 3D/4D quality and performance. Thus, core 3D/4D data models and open APIs need to be re-examined for this next-generation capability.
- Recommendation 4: Continue research in human factors, human cognition, and reducing cognitive burden with the focus on providing the optimal experience to users. This research should also include better understanding of the fact that users often have different specialized requirements for how they need to experience and interact with the same content.
- Recommendation 5: Research and development to develop approaches, best practices, algorithms, and applications that 1.) Reduce power consumption of the device at the edge and 2.) Provide additional processing capability at the edge. For example, investigate a Fog computing architecture that creates a local mini-cloud to run AI applications on real time sensor streams.
- Recommendation 6: Coordinated research and development activities to produce a shared, unified architecture vision and roadmap for achieving the desired level of interoperability and standards required for next-generation mixed reality experiences. This needs to be treated as a "living" document, and be incrementally updated as technology and requirements change. The scope of this architecture needs to cover the full expanse of end-to-end mixed reality production-operations lifecycles and workflows, and encompass all interoperability requirements and standards reflected in this report. This effort is crucial to achieving a cohesive unified capability across this mission space.

1.5. Prior-After Comparison

For many years, a variety of proprietary and consumer applications have provided basic mapping tools for users at the edge. A typical use case was a utility engineer out in the field that needed asset maps that they could then be updated based on observed properties. These tools were focused on static, 2-D maps with pre-programmed capabilities to add new content or update existing content. Once the user returned to their home office, the updates could then be uploaded and integrated into the master data store. Until updates to the master were completed, other users could not see the changes. Further, minimal 3D content was available for enhancing any visualization. Augmented Reality was not available. Mobile hardware capabilities simply could not support many of the emerging use cases demanded by users at the edge.

In the last few years, rapid technological changes in 3D content capture, big data analytics, geospatial information fusion, gaming, and hardware power have completely changed the landscape for what applications and capabilities can now be provided to users at the edge [3]. However, many of these advancements are uncoordinated, interoperability is lacking, and standards are often ignored.

Therefore, the goal of this CDS is to define a more holistic picture of the current state of technology and to help drive best practice use of current geospatial standards or to define requirements, architectures, and use cases for new geospatial/location interoperability standards.

1.6. Recommendations for Future OGC Work

Building on the results and the recommendations from the Mixed Reality to the Edge CDS, the OGC suggests the following strategy to achieve the goal of providing the right content at the right time for mixed reality data that can be easily shared with and among users at the edge.

Like other OGC Concept Development Studies (i.e., the Marine SDI CDS), the needs that can be identified for future work can be summarized as:

- 1. Provide stakeholders with appropriate access to the required geospatial data they need when they need it.
- 2. Allow different stakeholders, at different locations, to access the desired same content for different applications (multiple use and re-usability).
- 3. Allow for data exchange, especially for dynamic sensor data, in an appropriate, efficient and secure way.
- 4. To achieve these three objectives requires the continued and increasing use of OGC and other open standards. The use of standards identified in the RFI responses and workshops needs to be stress tested in an operational environment (see Pilot below).

Two other similar issues could be addressed in future work also:

- 1. The use of Mixed Reality in areas of limited bandwidth needs to be furthered explored. As noted in the Marine SDI CDS, it is possible that a significant portion of the base or core data can be prepared in advance and pre-loaded on mobile devices for field use. This can be accomplished by leveraging the GeoPackage standard. How this can be used to support the different user communities needs to be demonstrated and tested in real world exercises. In a possible future Pilot Project (below) and as as noted by the RFI Respondents concentrating on the lightweight data formats and symbology needed in low bandwidth environments is very important.
- 2. Careful attention to the correct application of metadata including provenance is needed. To support this requirement, proper metadata needs to be generated. As data is created, automated metadata generation techniques should displace costly manual techniques.

The proposed Pilot follows the OGC guidelines used in other programs, such as the Arctic Spatial Data Pilot (ArcticSDP). The goal is not to develop a different, new operational Mixed Reality to the Edge Spatial Data Infrastructure. Rather, the intent is to support the community by demonstrating the value of existing data and data infrastructures that already exist. The Pilot would provide a mechanism to pull the disparate sources together to demonstrate the current capability and issues concerning Mixed Reality to the Edge content processing workflows and provision.

The below steps that could be followed are documented and proven in the Arctic Spatial Data Pilot (ArticSDP):

Part 1: Assessment and Planning:

1. Project Initiation: Host key stakeholders to discuss the key issues identified in this CDS that have direct bearing on the form and structure of a Mixed Reality to the Edge content provision capability. Initiation includes outreach.

- 2. Conduct Stakeholder Survey: Survey organization interested in Mixed Reality to the Edge issues related to an infrastructure to better understand their needs and requirements. This would build on the work accomplished during the Mixed Reality CDS
- 3. Prepare Inventory and Assessment of Existing Mixed Reality to the Edge Portals, Applications, Data Sets and Databases. Inventory should cover formats, datums, metadata, standards used, etc.
- 4. Prepare an Implementation Strategy and unified Architecture. An Implementation Strategy should include a plan covering hosting and data sharing agreements combined with a security framework.
- 5. Execute a Pilot that demonstrates, in an operational environment, that the architecture can be implemented based on the identified strategies.

1.7. Document contributor contact points

All questions regarding this document should be directed to the editor or the contributors:

Contacts

Name	Organization
Carl Reed	Carl Reed & Associates
Tracey Birch	SOFWERX
Terry Idol	OGC

1.8. Foreword

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. The Open Geospatial Consortium shall not be held responsible for identifying any or all such patent rights.

Recipients of this document are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of the standard set forth in this document, and to provide supporting documentation.

Chapter 2. Acknowledgements

The OGC expresses its gratefulness to the sponsor of this Concept Development Study: USSOCOM. In addition, OGC wishes to thank SOFWERX for hosting the Mixed Reality Workshop. The OGC further wishes to express its gratitude to the following companies and organizations that provided excellent contributions in responding to the formal Request for Information that provide key content for this OGC Engineering Report.

- AGI-CESIUM
- Ball
- Clostra
- Cognitics
- Compusult
- ICT
- ImageMatters
- KaDSci
- Ordnance Survey
- PLANiT
- SensorUp
- University of Calgary
- US Army Engineer Research and Development Center and Cold Regions Research and Engineering Laboratory
- Army Geospatial Center
- US National Geospatial Intelligence Agency
- Virginia Modeling, Analysis, and Simulation Center (VMASC)

The OGC also thanks the following Panelists and their organizations for supporting the Mixed Reality Workshop:

- General Doug Brown USA (ret)
- John Simmins (EPRI)
- Chris Tucker (OGC Board)
- Mark Abrams (ExtremeGeo)
- Robert Austin (Enterprise GIS Tampa)
- Stu Bradin (COL, retired SOF)
- Talbot Brooks ((Delta state University)
- Patrick Cozzi (Cesium)
- David Finnegan (ERDC)
- David Graham (CAE)

- 1LT Lichtenwald / Tim Rodabaugh
- Dan Maxwell (KadSci)
- Joseph Mullins (US Futures Command)
- Harry Niedzwiadek (ImageMatters)
- Kumar Navulur (DigitalGlobe)
- George Percivall (OGC/CTO)
- Trent Tinker (Hexagon)

Chapter 3. Terms and definitions

The following is a list of terms with definitions that may help the reader better understand the content of this Engineering Report.

• 3D Mesh

The structural build of a 3D model consisting of polygons. 3D meshes use reference points in X, Y and Z axes to define shapes with height, width and depth.

• Cloud Computing

Paradigm for enabling network access to a scalable and elastic pool of shareable physical or virtual resources with self-service provisioning and administration ondemand.

• Edge Computing

Distributed computing in which processing and data storage takes place at or near the edge, where the nearness is defined by the system's requirements.

• Fog Computing

A term created by Cisco that refers to extending cloud computing to the edge of an enterprise's network. Also known as Edge Computing or fogging, fog computing facilitates the operation of compute, storage and networking services between end devices and cloud computing data centers.

• Mixed Reality

Mixed reality is the result of blending the physical world with the digital world.

• OBJ

A geometry definition file format first developed by Wavefront Technologies for its Advanced Visualizer animation package. The file format is freely and publicly available and has been adopted by other 3D graphics application vendors

• SLAM

Simultaneous localization and mapping (SLAM) is the computational problem of constructing or updating a map of an unknown environment while simultaneously keeping track of an agent's location within it.

3.1. Technical Tools, Applications, and Platforms

The following table provides information on the various technologies, tools, and platforms mentioned in the responses to the RFI. Mention of these technologies is informational only and is not an endorsement.

Name	Description
Blender	A free and open source 3D creation suite. It supports the entirety of the 3D pipeline—modeling, rigging, animation, simulation, rendering, compositing and motion tracking, even video editing and game creation.
Entwine	A data organization library for massive point clouds, designed to conquer datasets of trillions of points as well as desktop-scale point clouds.
FME	Safe Software data integration platform
GDAL	Geospatial Data Abstraction Layer. An Open Source translator library for raster and vector geospatial data formats.
OpenGL	Open Graphics Library is a cross-language, cross-platform application programming interface for rendering 2D and 3D vector graphics.
Matterport	3D camera and media platform that provides an end-to-end 3D media solution.
PDAL	Point Data Abstraction Library. A C++ BSD library for translating and manipulating point cloud data.
Plasio	A project by Uday Verma and Howard Butler that implements point cloud rendering capability in a browser.
Potree	A free open-source WebGL based point cloud renderer for large point clouds, developed at the Institute of Computer Graphics and Algorithms, TU Wien.

Name	Description
REVIT	Revit is building information modelling software for architects, landscape architects, structural engineers, MEP engineers, designers and contractors.
Snowball edge	Snowball Edge is a data transfer device with on-board storage and compute power that provides select AWS services.
Unity	For building high-quality, cross platform 3D and 2D games and deploying them across mobile, desktop, VR/AR, or console visualization platforms.
UnReal	Unreal Engine 4 is a suite of development tools made for working with real-time technology. Application use ranges from cinematic experiences to high-quality games across PC, console, mobile, VR and AR.
WebGL	WebGL is a JavaScript API for rendering interactive 2D and 3D graphics within any compatible web browser without the use of plug-ins.

3.2. Standards and Specifications Mentioned in the Responses

The following table provides information on the many standards and specifications mentioned in the responses to the RFI. For the purposes of this document, a document and/or technology is termed to be a standard if it has gone through a formal review and approval process in a Standards Development Organization (SDO). Examples of SDOs are the OGC, IETF, OASIS, and ISO. A document is termed a "specification" if that document has not been through an accredited review and approval process.

Name	Description
3DPS	OGC 3D Portrayal Service Interface standard.
3DTiles	An OGC Community Standard for streaming large heterogeneous 3D geospatial datasets.
ARML	The OGC Augemented Reality Markup Language (ARML) allows users to describe virtual objects in an Augmented Reality (AR) scene with their appearances and their anchors (a broader concept of a location) related to the real world.

Table 2. Standards mentioned

Name	Description
CDB	The OGC CDB Standard defines a standardized model and structure for a single, "versionable", virtual representation of the earth.
CityGML	OGC Standard that defines an open data model and XML-based format for the storage and exchange of virtual 3D city models.
GeoPackage	an open, non-proprietary, platform- independent and standards-based data format for geographic information implemented as a SQLite database container.
gLTF	(GL Transmission Format) is a royalty-free specification developed and maintained by Khronos for the efficient transmission and loading of 3D scenes and models by applications.
InDoorGML	OGC Standard that specifies an open data model and XML schema for indoor spatial information.
I3S	An OGC Community Standard for Scene Layer services and packages. A single I3S data set, referred to as a scene layer, is a container for arbitrarily large amounts of heterogeneously distributed 3D geographic data.
Layered Terrain Format (LTF)	A specification designed for efficient environment data interchange.
SensorThings	OGC Standard providing an open and unified framework to interconnect IoT sensing devices, data, and applications over the Web.
WCS	OGC Web Coverage Service Interface Standard specifies a simple HTTP interface for access to coverage data in forms that are useful for client-side analysis and rendering.
WFS	The OGC Web Feature Service Interface Standard defines a set of interfaces for accessing geographic information at the feature and feature property level over the Internet.

Name	Description
WMS	OGC Web Map Service. The WMS standard specifies a simple HTTP interface for requesting geo-registered map images from one or more distributed geospatial databases.
WPS	The OGC Web Processing Service Interface Standard provides rules for standardizing how inputs and outputs (requests and responses) for geospatial processing services, such as polygon overlay.

3.3. Abbreviated terms

NOTE: The following is a list of the abbreviated terms and the symbols necessary for understanding this document.

- AGC Army Geospatial Center
- AGE Army Geospatial Enterprise
- AI Artificial Intelligence
- AWS Amazon Web Services
- AR Augmented Reality
- BIM Building Information Model
- CDS Concept Development Study
- CRREL Cold Regions Research and Engineering Laboratory
- CRS Coordinate Reference System
- CSW OGC Catalog Service Web
- **DEM** Digital Elevation Model
- DDIL Denied, Disrupted, Intermittent, and Limited
- DL Deep Learning
- EO Earth Observation
- GPKG GeoPackage
- GRiD Geospatial Repository and Data Management System (ERDC)
- IARPA Intelligence Advanced Research Projects Activity
- JWICS Joint Worldwide Intelligence Communications System
- LAS LIDAR Data Exchange File
- LTF Layered Terrain Format
- ML Machine Learning
- MSTAR Manportable Surveillance and Target Acquisition Radar
- NFPA National Fire Protection Association

- **OWT** One World Terrain
- SAR Synthetic Aperture Radar
- SIPR Secret Internet Protocol Router
- SLAM Simultaneous localization and mapping
- SSGF Standard and Shareable Geospatial Foundation
- STE Synthetic Training Environment
- STRI Simulation Training and Instrumentation
- SWaP Size, weight and power
- VMASC Virginia Modeling, Analysis and Simulation Center
- VR Virtual Reality

Chapter 4. Overview

This OGC Engineering Report consists of the following main sections.

Section 1 provides an Executive Summary. This summary includes an overview of the Mixed Reality to the Edge (CDS). This section describes the situation prior to the testbed and discusses the requirements set by the sponsors.

Section 2 acknowledges the contributions of the RFI respondents, workshop panelists, and others involved in this CDS activity.

Section 3 provides a list of terms and definitions, a table of standards mentioned in the ER, and a list of abbreviations used in this ER.

Section 5 provides summaries of the characteristics of the organizations responding to the RFI and participating in the workshop.

Section 6 provides summaries of responses to the RFI Architecture Questions combined with workshop discussions.

Section 7 provides summaries of Responses to the Data Analytics Questions combined with workshop discussions.

Sections 8 provides summaries of Responses to the Scenarios and Use Cases Questions. The summaries are augmented with input from the CDS Workshop.

Section 9 provides summaries of Responses to the Operation and Organization Questions. The summaries are augmented with input from the CDS Workshop.

Section 10 provides summaries of Responses to the Technology and Applications Questions. The summaries are augmented with input from the CDS Workshop.

Section 11 provides summaries of responses to the Other Factors Questions. The summaries are augmented with input from the CDS Workshop.

Annex A provides clarifications and additional information content regarding responses to each of the questions posed in the RFI.

Chapter 5. RFI Respondent Characteristics

This section provides summary information regarding the respondents to the Mixed Reality at the Edge Request for Information. A list of the organizations that responded to the RFI is found in the Acknowledgements.

5.1. Type of Organization

Question 4.1.2: What organization are you affiliated with and what is your role dealing with 3D/4D data curation, analytics, and provision?

RFI respondents and workshop participants represent the following organizations and/or organization types.

- Entrepreneurial startup.
- Government-owned limited company.
- Private/Commercial companies
- Universities and affiliated research centers.
- US Federal Government (GeoInt community)
- US Army Research and Development
- US DoD

With regard to mixed reality at the edge applications and requirements, their roles include the following:

- Principal investigator gaming and simulation
- Company President help customers with technology transition
- SETA strategic investment and transition support to S&T and GRiD
- Lead software engineer
- Work OGC standards and help customers integrate OGC standards to evolve technology
- Lead team in the acquisition, processing, storage, distribution, and application of 3d terrain datasets for military applications.
- Chief Scientist focused on 2D/3D/4D geospatial architectures and solutions.
- Senior engineer working 2D/3D/4D geospatial architectures and solutions
- President and Chief Scientist for firm executing research relating to the collection, curation, and use of many types of data, including 3 and 4D data.
- President of company focused on ML and building measurement, BIM, CAD and AR integration.
- Research & Innovation Scientist
- Lidar Imagery Scientist Leading several 3D data initiatives

5.2. Provider or Consumer?

Question 4.1.3: Are you a provider or a consumer of technology, data, services, and so forth?

The majority of the organizations that responded to the RFI are both providers and consumers of 3D technology and 3D data. A few organizations are purely providers of 3D data and/or services. A few others are primarily providers of 3D services and data, including at the research level.

- Provider of high fidelity 3D terrain reconstructions using COTS drones and photogrammetric software.
- A provider of increasingly automated classification and segmentation processes that use machine learning.
- Largely a consumer of geospatial data for use in multiple forms of models and simulations as well as a researcher in approaches for advancing the state of the art and practice intended to improve the interoperability of collaborating information systems that provide command and control, intelligence, and simulation support to headquarters staffs and operational personnel across many domains.
- Provider of geospatial data and services but also a consumer of content from mobile mapping applications and volunteered geographic information.
- Technology provider of machine learning and deep learning solutions as well as network security.
- Consumer of data, bridging research initiatives to the analytic workforce.
- Provider of 3D data acting on behalf of the GeoINT Foundation community leader, NGA.
- As an end-to-end full lifecycle solution provider, both a producer and consumer of technology, data and services.
- A provider of technology and services.
- Both a provider and consumer of 3D technologies, data and services.
- Vast experience in taking 3D data, storing it in a GeoPackage, and rendering the content in an Android mobile app.
- Convert of 2D and 3D data, such as CDB shapefiles and CityGML, to produce GeoPackages containing 3DTiles.
- Provider of professional services providing solutions for measuring buildings and creating CAD and BIM of those buildings for a myriad client
- Provider of services and technology and also look at how to integrate with existing standards and data sets

5.3. Communities your organization Supports

Question 4.1.4: What user community(ies) do you work in or represent?

- Arctic research and community outreach
- Army (Synthetic Training Environment Cross Functional Team

- Army Research Lab
- Army Special Operations Forces
- Army Geospatial Center)
- Federal Bureau of Investigation
- First responders
- Internet technology companies
- Investment groups
- Local government
- Marginalized communities
- Marine Corps
- Naval Special Warfare
- National Geospatial Intelligence Agency (NGA)
- Outdoor leisure consumer community
- Smart cities
- US Border Patrol (USBP)
- US DoD
- US Department of Energy
- US Department of Homeland Security
- US Department Homeland Security First Responders' Group
- US Geological Survey (USGS)
- US Immigration and Custom Enforcement (ICE)
- US Intelligence Community
- US Office of Naval Research
- Utilities
- IoT (Internet of Things), arctic research and community outreach, smart cities, simulation and modeling communities

Chapter 6. Summary of responses to the Architecture Questions

This section of the Engineering Report summarizes the responses to each of the Architecture questions presented in the Request for Information (RFI). The summaries are augmented with input from the CDS Workshop.

6.1. Response to Question 4.2.1. Automated processes - general

Do you use any automated processes for analyzing, modelling, and/or visualizing 3D/4D content in your organization?

The majority of the respondents are using some level of automated or semi-automated processing of 3D content, including curation. A variety of tools are being used as part of the automated or semi-automated workflows. However, there were numerous statements about the current immature state of software and machine learning. Comments in that regard ranged from "we have a long way to go" to "we have developed our own toolset that we trust". In all cases, existing tools such as GDAL and PoTree, are used in all workflows.

Most of the automated and semi-automated workflows focus on:

- Photogrammetry
- Processing of point clouds including some level of feature extraction
- Generation of 3D content, such as buildings or other digital models
- The use of existing OGC standards, such as CDB, CityGML and GeoPackage, in these workflows.

Several respondents also pointed to the need to keep the human in the loop. There needs to be a symbiotic relationship between human and machine. In terms of generating fully curated data, AI can only go so far. A human operator is required in many cases to monitor the process and validate the final results.

Three organizations did not respond to this question.

Some of the detailed responses include:

- Using automated processes for image acquisition via an autonomous photogrammetric capture system, reconstruction with cloud-based COTS photogrammetry software, semi-autonomous segmentation, classification, and attribution of terrain features.
- Use LTF processing tools as well as other tools for generating and visualizing 3D data. These tools are also used for interoperability experimentation and analysis. Use extended mixed reality technologies to create multi-sensory (Smell, touch, site) user experiences.
- Using automated processes for generating digital surface models (DSMs) and semi-automated process for generating digital terrain models (DTMs). Building height values are generated through an automated procedure using DSM data. Current research projects are investigating

methods of manual and semi-automatic 3D building modelling at Level of Detail (LOD) 2 using optical imagery, planimetric building polygons, DSMs and DTMs.

- Have not yet discovered any automated processes that perform as well as those developed in house for specific applications. Certain ML/DL-specific libraries now assist with data representation so that low level curation is no longer quite as urgent.
- Currently leveraging several initiatives to derive 3D point cloud data from satellite imagery.
- Developed and host selected 3D decision tools and are exploring plans for additional tools. Applications accept data, as is, from varied providers and do not modify the data delivered.
- Automation is critical to making this technical means viable in many time-sensitive, missioncritical situations. Development is on an incremental path of increasing automation. But, just to be clear, there are critical roles for users to play too. For example, consider how difficult it is for a machine to autonomously validate the integrity, coherence, and usability of a 3D/4D model. [This is beyond the scope of modern AI. There is not yet have general intelligence and 3D/4D sense-making working on a machine.]
- Currently using services and automated processes to analyze, model and visualize 3D / 4D content. As part of the OGC Indoor Navigation pilot, implemented a custom-built process to retrieve CityGML, IndoorGML and LiDAR data from a CSW, convert it to 3DTiles and store the data in a GeoPackage. These data can then be visualized, and indoor routes can be planned, in a mobile or desktop-based preplanning client by a first responder.
- Developing innovations to increasingly automate the collection and creation of as-built BIM and further create a digital twin of the built world.

Tools and toolsets mentioned in the responses include:

VTAC Epic, GDAL, HTC Vive - A VR platform, LTF, Plasio, Potree, Unity, Unreal

6.2. Response to Question 4.2.2. Automated Processes for content integration

Do you use any automated processes for integrating 3D/4D content from multiple sources into AR/VR applications?

Responses to this question ranged from "we do not leverage AR/VR applications or have automated processes to deliver content" to "we are using automated processes to perform the integration process". Technical solutions range from custom (developed in house), to COTS to open source, or some combination thereof. In general, there is a sense that more research and technical capabilities are required to support more automated conversion and integration workflows with a strong trust and confidence in the provenance and quality of the final product. Three organizations did not respond to this question.

Detailed responses included:

• Using automated processes for integrating and converting 3D content from multiple sources to a desired file format and quality of mesh decimation prior to ingesting into AR/VR application. Currently developing automatic dataset conversion processes from 3D mesh to AR/VR compatible formats.

- Currently most of these efforts are accomplished with human in the loop construction concepts. Currently writing translators to do experimentation on the use of geospatial data in different formats in a federated fashion. In most cases 4D is executed using Unity and UnReal.
- Using Unity for integrating content into AR/VR apps but do not currently use any automated processes.
- Using FME to convert source DTM raster files into 3D meshes that can be subsequently ingested into 3D tools such as Blender or Unity (and potentially into VR/AR applications). Use FME to convert 2D geospatial data from native formats into GeoJSON, which is then imported into Unity (and thus transformed into '3D' objects) using a custom importer and used within custom VR applications. Have used Blender Python scripts to instantiate 3D models of street furniture based upon data prepared in FME. Have incorporated streamed data into Unity VR/AR applications, using the Mapbox Unity SDK.
- Have not yet discovered a model that allows full integration. However, increasingly sources of 3D/4D content provide open APIs that are useful in assisting in in-house integration (gdal.org).
- Continually coordinating with customers for serving community requested AR/VR solutions. The intention is to serve both the HUD 3.0 and ATAK communities for AR, and the VR communities via GRiD delivery of CDB and OBJ run-time executable data. GRID accepts input of 3D data of different formats converting on ingest as necessary.
- Development is on an incremental path of increasing automation, where multi-source processing threads contribute to a seamless mixed reality scene (e.g., AR on handset).
- Using several custom processes for integrating 3D / 4D content from various sources (weather, sensor, geospatial, LiDAR) into an Android mobile application. Have built a GeoPackage wizard into products that allows users to select data from various sources (OGC service, ArcGIS database, local file) and package them into a GeoPackage container.
- Some experience testing the OGC IoT standards using real-world scenarios defined by first responders. In this project, a scalable information system architecture combining the features of the Internet of Things and OGC Sensor Web Enablement is developed for first responders. Integrating AR/VR with real-time IoT data from multiple IoT systems is one of the research interests.

Technology used/mentioned in responses to this question include: Blender, FME, GDAL, MapBox, Unity, UnReal

6.3. Response to 4.2.3 - What are key technology components?

What do you think should be the key technology components (e.g., standards, networks, clients, web services, data storage) for more automated workflows for provision and use of 3D/4D geospatial content to a user at the edge?

There was consensus that all of the technology components identified in the question are important. As one respondent succinctly stated, "All of the above". Of all the technology components mentioned, standards were consistently mentioned has extremely to very important. For example, "If we were to prioritize the most important it would be in the area of standards." This focus on standards was made regardless of whether the respondent was speaking of integration, analysis, visualization, web services, data source, or APIs. A number of respondents mentioned point clouds and the need for better automated processing tools combined with network bandwidth and standards. Broader use of the cloud and cloud computing is another consistent technology thread. One organization did not respond to this question.

Detailed responses to this question included.

- Cloud-based storage, processing, segmentation/classification/attribution, dissemination.
- OGC compliant and open-source, engine-agnostic formats. Provision of OGC Compliant data and use of cloud and web services to seamlessly deliver geospatial content on the edge.
- For Mixed Reality at the edge, rich standards would allow for the construction of platform agnostic tools.
- Three key technologies are: Cloud data storage, High bandwidth and low latency network connections (e.g. 5G) for data upload and download,
- Automated workflows for data capture, feature extraction and feature delivery with low latency to the users at the edge.
- Standards-based point cloud data stream to open source data storage. In addition, a "simple" fully DL-based pipeline from PCD to rapid and accurate 3-D image segmentation could, in principle, be implemented at the cloud level and accessed by edge users. This technology is very, very close to feasible. The use of such a pipeline would be optional, and an open API would allow fast incorporation of data into ML/DL workflow and application development.
- A larger trade study that stress tests each component with the user requirements in mind to identify gaps and areas that require greater investment.
- If done prior to combat operations, network bandwidth should be considered but should not be essential for transferring geospatial data to or at the edge.
- Use of lightweight data formats and symbology are critical in a small bandwidth environment. Provisioning troops with foundation data prior to a mission start is helpful by mitigating a need to push large volumes of data around the battlespace. Snowballs at the edge for data storage is the current preferred solution.
- The focus should be on defining a unified, open 3D/4D interoperability framework, consisting of a core abstract model for 3D/4D and the open APIs for major functional components and processes.
- Key technology components for more automated workflows for provision and use of 3D / 4D geospatial content to a user at the edge would be 3D data representation standards and data storage standards, such as GeoPackage for offline data storage
- The proliferation of point cloud scanning (LIDAR and SLAM technologies) are all competing to create point clouds of building geometry. None of them are automatically creating BIM. Point cloud scanners, whether stationary or using SLAM, coupled with Machine Learning software hosted on a Fog network. The human computer (field technician) is the best present-day technology to recognize and make decisions to teach algorithms to automatically learn what certain clusters of point clouds represent in the real world. In other words, field technicians will identify objects such as light switches or doorknobs and teach a program to relate that cluster of point clouds to the real-world object.
- A data geosynchronisation capability that ensures that the latest (and best) geospatial content is

available at all times to the user at the edge. This is known as Delta Updates in OGC Testbed 15.

• In an AR context, integrating real-time information into the static 3D environment provides enhanced situational awareness. Integrating both sensing and actuating (tasking) will empower users in the AR environment to interact with the physical world (actuating/tasking) and see changes through real-time sensing.

Standards and technology platforms mentioned in this section: GeoPackage, OpenGL, SensorThings, SLAM, Snowball Edge, WebGL

6.4. Response to 4.2.4 - Controlled Vocabularies

How important is the use of controlled vocabularies (e.g. ontologies) in automated workflows especially in the use of machine learning as related to automated processing and curation of 3D/4D content?

The consensus is that controlled vocabularies are critical to both automated workflows and machine/deep learning. Only one response stated that controlled vocabularies are not necessary. Most respondents view controlled vocabularies as critical or crucial to developing and deploying automated workflows integrated with ML/DL. However, the use cases varied in terms of why controlled vocabularies are important. The more detailed responses below show the various (but related) use cases. There is also an undercurrent in the responses that developing and maintaining controlled vocabularies is a hard task. Four organizations did not respond to this question.

Detailed responses to this question included:

- Crucial. These shared vocabularies allow for deconfliction during processing and curation of data. We will be using Standardized Sharable Geospatial Foundation (SSGF) for DoD and OGC users, as well standardized terminology developed for the Army's Synthetic Training Environment (STE) Cross Functional Team
- In the near term this is essential to automating interoperability. Without this discipline, serious inconsistencies in data will surface at time of use; which is the worst possible time for that to occur. It presents great risk to tactical operators.
- Ontological classes linked to data concepts through semantic mapping is extremely useful for machines in interpreting heterogeneous big data.
- We do not believe that these are currently critical. Indeed, they are quite distinct from machine learning, algorithmic approaches. However, they may have a future role.
- Very valuable for enabling machines to communicate with each other. For example, when we rely on training data for various ML approaches, having a consistent body of curated content with strictly enforced ontologies determines the success of the model.
- This has not been part of the organization's regular S&T work space (yet). However, this makes sense. At a minimum, approaches like the use of ontologies should help to standardize the output of 3D curated information.
- Controlled vocabularies are critical to machine understanding, and therefore to automation. This needs to be strategically introduced, first on the production side, and then incrementally introduced into DDIL settings. Semantic integrity needs to be baked into 3d/4d workflows from inception.

- The use of controlled vocabularies is of paramount importance in automated workflows, especially in the use of machine learning, and for packaging, sharing and deployment of ML models in "moving the process to the data" scenarios.
- In order to make the geospatial data AI or automated workflow ready, controlled vocabularies are critical. This includes not only relatively static geospatial information (building models, etc.), but also the observed properties and unit of measurement in an IoT context.

6.5. Responses to 4.2.5: Provision of interoperable, automated 3D workflows

What do you think is the best way to provide interoperable, automated 3D workflows that support the user at the edge?

A number of detailed responses were provided for the question. There were a range of suggestions but many of the respondents mentioned the need for 1.) consistent 3D models, 2.) use and enhance existing standards, such as the OGC GeoPackage standard and 3.) ease of use for any applications and tool sets provided to the use at the edge.

In general, the best way to provide interoperable, automated 3D workflows that support the user at the edge is to minimize the number of steps required for the user to collect and upload/transfer raw photogrammetric data to an automated terrain development pipeline via a web service and/or cloud repository. The uploaded raw data is processed into an optimal format and attribution on the cloud prior to being exported to the end user at the edge as an update or patch they can download/transfer to their local device whenever internet connectivity is viable, or a peripheral containing processed terrain data is otherwise physically provided to transfer data onto their device.

The following are additional detailed responses for recommendations and identification of issues related to this question.

- Workflow in this context is a confusing term. The user at the edge should be able to intuitively consume 3D and 4D data (models) as part of their workflow vs creation of a special workflow optimized for the 3d data. Our vision would be that the 3D & 4D representations be immediately available and integrated into the performance of their primary tasks. The technical infrastructure, data models, and Concept of Operations (CONOPS) to do this well do not yet exist. Developing them will require a period of investment in experimentation. Moreover, given the continuing advances in technology, any CONOPS that is developed should accommodate change.
- Recommend that 'capture once, use many times' philosophy for 3D data to ensure consistency and reduce costs.
- Recent developments in Deep Learning (DL) have provided roadmaps for direct "data to DL" without intermediate and time-consuming data preprocessing. Such pipelines would benefit greatly from automated (DL-based?) ontological parsing prior to insertion into data repositories, removing end DL-development of the burden of data preprocessing in the form of pretraining for basic semantic classification.
- We need a larger volume of curated data for algorithm developers to test with. Then, we could use a sandbox that allows developers to plugin to the data in an easier way and run on new

data. Finally, we need automated methods to test the outputs to report back to developers and assist investors in making better decisions on who is making progress and should continue to receive more funding.

- Provide capabilities at the edge in a physical environment that are not starved for SWaP. More specifically, we envision a family of GRiD instances wherein GRiD Environment lives within the NGA cloud domain, GRiD Tactical will exist at the BCT echelon, and GRiD X will be in place at the tactical edge to support forward actions. Furthermore, direct connectivity to the myriad organic collectors such as Unmanned Aircraft Systems (UAS), mobile, Unattended Ground Sensors (UGS), and Intelligence, Surveillance, Reconnaissance (ISR) Sensors, as well as connectivity to relevant Intel Broadcast System (IBS) feeds are crucial for near real-time Situational Awareness.
- Must have tools and processes that converge on coherent scenes that can be packaged and disseminated to DDIL settings. So, for example, OGC's GeoPackage encoding standard needs to be extended to provide more comprehensive 3D support.
- The best way to provide interoperable, automated 3D workflows that support the user at the edge is to package large 3D / 4D content into a robust 3D data model such as 3DTiles, package this data into a container, such as GeoPackage, and allow offline syncing or online download of this data to the device which is used at the edge. Smaller data payloads, such as status and real time situational data, can be streamed to the user based on network availability.
- Equipping users at the edge with FOG computing power will be necessary to meet user at the edge processing requirements as well as support AI based applications.

Standards mentioned in the responses to this question include 3DTiles, GeoPackage and CDB.

Two organizations did not respond to this question.

6.6. Responses to 4.2.6 - Provenance and metadata

How important is provenance and metadata in these automated workflows?

Overwhelmingly, the consensus is that capturing, curating, and using provenance and metadata information in automated workflows is critical. The reasons for the consensus vary but in general focused on trust, using the right data at the right time, determining whether data is fit for purpose, calibrating machine learning models, data reuse, and tracking uncertainty and error propagation. Four organizations did not respond to this question.

Detailed responses to this question included:

- Provenance matters in the classification, segmentation, and attribution of geospatial data, as each area of interest provides unique features that must be processed. Metadata, however, is vital for the categorization of terrain data.
- If we are thinking about system interoperability then we should be platform agnostic e.g. we do not care how the information was created. From a trust standpoint, understanding the quality (accuracy, bias, and timeliness) is absolutely essential.
- Very. Trust in when, where and how data has been captured and generated is vital to assure confidence in automated processes.

- Provenance and metadata provide a basis for comparability and repeatability of ML experiments. Extracting, storing and managing metadata and provenance information of common artifacts of machine learning (ML) development is of great importance, particularly when curating datasets, and assisting with prediction and evaluation during training.
- Extremely important. It is critical we understand the inputs in a way that helps inform which input data works best. When we lack appropriate provenance and metadata we cannot make more informed decisions on data (input) quality. Similarly, outputs must keep data on the associated processing pipeline to inform the end user.
- Metadata is critical for tracking pedigree, conflation, uncertainty and error analysis. We have witnessed its absence over the years preventing our ability to carry out these critical functions. Additional provenance is recommended to capture all phases of processing pre-product delivery. Metadata we can capture we retain.
- Critical to integrity of automated processes.
- Provenance and metadata are crucial to record items, such as where the data originated, who created the data, etc., and to outline the historical timeline of the data. Metadata needs to be updated at each stage of the workflow so that there is accountability and traceability throughout the workflow.

6.7. Responses to 4.2.7 - Provenance and the end user.

How important is it for the end user to have access to provenance and metadata information?

Overwhelmingly, the consensus is that providing provenance and metadata information to end users is important. However, the additional consensus is that the systems and applications providing the information to the end user should filter the metadata and provide only those elements of importance to the end user at the edge. This is to prevent information overload. Four organizations did not respond to this question.

Detailed responses to this question included:

- Users on the edge may be interested in some types of metadata, though not most. For example, they may want to know time and resolution metadata, though probably not sensor type, etc.
- Provide metadata access to end users but in most edge situations we would want automation to assume most of the burden associated with using metadata.
- Key importance that end users can trust they data they are using, potentially in mission-critical contexts.
- Analysts and researchers have a high need for access to provenance and metadata information. The former (analysts) rely on such information in order to make decisions regarding the relative importance or veracity of their data.
- Some advanced users might be looking at the source data provenance and metadata, but most end users perhaps could use simple color codes: If the data is authoritative and trusted it could be labeled "green" while less trusted are orange and unknown datasets "red".
- It is useful for users in data filtering / data discovery. It would allow developers to calculate error and uncertainty, and that in turn provides users with a level of comfort on how their mission can best be accomplished.

- Critical to both producers and users when it comes to understanding and resolving data fusion/conflation/integrity issues.
- The end user needs to be able to access this provenance and timeline information to be able to validate the source of the data, identify the changes to the data and access the current status of the data.

6.8. Responses to 4.2.8 - Does your organization have a ML/AI plan

Does your organization have an ML/AI plan for geospatial data analytics? If so, please provide a brief description if you are able.

Again, there were detailed responses to this question. The majority of organizations responding to this question either currently have some ML/DL capabilities or are planning on developing and deploying these capabilities in the near future. A range of applications were specified, including:

- Line-of-sight analysis and automatic routing calculations based on line-of-sight threat levels;
- Developing an "AI sandbox" for users to determine optimal mission rehearsal strategies.
- Identify objects of specific tactical and actionable interest to the Army.
- Primary focus is on producing and exploiting object-based and activity-based intelligence.
- Trajectory mining: analyzing moving feature trajectories with data mining algorithms in order to clean noises, detect anomalies, segment trajectories, assign meanings, detect stay points and co-movements, etc.
- IoT data stream analytics: aggregate, filter, enrich, and process real-time geospatial data streams from multiple IoT systems.

Special Note: The Ordnance Survey of Great Britain (OS) has a major commitment to ML/DL. OS has various teams which are developing geospatial data analytics for various aspects of remote sensing, feature extraction and change intelligence. This is partly driven by a requirement to automate and improve consistency in conventionally manual flowlines, and also to generate data that is implicit in raw observational data but is too costly to extract manually or semi-automatically. Their aim is to use machine learning to automate data extraction and also to use deep learning to derive meaningful geospatial semantic relationships between topographic objects. They have a trained a deep network to infer and discover new ways of describing the landscape. This has culminated in a program of testing their networks for new features they can help them identify, such as roof attribution. The OS is now experimenting with different methods of scaling the deployment of their deep learning models. In a recent project they trialed multiple methods of deploying their deep learning models at national scale and now have a validated methodology for creating nationally-consistent datasets as well as a robust rules-based approach to change detection in geospatial data that is close to being part of their production process.

There is an Machine Learning thread in the OGC Testbed-15 activity. One aspect of this work is to provide an ML model to assess, rank and catalogue geospatial data from various online data providers. This machine learning model will allow a user to create a CSW catalog of content based on provided metadata attributes. The ML model will determine the best fitting services based on

the metadata provided and populate a new CSW catalog accordingly.

Two organizations did not respond to this question.

6.9. Responses to 4.2.9 - Geospatial Standards

Currently, what are the key geospatial standards you use to access and disseminate 3D geospatial content and/or services to users at the edge?

In line with the consensus response to 4.2.3 that standards are critical to successful automated workflows, machine learning, and content integration, the responses to this question focused on which standards and standards organizations are important to supporting users at the edge.

The Open Geospatial Consortium (OGC) was the most mentioned Standards Development Organization (SDO). The International Organization for Standardization (ISO) Technical Committee (TC) 211 and the International Hydrographic Organization (IHO) were mentioned. There were also numerous open source and de-facto industry specifications mentioned. The latter specifications were most focused on point cloud processing and 3D content visualization. A mix of content delivery standards was suggested. These are standards ranging from providing simple map images via some network (cloud) to the end user to streaming and/or packaging processed 3D content to the user at the edge to providing a self-described container of geospatial content to the user at the edge. One response stated that the industry needs guidance on what are the best standards to implement in specific environments including guidance for maximizing performance.

One organization did not respond to this question and one responded that theywe do not provide 3D content to the user at the edge.

Some additional detailed responses are provided below.

- GeoPackage is the primary standard combined with the LTF specification used to get data to the edge operationally and in our experimentation. Additionally, WFS and WMS are in use for terrain data distribution among our collaborators.
- For point clouds (.bpf,.las,.laz, entwine point tiles) DEMs (dted, GeoTIFF, and WMS/WCS), for vecors (.shp, Geojson, little vector tiles, and WFS) mesh (3D tiles, .obj), 3D models (COLLADA, OpenFLT), and for multi-format GeoPackage and CDB.
- WMS, WFS, Entwine. ERDC is pursuing Level of Detail (LOD) mesh representations for 3D data to facilitate dissemination.
- OGC models and web services. We advocate continued enhancement and evolution of this body of work.
- Catalog Services for Web, GeoPackage 3DTiles for GeoPackage, Related Tables for GeoPackage, WPS services to interact with GeoPackage content, 3DPS services to stream GeoPackage 3DTiles, i3DM, B3DM. Compusult also processes data from the following standards: LiDAR, IndoorGML and CityGML, 3DTiles
- BIM/IFC, REVIT, Matterport

Standards mentioned in response to this question are: 3DTiles, Catalog Service, CDB, CityGML, COLLADA, GeoPackage, GeoJSON, GeoTIFF, I3S, IndoorGML, LAS, LAZ, OpenFLT, Shapefiles, WCS,

WFS, WMS, WPS

Chapter 7. Summary of Responses to the Data Analytics Questions

This section of the Engineering Report summarizes the responses to each of the Data Analytics questions presented in the Request for Information (RFI). The summaries are augmented with input from the CDS Workshop.

7.1. Responses to 4.3.1: What 3D/4D data does your organization already use?

There is a wide range in responses to this question. The 3d/4d data in use currently include:

- Photogrammetric data via UAS and satellite captures
- Interior photogrammetric data
- Satellite imagery
- LIDAR data
- Thermal/infrared data
- Building interiors and spaces
- Point clouds, 3D meshes, 3D city models and 3D computer graphics/gaming engine models.
- Synthetic Aperture Radar (SAR)/ Manportable Surveillance and Target Acquisition Radar (MSTAR)
- Use a variety of 2D, 3D and 4D geospatial data in our products.
- Building Geometry
- Wearable, environmental and industrial sensor data streams
- Rail car movement data from GNSS devices

Some documented current projects include:

- A semi-autonomous robot, which uses GPS and LiDAR data to dynamically build 3D maps that enable the robot to navigate to location waypoints while avoiding obstacles in indoor and outdoor environments. Additionally, the LiDAR data to create 3DTiles is processed and stored within a GeoPackage to be viewed on a mobile or desktop client.
- Created 3D data to perform line of sight analysis to optimize the placement of 5G small cells, and to support the CityVerve Internet of Things project in Manchester.
- Mobile client utilizes OpenGL to render 3D maps and 3D models based on formats such as DEM, OBJ, COLLADA, LAS, GLTF and 3DTiles.

One issue stated is that indoor standards and capabilities need a lot of work to become operationally viable. Yet, both are important in terms of improving future mixed reality driven operations. Standards mentioned for encoding or communicating these data include BIM, CityGML, Collada, IFC, IndoorGML, and LAS.

7.2. Responses to 4.3.2: 3D/4D geospatial content and ML/AI analytics

What 3D/4D geospatial content do you regularly use whose value would be enhanced by the use of ML/AI analytics?

The majority of organization that responded to the RFI or attended the workshop either use ML/DL or are planning on using ML/DL in their 3D/4D processing and analytics workflows. There is general agreement that ML/DL can enhance the value of 3D sources such as LiDAR and Full Motion Video (FMV). There is also agreement that ML/DL can decrease time from capture to provision to the user at the edge. ML/DL capabilities are applicable to backend server processes, workflows between servers and the end user, and finally local application at the edge, such as in a mesh computing scenario. Three organizations did not respond to this question.

Current and/or proposed activities for using ML/DL with 3D/4D applications and workflows include:

- Photogrammetric terrain datasets would be enhanced by the use of ML/AI. ML/AI processes are applicable for the Classification, Segmentation, and Attribution content of 3d datasets.
- ML/AI analytics could help correct for the margin of error between data sources during fusion and eventually when merging multiple areas of interest together into larger datasets.
- ML/AI analytics-based workflows such as change detection, path planning, and more abstract geospatially informed social analytic tasks are envisioned.
- There are two classes of activities that should be considering when addressing machine learning and AI with 3D and 4D data. First, are tasks associate with managing the data. These are relevant from acquisition through to retirement. Second is the use ML/AI in support operational intelligence and C2 tasks as an assistant to planners and decision makers.
- The change detection approach relies on height data (DSM/DTM) to identify real-world features such as types of vegetation, and there are future plans to enhance existing approaches with image-based machine learning to working with height data. Currently have not adopted any deep learning practices with regards to 3D models or point clouds. There is experimentation with using DSMs and using 2.5D height data as training data with some success. Research into training with 2.5D datasets and the benefits of using multi-modal approaches in deep learning algorithms will continue.
- All 3D/4D data conversions have significant room for improvement using ML/AI. Additionally, LiDAR collectors especially single photon sensitive, could leverage these techniques for noise filtering and conversion to models.
- The use of ML / AI analytics would enhance the value of 3D and 4D geospatial content in many areas, such as classification, ranking, object detection and semantic segmentation. ML can be used to process imagery to determine and classify the features within. ML can also be used to find 3D content relevant to a mission given a set of parameters.

Three organizations did not respond to this question.

7.3. Responses to 4.3.3: What 3D formats/encodings do you use?

What 3d geospatial storage formats or encodings do you use?

The following formats, encodings, and models were mentioned in the responses to this question.

• 3D Tiles, BIM/IFC, CityGML, COLLADA, DWG (Autodesk), FBX (Autodesk Filmbox), .gltf (Khronos), GeoPackage, GeoTIFF, JPG2000, LAS for point clouds, LTF, NITF, .obj (Wavefront), OGC CDB, OGC I3S, OGC GML, OGC KML, OpenFlight, ShapeFiles, X3D, Zip

7.4. Responses to 4.3.4: What interoperability capabilities are missing?

What needed data sharing (interoperability) capabilities are missing and should be made available, revised, or developed?

Responses to this question tended to be quite detailed. Even with existing standards and models for sharing geospatial data, a concerted community standards effort is needed to bring all important aspects together in support of mixed reality to the edge. In light of industry's quest for greater automation and higher quality 3D/4D data and 3D/4D experiences, there is a belief that end-to-end production, exploitation, and dissemination processes need to be examined in terms of enabling the desired new automated workflows required for creating, integrating, and sharing 3D/4D objects, scenes, user-views, and embedded autonomous analytics. The present environment is plagued with many disjoint models and tools, with users manually passing outdated file schemas around. The industry is sustaining a lot of non-unified legacy, which is fine, but to make progress there is need to abstract away from legacy approaches to get to the next level of capability. A fresh new "3D/4D capability layer" is needed so that we can incrementally and progressively move towards more "automated workflows" and improved 3D/4D quality and performance. Thus, core 3D/4D data models and open APIs need to be re-examined for this next-generation capability. One organization did respond to this question.

A variety of issues and requirements were highlighted. These included:

- The need for lightweight but consistent metadata including provenance is available for each 3D/4D data resource.
- Using OGC-compliant data and standardized, engine-agnostic formats will support this effort.
- Exchange data models that supports diverse applications without being bloated with additional fields and attributes irrelevant to a specific data exchange context. For example, a way to exchange data from a database developed for infrastructure modeling to one that will assist firefighters to optimally save structures.
- The ability to "geosynchronize" data stores.
- Semantic interoperability.
- An enhanced edit capability which encompasses all 3D-exchangeable formats, such as I3S, 3DTiles.

- The streaming of mesh/polygon model formats requires additional work. Need DEMS and polygon/mesh models that store versions (provenance) and metadata (accuracy). GeoPackage tries to solve some of these issues and is mobile friendly. However, this is all primarily one way (to the user). The two-way capability is required (back to provider).
- Indoor mapping standards are inadequate for enterprise type data initiatives currently.
- As application drives format, each application benefits from data and metadata presented in certain ways to accentuate speed and/or accuracy. Information is saved in a manner that is as neutral as possible to be most agile on transmission / export.
- Best practices of integrating IoT (real-time data streams) with other OGC standards

Additional Note: An interesting observation is that everyone should not need to use the same format, even though they are using the same data. There is acceptance to the idea that legacy systems and formats continue to exist.

From a standards perspective, the largest gap in terms of supporting 3D to the edge, are the 3D/4D extensions needed for GeoPackage. The highest immediate gains and returns on investment are best realized there. The extensions to GeoPackage should address: (1) core (unified) model upgrades, (2) open portrayal framework considerations, (3) adding semantics to 3D/4D models to support machine understanding and automation, (4) coping with provenance, and (5) open APIs for core 3D/4D services. This could be achieved in phased increments.

7.5. Responses to 4.3.5: Fit for use?

Is the data you have access to "fit for use"? Available in the formats you require? Are the datasets updated in time intervals that meets your needs?

Responses to this question ranged from "Yes" to "Sometimes" to not often. Typically, what is true is that the fit for use data concept has not been implemented in its entirety to accommodate users across all the varied mission space. The respondents then proceeded to identify the types of issues they encounter when accessing and processing 3D/4D content for their workflows. Four organizations did not respond to this question.

Detail responses to this question included:

- Yes, the data is fit for use and available in the required formats, and the time intervals are sponsor driven.
- Cannot directly comment as 3D/4D data are not being used in an operational context in our organization. Current research projects have multiple requirements, but inevitably involve a significant effort in extracting, cleansing and transforming 3D data prior to use.
- Although 3D data that is "open source" often is available in formats that can be used, it can be difficult to actually acquire and there are few user-unfriendly Graphical User Interfaces (GUIs).
- Lack of rapid and intelligent indexing makes it difficult to discover data that is actually useful (Discovery Catalogues)
- Data is currently too large to disseminate as a 3D product both (especially with bandwidth limited customers).

- Registration issues between datasets are common.
- Breaks in topo to bathy presents issues in littoral areas.
- Airborne collection often outdated but can still be useful. Not updated frequently enough (currency or geographically) to satisfy many users.
- The data accessed is not always in the format required, so custom data conversion utilities to convert these data are built as needed.

In general, fitness for use will dramatically improve as new automated workflows gradually displace (or at least augment) existing brittle, disjoint capabilities. For example, in addition to asking: "What formats do you require?" the discussion needs to also address: "What core models and open APIs are needed to implement this automated workflow (or constituent workflow segment)?" In short, a new generation of interoperable middleware needs to come along. Fitness for use is a very complex idea that still requires significant research. (As they say: "The devil is in the detail.")

7.6. Responses to 4.3.6: Are there adequate tools?

Are there adequate tools for your analysis, dissemination, and visualization of the 3D/4D data?

While two organizations stated that there are adequate tools, the majority of responses clearly concluded that "we have only scratched the surface on . . . tools". Further, even if the tools are adequate for the organizations 3D workflows and/or customer base those organizations stated that there is room for considerable improvement. Suggestions included improved speed, lighter weight representation/file size, with maintained or enhanced cognitive understanding coming in the future. A number of organizations also identified gaps to the available toolsets and applications baseline, such as the lack of a geosynchronization capability. Another set of comments related to the fact that quite often, key functionality is only available on desktop servers and not on the mobile assets required by users at the edge. Two organizations did not respond to this question

Detailed responses to this question included:

- Yes, but more tools are being developed.
- Most jobs get accomplished, especially if they are high priority. That said, improved tools would provide additional capability that is likely to reduce risk and increase the likelihood of operational success.
- Currently there are some very basic tools available for analysis of 3D data, but they are not adequate for most needs. Dissemination more sophisticated tools such as Skyline and virtualCity PUBLISHER are now available for disseminating 3D data via the web. Visualisation there are several well-developed tools in this space, see answers to question 4.2.2.
- Yes, there generally are adequate tools for analysis, dissemination and visualization of the 3D / 4D data at the level as required. If those tools do not exist, custom utilities are developed.
- Needs to evolve. Many sensors/platforms producing point clouds and digital imagery. None of them are automatically creating BIM. SLAM, coupled with Machine Learning software, needs to evolve. Still need heavy technician involvement. These tools appear to be very limited currently to objects such as pipes, corners of walls and floors and other easily identifiable geometric

7.7. Responses to 4.3.7: Experience level of users

Are the tools or data you have only accessible to limited, experienced people, or are they also accessible by general populations?

Responses to this question ranged from requiring only minimal training to largely experienced people with special skills. The range of responses is dependent on the type of processing, analytics, or visualizations required. For example, training DL based applications for specific types of feature recognition or patterns requires skilled technicians with domain knowledge. However, there was agreement that tools and applications used at the edge needed to be easy to use with light training requirements. One key differentiator is security: Access may be restricted to appropriate security domain (unclassified, SIPR, and JWICS) clearance. Two organizations did not respond to this question.

A few other provided observations include:

- 3D tools and data are currently only accessible to specialist research teams.
- You can buy many of the tools commercially. However, they are primarily used by government/DOD customers because of the costs. Open source or GOTS tools, especially for server based applications, are critical.
- Potree is one of the best web viewers of point cloud data available in my opinion and fits this license model. We see users coalescing around the TAK ecosystem.

7.8. Responses to 4.3.8: Use of domain models.

Do you use conceptual or domain models, and if so, how?

The majority of the organizations stated that they did use conceptual models in some aspects of their work with 3D/4D content. Five organizations responded "No" or did not respond.

Specific stated reasons are:

- In conjunction with the Army's STE One World Terrain (OWT) effort, there is work in devising a data model that will be rooted in 3D representation of geospatial objects for use in modeling and simulation. This will leverage existing sources and specifications, but intend the format to be open and modifiable by existing 3D authoring tools (3dsMax, Maya, Blender).
- Using conceptual models of the data lifecycle in the development and evaluation of new processes as well as the development of and description of best practices in the management of geospatial data.
- This depends on the application. Conceptual modeling may be required to reduce native complexity of certain problems in order to optimally apply algorithmic solutions to separate parts, prior to reintegration.
- Uniform models and model semantics are central in our work, and key to enabling higher levels of automation.

• Use domain modelling in what is called a portfolio. Portfolios allow a user to gather and organize 2D and 3D data in a format relevant to the mission or survey being conducted.

Conceptual models mentioned in the responses include BIM and CDB.

Chapter 8. Summary of Responses to the Scenarios and Use Cases Questions

This section of the Engineering Report summarizes the responses to each of the Scenarios and Use Case questions presented in the RFI. The summaries are augmented with input from the CDS Workshop. Before the details of the responses to the questions in this section are provide, one respondent, Planit, provided an excellent use case with a scenario, agents, and users.

8.1. Planit provided use case

Presently, first responders arrive at a building and search for the little white box mounted on the wall inside the main lobby. They open the box and take out a 30-page document filled with information about the building. There will be floor plans showing where fire equipment is located. This document also provides information about which occupants may be in need of assistance to vacate. The document should also describe where hazardous materials may be stored at the time the document was prepared. Contrast this with calling up a building address on your tablet on the way to an emergency.

- Open up the BIM Platform
- View the 3D image of the building
- Locate the Siamese connections, the nearest fire hydrants
- Locate the fires exits and stairwells
- Check the location of the fire as indicated by the sensors monitoring smoke and heat and pinpoint where the fire is exactly.
- Run a smoke propagation program to determine which exits need to be closed down or used.
- Send an alert to all occupants in and around the fire and direct them to the proper exit.
- Run an occupant evacuation program to understand and direct which exists/stairwells to use and notify.
- Set off fire extinguishers in critical locations to maximize fire suppression and minimize damage.
- Virtually enter the building and program in the most efficient route to the fire. Send out to all first responders.

8.2. Responses to 4.4.1: Key Workflows

What do you see as key workflow use cases for automated 3D/4D analytics?

Responses to this question are wide ranging, likely due to the different communities and domains represented by the responding organizations and workshop participants. From the data side, ongoing maintenance of "foundation data" and associate products combined with ability to quickly and "easily" integrate and update based on real time sensor and other sources is highly important. This includes integration of content from new collection technologies. Further, the ability to provide alerts to users at the edge to changes in situation, content updates, and so on is considered

important. Alerts include critical alerts such as change in wind direction at a wildfire.

Key analytical processes and/or applications include: change detection, line-of-sight, route planning, 3D pathfinding (especially indoors), point-of-need urgency identification, local conflation via update, and local situational awareness. Other suggestions are more automated processes for capture of buildings Indoors and outdoors) such as SLAM and the use of ML/DL in these processes.

Some additional detailed responses include.

- Real-time 3D mapping with user-enabled revisioning; object recognition in low Signal-to-Noise Ratio (SNR); spatio-temporal support for event prediction/location based on collaborative "recommender" engines, environmental changes, infrastructure development.
- From input data (on a server/from a service) a user could request a 3d model to be generated. Depending on bandwidth, they could receive that model in a browser/mobile device and ask questions of the data (measurements, viewsheds, line of sight, etc.). Alternatively, the user could ask that question of the 3d model on the server and only receive the answer in the client (versus calculating on the client).
- A valid workflow would be to take 3D data, such as LiDAR, from the raw form, create tiled 3D content from the raw data and store these data in a GeoPackage. The mobile and web clients would then render the GeoPackage either locally or through an OGC 3DPS service. Users can then capture additional data to enhance the content in the GeoPackage and sync with other users and the main system.
- There are various challenges for CDB use-cases such as the size of the Synthetic Environment storage is increasing by massive amounts of data, dynamic simulators, and streams of real time sensor data.

Two organizations did not respond.

8.3. Responses to 4.4.2: Key use cases for provision of 3D/4D content

What do you see as key use cases involving provision of 3D/4D content and other information assets to the user at the edge?

A key use cases is the provision of the best pre-populated geospatial content to users before they "move to the edge". Some of the 3D/4D content could be generated on demand and communicated while they are in transit. There are variety of provision methods to deliver 3D/4D content to the user at the edge, including but not limited to: Intermittent internet connections, dismounted external drive transfers, using drones to deliver physical storage peripherals and/or as internet hotspots, satellites, etc. Bandwidth will determine the CONOP afterward. In normal bandwidth use cases, detailed textured data could push to the edge over browsers or mobile devices. This would also allow for data collection in situ to push back to the data brokers. In limited bandwidth situations, additional 3D data may be generated on scene but only shared locally with teams as mesh (or compressed) services in real-time. These data should be pushed back when connections are available. In no bandwidth scenarios, the 3D data can provide a foundational layer that users can markup with relevant content that can sync upon reconnection of comms networks. ML/AI based applications that work with sensors at the edge will increase in the future.

Additional use cases provided include:

- Rehearsals / walk throughs informed by current environmental data. Real time alerts and decision support.
- Three key use-case are: 1) Emergency response and rescue: real-time situational awareness 2) Emergency scenario planning and training 3) Navigation, particularly in complex geographies such as airports and hospitals.
- Rapid access to a standard data model, regardless of provenance, that can quickly be wrapped in application-specific interfaces, if required, or used raw where applicable (many ML/DL algorithms).
- Rapid intersection (co-registration) of 3D data with other INT sources, then conversion to easily transmitted and understood information to the edge users displays (HUD 3.0, ATAK, etc.).
- See the firefighter use case above. Other use cases: Shooter at a University campus, bomb alert in the basement of a building, water pipe burst on the top floor of an apartment, gas leak at an industrial complex.
- There are use cases for synchronizing CDB or GeoPackage data where peer-to-peer synchronization across disconnected devices is desired. In these use cases, each mobile device needs its own copy of the GeoPackage data. However, at the same time, the user may wish to share additional feature data and/or changes to the feature data, without a centralized server to manage the synchronization.

One organization did not respond.

Chapter 9. Summary of Responses to the Operation and Organization Questions

This section of the Engineering Report summarizes the responses to each of the Operation and Organization questions presented in the RFI. The summaries are augmented with input from the CDS Workshop.

9.1. Responses to 4.5.1: Policy, organizational, and administrative challenges

What policy, organizational, and administrative challenges do you have that must be addressed to improve provision of curated 3d/4d content to the user at the edge?

Responses to this question can be summarized by 1.) Greater emphasis within programs of record on shared and consistent standards for core models and open APIs 2.) the need for enhanced curation and metadata, 3.) a governance structure, and 4.) through the use of standards a more "permissive" environment for using open source data, crowd sourcing, cloud computing, and other data/technologies as appropriate. Five organizations did not respond or stated "No comment" to this question.

Additional details provided include:

- The use of commercial cloud solutions to scale processing abilities.
- Currently the US Government in particular applies different 3 and 4D data standards for similar (sometimes identical) use cases. In most cases the reason for this that the kind of interoperability leaders are demanding today were only imagined by research scientists.
- Future wish list: Domain enabled 3D/4D standards-based workflows with automatic curation depending on user specification.
- Given the volume of data, a mix of crowd sourcing and curation is needed. Exiting models such as Wikipedia or Google Maps can be looked to as examples.
- Obvious network challenges (classified vs unclassified) and data classification issues. Limited curated/labeled data especially in 3D-fewer standards or no standards for 3d labeled data. Insufficient synthetic data generators for these applications. Overall, there is a large divide between the terrain side of the DOD and the modeling/simulation world (which cares less about real coordinate systems, accuracy, etc.). As a result, different standards and formats develop in each.
- Collect ephemeris metadata provided with 3D data to workflows to allow for the calculation of uncertainty and error.
- We need new guidelines and governance framework for "Mixed Reality to the Edge." These policies need to be well-calibrated on the practical limits set by the state of technology and data.

9.2. Responses to 4.5.2: Unique Needs

Are there unique needs that need to be considered at various levels of operations (local, state, regional, national, international levels, and by various players (government, commercial, NGO, academia/research)?

While the majority of organizations responding to this question states "Yes", the responses suggested differing viewpoints and solutions to the "unique needs" aspect of the question. Issues such as classified versus unclassified and implications to more general access, different domain models, and technologies were suggested. More specifically, different missions demand different storage plans, analysis tools, dissemination policies, and time latencies were mentioned. And these different missions occur across varied levels of operation, as well as across a variety of players. This variety dictates a need for unique CONOPs and implementations, with each use case likely to be different. However, throughout the majority of the responses, there was the suggestion that the consistent open standards for all providers and consumers would ameliorate the unique needs at various levels of government. Four organizations did not respond to this question.

More detailed responses to this question included:

- For example, first responders look to the National Fire Protection Association (NFPA) for their guidance on standards and there is little to no connection between NFPA and other standards bodies like OGC that focus on geospatial data. This leads to inconsistences and interoperability failures. Think about National Information Exchange Model (NIEM), and of the data models as an example of disconnects.
- Levels differ largely based on domain, rather than organization type. For example, at all levels, infrastructure and socio-economic models are required, while local/state/regional concerns regarding disaster response and terrorism intersect with national and international military and security concerns. So the unique needs are more domain-specific than organizationally dependent, at least at the algorithmic level.
- Given the sensitivity of national defense work, devices with cameras, sensors, and constant connectivity need to be carefully vetted prior to use by the military.
- While the enterprise data challenges for very large enterprises (e.g. NGA) are much different than academic communities, there is no reason both cannot leverage the same technology.
- Greater emphasis on the use of cloud-enabled, technology-independent, open APIs for sharing 3D/4D artifacts/data streams. This in turn enables a greater seamless ecosystem of producers and consumers. Research also needs to explore data sensitivity, labeling, encryption, and expiration issues.

Chapter 10. Summary of Responses to the Technology and Applications Questions

This section of the Engineering Report summarizes the responses to each of the Technology and Applications questions presented in the RFI. The summaries are augmented with input from the CDS Workshop.

10.1. Responses to 4.6.1: Portals?

Are you aware of national, regional or topical portals that can be used to support your automated 3D/4D workflows? How might they be improved?

This question did not receive many responses. Five organizations did not provide a response. Five responded with Yes and three with No. The most definitive response was "Yes, www.geoplatform.gov (and similar major portals) would directly benefit from such 3d/4D information products and analytic services".

Some specific responses to this question included:

- Awareness of several regional and topical portals for cloud-based storage of data to support inhouse automated 3D workflow. They could be improved by incorporating capabilities to ingest and convert raw data to a desired format and could also improve the way in which they associate 4D data for any given area of interest collected multiple times.
- Working on one at NGA, but it is in the early stages. IARPA is also heavily involved in this effort with their CORE3D program.
- Army users have access to many geospatial data types through an IC portal but it is not easy navigated, not a one stop entry, does not allow for ready complementary geospatial services, and is not unclassified. Assessment Limited utility.
- Evolution of GRiD connectivity to IC data bases across several agencies that house geospatial content of potential value.

Other suggestions included using standards-based access and toolsets such as GDAL. Direct access to datasets (no intermediate steps) is desirable. Standardization as described previously (at best), or continued GDAL development toward coupling to native data portals is recommended.

10.2. Responses to 4.6.2: What other capabilities need to be developed?

What other type of applications, tools, and services do you believe should be developed or integrated that enhance the currency and value of provision of 3D/4D content to the user at the edge?

Responses to this question tended to focus on three key functional areas: 1.) the ability to provide delta updates (geosynch) in a DDIL environment 2.) enhanced ML/DL workflows to facilitate the processing, delivery, and visualization of enhanced 3D/4D content to users at the edge and 3.) new

and enhanced tools for field data collection and processing of 3D content, especially indoors. These focus areas included recommendations on "lightweight" formats and compression. There was also discussion on being able to train and use ML/DL based workflows in the field. Several responses suggested the need for easier access to national and other open source data repositories. Three organizations did not respond.

Some more detailed responses are below.

- Applications and infrastructure that enable rapid editing and publishing of 3D/4D content for users at the edge would be highly valuable.
- A single portal login for all national open source data with links to regional and local data. ML/DL-driven recommendations (if user-selected) for rapid discovery and delivery of appropriate data formats.
- By in large, the industry has focused work on interoperability standards and models for visual experiences (simulation). The community needs to be thinking about a future with increasing tactical autonomy, where machines are crunching on 3D/4D models to assist humans.
- The human computer (field technician) is the best present-day technology to recognize and make decisions to teach algorithms to automatically learn what certain clusters of point clouds represent in the real world. In other words, field technicians will identify objects such as light switches or doorknobs and teach a program to relate that cluster of point clouds to the real-world object.
- Investigate into and the development of an extension to GeoPackage that enables synchronization of feature data that follows the git model. Using this model, applications can be developed to synchronize GeoPackage feature data in a decentralized environment, in a manner that can be transparent to existing applications when desired.

Chapter 11. Summary of responses to the Other Factors Questions

This section of the Engineering Report summarizes the responses to each of the Other Factors questions presented in the RFI. The summaries are augmented with input from the CDS Workshop.

11.1. Responses to 4.7.1: What sensors and sensor capabilities?

What sensor inputs, computing architectures, communications networks, and/or other technologies would enable additional 3D/4D capabilities at the edge?

General: ALL sensor observations and ALL types of derived location-based information are candidates for consideration as shared 3D/4D that may be needed at the edge. For example, IoT provides an important chapter in this discussion.

- Sensor inputs: Source data: LIDAR and Infrared, Multispectral Imagery (MSI), and Radar from satellites, airborne and ground based sensors. Real time observations from a variety of sensors, including those onboard a mobile device. These sensor inputs are required by AR applications. Input from other sensors, such as gas sensors and infrared sensors, could be used for public safety and to assist first responders in search and recovery operations. IoT standards and technology could allow easy integration of real time sensor inputs would bring additional 3D / 4D visual immersive mapping capabilities to users at the edge in a variety of applications. All this means that technologies enabling fusion of data from multiple sources would be very useful. There is also the requirement for algorithms that evaluate the "importance" of a new dataset and whether it (part, all, none) should update existing holdings. 3D version control. Interior space mapping.
- Computing Architectures: Processing with virtual machines on the cloud (Azure instance or Amazon Snowballs). Need to research GPU-based computing at the edge. GPU-based computing as a trend in high performance computing can be investigated to be adopted for 3D/4D capabilities at the edge. The GPU is used together with a CPU to accelerate computation and visualization. One example is the newly announced NVIDIA Jetson Nano platform offering 472 GFLOPs computing power at the edge. It also has the GPIO ports that can connect to IoT sensors. As more and more chips and boards are available for the development community, we expect to see more GPU-based computing devices at the edge, and many of them will be collecting or used for rendering 3D geospatial information and integrating with real-time camera feeds, and insitu sensing feeds.
- Communications Networks: Cloud-based repositories, version control. All solutions at the edge need to be developed to support both connected and disconnected operations. Blue force operations need to remain fully operational and not deterred by adversarial intervention. There is a requirement for a secure means of communication with low latency and high bandwidth. On demand wide-band P2P, D2D self-healing networking capabilities to out-maneuver communications interruptions due to natural and malicious environments. Mesh networks to support local data transmission.
- Issue: One of the key challenges with data and visualization at the edge is localization. If the

location and orientation of the user is not known with high precision, then the data cannot be accurately presented. Also consider denied environments where pervasive technologies like GPS may not be available. How do software and algorithms accurately determine a precise location and orientation from sensor fusion if one or more sensor inputs become degraded or unavailable?

Software mentioned in the responses: SLAM, PostgreSQL, Post GIS, other Geospatial Open Source software such as PDAL, Docker containers, and LAS libraries).

11.2. Responses to 4.7.2: Effects of a DDIL environment

Does a DDIL environment change your workflow for 3D/4D information at the edge?

There was an interesting range of responses to this question from DDIL does not change acquisition and processing, to only sharing of data, to DDIL changes everything. One consistent response is the need to use mesh networks. One interesting note on mesh implementations is that it is not possible to make promises with respect to anyone's consistent delivery of applications to the edge. Other suggestions are to use data caching, use symbol encodings and content streaming that significantly reduce transmission impact, and analytical services provisioned at the edge need to be designed with DDIL as a primary design criterion. However, there are scenarios that bandwidth should not determine the workflow, but instead the required information by the end user. This required information does not necessarily correlate with larger data (or bandwidth). Maximize processing at the source of the data collection

Five organizations did not comment on this question.

11.3. Responses to 4.7.3: Other success factor?

What other success factors or considerations do you see as needed for successful automated or semi-automated 3D/4D workflows?

There are many success factors that are dependent upon addressing technological gaps in existing automated or semi-automated 3D/4D workflows. These technological gaps include: cloud computing, ML techniques, standardized data models, and better training data. An acknowledgment of the criticality of automated workflows to reduce these gaps at all levels of government is required. As such we need to think about the resources and prioritize the highest priority items, such as which algorithms are most efficient, where is the utility curve in data produced fully automated etc. Four organizations did not response to this question.

Some other research areas might be:

- An increase in delivery speeds of geospatial content to data-scientists and AI engineers would lead to much more rapid development of both life-enhancing and life-saving technologies (possibly an order of magnitude).
- Use of 3D measurements during Machine Learning, vice only 2D, to improve probability of ID and reduce false alarms.
- Use of high speed on-platform ML computing for speed in getting answers to the ground users

- Careful selection of actionable targets/objects of interest for Machine Learning Identification.
- Analytical Processing against a smaller subset of targets would be accelerated.
- Positive ID of a smaller subset of targets would be improved.
- Transmission / bandwidth use would be limited to actionable targets; users would only be sent targets that warrant their immediate attention.

One issue is that there will likely be strong incentives for commercial entities to begin cataloging and geo-referencing large amounts of 3D geo-referenced data which can lead to fragmentation, proprietary formats, and closed repositories. Consumer adoption, historically, tends to gravitate toward commercial sources. Therefore, intensive standards activity is required to ensure that all geospatial content - regardless of source - is needed to support users at the edge

Appendix A: Revision History

Table 3. Revision History

Date	Editor	Release	Primary clauses modified	Descriptions
February 15, 2019	C. Reed	.1	All	initial version
May 2019	C. Reed	.2	All	Develop summaries for each question
June 2019	C. Reed	.3	All	Replace all "we" and so on usages
July 2019	C. Reed	.4	All	Final edits in prep for publication
August 2019	C. Reed	.5	All	Incorporate all of Gobe's suggested edits.

Appendix B: Bibliography

- 1. Milgram, P., Kishino, F.: A taxonomy of mixed reality visual displays. IEICE Transactions on Information Systems. 77, 321–1329 (1994).
- 2. Microsoft: What is Mixed Reality?, https://docs.microsoft.com/en-us/windows/mixed-reality/ mixed-reality, (2018).
- 3. Percivall, G., Reed, C., Simonis, I., Lieberman, J., Ramage, S.: Big Geospatial Data an OGC White Paper. OGC 16-131r2, Open Geospatial Consortium, http://docs.opengeospatial.org/wp/16-131r2/ 16-131r2.html (2017).