



BUILDING CLOUD-NATIVE GEOSPATIAL STANDARDS

CHRIS HOLMES EXPLAINS HOW THE INDUSTRY CAN MAKE THE CLOUD-NATIVE GEOSPATIAL VISION A REALITY

In 2021, I was temporarily appointed as OGC's first 'visiting fellow'. My time in that role was mostly spent investigating a riff on a question I posed five years ago in a 'Cloud-Native Geospatial' blog series: 'What would geospatial standards look like if they were built for the cloud?'

Previous to my appointment, I'd spent much of my time focused on two core parts of the cloud-native transformation: Cloud Optimised GeoTIFFs (COG) and SpatioTemporal Asset Catalogues (STAC). Early adoption of these formats has mostly centred on multi-spectral satellite imagery. So, in my time as an OGC visiting fellow, I was excited to take the time to look at the entire geospatial landscape, not just imagery, and the potential for OGC to play the key leadership role in making the cloud-native geospatial vision a reality.

I found that OGC's existing standards development work could easily evolve to align the industry on cloud-native geospatial architectures. Indeed, there is no organisation better situated to make it a reality than OGC: it is already trusted by every government as the steward of geospatial standards and has the largest community of geospatial

experts collaborating across commercial, non-profit, government, and academia.

The cloud-native geospatial vision

The OGC's mission is to 'make location information Findable, Accessible, Interoperable and Reusable (FAIR)'. Cloud-native geospatial shares the exact same goal, but leverages the cloud to radically simplify the effort needed to make geospatial data FAIR.

Instead of forcing data providers to stand up, maintain and scale their own APIs, it should be as simple as uploading the right cloud-native geospatial format and metadata to a cloud. All the APIs and scalability come from the cloud itself, enabling geospatial to ride the continuous waves of innovation in the broader IT world instead of continually playing catch-up.


A core aim of cloud-native geospatial is to decrease the burden on data providers and in turn enable far more geospatial data to be FAIR. The only cost that providers should need to pay is for the cloud storage, which continues to fall. If core data is hosted on the cloud, then general cloud-native technologies enable the cost equation to be flipped on its head, with the users of

the data paying for the computation they do and potentially the egress costs.

Things get really exciting when thinking about a whole new class of cloud-native geospatial tools that can layer on top of the core FAIR data. Google Earth Engine (GEE) has been operating on the cloud for years, albeit traditionally as a walled garden (though GEE does now support COG registration). In the cloud-native geospatial vision, however, any data on any cloud could be used by GEE or any other cloud-native tools.

Crucially, any new cloud-scale compute tool wouldn't need to build up its own data catalogue, as it could just access the same cloud-native geo formats that other tools use. Having a suite of cloud-native geospatial tools coupled with cheap data-hosting then enables a much longer tail of geospatial data to be FAIR, as smaller organisations that have valuable information – but not the wherewithal to run servers – can embrace cloud-native geospatial.

Access to all the world's information in one connected computing paradigm, combined with infinite scale computation should, in turn, usher in a whole new wave of innovative tools that move beyond



traditional geospatial analysis to find and understand broader patterns. By making geospatial information cloud-native, then, the line between geospatial and non-geospatial information will blur, greatly magnifying the potential impact of geospatial insight – but that’s worth a whole other article on its own.

Towards a cloud-native geospatial standards baseline

So how do we actually make progress towards this vision? The core standards are much closer than you might expect, with many potential standards already (being) developed, such as: COG for core raster; Zarr for multi-dimensional raster; GeoParquet and FlatGeoBuf for core vector; COPC for point cloud; OGC API – Features ‘Collection’ and STAC Collection constructs for collection and dataset metadata; and STAC for granular/scene level/asset metadata. Indeed, early adopters of these and other standards showcased their application and impact during the ‘Cloud-Native Geospatial Outreach Event’ organised by OGC in April, the recordings of which are publicly available online.

But it must be emphasised that achieving this vision will take far more work from the geospatial industry than just releasing some standards. We need a sustained effort to bring every piece of location information to the cloud in standard

formats, to update every tool to be able to work with it, and to build – together – a whole new class of next-generation tools that show the power of having petabytes of information about the world in one place.

This means a very solid baseline foundation to build on, enabling layers and layers of innovation on top. But this standards baseline must also be adaptable to the overall technology landscape (like the shift from XML to JSON and to whatever will come next). The key to this is to build small pieces that are loosely coupled, with few moving parts, and really focusing on the truly geospatial components. This approach is one of radical simplicity, getting the core atomic units right to enable unimagined innovation on top.

We are potentially close to this baseline, and I’m glad we have an organisation like OGC to bring together disparate communities of data providers, developers, domain experts, and users to collaborate. I believe that if we work together to realise the solid foundation and build out the interoperable ecosystem of tools around it, then the power of geospatial information will be available to all. Those of us who work in the field understand its power, and if we can build simple cloud-native interoperability that makes it easy for anyone to access our data, then the impact on the world will be immeasurable.

Chris Holmes is an OGC visiting fellow, and open fellow and VP product/strategy at Planet (www.planet.com)