

OPEN DATA FOR RESILIENCE INITIATIVE
& GEONODE

A CASE STUDY ON INSTITUTIONAL INVESTMENTS IN OPEN SOURCE

© 2017 International Bank for Reconstruction and Development /
International Development Association or

The World Bank
1818 H Street NW
Washington DC 20433
(202) 473-1000
www.worldbank.org

This work is a product of independent consultants working in collaboration with The World Bank. The findings, interpretations, and conclusions expressed in this work do not necessarily reflect the views of The World Bank, its Board of Executive Directors, or the governments they represent.

The World Bank does not guarantee the accuracy of the data included in this work. The boundaries, colors, denominations, and other information shown on any map in this work do not imply any judgment on the part of The World Bank concerning the legal status of any territory or the endorsement or acceptance of such boundaries.

Rights & Permissions

The World Bank supports the free online communication and exchange of knowledge as the most effective way of ensuring that the fruits of research, economic and sector work, and development practice are made widely available, read, and built upon. It is therefore committed to open access, which, for authors, enables the widest possible dissemination of their findings and, for researchers, readers, and users, increases their ability to discover pertinent information.

The material in this work is made available under a Creative Commons 3.0 By IGO License. You are encouraged to share and adapt this content for any purpose, including commercial use, as long as full attribution to this work is given. More information about this license can be obtained at: <http://creativecommons.org/licenses/by/3.0/igo/>

Any queries on rights and licenses, including subsidiary rights, should be addressed to the Office of the Publisher, The World Bank, 1818 H Street NW, Washington, DC 20433, USA; fax: 202-522-2422; e-mail: pubrights@worldbank.org.

Open Data for Resilience Initiative & GeoNode: A Case Study on Institutional Investments in Open Source

Attribution

Please cite the work as follows:

GFDRR. 2017. *Open Data for Resilience Initiative & GeoNode: A Case Study on Institutional Investments in Open Source*. Washington, DC: GFDRR. License: Creative Commons Attribution CC BY 3.0.

Acknowledgements

This report was written by Karl Fogel and Cecilia Donnelly, with design work by Felicia Doolin, all of whom are part of Open Tech Strategies, LLC.

We thank all of the interviewees for giving so much of their valuable time, and for their willingness to follow up with further information by email and even to do a second interview when asked: Sebastian Benthall, Scott Clark, Simone Dalmasso, Andy Dearing, Francis Ghesquiere, Stuart Gill, David Hemphill, Chris Holmes, Bishwa Raj Pandey, Francesco Stompanato, Chris Tucker, and Ben Wyss. Camille Acey, Kenneth Bossung, Rolando Peñate, and Eddie Pickle responded promptly and helpfully to requests for information about GeoNode's early history.

Jeffrey Johnson and Ariel Núñez gave multiple interviews, reviewed some of the data, particularly regarding development activity, and sent by email a tremendous amount of useful information that we could not have discovered on our own.

Vivien Deparday and Robert Soden not only graciously made themselves available for multiple conversations, provided introductions, and sent followup information, they also reviewed several draft versions of this report and gave detailed feedback. Alanna Simpson provided the support to make this publication possible as well as carried out essential reviews. Shanika Hettige also helped facilitate this review process. The report benefitted greatly from their attention; any errors or omissions that remain are solely the responsibility of the authors.

Funding for this report was provided by the GFDRR Innovation Lab.

TABLE OF CONTENTS

EXECUTIVE SUMMARY	5
1 INTRODUCTION	8
2 WHAT GEONODE DOES	10
3 BACKGROUND AND HISTORY	13
4 OPENDRI'S INVESTMENT IN GEONODE	19
5 RETURN ON INVESTMENT	28
6 EMERGENT BEST PRACTICES	42
CONCLUSION	44

EXECUTIVE SUMMARY

Starting in 2009, the Global Facility for Disaster Reduction and Recovery (GFDRR) and its partners developed GeoNode: web-based, open source software that enables organizations to easily create catalogs of geospatial data, and that allows users to access, share, and visualize that data. Today, GeoNode is a public good relied on by hundreds of organizations around the world, and which receives a continuously increasing investment from existing and new partners. These partners form the core of a thriving, mutually beneficial ecosystem of users and contributors — an ecosystem that includes NGOs, government agencies from a variety of countries, commercial participants, and motivated individuals.

GFDRR's direct and in-kind investment in GeoNode over the past six and a half years has been in the range of \$1.0–\$1.5 million USD. Partners have also made significant investments in GeoNode; a conservative estimate of these partner investments comes to approximately \$2 million USD over the same time period.

GFDRR's investment in GeoNode would be a reasonable amount even viewed strictly as a software development cost: the GeoNode software today represents an approximately 200% return on investment in terms of code written, since the current GeoNode project would most likely have cost \$2.0–3.0 million USD if GFDRR had produced it alone as proprietary software, without building an open source community around the codebase. (The cost of licensing and configuring a commercial “off-the-shelf” proprietary solution would have been even greater, as the total cost would grow directly with the number of installations, while offering less long-term flexibility to meet the evolving needs of GFDRR and its partners.)

However, the resultant software code is only part of the story. As this report examines in detail, the key to GeoNode’s long-term success is that GFDRR and its partners have structured the project in ways that encourage participation by others who have similar needs, creating a self-sustaining open source community that functions independently of the continued presence of any particular long-term sponsor.

In particular, GFDRR followed these principles:

1. Simultaneously contract out and hire internally:

Outside developers increase the commercial viability and “social surface area” of the software project, while internal staff both contribute to developing the software and provide natural day-to-day oversight of the outside contractors.

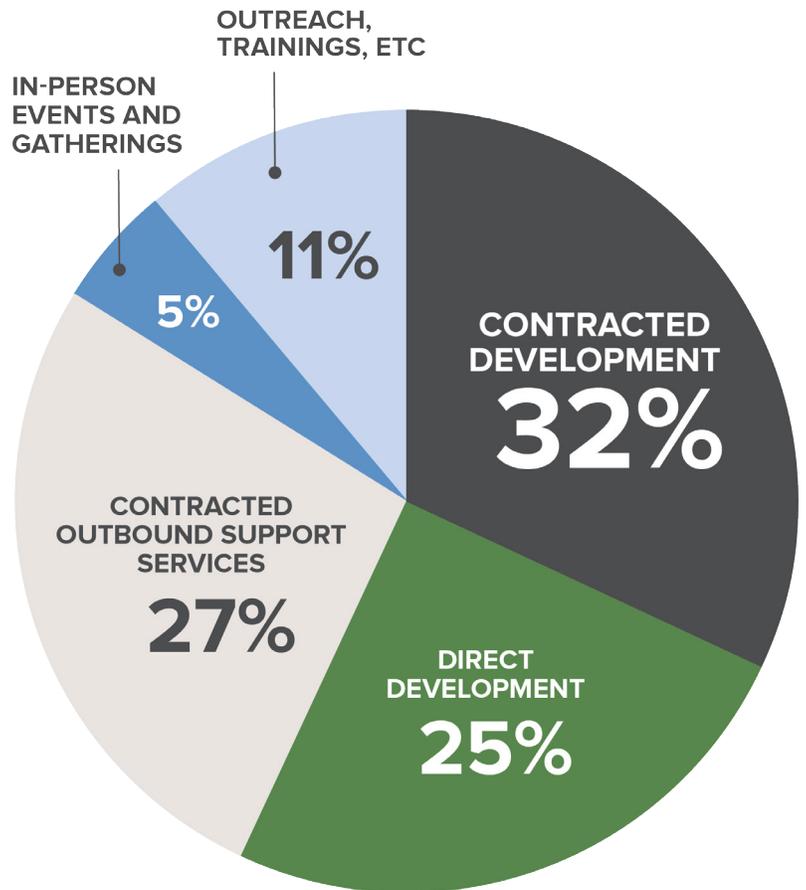
2. Sponsor in-person events:

Partners met, learned, and collaborated at these events much more effectively than they could have if they had worked together only remotely.

3. Create partnerships: GFDRR used staff time and connections to bring in peer institutions, which then invested in GeoNode themselves.

4. Train users: GFDRR encouraged client countries to deploy GeoNode, and invested in these deployments by allocating some staff time for training. Note: the numbers presented here do not include in country training, only trainings at relevant events.

Beyond the considerable technical success of the GeoNode software, this open source ecosystem itself represents a significant return on investment for GFDRR and its partners.



With GeoNode now a self-sustaining project, GFDRR is able to reap the ongoing benefits of its continuing software development while shouldering very little of the costs. Instead, GFDRR is able to invest in building capacity of partner countries and user communities to deploy, maintain, and use GeoNode platforms. Furthermore, GFDRR — along with everyone else who uses GeoNode — receives the additional benefit of having a place in which to interact with other organizations who have similar needs. The GeoNode project is a living repository of best practices for geospatial data, a forum in which to find highly qualified geospatial specialists around the world, and a place to discover new collaborators in gathering, managing, sharing, and using geospatial data. The software itself is now easy enough to use that people from local governments or universities can set up their own instances of GeoNode without assistance from or involvement by GFDRR.

This steady growth in number and diversity of participants is a hallmark of a flourishing open source project, and the final section of this report discusses some elements of GFDRR's approach to GeoNode that could be applied to other projects to achieve similar results. These best practices include, among others:

- **Run as an open source project from the very beginning**
- **Engage other organizations commercially**
- **Focus on communications and evangelism early**
- **Find and encourage the right partners**
- **Invest in collaboration infrastructure**
- **Hold events and sponsor attendance**
- **Use funding choices as a signal to peer institutions**
- **Improve user experience to attract new users**
- **Change the nature of your investment as needed**

GeoNode's future as a public good seems secured. It is now used and maintained by hundreds of organizations — governmental, non-profit, and commercial — and GFDRR can expect to benefit from the project for many years to come.

1

INTRODUCTION

GeoNode is a web catalog of geospatial data that allows users to share, access and visualize geospatial data.

GFDRR's Open Data for Resilience Initiative (OpenDRI) began supporting the GeoNode project in 2009 and continues that support to the present. Over the past seven years, OpenDRI's involvement has helped GeoNode grow, in both technical and organizational terms, to become a major geospatial data platform with many different applications and stakeholders. OpenDRI has been particularly focused on GeoNode's applicability to disaster risk management, but GeoNode is now a public good that addresses many different applications.

This report examines the history of the GeoNode software project from its inception, tracing how GFDRR contributed to the project's success. We look closely at the technical and the social aspects of creating and participating in an open source ecosystem, paying particular attention to how OpenDRI's investment strategy encouraged the arrival of outside institutional investment, by both non-profit and for-profit organizations.

By choosing an open source development strategy, and successfully creating a self-sustaining open source community, GFDRR obtained a much greater return on investment than it would have otherwise.

GFDRR and The World Bank

The Global Facility for Disaster Reduction and Recovery (GFDRR) is a global partnership that helps developing countries better understand and reduce their vulnerabilities to natural hazards and adapt to climate change. Working with over 400 local, national, regional, and international partners, GFDRR provides grant financing, technical assistance, training and knowledge sharing activities to mainstream disaster and climate risk management in policies and strategies. GFDRR is managed by the World Bank and many of OpenDRI's in-country programs and GeoNode deployments were carried out in partnership with World Bank teams.

SOURCES AND METHODS

We produced this report using a combination of quantitative and qualitative techniques:

- Phone interviews with 15 people, from approximately 10 different organizations, conducted in March–July of 2016. Most interviews lasted between one and two hours; some interviewees also responded to followup questions in writing, and a few did a second interview.
- Quantitative data from public sources such as the project’s discussion forum archives, issue ticket tracker, source code repositories, etc¹. We retrieved this data through a combination of manual and programmatic means. Details of the tools we used and the data gathered are given at <https://github.com/OpenTechStrategies/geonode-report>.
- Quantitative data from non-public sources, mainly to determine the amounts of money spent at various points in the project. In these cases, we have tried to check each claim with at least two sources. In most cases such cross-checking was possible, though occasionally there was only one source with the requisite knowledge. However, although sources did not always agree on matters of opinion, we encountered no disagreement on matters of fact, and are confident that the factual data provided to us was accurate.
- Qualitative examination of interactions between participants, of issue tracker activity and discussion forum activity, of involvement from partner organizations, etc. Among other things, such information serves as a rough proxy for various ways in which the GeoNode project has evolved over time: change in user base size, change in development community size and demographics, change in project priorities, technical progress, and growth of usage categories. We relied on these analyses the most in sections “5. Return on investment” and “6. Emergent Best Practices”.

1. Throughout the report, we will explain technical jargon where it is first substantively used.

2

WHAT GEONODE DOES

Note: Although this report is not about GeoNode’s technical capabilities per se, it will be much easier to understand if one has basic familiarity with what GeoNode does and what kinds of problems it solves. Readers already familiar with GeoNode can skip this section.

GeoNode is a web-based application for cataloguing, displaying and interacting with geospatial data — data that is associated with geographic locations. An organization typically starts using GeoNode by downloading a copy of the software (which is distributed for free from GeoNode.org) and installing it on a web server either locally or “in the cloud”. Alternatively, the organization may hire a vendor that provides commercial support for GeoNode to perform those steps.

The organization then uploads various files of geospatial data to the new GeoNode server, as data layers. These layers are made up of different kinds of data, associated metadata, documents, permissions, and sometimes maps. GeoNode catalogues and, when appropriate, displays that data visually to web browsers, as selectable layers or “overlays” superimposed on a base map. The base map can come from any digital map provider, such as Google Maps, OpenStreetMap.org, or Microsoft Bing Maps.

Each user of the site can toggle individual data layers on or off, in various combinations, to view relationships between chosen data sets —

say, showing where outbreaks of a particular disease happened versus where health clinics are located — without being distracted by extraneous information. Users can also manipulate the data directly in their browser, via GeoNode’s user interface, and can even save their changes permanently on the server when they have authorization to do so. Naturally, a user can also download an image file or print a PDF of anything displayed to them in GeoNode. Moreover, GeoNode acts as a platform on which users can develop their own tools and analysis.

The organization that manages a particular instance of GeoNode can also authorize others — including people outside the organization — to upload data and create new layers and maps themselves, using web-browser-based controls and forms that do not require great expertise to operate. All the data is in standard formats that are widely used across the geospatial information community, principally a file format called “Shapefile” plus various well-standardized image and metadata formats. Other software tools know how to generate these formats: for example, there are tools to convert information from Microsoft Excel-style spreadsheets to Shapefile format.

We cannot directly show real-time interaction with GeoNode in a printed report, of course, but Figures 1 and 2 give an idea of the software's basic use:

FIGURE 1: A GeoNode map for the World Food Programme (the WFP, an early adopter of GeoNode), showing food insecurity classifications in West Africa, with several layers of contextual data. In the left pane, under “Overlays”, the checked boxes indicate which data layers are displayed in the map on the right. By hovering the mouse pointer over a layer’s name, we get a popup window describing that layer in more detail: “This map represents, the current (March - May 2016) food security situation over RBD [Regional Bureau Dakar] countries. Cameroon Data : Extrême Nord, EFSA September 2015. Central African Republic Data : IPC, December 2015. Other Countries Data: Cadre Harmonisé, March-May 2016.” The text in the popup window comes from the data uploaded to create that layer.

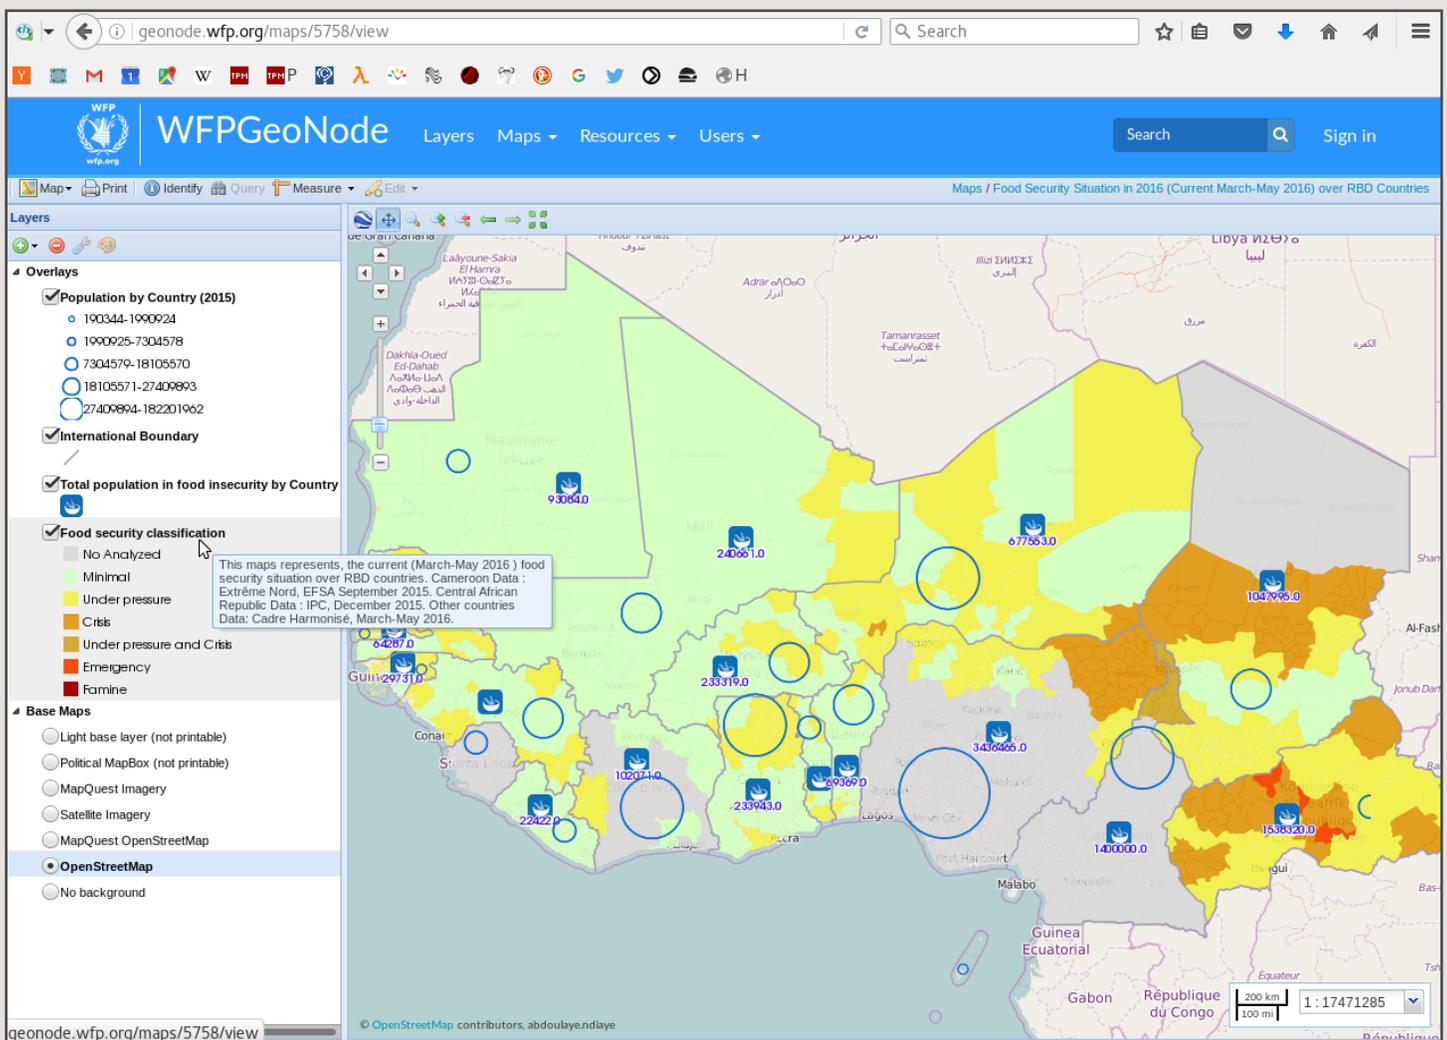
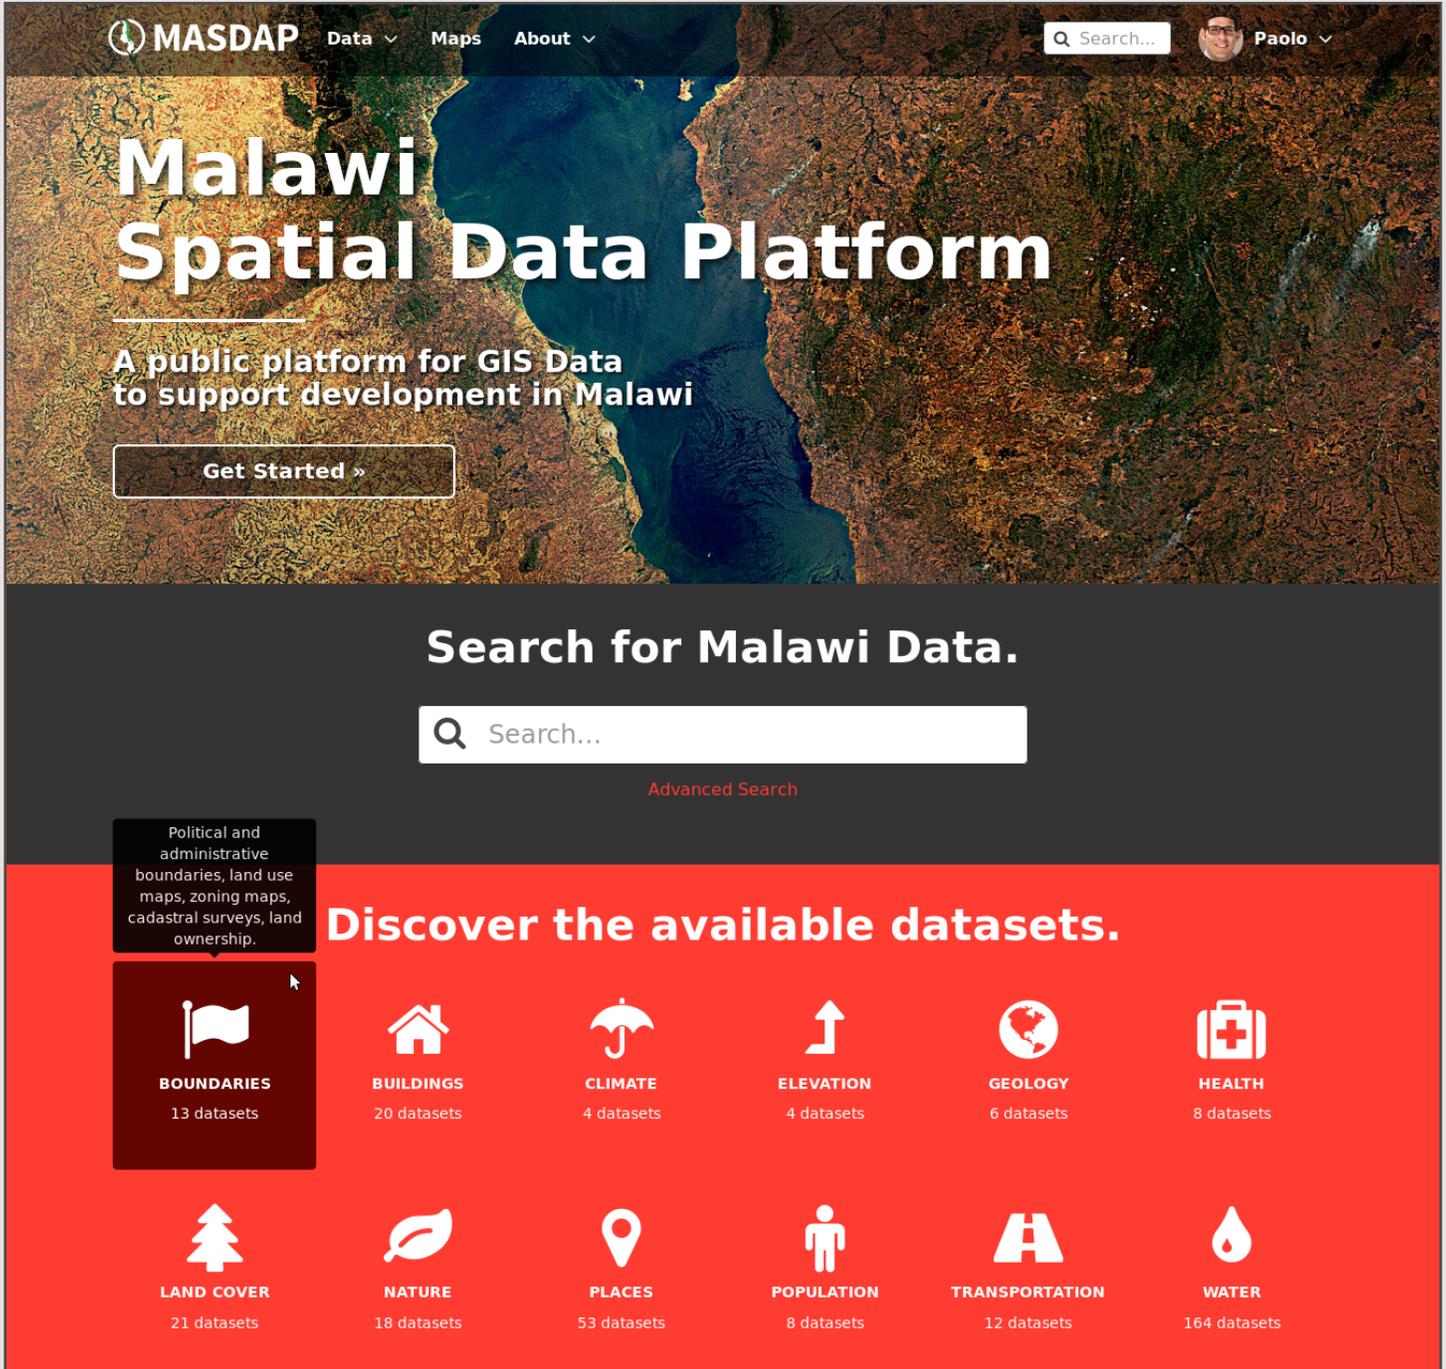


FIGURE 2: A geospatial data catalog for the Malawi Spatial Data Platform (MASDAP), showing how GeoNode can be used not just to overlay datasets, but to group them into categories that can be selected among interactively. Hovering the mouse pointer over a category name produces a popup window describing that category in more detail. In this example, the description for the 13 boundary-related mappable datasets is shown: “Political and administrative boundaries, land use maps, zoning maps, cadastral surveys, land ownership.”



3

BACKGROUND AND HISTORY

The GeoNode open source software project today is a collaborative endeavor involving more than ten separate organizations and private sector entities. Later in this report we will examine in more detail exactly what it means that GeoNode is an open source project, because the dynamics of open source software are key to understanding how GFDRR contributed to GeoNode. In the broadest sense GeoNode can be viewed as a communally-maintained technological solution to a set of common, overlapping problems shared by a number of different organizations.

The GeoNode open source software project today is a collaborative endeavor involving more than ten separate organizations and private sector entities and many individuals. Later in this report we will examine in more detail exactly what it means that GeoNode is an open source project, because the dynamics of open source software are key to understanding how GFDRR contributed to GeoNode. In the broadest sense GeoNode can be viewed as a communally-maintained technological solution to a set of common, overlapping problems shared by a number of different organizations.

GeoNode originated as a response to a fundamental problem faced by GFDRR and its partners: effective risk assessment and disaster recovery efforts can only be as good as the data that informs them, and the uncertain quality of geospatial data available at the time of GFDRR's founding in late 2006 was hampering their work. In particular, the Central American Probabilistic Risk Assessment (CAPRA) program suffered from lack of high-quality, recently-updated, standards-compliant geospatial data.

2. "Metadata" is machine-readable labeling that accompanies a dataset and specifies exactly what the dataset is, when it was created, what its authoritative source is, how it relates to other datasets, etc. With geospatial information, clear and standard metadata is especially important because geospatial data sets are most useful when they are layered on top of one another within a common base map, to provide a multi-factor view of the situation on the ground.

The problem was exacerbated by a mutually reinforcing interaction between inadequate tools and the human responses provoked by having inadequate tools. In the absence of easy-to-use technology for sharing and managing geospatial information, people fall back on ad hoc solutions — for example, sending data via physical CD-ROM — that don't always have consistent standards for descriptive metadata². Lack of dependable metadata in turn causes people to be less enthusiastic to share data in the first place, because each act of sharing may just result in further questions about the data, leading to a possibly

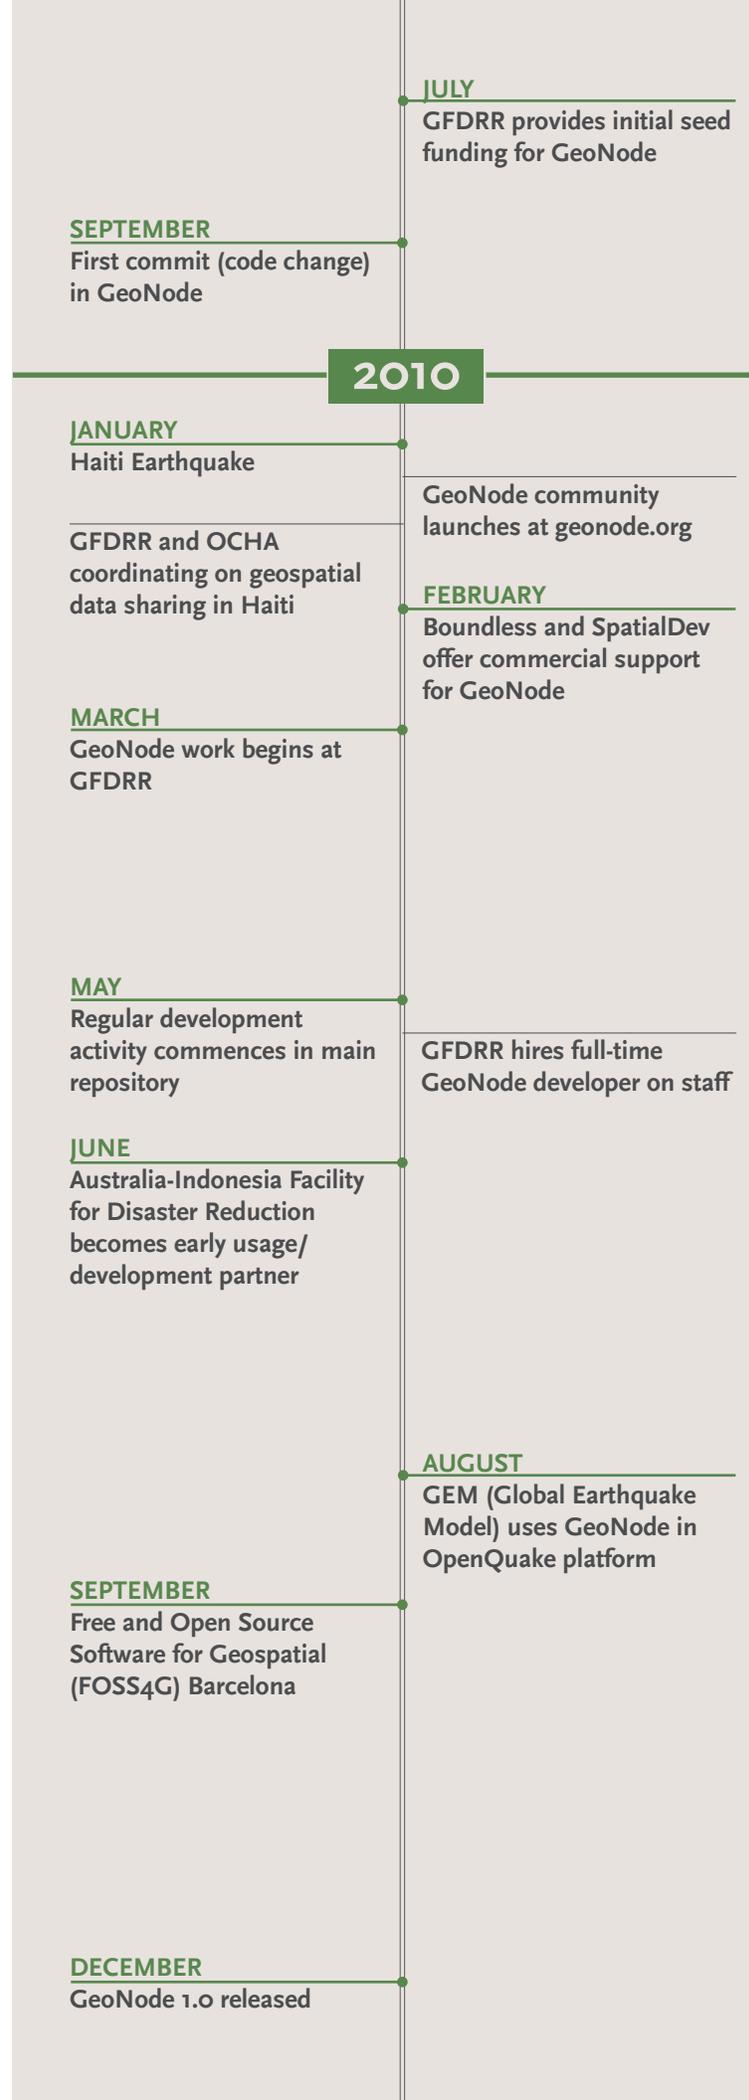


time-consuming and distracting cycle of questions and answers that, under better circumstances, would have been handled automatically by machines. Worse, sometimes the data is never used or even lost altogether as a consequence of there being no reliable place to store and share it. This loss represents a waste of the resources which were used to collect and organize that data. Finally, for people who are inclined to err on the side of holding information close anyway, the technical difficulty of sharing gives them an excuse for indulging that inclination. It's much easier to be reticent about sharing data when there is no single, agreed-on answer to the question "How do I share?"

As Francis Ghesquiere, head of GFDRR, commented retrospectively: "People were protective of their data — information is power — but in any case, even if they agreed to share the data, the technology didn't exist for them to transfer it to us or other partners. It could be 'Here's my CD' with no metadata, no context, no standards: information that cost a lot of money to collect but was useless. We needed a tool to allow people to document data, hold it, and then maybe, eventually share it with other people within their department or to the general public."

GFDRR staff were aware of the benefits that better tools could bring, and in early 2009, an opportunity came along to do something about it. GFDRR and OpenGeo (a geospatial company now known as Boundless Geo) were able to identify a common set of problems and challenges around geospatial information, and the direction of web-based geospatial systems generally. Soon GFDRR and OpenGeo decided to collaborate on what would become GeoNode. Both organizations were familiar with the problems of managing geospatial data: the OpenGeo team brought deep technical expertise, while GFDRR brought a clear understanding of on-the-ground needs and a worldwide network of contacts — both of which proved to be essential early contributions to the project. GFDRR knew by this point that its mission of risk assessment and disaster risk reduction would be well-served by improved geospatial data management, and the initial group at GFDRR (not yet called OpenDRI) was able to find approximately \$30,000 of seed funding to get GeoNode under way.

The Haiti earthquake of 12 January 2010 provided an unfortunate confirmation of their belief that GeoNode was an idea whose time had come. The aftermath of the quake required coordination among many different organizations,





including some who do not normally engage in closely-coupled temporal and spatial cooperation on a tight schedule (such as the World Bank and OCHA, the United Nations Office for the Coordination of Humanitarian Affairs). Many of these organizations were both producers and consumers of geospatial data. As they continued working on the ground and leading the coordination efforts the need for better data management tools became ever more clear. GeoNode development had begun just in time to provide a solution, and, as one person involved observed, “people are more open to sharing when there's a disaster”.

The large amount of data resulting from the earthquake in Haiti was not the only driver of GeoNode progress. Around the middle of 2010, the GFDRR team began working with Global Earthquake Model (GEM) and, a bit later, the Australia-Indonesia Facility for Disaster Reduction (AIFDR), both of which saw obvious benefits to better management of geospatial data and became early adopters of GeoNode.

The involvement of institutional partners and the increasing use of GeoNode in the field encouraged GFDRR to increase its investment. This involved hiring a full-time developer as a member of the GFDRR staff, and by procuring a development continuation contract, which was competitively bid and again won by OpenGeo³. (The nature of GFDRR's investment, and the results, are examined in detail later, in the section “OpenDRI's Investment in GeoNode”.)

From this point in mid-to-late 2010, the history of the GeoNode project is mostly a straightforward, even classic, example of open source investment: a founding sponsor, a few well-chosen early partners collaborating to make significant contributions, and eventually a broad ecosystem that includes more lightly-involved people and organizations who participate through conversations, issue reports, feature suggestions, and by simply using GeoNode and spreading the word.

The early partners, beyond GFDRR and OpenGeo, included: the Australia-Indonesia Facility for Disaster Reduction (AIFDR), the MapStory Foundation, the GEM (Global Earthquake Model) Foundation, the Harvard Center for Geographic Analysis (through their WorldMap project), NASA (through the SERVIR project), Information Technology for Humanitarian Assistance, Cooperation and

3. OpenGeo, where several key GeoNode developers worked, later changed its name to Boundless Geo (<http://boundlessgeo.com/>).

2011

FEBRUARY

Harvard Center for Geographic Analysis uses GeoNode for WorldMap project

InaSAFE extends GeoNode

First code contribution from a GeoNode newcomer

JULY

OpenDRI founded

AUGUST

MapStory project starts development work, based on GeoNode

FOSS4GNA: Denver, CO

MARCH

GFDRR hires OpenDRI Lead

MAY

GEM builds Faulted Earth from GeoNode

First GeoNode roadmapping summit

World Food Programme hires full-time GeoNode developer

SEPTEMBER

GFDRR initiates second major investment in GeoNode: overall user experience refresh, to encourage buy-in by new partners and users.

2012

FEBRUARY

GeoNode 1.1 released

GeoNode code sprint

APRIL

MapStory reaches alpha release

FOSS4GNA Washington, DC

JULY

Understanding Risk Forum: Capetown

WorldMap released in beta

Action (ITHACA), and the World Food Programme (WFP).

These partnerships will be discussed in more detail in the “Methodology of investment” section below, but in brief: GFDRR staff was instrumental in attracting and cultivating relationships with partners who brought both material resources and credibility, both of which were in turn important to the growth and widespread adoption of GeoNode.

One way in which GFDRR encouraged partnerships was by hosting code sprints and roadmapping summits and by presenting and running trainings about GeoNode at geospatial gatherings around the world (especially the FOSS4G, FOSS4GNA, and Understanding Risk conferences). These sprints and summits determined the direction of development work collaboratively, giving partners the opportunity to have their needs heard and discussed. In order for partners to invest deeply in an open source project, they have to have the sense that their needs and contributions can influence the project’s direction.

Ben Wyss was a longtime developer of the GEM Foundation’s OpenQuake platform for assessing seismic risk, which is built on GeoNode. He notes that these in-person meetings at conferences and code sprints were crucial for the GEM team to learn about GeoNode and also to meet the core developers. Wyss emphasized

“People were protective of their data—information is power—but in any case, even if they agreed to share the data, the technology didn't exist for them to transfer it to us or other partners. [...] We needed a tool to allow people to document data, hold it, and then maybe, eventually share it with other people within their department or to the general public.”

FRANCIS GHESQUIERE

OCTOBER

GeoNode 1.2 released

GFDRR Hires OpenDRI Deployment Lead to oversee in-country partnerships

NOVEMBER

U.S. Army Corps of Engineers begins ROGUE development, extending GeoNode UI

2013

JANUARY

LMN Solutions' involvement in ROGUE builds on MapStory/GeoNode work, leading to GeoGit

APRIL

GeoNode deployed in Belize

JUNE

GeoNode Code Sprint: Alexandria, VA, USA

SEPTEMBER

FOSS4G Nottingham

2014

JANUARY

GeoNode roadmapping summit in DC

APRIL

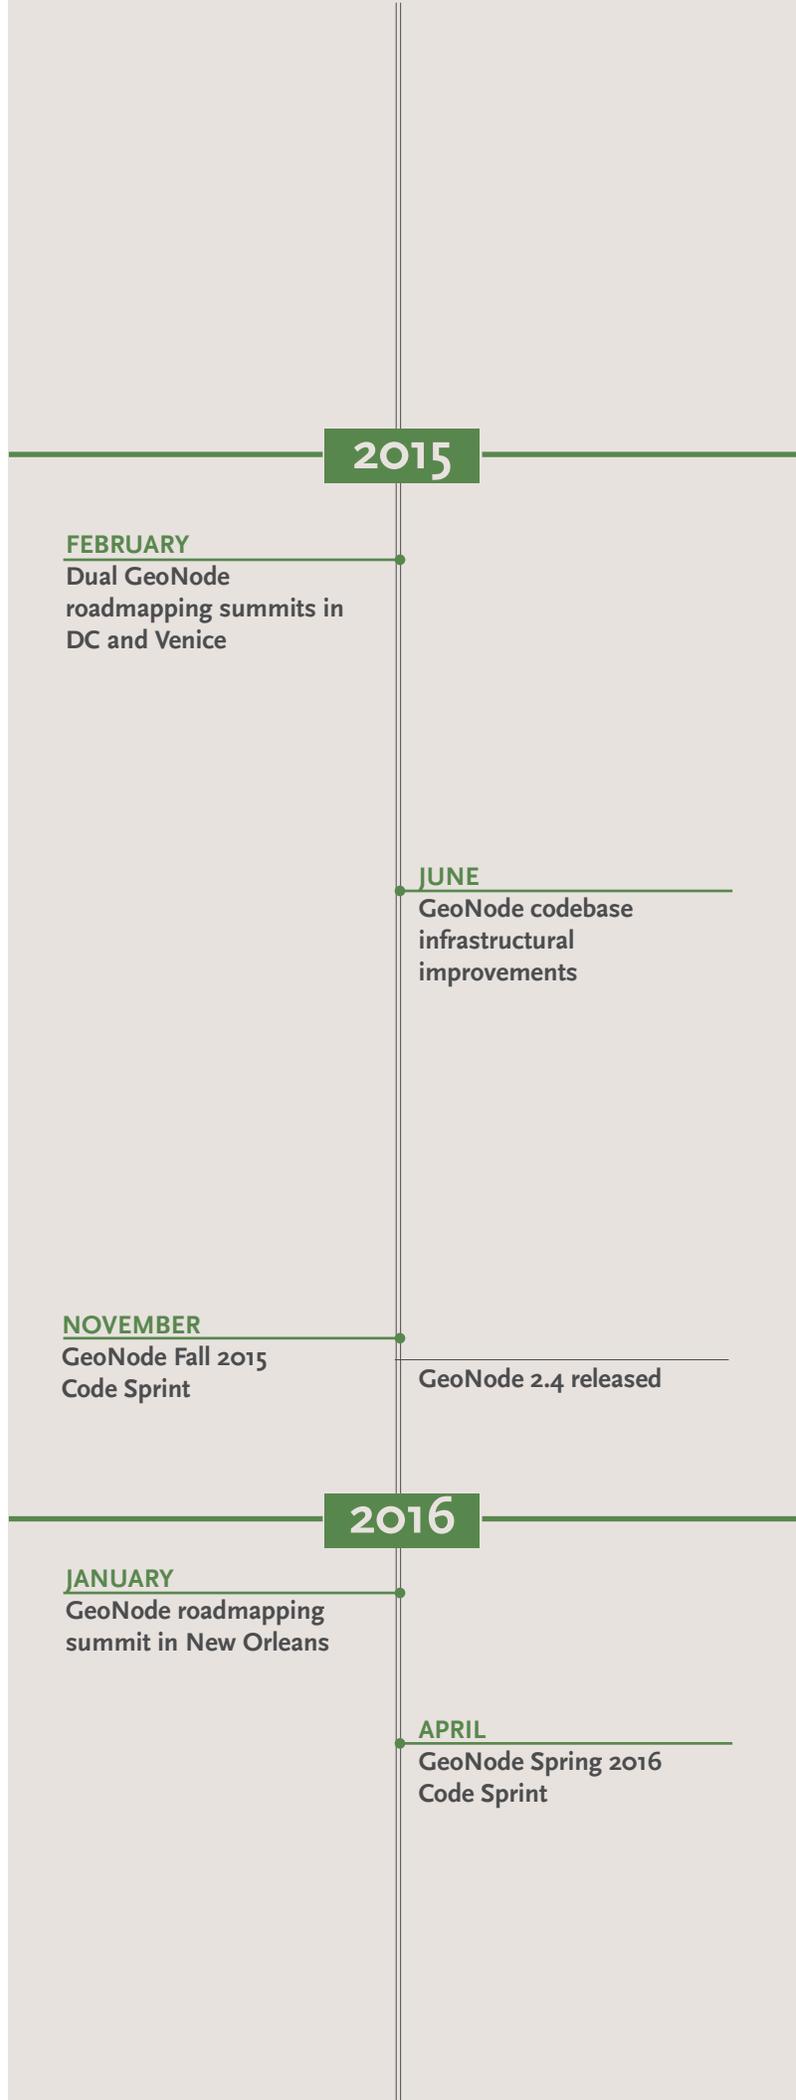
GeoNode 2.0 released

JUNE

Understanding Risk Forum: London

that GFDRR’s hiring of a full-time developer in 2010 was “a huge resource” for its partners. These personal relationships were made possible by frequent meetings, initially sponsored by GFDRR and partners, and were key for the collaborative work that these two teams, among others, did on GeoNode.

This collaborative decision-making between partners has been made possible by frequent meetings from 2011 to the present, with the most recent summit taking place in January 2016. More organizations, like USAID, NOAA, and the World Agroforestry Centre (ICRAF) have begun to participate over the years. The increasing size and diversity of these roadmapping events is a testament to GeoNode’s continuing growth. As will be shown in more detail later in this report, the size, institutional diversity, developer base, and user base of the GeoNode community are now much greater than in the first years of the project, but are also a direct result of the decisions and investments made in those first years.



4

OPENDRI'S INVESTMENT IN GEONODE

A BRIEF INTRODUCTION TO OPEN SOURCE SOFTWARE PROJECTS

Note: To understand what OpenDRI did with GeoNode, and why it was successful, some knowledge of how open source projects work is necessary. Readers who are already familiar with open source projects and how they operate can probably skip over this next subsection.

An open source software project is a voluntary association of people and organizations who have decided to pool their resources and work together on a single, shared copy of a piece of software. The stipulation that this is “voluntary” is an important element: in all open source projects, the software code is placed online under a license that allows anyone to copy, use, modify, and redistribute (verbatim or with modifications) the software, for any purpose, for free. Crucially, the license bestows that permission on everyone in the world, not just on the parties who happen to collaborate on a particular copy of the code.

That absence of monopoly control by any single agent has profound effects on the way open source projects operate. Because any participating party, or consortium of parties, could, in theory, take a copy of the code and go off in their own direction with it⁴, an open source project that persists over a long period of time is one that has by definition found ways to accommodate differences and disagreements among its members.

In projects like GeoNode, that involve participants who have mission-driven and in some cases commercial⁵

interests in the code, these disagreements tend to arise with a frequency and severity proportional to the differences between the goals of the participating organizations. Yet such disagreement is not necessarily destabilizing — indeed, it can be the opposite: it often provides a motivation for the project to develop long-term processes of compromise and cooperation, as each participant realizes that they are still stronger together than any of them would be alone.

It would have been prohibitively burdensome for any of the participants, on their own, to create and maintain software with equivalent functionality to GeoNode. To give a rough but basically accurate model: if 10 groups collaborate on a piece of open source software, then (all other things being equal) each group will contribute 10% of the effort and, in the ideal case, each will see a 9× return-on-investment. But suppose that each participant disagrees with, or has no use for, say half of the work of each other participant — a pretty extreme, though not unheard-of, degree of goal divergence. That still would result in a roughly 4.5× ROI for each participant. Of course, this idealized example leaves out some edge costs: there is technical overhead in adjusting one’s own usage of the software to accommodate design decisions that resulted from other parties’ needs, and there is communications overhead for participating in the project community. But on the whole, the decision to stay with the collective project is usually an easy one for most participants, even in cases where significant technical disagreement persists for a long time.

For the institutional members of the GeoNode community, the decision to concentrate resources on the shared group project — that is, the decision to not fork — has always been not merely financial, but informational as well.

4. An event known as “forking”, or as a “social fork” of the project. (Note that this sense of “fork” is not to be confused with the more colloquial and casual use of the word “fork” on the popular code hosting site GitHub, where it is just used to mean “make a temporary clone of the code for development purposes”, and has no implication of disagreement about project direction.)

5. There is no contradiction between “commercial” and “open source”, and many open source projects are the basis of significant revenue streams for companies that participate in them. The opposite of “open source” is “proprietary”, i.e., under monopoly control by one party; commerciality per se does not require monopoly.

Participation in an open source project is a very effective source of perspective about one's own usage of the software. Many of the individual participants understood that with a new tool like GeoNode, their best source of information about their own potential future needs would be to look at other organizations' current uses and needs. That is, the ideas and feature suggestions coming from other partners were often a good preview of one's own likely future directions. Informationally, as well as financially, it pays to stay where the action is rather than strike out on one's own.

HOW GFDRR AND OPENDRI INVESTED IN GEONODE

GFDRR's investment in GeoNode was not merely a total amount of money expended over a period of time. The structure of the investment — the specific ways in which GFDRR deployed money and attention in the project — was also key to GeoNode's success.

To understand that structure, it may be useful to start with an overview of the investments made, indexed by date and amount of expenditure, which we provide in a chronological list below. Note also that some of these investments can be measured only imprecisely. We can say that part of a given investment is related to GeoNode improvement, but would quickly become entangled in definitional questions were we to try to make the measurement precise. Without going into those precise definitions, our purpose is to show the approximate size of GFDRR's investment in GeoNode since the inception of the project, to explain the shape or structure of that investment (which is just as important as the total amount), and to examine the returns GFDRR obtained for its commitment.

With all of those caveats laid out, here is a list of the major, quantifiable investments GFDRR has made in GeoNode; all figures are in unadjusted U.S. dollars:

MID-2009

\$30,000 initial development contract.

2009—PRESENT

About \$580,000 worth of GFDRR staff time. This includes a full-time GeoNode developer on staff,

plus various technical, documentation, training, communications, and outreach investments from other staff members.

SEP 2010

\$50,000 contract to develop and support HaitiData.org.

JUNE 2010

\$20,000 worth of core development as part of a larger contract to develop GeoNode for Pacific Catastrophe Risk Assessment & Financing Initiative (PCRAFI) at the South Pacific Applied Geoscience Commission (SOPAC).

MAY 2011—PRESENT

About \$20,000 in financial and in-kind support for roadmapping summits, code sprints, and GeoNode gatherings at conferences (does not include major in-person events already mentioned elsewhere in this list).

LATE 2011 / EARLY 2012

\$240,000 contract for improvement of GeoNode user interface and user experience.

JULY 2012

About \$25,000, to fund around 25 participants to attend Understanding Risk conference and attend GeoNode training workshop.

2012—PRESENT

About \$250,000, for various smaller-scale contracts in many locations, usually for deployment and related services. (See further explanation below about these.)

JULY 2014

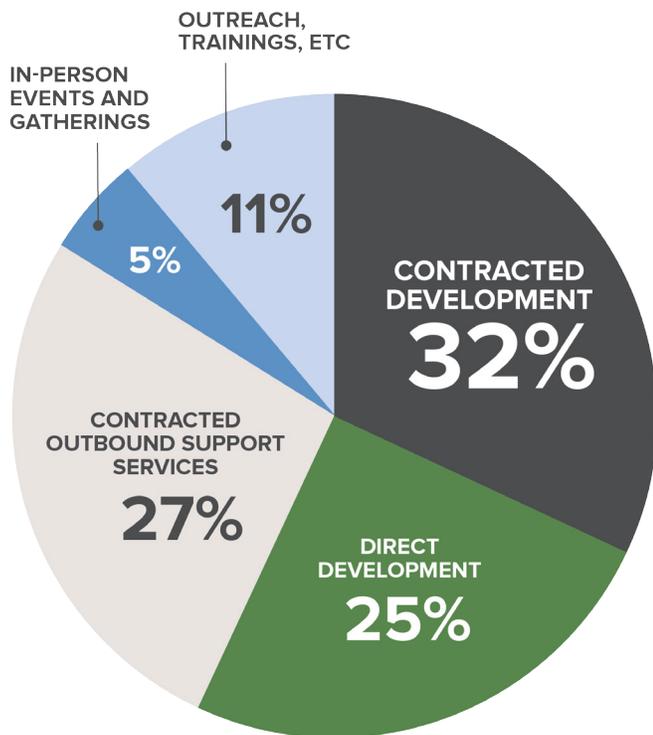
About \$25,000 to fund around 25 participants to attend Understanding Risk conference and attend GeoNode training workshop, similarly to UR2012.

2014–PRESENT

\$175,000 for core GeoNode development in a larger contract that is half for core development and half for country support.

STRUCTURE OF TOTAL INVESTMENT

Overall, GFDRR's investment in GeoNode, over approximately **six and a half years** from mid 2009 to early 2016, is in the range of **\$1.0–\$1.5 million USD**:



The proportions above are necessarily estimates, and the overall expenditure is an estimated range, because measurements of staff time are naturally imprecise and reliant on subjective recollection, and because some GFDRR activities that supported GeoNode development also served other related purposes simultaneously. Note: the numbers presented here do not include in-country training, only trainings at relevant events.

METHODOLOGY OF INVESTMENT

GFDRR invested in GeoNode in several different ways:

- Hiring contractors to do development work;
- Hiring individuals directly as software developers;
- Using its influence and credibility to form partnerships with other organizations that became users and supporters of the software;
- Allocating staff time toward training, deploying, and community management;
- Encouraging the deployment and use of GeoNode in client countries.

The way that GFDRR made these investments ensured that each one reinforced the effectiveness of the others. Instead of simply hiring a single company, giving them software requirements, and sending them off to build a product (open source or not), GFDRR purposely built an ecosystem of vendors and users.

Choosing the right early contractors

Initially, this ecosystem took the form of a few partners, among whom GFDRR was the de facto leader, and both GFDRR and the partners hired software developers internally and contracted with external companies.

GFDRR started its investment in GeoNode by procuring a software development contract that deliberately required both geospatial expertise and open source community-building expertise. The company eventually selected through a competitive process, OpenGeo, was a strategically fortunate and favourable choice: they were already experienced in open source development and very well-connected in the geospatial web community. By having OpenGeo developers start work on GeoNode, instead of using in-house or directly-managed contract developers, GFDRR created early socialization of the software among likely early developers and adopters for free.

Contracting to OpenGeo, an experienced software development firm, also saved GFDRR from becoming a development shop. Instead of hiring the several internal developers that would have been needed to build GeoNode, GFDRR invested in one or two people who could supervise the process and make sure that proper open source practices were followed. This had the added

benefit of ensuring that the software was usably open source, because GFDRR staff collaborated with OpenGeo developers via public, open source channels (see “Run the project in the open from the start,” in section 6). OpenGeo also had a motivation to start thinking about potential commercial demand for GeoNode as a business opportunity for themselves.

That last element is key, because open source projects are healthiest when different contributing organizations work together while each brings its own objectives. OpenGeo’s business needs were different from OpenDRI’s goals around risk management, making GeoNode as a whole a stronger and more well-rounded platform which could attract even more users.

This broadening is a general principle: the more organizations that are involved, the more people gain expertise in the software, and the more of them have incentives to improve the software and make commercial offerings based on it. Then more organizations see that the software is viable and can support commercial offerings, and they get involved too.

The most recent history of GeoNode confirms this pattern unmistakably: as we will see later on, some U.S. federal government agencies have recently adopted GeoNode in such a committed way that they are now making investments as large or possibly larger than GFDRR’s total investment to date.

Having a full-time developer on staff

Hiring a full-time developer at GFDRR was an important early move. By having an in-house developer closely involved in the project, GFDRR increased the expected return on its existing investment with OpenGeo (and, as we will see later in Figure 4, this increase is visible in the productivity of the OpenGeo developers in the period after this developer joined).

Having in-house resources both reduced risk, by ensuring that technical knowledge would be maintained in multiple organizations, and increased the effectiveness of the OpenGeo developers, by giving them a full-time customer-side partner to check assumptions with, react to, vet technical and design decisions with, etc.

This hire highlights the importance of investment structure in another way. GFDRR’s early socialization of

GEONODE ORGANIZATIONAL PARTNERSHIPS



2010–2011 PARTNERS

This and all partner lists are non-exhaustive, and are meant only to give an idea of the range of institutions that have become involved in GeoNode over the years.

- 📍 The World Bank
- 📍 OpenGeo
- 📍 Australia Indonesia Facility For Disaster Reduction (AIFDR)
- 📍 MapStory
- 📍 Global Earthquake Model (GEM) Foundation
- 📍 Harvard WorldMap
- 📍 ROGUE (U.S. Army Corps of Engineers)
- 📍 South Pacific Applied Geoscience Commission (SOPAC)
- 📍 SERVIR (U.S. National Aeronautics and Space Administration / NASA)
- 📍 Regional Centre For Mapping of Resources For Development (RCMRD, Kenya).
- 📍 Information Technology for Humanitarian Assistance Cooperation and Action (ITHACA, Italy)
- 📍 World Food Programme (WFP)



World base map image, here and in subsequent maps, from CARTO.

GeoNode made qualified developers, including the person ultimately hired for this role, interested in this work. Developers working at companies in the geospatial space knew about GeoNode because it was open source and because GFDRR had made efforts to circulate knowledge of it in the geospatial community by hiring contractors instead of keeping development in-house.

Forming partnerships

Outside of GFDRR, the ecosystem was growing through dedicated outreach work by the OpenDRI team. MapStory, the GEM Foundation, AIFDR, Harvard, ITHACA, and the Army Corps of Engineers all became early partners, in a kind of snowball-effect accumulation that was helped along by the credibility GFDRR’s early involvement brought to the project. GEM, which eventually created the OpenQuake platform as a wrapper around GeoNode, was brought in through the efforts of the GFDRR team.

MapStory joined GeoNode — and ultimately became significant contributors, albeit with a long period of technical divergence due to having different needs — through a connection with the OpenGeo team, and then introduced the Army Corps of Engineers to GeoNode, which led to the Corps’ ROGUE (Rapid Open Geospatial User-Driven Enterprise) project being based on GeoNode.

Inter-developer connections, community management, and the spread of GeoNode knowledge

Personal relationships between developers and users at different organizations led to greater resilience for the project as a whole. Individuals at the initiating organizations naturally came with their own personal and professional networks. In GeoNode’s development, these personal connections in the geospatial community were crucial for early uptake, and were formed and strengthened through the developer meetings discussed in the “Background and History” section.

As the project progressed, developers would sometimes move to different organizations, carrying GeoNode knowledge with them and drawing the new organization more deeply into the project. This is a very common pattern in open source, since a developer’s position and influence in a project are independent of their employer.

For example, one of the most active long-term GeoNode developers, Simone Dalmasso, started working with GeoNode at ITHACA in 2010, while getting his Ph.D.

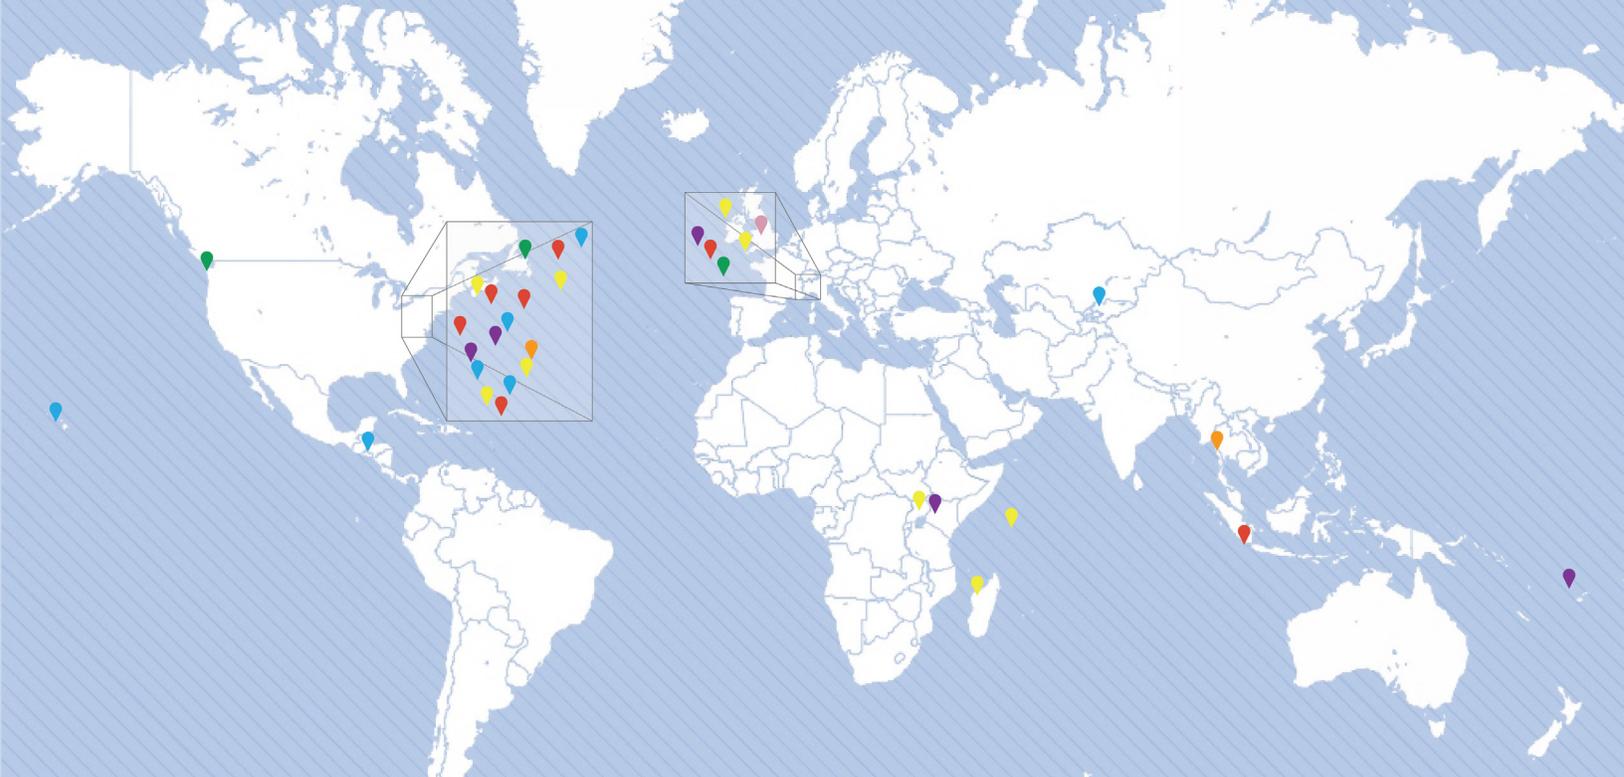


2012–2013 PARTNER ADDITIONS

This and all partner lists are non-exhaustive, and are meant only to give an idea of the range of institutions that have become involved in GeoNode over the years.

-  Spatial Development International (SpatialDev)
-  Eldarion
-  European Commission Joint Research Centre (JRC)
-  Asian Disaster Preparedness Center (ADPC)
-  LMN Solutions





in Turin, Italy (eventually he would become the release manager for GeoNode 2.4 release, in 2015). In 2011 he moved to the World Food Programme, where his work on GeoNode intensified, and he later went to the European Commission as a senior technical projects manager, where his work continues to involve GeoNode.

Moving from organization to organization, Dalmaso carried a tremendous amount of GeoNode knowledge and influence with him. Meanwhile, his departure from ITHACA and later from WFP did not signal the end of those organizations' involvement in GeoNode: in both places he was succeeded by other developers (Paolo Pasquali and Paolo Corti respectively; the latter would himself later move to Harvard University Center for Geographic Analysis).

Each of these organizations contributed money and time to the GeoNode project based on their own needs, expanding the software's capabilities and audience. In many cases they built software products that depended on GeoNode, contributing back to the core project in order to meet the requirements of their own applications or platforms (this is also a typical pattern in open source). Many of these organizations did not have a formal relationship with GFDRR, beyond their participation in the GeoNode project, but none was necessary to collaborate in technical discussions and in the work they needed done.

GFDRR's GeoNode-related investments in communications

2014–2016 NEW PARTNERS

This and all partner lists are non-exhaustive, and are meant only to give an idea of the range of institutions that have become involved in GeoNode over the years.

-  Comisión Permanente de Contingencias (Permanent Contingency Commission / COPECO, Honduras)
-  Humanitarian Information Unit (HIU, U.S. State Department)
-  Marine Civil Information Management System (MARCIMS, U.S. Marine Corps)
-  National Geospatial-Intelligence Agency (U.S. NGA)
-  Office of the Secretary of Defense (U.S.)
-  Pacific Disaster Center
-  Central Asian Institute for Applied Geosciences (CAIAG, Kyrgyzstan)
-  National Research Council: Institute of Marine Sciences (Italy)
-  World Agroforestry Centre (ICRAF)
-  Massachusetts Institute of Technology (MIT, U.S.)
-  National Oceanic and Atmospheric Administration Center for Weather and Climate Prediction (NOAA NCWCP, U.S. Department of Commerce)
-  Politecnico di Milano (Polytechnic University of Milan, Italy)
-  Humanitarian Data Exchange (HDX, United Nations Office for the Coordination of Humanitarian Affairs)
-  Agency for International Development (U.S. AID)
-  allSpatial
-  HABAKA Innovation Hub (Madagascar)
-  GESP ("Gestione, Elaborazione, Studio, Pianificazione", Italy)

and outreach to these partners sometimes had major effects on them beyond just GeoNode. According to one source, the GEM Foundation ultimately became a leading open source seismic hazard and risk assessment initiative. Open source became a major part of their mission, and thus GFDRR's investment in relationships and partnerships around GeoNode ended up having a broader effect on the open source risk assessment community.

Expansion of in-house team

GFDRR continued to hire internally for the project in 2011-2012, especially once OpenDRI was officially formed. New members joined the team, some of whom had technical experience, but much of their work was focused on scaling up usage and deployments of GeoNode, as well as continuing the outreach and communications that the existing GFDRR team had been doing. The new OpenDRI team had responsibilities beyond GeoNode, but some of the team members eventually became active in GeoNode mailing lists and occasionally even in direct development, which is a sign of how much the GeoNode project had become a common gathering place for people engaged in a range of related activities.

IN-PERSON EVENTS

Beyond these internal hires and continuing GFDRR contracts with outside vendors, GFDRR invested directly in the growth and care of the community around GeoNode through in-person events. By hosting and subsidizing road-mapping summits and code sprints, GFDRR strengthened relationships between community members and made much of the work by partners possible. Since most open source development happens remotely, GeoNode collaborators didn't always have a chance to build trust and personal relationships. Coming to these events and meeting in person allowed them to do so, helping them work closely together afterward. See Ben Wyss' comments about the importance of in-person events at the end of the "Background and History" section.

A close look at one of these events, the first GeoNode Roadmapping Summit and Sprint (May 2011)⁶, shows the effect that a well-timed and well-planned stakeholder gathering can have on a project.

The Summit was carefully structured. The first day, each organization gave a short presentation on how they were using GeoNode, and what their goals were for the

coming year. Although there was some consistent overall vision, it was also clear that different institutions had very different priorities, and there was "some tension in the room" according to one participant. Then, in the evening, everyone went out together and were relieved to be able to just relax and interact socially for a while, without presentations and roadmaps and technical worries. When they reconvened the next day, they brainstormed on development issues as a group, and finally held a project prioritization bidding exercise, a consensus-oriented process in which everyone could attach up-votes or down-votes to weight the various proposals.

At the end of that process, the proposal that turned out to have the most up-votes was not about developing a particular feature, but rather a proposal to make the foundation of GeoNode more stable and reliable — to "pay back technical debt", in the language of software development⁷ — so that future work on GeoNode would be easier.

Thanks in part to the attendees' familiarity with each other, from the previous day's meeting and that evening's socialization, they now had enough mutual trust to make concrete commitments to materially support that goal, which turned out to be very good for GeoNode and thus, ultimately, for their various organizational goals as well. One attendee, Sebastian Benthall, the GeoNode Program Manager at OpenGeo for a couple of years in the early history of the project, summed up the importance of the event's structure this way:

It's not just that we held an in person event. We scaffolded a procedure for collective decision making. ... What I thought was noteworthy was that we brought collective multi-stakeholder action up one level of hierarchy to the funders themselves. And then only collectively did they arrive at the one intangible thing they all really needed: a "rock solid" core.

Rolando Peñate, who designed the consensus exercise, said:

The event was pretty intentionally structured around the stakeholders getting together to discuss ideas, make the case for their priorities, and then partake in a formal consensus exercise to weight the proposals. It's worth noting that the results of this exercise later served as a "menu" against which stakeholders would contract future work.

6. <http://geonode.org/blog/2012/02/13/sprint/>

7. See https://en.wikipedia.org/wiki/Technical_debt for more on the concept of "technical debt".

PROMOTING A CULTURE OF DOCUMENTATION

An important component of GFDRR's investment strategy was to establish a culture in which documentation improvements are as important as software code improvements. It is well-known in the open source world that documentation is crucial for making the software easy for new users to try out and onboarding new contributors. Many people will start using a piece of software based only on the project's website and documentation — if these are out of date or lacking in detail, these potential users and contributors may well move on to a different project. Moreover, good documentation saves developer time by reducing the number of basic questions they need to answer. When they are asked simple questions, they can simply point the person to the relevant section of the "docs." Finally, and this may be especially relevant for GeoNode, clear and complete documentation serves as an excellent starting point for in-person trainings.

For example, in preparation for the "Advanced Spatial Data Management Training" event⁸ held in Trinidad in February 2013, GFDRR made a deliberate preparatory investment in documentation infrastructure. This meant not only finding and filling gaps in GeoNode's documentation itself, but setting up a dedicated documentation site (<https://docs.geonode.org/>), a style guide, and a polished workflow to make it easier for new contributors to participate in improving the documentation⁹. The short-term goal of this investment was to improve the project's documentation before the Trinidad workshop, but the long-term goal was to foster a culture in which people who contribute new features are expected to include documentation along with their code.

IN-KIND AND INCIDENTAL INVESTMENT

The GFDRR team also encouraged the World Bank's client countries to use GeoNode. In these areas, GFDRR hired contractors to produce data, and used the data collected in a GeoNode as a measure of whether those companies were fulfilling their obligations. World Bank program managers and the GFDRR/OpenDRI team helped those countries with installing and configuring GeoNode, and sometimes with loading their initial data. These investments were usually made for a particular in-country project, and we have not counted them as direct investments in GeoNode

here. However, these efforts still had the effect of increasing GeoNode's mindshare among those who work with data for disaster risk reduction, and in some cases (see Bishwa Raj Pandey's comments about Belize later in this report) that mindshare became self-sustaining instead of being driven by GFDRR.

GFDRR sometimes provided funding for these country- and region-specific GeoNode installations through a contract with a local provider of geospatial technology services. We do not count all of these contracts here, because not all of them resulted in direct and measurable benefit to the GeoNode project itself. We have counted some of the most important ones, however, for example contracts with AllSpatial and Habaka to set up GeoNode in Mauritius-Seychelles and Madagascar respectively.

In general such contracts have a positive long-term effect on Geonode: by increasing awareness and spreading expertise in GeoNode, they helped create an ecosystem of demand and exchange that supports the project to this day. And when a contract involves a third-party service provider, it has the secondary positive effect that now GeoNode knowledge resides in two new places: the organization for which the instance was deployed, and the organization that did the deploying. These effects will be discussed further in the section "Return on investment".

8. This event also illustrated opportunistic partnership building: it was a collaboration between the University of the West Indies (UWI) in Trinidad and Tobago, GFDRR, and the World Bank's Latin America and Caribbean Region Disaster Risk Management and Urban Development (LCSDU) Unit.

9. See http://docs.geonode.org/en/master/organizational/contribute/documentation_guidelines.html for details about the documentation system, which uses "Restructured Text" as its master input format and a tool called Sphinx to generate output in multiple formats (web pages, PDF, and so on) with automatic cross-referencing, indexing, etc.

“The event was pretty intentionally structured around the stakeholders getting together to discuss ideas, make the case for their priorities, and then partake in a formal consensus exercise to weight the proposals. It's worth noting that the results of this exercise later served as a 'menu' against which stakeholders would contract future work.”

ROLANDO PEÑATE

5

RETURN ON INVESTMENT

We conservatively estimate GFDRR’s return on investment (ROI) for GeoNode to have been at least 200% — with an emphasis on “conservative”. As we will show below, the long-term and indirect ROI is much greater than that, and continues to grow, thanks to significant new investments from non-GFDRR partners.

To evaluate GFDRR’s decision to invest in GeoNode as an open source project in 2009, we must look at the options available at the time.

GFDRR and its clients needed a tool for managing geospatial data on the web. There was no open source tool that really met the need: the ones that existed at the time required a high degree of expertise to set up and use, and offered only a subset of the features GeoNode would eventually offer. There was likewise no proprietary tool that had all the required features, though with various workarounds and adjustments some of them could probably have done the job for a while. But in any case, the proprietary tools came with their own set of problems: licensing fees, a monopolized vendor landscape, limits on possible customizations, and limits on how involved customers can be in development and knowledge exchange. By contrast, while no open source tool existed with GeoNode’s set of features, the development team knew that they could take advantage of existing open source components to build GeoNode more quickly and easily than if they were starting “from scratch,” as discussed below in the section titled “deep extensibility.”

The GFDRR team’s decision to commit to GeoNode as an open source project rested on two key insights. One, the team knew that they had a fairly good understanding of their immediate, near-term technical needs for geospatial

data management, and exactly how currently available software fell short in meeting those needs. Two, they also knew what they did not know — namely, what their needs and their clients’ needs would be in the medium- to long-term future, except in general terms. GFDRR works around the world with many countries and partner organizations, and sees variation over time in the financial and technical resources available to its clients and partners, as well as in the risks and disasters they face.

For example, data management in the aftermath of the Haiti earthquake of January 2010 was a major catalyst of GeoNode development, and <http://HaitiData.org/> was one of the first production deployments of a GeoNode server. But four years later, a GeoNode site¹⁰ (<https://ebolageonode.org/>) would be deployed to manage data related to the Ebola outbreak in West Africa — a very different type of disaster, and one that crossed far more national boundaries; more recently, GeoNode has been chosen by the International Fund for Animal Welfare and the Kenya Wildlife Service to help coordinate counter-poaching efforts¹¹.

The best way to deal with a wide variety of clients and needs is to look for what they all have in common — the core elements of geospatial data management, that are likely to remain constant across all uses — and offer a system that meets those core needs while also providing well-documented mechanisms for future extension and

10. EbolaGeonode.org is being migrated to the Humanitarian Data Exchange at <https://data.hdx.nwlabs.org/> as of this writing.

11. <http://americangeo.org/news-and-information/ifaws-tenboma-project-adopts-opengeo-in-support-of-innovative-counter-poaching-initiative/>



Coding Events 2010–2011: FOSS4G (Barcelona, Spain), Risiko / Risk-in-a-Box Workshop (Jakarta, Indonesia), GeoNode Roadmapping Summit (Washington, DC).



Coding Events 2012–2013: FOSS4GNA (Denver, Colorado, USA), Understanding Risk Forum (Cape Town, South Africa), GeoNode Code Sprint (Cambridge, Massachusetts, USA), FOSS4G (Nottingham, England, UK), GeoNode Code Sprint (Alexandria, Virginia, USA).

enhancement, so that the system can evolve naturally with usage.

Ideally, this means extensibility on two levels:

On-the-fly extensibility: This is extensibility in day-to-day usage. In GeoNode, users can upload data and create custom maps and layers, to manipulate information and share the results with others. More sophisticated users can also connect to a GeoNode site programmatically (i.e., using GeoNode’s “application programming interface” or API), allowing them to automate how they upload, search, customize, and share data. In other words, GeoNode can become an extension of another computer program, as long as that program wants to use features that already exist in GeoNode.

Deep extensibility: This is the heart of what an open source project is about: a technical architecture and a social environment that encourage people to get involved in improving the system as a whole, so that it can grow new features organically to fit real-world usage. With deep extensibility, even if the first version of GeoNode didn’t

have certain features that would later turn out to be needed, there would at least be a path by which motivated parties could provide those features, thus benefitting everyone. It is an especially powerful approach when combined with an open-data ecosystem, because the data provides a common focal point around which people with different goals can, eventually, agree on how to design and implement new features.

The GFDRR team understood the benefits of both kinds of extensibility, and knew that not only did no proprietary system fully meet their needs, but even if one came close, it would still have both less on-the-fly extensibility and less deep extensibility than GeoNode had the potential to achieve¹². Thus the choice they really faced was:

Given the present technical needs of GFDRR and its clients, and likely future needs, what are the relative merits of running a purely product-oriented software development process to create tools to address those needs, versus the more complex task of creating an open source software system and a community around it?

12. One interviewee, a well known and widely experienced geospatial developer, said of government agencies now switching to GeoNode: “They don’t get any flexibility to customize [the proprietary software they had been using]... they can’t get [the vendor] to customize it for them. They get a lot more flexibility with open source and GeoNode”.

Note that “more complex” does not necessarily mean more expensive, in terms of direct costs. It is possible to roughly calculate the hypothetical cost of the road not taken, that is, how much it would have cost GFDRR on its own to bring a piece of software of GeoNode’s scope and functionality from conception to its present maturity. The answer is approximately \$2.0 – \$3.0 million¹³ USD, whereas the amount GFDRR actually spent is \$1.0 – \$1.5 million.

So if the goal was simply to get GeoNode-like functionality, GFDRR achieved that in a fairly thrifty way. In addition to benefitting from the early involvement of key partners, GeoNode also saved time and money by integrating itself closely with existing open source technologies, including but not limited to geospatial data tools.

While no software today is written fully from scratch, GeoNode stands out for the extreme extent to which it took advantage of already-available open source platforms and code libraries, and thus avoided writing unnecessary new code. Open source platforms like Django, GeoServer, Bootstrap, jQuery, PostGIS, Geospatial Python, OpenLayers and GeoExt are not widely known outside programming communities, but they represent many millions of dollars of investment, and any experienced developer, on learning that they are the foundation on which GeoNode is built, will immediately grasp how much time and development effort they saved the project.

INVESTING IN A COMMUNITY, NOT JUST IN SOFTWARE

Building just a software product was never GFDRR’s only goal. Their broader goal was to have a system that offered GFDRR and its partners deep extensibility, along with a rich worldwide network of commercial and informal support. The combination of open source extensibility and widely-available expertise meant that countries and other institutional users of GeoNode would not be locked into either GFDRR nor into a monopolistic single vendor to obtain expert help in taking advantage of that extensibility.

Another way to put it is that by making GeoNode open source, GFDRR successfully avoided being in the software business. Instead, it positioned itself as a participant in an open source project, which meant that it could later change the size or nature of its investment without endangering the project as a whole, and without endangering the

availability of third-party support to World Bank clients who became reliant on GeoNode.

GFDRR may have initiated the GeoNode project, but at this point they are just another partner, making new investments based strictly on whether those investments are an efficient way to achieve their own specific goals, rather than out of any existential concern about the project, which is now self-sustaining.

In fact, the largest current investment in GeoNode development comes not from GFDRR but from a consortium of U.S. government agencies who are using GeoNode for purposes related to, but not exactly the same as, those that GFDRR originally started the project for. The tremendous amount of core development work they are doing — well over \$1 million USD worth — will benefit GFDRR and its partners as well, and this continued investment comes at no cost to GFDRR. It is, rather, a return on GFDRR’s original investment.

As one of the developers working for that consortium said: “GFDRR laid the foundation and now [these agencies] have something to invest in... The project has reached some level of stability and maturity where these organizations feel like it’s the obvious place to build new features.”

THE ROI OF AN OPEN SOURCE PROJECT

Given that the GFDRR team understood the potential benefits of running GeoNode as an open source project, and chose that course, how successful were they in implementing it?

There are a couple of ways to evaluate the success of an open source investment. The narrower way is to look strictly at cost savings, that is, at expenses foregone because they were covered by partner investments. As we saw earlier, GFDRR did quite well in this regard, with approximately a 1x ROI, i.e., 200% — it got about twice as much as it paid for, in terms of raw code development. This is the maximally conservative, quantifiable portion of the ROI: we base it on standard rates of software development, and on a survey of other institutions’ past contribution to and reliance on GeoNode. This estimate does not even include the current wave of new investment that is being poured into GeoNode today from non-GFDRR sources,

30 ¹³. This is based on an analysis of the development history of the GeoNode code repository history. We looked at the ratio between code contributions known to have been funded, directly or indirectly, by GFDRR and those that came from other sources, including but limited to the partner organizations mentioned in this report.

“GFDRR laid the foundation and now [these agencies] have something to invest in...The project has reached some level of stability and maturity where these organizations feel like it's the obvious place to build new features.”

except to treat it as an affirmation that past investments achieved their intended results and were not wasted.

Thus, looked at purely in financial terms, open source was a good choice.

But a broader and more useful way to evaluate success in open source projects is to look at the community created by the investment, and at the benefits that community continues to bring to the original investor as well as to its other members¹⁴.

There is as yet no standardized framework for that kind of evaluation; the most important aspects of ROI are also those least amenable to easy quantification. However, there are some proxy measurements we can use. Some of the best to look at are growth in the supply of commercial support options, certain trends in the types of interactions happening in the project's development forums, and growth in the participation rate in community's discussion forums — these latter two may be summed up as growth in “community complexity” or in “community depth”.

GROWTH IN COMMERCIAL SUPPORT OPTIONS

Because GeoNode is open source software, there is no centralized registry listing all known GeoNode instances in the world — GeoNode instances do not automatically “phone home” back to their original supplier to report on how they are being used, nor even to report on the mere

fact that they are being used, because that would violate the default privacy expectations of those who set up the instances¹⁵.

However, a good proxy for relative growth in the number of GeoNode installations is to look at the number of commercial entities offering support for GeoNode — customization, deployment assistance, maintenance and operations, training, etc. The appearance of such entities not only tracks the growth of a market, it also signifies increased stability in the GeoNode development ecosystem, since each entity is a reservoir of expertise and development ability that is, at least in part, economically independent of the other organizations involved in maintaining GeoNode. That is, each new commercial actor is one more party who has an interest in GeoNode flourishing: they would not have made the investment in acquiring expertise unless GeoNode looked like a good bet, and now that they have made the investment, they are unlikely to let it go to waste.

In 2009 — and, speaking roughly, even through 2010-2012 — there was, more or less, only one commercial outfit that did significant business based on GeoNode: OpenGeo (today named Boundless Geo). There were a few independent contractors here and there, but to the extent they worked on GeoNode, most of them did so as subcontractors to OpenGeo.

In 2016, by contrast, two GeoNode developers named 17 sources of commercial support for GeoNode off the top of

14. In other words, the project's Social Return on Investment (SROI) — see https://en.wikipedia.org/wiki/Social_return_on_investment. Although out of scope for this report, a general framework for measuring Open Source SROI would be a valuable tool, and the beginnings of such a framework are implicit in our evaluation here. If some of GFDRR's partners are also starting to create open source projects, as appears anecdotally to be the case (based on our interviews for this report), it might be worth exploring the creation of some shared standards for Open Source SROI.

15. Upcoming versions of GeoNode will have an opt-in “phone home” feature, so that the GeoNode project can track where GeoNode is being used in those cases where the managers of the instance explicitly agree to that information being shared. This will make gathering statistics about GeoNode usage much easier in the future; however, this feature had not even been released at the time of the writing of this report.

Commercial Support Providers for GeoNode

2009—Approximately 2011

OpenGeo

As of June 2016

Gesp (Italy)
AllSpatial (Seychelles)
LTS International (UK/Indonesia/Malawi/Kenya)
Prominent Edge (USA)
Piensa (location unknown; team fluent in Spanish and English)
Terranodo (USA/Colombia/EU)
AcuGIS (USA)
Kartoza (South Africa)
Geo Enviro Omega (Indonesia)
Applied GeoSolutions (USA)
Boundless (USA; formerly OpenGeo)
Agrisoft (Indonesia)
Team Geode (Belgium)
Geocent (USA)
Quantitative Engineering Design (USA)
GeoSolutions (Italy)
ADPC (Asia/Pacific; not strictly commercial but cited by developers as a source for GeoNode development expertise)
Structured Data Systems Limited (SDSL) (Bangladesh)
GeoEdge (Sri Lanka)
iMMAP (Afghanistan / US)
RMSI (India)

FIGURE 3: Commercial Support Providers for GeoNode.

Disclaimer: The 2016 list is ad hoc and incomplete, being biased toward companies that have interacted with GeoNode developers and/or GFDRR.

their heads, each emphasizing that these represented only a "partial list."¹⁶ To these we have added a few more from other sources, with the total — but still incomplete — list shown on the right side of Figure 3.

Furthermore, the sizes and natures of these companies vary widely. Some are established geospatial technology services firms that long predate GeoNode, and who likely decided to enter into GeoNode support only after the software had become a clearly-established player in their industry. Others are relatively new startups founded by small groups of GeoNode developers, who realized that their GeoNode experience gave them a natural advantage in a growing market, and decided to seize the opportunity. Still others lie somewhere on a spectrum between those two, or are something else entirely, for example a data analysis shop with an open source background that, we

may guess, probably saw GeoNode as a route to handling more geospatial contracts.

The number and diversity of commercial support offerers available today, compared to the first few years of the project, suggests quite large growth in the number of current and potential customers, and is a clear sign that GFDRR's initial investment helped to reveal a previously-unsatisfied pool of demand for what GeoNode offers.

GROWTH IN DEVELOPER AND USER PARTICIPATION, AND IN COMMUNITY COMPLEXITY

Another good proxy measurement is to look at where people devote their time and attention. In general, people would not expend sustained attention on something unless

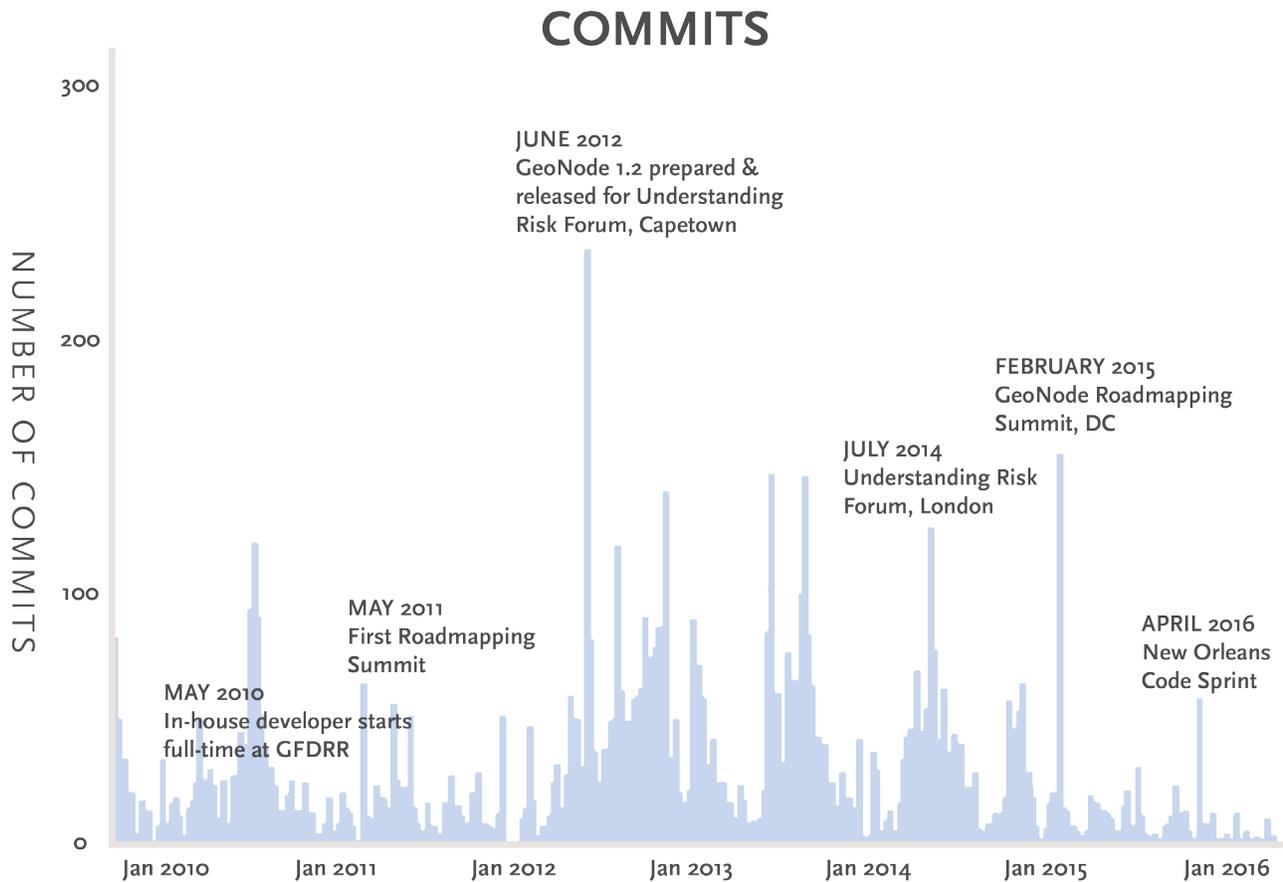


FIGURE 4: Code changes (“commits”) per week. A commit represents one unified code change submitted by a developer. A commit may be a simple issue fix, or one step in a multi-part feature addition, or a documentation update, etc. In general, the rate of commits is a good rough proxy for overall project activity. This chart shows that commit activity increases in response to institutional investment; the selection of events labeled here are all ones that GFDRR was involved in. Most notable is the fact that the sustained spike in commit activity that started right after GFDRR hired a full-time GeoNode developer in May 2010 is mostly not made up of commits by that developer. In other words, the spike consists of other developers reacting to the GFDRR developer’s arrival by doing more work themselves.

it were worth their while to do so. Fortunately, the nature of open source projects is such that attention is often quite visible: we can see who posts to the discussion forums and when, what they’re asking, and who responds.

For a significant span of GeoNode’s history, the primary discussion forums were in Google Groups, so we can even see relative changes in how many people viewed the conversations. We can also look at activity in the project’s issue tracker (explained below), and at changes in how developers respond to activity by other developers.

All of these things represent investments of attention. Wherever we see other parties increasing their investment, correlating with and slightly preceded by an investment by GFDRR, there is a good chance that we are seeing a cause-effect relationship — a return on investment.

Other parties spend more (attention or money) when they see GFDRR spend more. Thus, the trick for GFDRR was to make sure to do their spending in a way that was visible to the community and that sent signals about where it would be most productive for attention to coalesce.

DRIVING THE RATE OF CODE IMPROVEMENTS

Figure 4 shows some of these effects by looking at the rate of code improvements¹⁷.

The spike around May 2010 is right when GFDRR hired a full-time developer, but note that the majority of the code changes inside the spike are not by this developer; in fact, only about 15% were by him. It appears instead that this hiring was a catalyst for involving others.

17. Known as “commits” among programmers.

Open source developers typically become more productive when those around them become more productive, and GFDRR’s in-house developer had just begun working on GeoNode full time — a significant increase in daily development energy, although a few developers at OpenGeo had already been working on the project for a while by that point. Most of the commits in that spike are from OpenGeo developers. (Note also that this period overlaps with the run-up to GeoNode’s first official release: in August 2010 GeoNode entered beta-testing for its 1.0 release, with the production 1.0 release coming in December 2010.)

Other spikes correlate with GFDRR’s investments in hosting code sprints and roadmapping summits, where attendees are motivated by the presence of other attendees and the opportunity to form personal bonds through shared work.

IMPROVED RESPONSIVENESS TO USER FEEDBACK

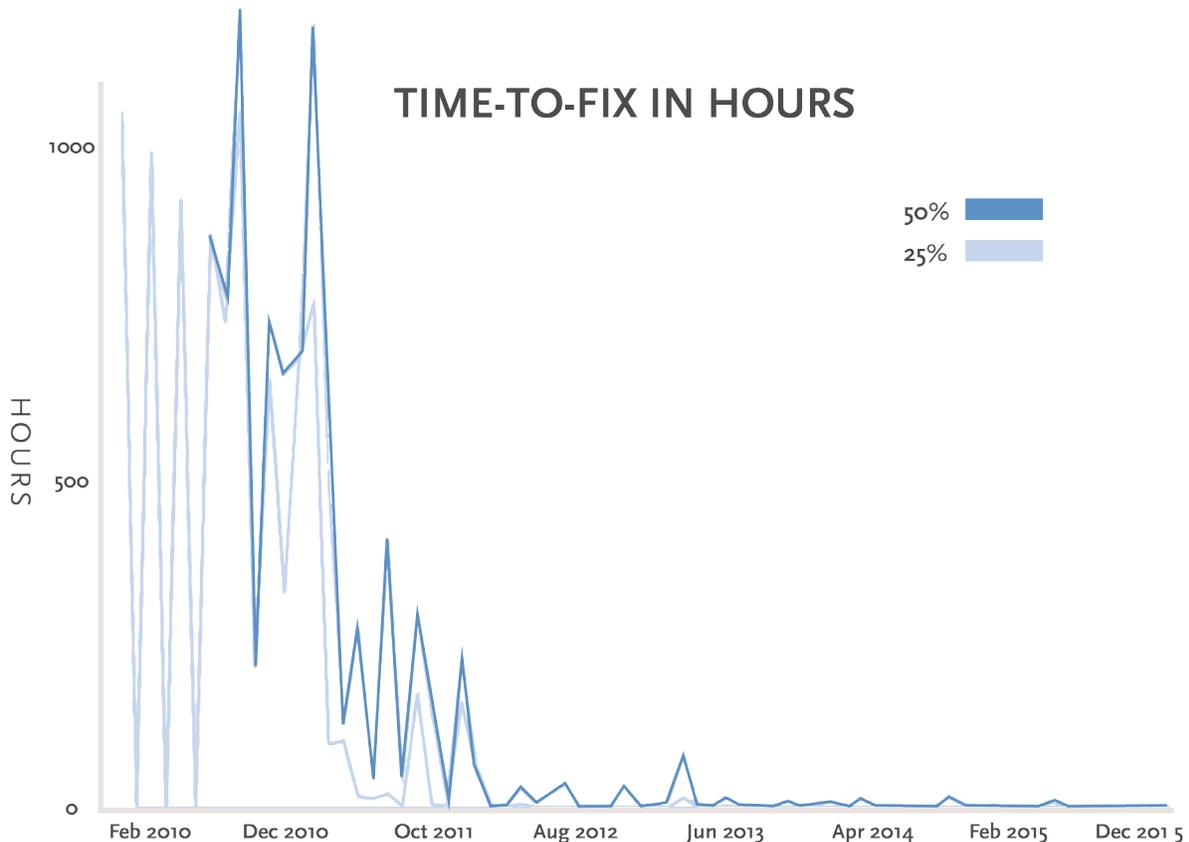
We also looked at changes in the time-to-fix rate for issue reports. In GeoNode, as in most open source projects,

problems in the software are filed as tickets in a ticketing system, colloquially known as an “issue tracker” or “bug tracker” (in this report we use the word “issue,” which is synonymous to “bug”). Each issue has its own ticket, which has a unique number and contains all the information the user -- that is, the reporter -- can provide to the developers about how to demonstrate the issue and diagnose its causes.

A ticket often records a dialogue between users and developers as it moves through a sequence of well-defined states: e.g., “new” > “verified” > “fixed” > “resolved”. By analyzing data from the issue tracker’s archives, in particular how long it takes a ticket to go from “new” to “resolved”, it is thus possible to get a statistical overview of the community’s responsiveness to issue reports.

In Figure 5, the GeoNode project shows a very healthy trend, in which issue tickets that can be handled right away are fixed with increasing promptness over time. For this trend to occur in a project which is getting an increasing number of users and deployments over time, as GeoNode certainly was during the period shown above, it

FIGURE 5: Time-to-fix in hours for the quickest-closed 50% and 25% of issue tickets. As the project became more mature, its processes for handling incoming issue reports improved. The trends are similar for the other closed tickets not represented here, but for demonstrating community responsiveness, it makes sense to look at the most-quickly-closed half of issues, since there are entire classes of issue reports that by their nature cannot be handled promptly, such as feature requests, edge-case issues that are difficult to diagnose and affect few users, etc.



GEONODE USERS GOOGLE GROUP POSTS

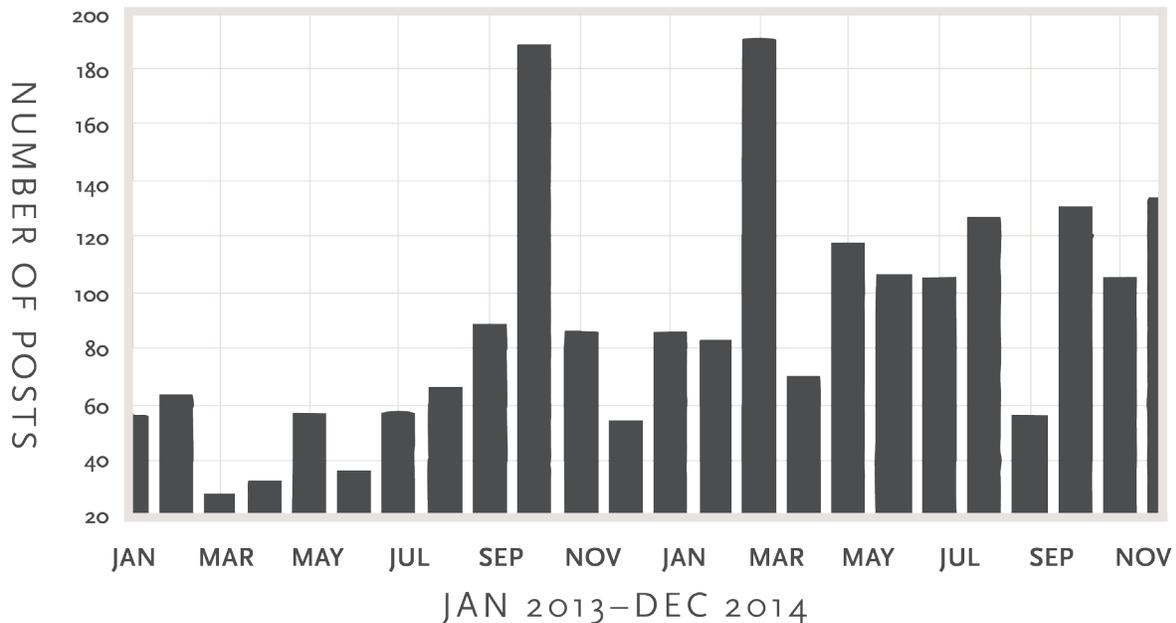


FIGURE 6: Posts to the GeoNode Users Google Group, over the two years 2013 and 2014. Aside from the general upward reflecting GeoNode’s growth, the interesting features are the peaks, which correspond to GeoNode release activity and to in-person events. For example, the sharp increase in traffic in late September 2013 coincides with the FOSS4G Nottingham; the one in March 2014 with the final beta-testing period of GeoNode 2.0, a major release of GeoNode. These trends should be compared with the views counts for the same Google Group, shown in Figure 7 (please note the difference in vertical scaling — vertical distance counts for much more in the views chart). Both charts conclude in December 2014 because the GeoNode mailing list forums left Google Groups at the beginning of 2015 and are now hosted at the Open Source Geospatial Foundation.

means that, at least as far as the issue tracker is concerned, developer attention is generally keeping pace with user involvement, and even getting ahead of the curve a bit. Otherwise the time-to-fix rate for issues would have either remained constant or gone up -- which is exactly what can be seen in some other open source projects that have widespread adoption but whose developer communities have not grown proportionally to that adoption.

The developer responsiveness represented in Figure 5 is a relatively smooth and (so far) permanent improvement: it implies a project community that is functioning increasingly well over time.

This too is a return on GFDRR’s investment, but the word “return” may mislead by implying a one-time return. The returns here are qualitatively different from returns on financial investments. There is no precise way to calculate the value to GFDRR and its clients of having a healthy and long-lasting project community, one that that GFDRR can

always step into or away from depending on their needs at the moment, without wondering whether the community will be there next year. It is the difference between planting a crop once for later harvest, and making a permanent increase in the size and fertility of the entire field.

GROWTH IN COMMUNITY PARTICIPATION AND ATTENTION

Another proxy measure of the community’s attention is activity on the GeoNode mailing lists¹⁸. This is a bit hard to calculate for GeoNode because the project has migrated between at least three mailing list services from 2009 to 2015, with some periods of list overlap. However, there is a clear two-year period comprising all of 2013 and 2014 when the GeoNode Users mailing list was consistently hosted on one Google Group.

Figures 6 and 7 show the levels of posting and viewing activity, respectively, during that period.

18. Mailing lists are typically used in open source projects to coordinate activity and discuss upcoming changes. For many open source projects, especially those that began more than two or three years ago, the mailing list is the primary place where developers and users interact with one another.

GEONODE USERS GOOGLE GROUP VIEWS

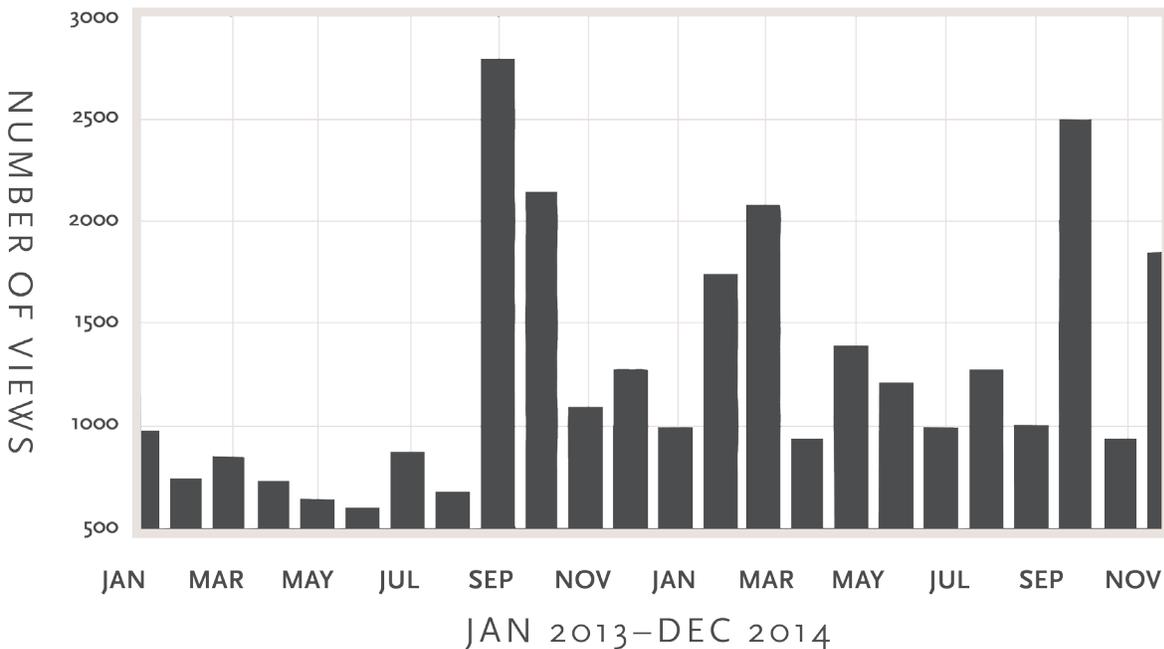


FIGURE 7: Views on the same GeoNode Users Google Group, again over the two years 2013 and 2014. Compared with posts (Figure 6, and note the difference in vertical scale), the increase in views is much sharper when it starts, and is more sustained afterwards. In other words, coinciding with the 2.0 release of GeoNode, a lot of people became interested and stayed interested in the project. Note that this chart only counts views that came in through the Google Groups web interface, not views by people who read the group via email-based subscription — so it is useful for seeing relative growth but probably undercounts the absolute numbers. (It is also likely that a higher proportion of developers than users subscribed via email, if GeoNode developers behave like developers in other open source projects, so the data probably reflects growth in user views more strongly than it reflects growth in developer views — which for our purpose here is a desirable bias anyway.)

Both charts show a permanent step up in activity in the second half of 2013 — when GeoNode 2.0 was finishing beta-testing and being released, and GeoNode was being distributed (on physical media) to around 1000 people at the FOSS4G conference in Nottingham, England.

However, looking at the difference between the two charts is revealing: the view rate showed a sharper and more sustained increase: relatively suddenly, a whole lot of people got interested in GeoNode and stayed interested. Even those not posting were still giving the project their attention.

This increase in interest was actually anticipated by the GFDRR team. GFDRR’s in-house developer posted¹⁹ to the Users group on 13 Sep 2013, saying, “Let’s prepare to help out new users in case the mailing list starts to get more traffic than usual.” The post indicates that the core development team, including GFDRR staff, knew the value of all those viewers and took care to keep the list valuable for them.

In general, GFDRR made a very deliberate investment in setting up a self-sustaining community on the mailing list. Although we did not count it separately when listing investments earlier, some fraction of GFDRR staff time was always devoted to reading and responding to user questions.

Even a cursory browse through the archives of the GeoNode Users mailing list, whether as a Google Group or in or any of its other instantiations, shows that this investment paid off. A great deal of the forum’s activity is users helping other users, often passing along advice that originally came from one of the core developers at OpenGeo or at GFDRR (see the sidebar “Examples of mailing-list interactions”).

The overall result of hundreds of this kind of interaction happening over time is that a finely tuned and socially maintained knowledge base is built up, archived and searchable, waiting for new GeoNode users to find in it the answers to their questions. As the project matures, less and less of this knowledge base needs to be produced by

EXAMPLES OF MAILING LIST INTERACTIONS

Two examples of how GFDRR staff thoughtfully interacted with the community on the GeoNode mailing list.

In general, GFDRR staff showed a keen understanding of the “long tail” value of the mailing list archive. An answer posted today might show up much later in someone’s search results when they’re looking around on the Internet for help solving some problem in GeoNode. Here we show two examples of the kinds of interactions that lead to a project’s mailing list archives becoming a valuable community resource.

1. In the conversation “Archivos de Estilo , SLD”²⁰, started in Spanish by Sergio Arnal on 10 August 2013, GFDRR’s in-house developer responded in English to Arnal’s technical question, and then in a postscript explained — in Spanish — that he had used English so others who saw the conversation could benefit from his answer:

PS: Te respondo en inglés para que otras personas puedan beneficiarse de la respuesta :)

Often, people who have solved a problem will then do an archive search themselves, to see who else has had the same problem in the past, and then post their solution as a followup, both to offer belated help and to make sure that the question and answer show up together in future Internet searches by others. Something like this seems to have happened here: more than a month later, on 17 September 2013, a third person made a new follow-up post, in Spanish, with some very specific technical suggestions for Sergio Arnal to try.

2. In the conversation “Error upgrading to Geonode 1.1”²¹ started on 15 February 2012, user “Thomas” (who appears to be with an Italian company that does some geospatial data work — a sign of GeoNode’s success in getting commercial adoption) asks a question about an installation problem, saying that he has searched the Internet for answers and citing one specific earlier post from the previous May on the same GeoNode Users mailing list.

Thomas receives a response that same day from Joe Larson, who in turn points to a post of his own from just a week before, and speculates that Thomas is perhaps having the same problem that Joe’s earlier post describes. Joe then quotes some text from an email exchange he had had earlier with a GeoNode developer that offers some tips for diagnosing this kind of problem — at which point a different developer (David Winslow, at OpenGeo not at GFDRR) joins the conversation, asking Thomas and Joe for some more technical details and offering further tips on how to trace the problem to its cause.

20. <https://groups.google.com/forum/#!topic/geonode-users/yOioVQfQjok>

21. <http://librelist.com/browser/geonode/2012/2/15/error-upgrading-to-geonode-1-1/#d9b53e12881e97dbd90b3f1feb810e67>

DEVELOPER RESPONSIVENESS

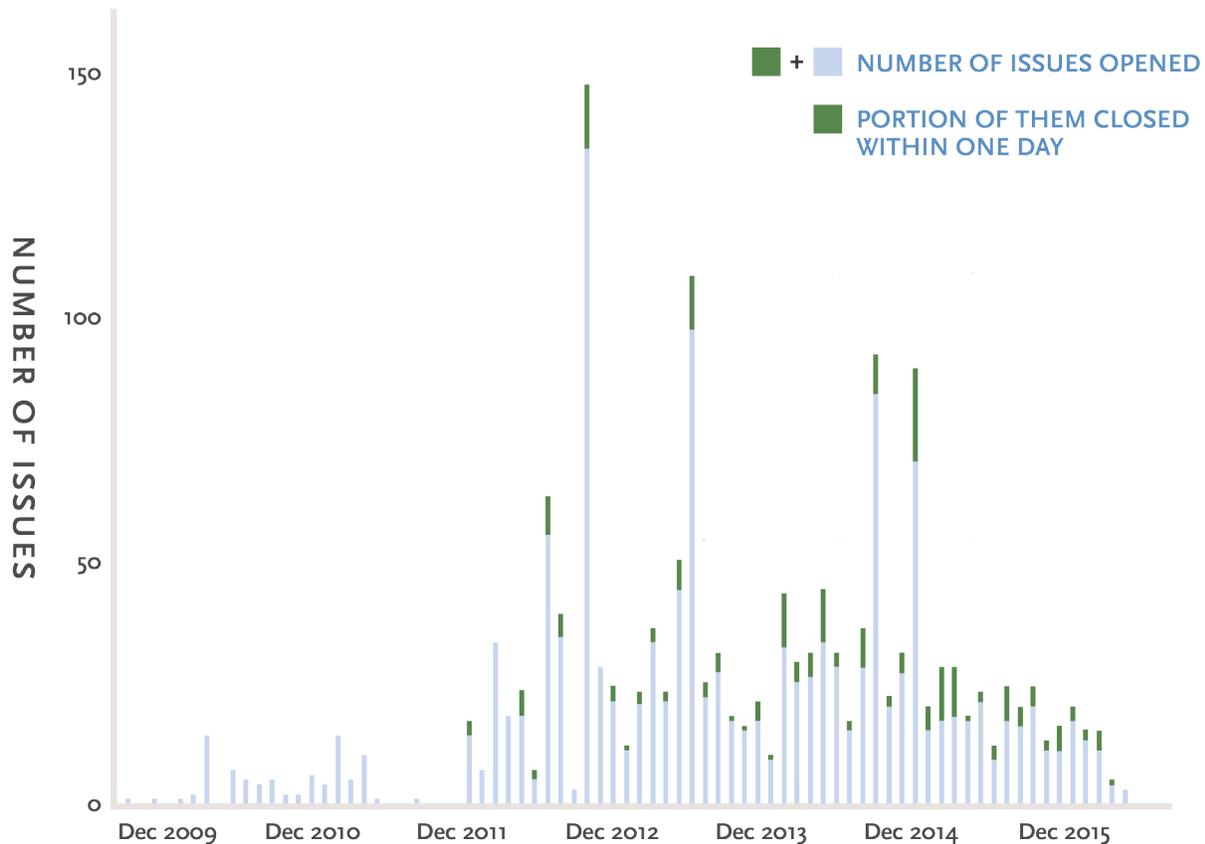


FIGURE 8: Issues opened each month, and the portion of those issues closed within one day. Similarly to Figure 5, this shows the project generally improving at responsiveness to user reports over time.

core developers; the users increasingly supply it themselves (often by expanding on things they learned earlier from developers), in a kind of long-term gift economy.

This dynamic is very familiar in open source projects, and achieving it, as the GeoNode project has done, must be counted as a significant return on investment for GFDRR, one that will continue to pay dividends for GeoNode users — including GFDRR itself — for many years to come.

GROWTH IN GEONODE ADOPTION

Growth in GeoNode adoption is also a return on investment — in particular, countries and organizations adopting GeoNode on their own, as opposed to doing so with external encouragement and funding.

This is harder to quantify post facto, but is apparent to those closest to geospatial data management projects in the field.

Bishwa Raj Pandey, a disaster risk management information specialist at the World Bank who was very active in Latin America and the Caribbean during the time covered by this report, told us that in that region there are now 15 countries using GeoNode, either because the World Bank introduced them to it (as part of data management in a project plan) or because they themselves were interested.

Pandey described the return on this investment as “huge” and gave the example of Belize, where the World Bank made only a small investment. According to Pandey, there was not only a lot of user uptake of the system in Belize — intrinsic uptake, as opposed to the World Bank coming in and recommending GeoNode usage — but experience with GeoNode also provided a motivation to use open source development for other projects.

USER ENGAGEMENT

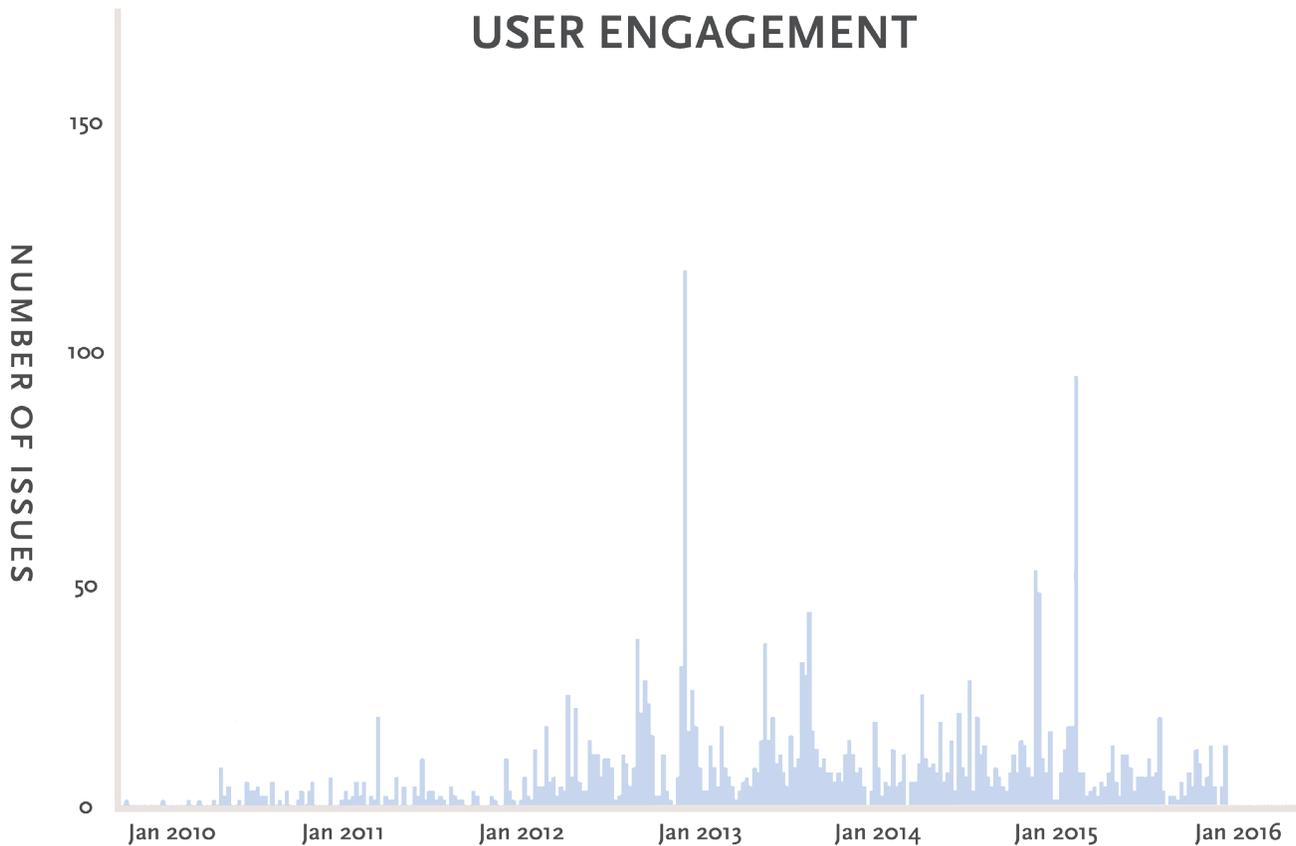


FIGURE 9: Rate of issues (bugs) being opened per week, across the history of the GeoNode project. Issue ticket filings are a non-linear proxy for growth in users and developers. As a project leaves its early stages and enters mature usage and development, there is an initial increase in the rate of issue ticket filings, because more people are exposed to the software and therefore encounter more new issues. The key word is “new”, however: once the project’s issue tracker database is well-seeded, the probability that a particular issue encountered by any given user or developer is new goes down — most issues are ones someone else has already encountered and filed. The new reporter may sometimes add a comment to an existing issue ticket, or often simply track the ticket but leave no comment nor additional information. Therefore, to see the rate of new filings remain fairly steady, as it does here — except for event-related spikes — is an indication of continued growth in user and developer base, because if the set of involved people were remaining static, then over time they would have fewer and fewer previously-unencountered issues to report, and we would see the filing rate going down.

Pandey remarked that GeoNode had changed the whole concept of how data is managed in some of the countries he was familiar with — that now people say “oh, we have GeoNode”, meaning, in Pandey’s words, “we have a system where instead of going to each department of the country to find the data on a USB drive, now we have a place where you can go and get the data.”

He also pointed out how GeoNode makes OpenDRI’s job easier. For one thing, it is used by GFDRR to see whether contractors in client countries are fulfilling their data-collection obligations, because it’s easy to check if the data is in the relevant GeoNode server or not. For another, GeoNode’s record-keeping features will note a spike in usage where there is sudden need.

For example, when Tropical Storm Erika damaged infrastructure in Dominica in 2015, suddenly many people

were looking at data related to Dominica. Some of them had data to contribute, too, and the World Bank was able to encourage them to work with the government because the government was already using GeoNode. Pandey noted similar effects in Saint Lucia and Saint Vincent, and made the general observation that countries are using GeoNode as a central repository, and especially trying to use it during disasters to guide response.

OTHER INDIRECT ROI

In addition to the returns described above, we may count two other indirect but significant returns on investment for GFDRR: lowered barrier-to-entry for clients, and the building of local capacity.

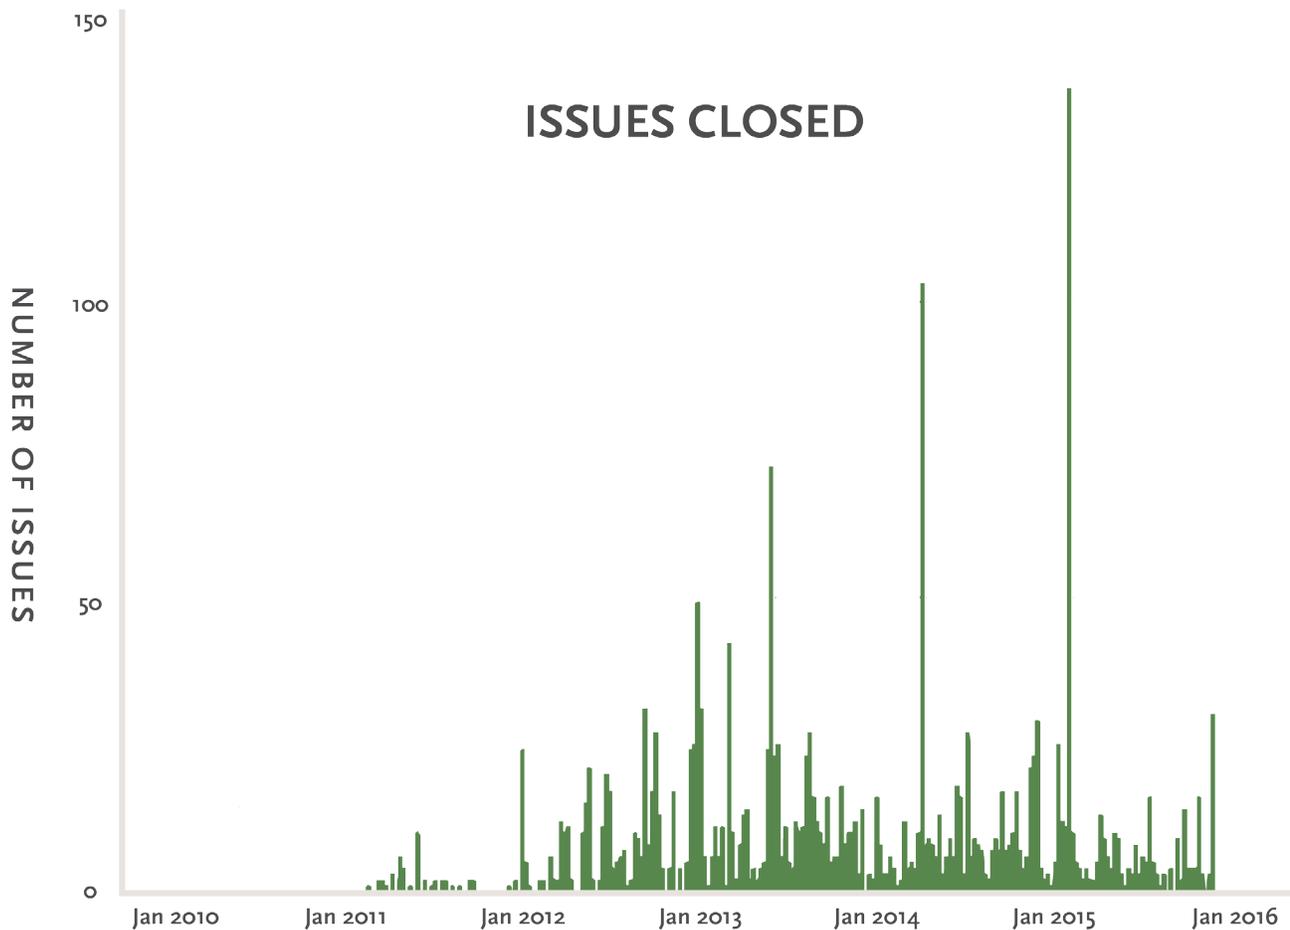


FIGURE 10: Closing issues is a sign of a responsive development team. As this chart shows, in general the GeoNode team maintains steady attention to clearing out the issue database, with spikes of increased attention that (based on our timeline of events, see pages 9-14) appear to be related to in-person gatherings and code sprints, which have a well-known effect of motivating developers to “clean house”.

Lowering barrier-to-entry

GeoNode significantly reduces the barrier-to-entry for client countries launching geospatial projects. GeoNode itself is universally available at no charge, of course, and GeoNode expertise is increasingly widely available as well. Even in resource-challenged environments it is now realistic for local institutions to deploy GeoNode with little or no GFDRR assistance, both for their own purposes and to collaborate with GFDRR.

Building local capacity

GeoNode empowers countries, and local universities and NGOs, to autonomously maintain and customize systems that meet their own geospatial needs, and to develop internal expertise in the course of doing so. Furthermore, it

provides opportunities for in-country companies to expand their business by providing support and development locally — GFDRR observers have seen this happen in Bangladesh, Sri Lanka, and Madagascar, for example — and it provides opportunities for local innovation, some of which then makes its way back to the core (global) GeoNode project.

THE CREATION OF A PUBLIC GOOD

The return on investment that is most difficult to quantify, and yet that is perhaps most important in the long run, is that GeoNode has enabled hundreds, perhaps thousands, of people and organizations around the world to do things more easily and to to them together.

The fact that some of those people and organizations have missions that overlap with the mission of GFDRR is merely the near edge of a large and unpredictable ecosystem. Robert Soden of GFDRR observed in mid-2016: “There are hundreds of GeoNodes now, and new ones springing up all the time. It’s become completely impossible to keep track of...”

These are just a few examples of recent GeoNode installations that were done entirely independently of GFDRR:

- <http://geoinfo.iplaneg.net/> – The Institute of Planning, Statistics and Geography in Guanajato, a small-medium size city in Mexico, deployed a fairly straightforward, “out-of-the-box” GeoNode.
- <https://geonode.senasa.gov.ar/> – The National Directorate for Vegetal Protection in Argentina set up GeoNode to track plague prevention programs, with a customized homepage.
- <http://asdc.immap.org/> – USAID deployed GeoNode in Afghanistan, with some custom filtering and statistics display features.

New people are now installing their own GeoNode instance on a weekly and sometimes near-daily basis, as can be seen from watching mailing list activity²², where questions often appear from people who are setting up GeoNode for a university, a local government, or with a private firm — on their own, without any involvement by GFDRR, as the

software itself is now well documented and easy to deploy

Any time an improvement is made in GeoNode by any of its contributing partners, all of these many people and institutions benefit. GeoNode is in this sense a classic public good, but this dynamic also represents an extended return on investment for GFDRR. New users provide feedback, and when similar feedback is received from multiple sources, the GeoNode maintenance team knows it represents a common user need, one likely shared by many of GFDRR’s partners. Every time the maintainers fix or improve something in response to that feedback, everyone benefits — which in turn draws in more users, in a snowball effect.

And, of course, some percentage of those new users will themselves go on to become active participants in GeoNode maintenance, by contributing code, testing for issues, writing documentation, etc.

Just as important as the data provided by all these GeoNode sites are the conceptual and cultural effects of GeoNode’s spread. Risk reduction experts, disaster response specialists, and scholars and researchers are becoming accustomed to using modern web-based geospatial techniques in their daily workflow. By being available as a public good, GeoNode has played a significant role in training people and setting expectations for smooth technical collaboration on geospatial data, which is a long-term benefit to GFDRR and its partners.

22. See the GeoNode Users mailing list archives: <http://lists.osgeo.org/pipermail/geonode-users/>

Coding Events 2014–2016: Understanding Risk Forum (London, England, UK), GeoNode Roadmapping Summit 2014 (Washington, DC, USA), GeoNode Roadmapping Summit 2015 (Washington, DC, USA and Venice, Italy), GeoNode Roadmapping Summit 2016 (New Orleans, Louisiana, USA).



6

EMERGENT BEST PRACTICES

The OpenDRI / GeoNode collaboration is an example of successful institutional investment in an open source project, and it is worth stepping back and attempting to draw some conclusions about what aspects of OpenDRI’s strategy with GeoNode might be generally recommended for new open source projects started by GFDRR or other similar institutions.

The recommendations below are based primarily on our study of the GeoNode project, but are also influenced by what we have observed — both positively and negatively — at other non-profit organizations and government agencies that have helped launch open source projects.

1. RUN THE PROJECT IN THE OPEN FROM THE START

One thing GFDRR got right (that many organizations don’t) is that they didn’t wait until they had working software to start behaving like an open source project. Instead, they ran the project in the open from the very beginning. This sends a signal to potential collaborators that the door is really open, and it ensures that no behind-closed-doors development habits arise among the initial development team.

It also ensures that there is a public record of the evolution of the code, and of all the design decisions made from beginning of the project. This helps other participants who join the project later: when they want to understand why a decision was made a certain way, there is a trail they can follow to find out the answer.

2. ENGAGE OTHER ORGANIZATIONS COMMERCIALY TO DIVERSIFY INVESTMENT

GFDRR saw commercial involvement in the project as desirable, because having software development shops with a commercial interest in GeoNode’s success is a good way to ensure the project stays healthy. To this

end, GFDRR contracted with multiple companies for development work. A developer who worked at one of these companies at the time told us that they soon had contracts with other, non-GFDRR-related, customers for GeoNode work.

Beyond GFDRR’s contracts with companies for coding, the team also worked with universities and other organizations to build GeoNode, as discussed above. Local providers grew up in many different countries to support their GeoNode installations. The overall result is that the global supply of GeoNode expertise is not housed under one management hierarchy, which is good, and yet the disadvantages one might expect from such decentralization are mostly absent, because the open source project itself provides a place for that expertise to congregate and circulate.

3. EVANGELISM & ACTIVE COMMUNICATIONS MAKE A DIFFERENCE, ESPECIALLY IN THE EARLY STAGES

A few of the interviewees mentioned that sustained early efforts by the OpenDRI team to simply explain and champion GeoNode, both inside GFDRR and among its clients and partners, were crucial to getting mindshare in the early days of the project. Interestingly, one of the more technical interviewees recalled being surprised at the amount of socializing work that the team was putting in, and said he did not appreciate until later that this work had been as important in securing GeoNode’s

future as the technical discussions and software development on which he spent most of his time.

4. FIND & ENCOURAGE THE RIGHT PARTNERS

The importance of this to GeoNode is described in the section “Methodology of investment”. We would only add here that it is important to keep the previous point in mind as well: some persistence in advocacy may be needed before partners understand what they have to gain by joining the project.

The organization — really, the people — that started the project know it the best and are already committed to it. They do not need to have explained to them what problems it is going to solve and how. But for everyone else, including some potentially very good partners, the project may be just one more thing making a claim on their attention. Finding the right partners is important, but you may also need to patiently encourage them.

5. INVEST IN INFRASTRUCTURE AND PROCESS, TO AMPLIFY OTHERS’ WORK

Although it is tempting, in a software project, to focus organizational investment directly in writing code, this is not always the best allocation of resources. When an organization has a clear overview of the project’s needs, and is able to spot impediments to new partners and contributors participating, the best use of that organization’s resources is sometimes to invest in improved infrastructure and process, benefitting the project by amplifying everyone else’s work. An example of this is examined in the subsection “Promoting a culture of documentation” on page 26.

6. HOLD EVENTS, AND HELP PEOPLE ATTEND THEM

Every interviewee who had hands-on experience working on GeoNode mentioned the importance of in-person events. Even a brief face-to-face acquaintance causes people to be more willing to make efforts for each other for a long time afterward, to compromise, to engage in constructive discussions, and to trust each other, particularly once collaborating remotely again. See the subsection “In-person events” on page 25 for a detailed example.

7. INITIAL INJECTION OF FUNDING IS BOTH AN INVESTMENT AND A SIGNAL

An organization’s initial commitment of funding is not only important in terms of the technical results it

produces (e.g., lines of code written), but in terms of the signal it sends to potential partners. Committing funding means also committing your institution’s credibility; depending on the institution, that can be quite influential, as it was with GFDRR and GeoNode.

8. TO EXPAND THE RANGE OF USERS, INVEST IN IMPROVED USER EXPERIENCE

The earliest partners in a project will often tolerate a less polished overall user experience than would be expected from a mature project — they knew they were getting in on the ground floor, while the building was still under construction. But it is important not to let developer acclimatization to that initial user experience become complacency. It is all the people who try the project once and walk away who hold the key to future expansion of user base.

Aware of this, GFDRR made a significant investment, from late 2011 through 2012, in user interface improvements to GeoNode, as part of a deliberate — and successful — attempt to encourage buy-in by a wider range of people and organizations. This kind of investment is sometimes harder to justify, on its face, than technical work that leads to direct benefits for the current users of the project. However, in the long run, it is important to do, because of the broader user base that results.

9. THE NATURE OF YOUR INVESTMENT CAN CHANGE OVER TIME

It is normal for the nature and size of an organization’s investment in an open source project to change over time. Neither the project nor the organization remain static, and even a founding organization may find eventually that it is able to get the benefits it needs with either a lower investment or simply a different kind of investment than it had been making in the past, especially after the project has attracted many participants and become effectively self-sustaining.

Furthermore, a reduction or change in role now does not rule out further change later: an organization may reduce its involvement a project for a while, only to come back later and become highly involved again. It just depends on what the organization’s goals are. Other participants in the project will understand this.

7

CONCLUSION

For all the returns already realized on GFDRR's investment in GeoNode, the most important ones may still be in the future. GeoNode is today a flourishing open source project, with increasing uptake by organizations that will be able to support its continued maintenance and development independently of GFDRR. OpenDRI has contributed to the creation of a public good, but a public good of a very particular kind: GeoNode is both a software tool and a software project, and as a project it serves as a gathering place where inter-organizational cooperation can happen with a minimum of bureaucracy and with immediately tangible results.

The institutions listed earlier in the section "Methodology of investment" were already a diverse group, although most of them had in common at least a shared interest in risk assessment and disaster recovery. The institutions involved in GeoNode six years from now are likely to be both more numerous and far more diverse, and this is perhaps the most significant long-term return GFDRR will see on its investment in GeoNode. The presence of those institutions means that as GFDRR finds itself dealing with new kinds of data, new clients, and changes in mission scope or tactics, it can take advantage of ever-increasing capacity in the increasingly complex and interdependent project it helped create. GeoNode was begun by GFDRR, but it has now grown into a public good that will remain useful for many organizations for years to come.

