

Geodata in the Cloud

Reviewed paper

Peter Riisager (Peter.Riisager@atkinsglobal.com), Morten Bødtkjer (Morten.Boedtkjer@atkinsglobal.com), Jan Jørgensen (Jan.Joergensen@atkinsglobal.com), Thomas Bergstedt (Thomas.Bergstedt@atkinsglobal.com), Carsten Helsted (Carsten.Helsted@atkinsglobal.com), alle ATKINS, Danmark samt Peter Herzberg (phol@vd.dk), Vejdirektoratet

Abstract

The cloud offers great advantages for internet hosting of applications and data. In this paper, we review and discuss some of these advantages in relation to a new cloud-hosted Web-GIS application for the Danish Road Directorate, serving live traffic geodata to Danish road users. The new traffic map represents an example of geodata served from a cloud platform offering a fast, scalable, and stable Web-GIS application, in a solution that is cost-efficient compared to what would otherwise be possible based on dedicated servers. The paper shortly discuss some of the considerations that went into the design of the solution, touching also on important issues related to cloud hosting.

Keywords: Geodata, GIS, Web, Cloud Computing, Live Traffic Data.

1. Introduction

Geodata is defined as data with a spatial aspect or component, i.e. data tied to a geographic location. Geodata are crucial for modern societies, especially following the near explosive growth in the use and importance of location-aware devices. The United Nations (2011) has recently underlined the significance of geodata in modern societies, stating that: "Building infrastructure for the gathering, validation, compilation and dissemination of geospatial information has become as important to countries as the building of roads and telecommunications networks and the provision of other basic services". Geodata are estimated to represent as much as 80 percent of all data collected, stored, and maintained by local governments (Ohio Department of Administrative Services, 2011).

The obvious and most straightforward way to convey geodata to, and between humans is by the use of maps. Although maps are abstract generalizations of reality, they are easily understood and appreciated by most people. The deep-seated human capacity to understand and read maps is demonstrated by the fact, that the oldest known land-map, the 6200 BC Catal Huyuk map, appear to predate the oldest known writing system by ca. 2000 years (Dodge et al., 2011). For online web communication, the written word came before maps, with the first ever website published in 1991 serving text that explains the World Wide Web project (World Wide Web Consortium, 2015). However, on-demand web maps began to appear shortly thereafter in the mid-1990s, following the introduction of the Mosaic web browser in 1993. Given that web browsers -and online mapping technologies continually improve, it is no surprise that online mapping applications, hereafter named Web-GIS applications, are becoming more and more prevalent on the internet (Peterson, 2014). For private corporations as well as government agencies, Web-GIS represent an excellent platform to disseminate geodata in an interactive manner that engages and empowers the public. This paper discusses the new Danish traffic map, which represent a typical example of a Web-GIS application. The aim of the traffic map is to communicate geodata to the public through state-of-the-art online technologies, including cloud-computing technologies.

Cloud computing has recently emerged as a compelling new paradigm for managing and delivering data and services over the internet. Cloud computing offers powerful technologies to host fast, scalable and stable internet applications, without the need to maintain expensive hardware, dedicated space, and software (Chou et al., 2015). The cloud offers virtually unlimited resources in terms of storage and computing power, and based on the chosen cloud solution service, it scales transparently and in a semi-automated manner, while offering up-to-date underlying technologies. Moreover, its pay-per-use delivery model has the potential to significantly reduce traditional IT infrastructure costs. Geodata -and Web-GIS applications are naturally taking advantage of the emerging and ever-expanding cloud computing technologies. A long list of projects have successfully implemented geodata services in the cloud (Pross and Caumont, 2014) as well as Web-GIS online mapping applications (MapCentia, 2015; CartoDB, 2015). In this paper, we review "Geodata in the Cloud", and discuss some of the advantages as well as potential caveats related to the new cloud-hosted Web-GIS solution developed for the Danish Road Directorate. To the best of our knowledge, the map represents the first major Danish state-funded geodata project that utilizes public cloud technologies.

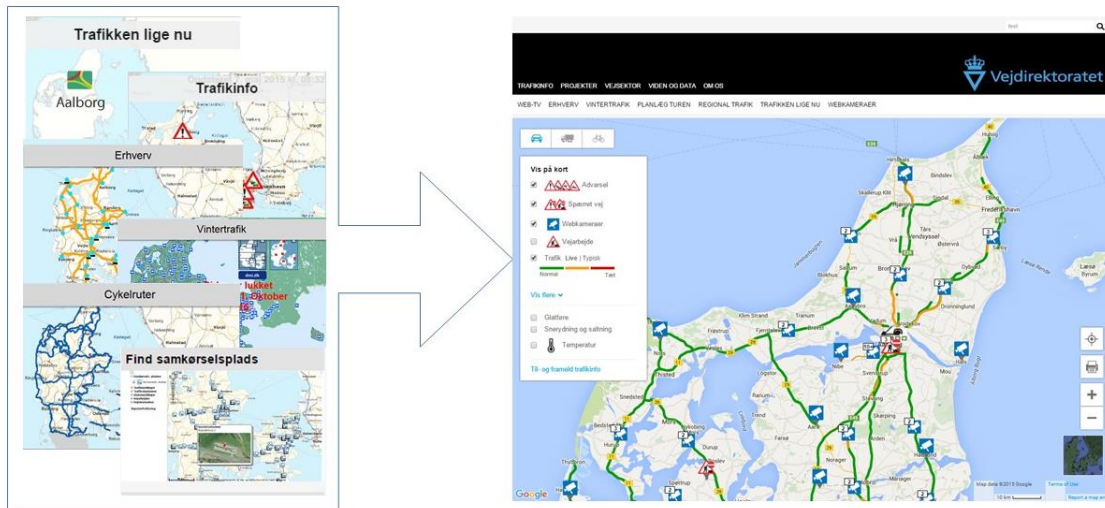
2. The Case: Live and Interactive Online Map for the Danish Road Directorate

The Danish Road Directorate (DRD) is responsible for the state-owned roads in Denmark. Besides the physical maintenance of road and road structures, the DRD is also responsible for the dissemination of traffic information to the public, in order to ensure a smooth and safe handling of the traffic. Disseminating traffic-information to the public through internet technologies is becoming increasingly important as traffic congestion continues to grow, and the road users are demanding up-to-date information on their location-aware mobile devices.

Through the last few years, the main DRD webpage (<http://vejdirektoratet.dk/>) has had about eight million visitors yearly. Roughly, two thirds of these visitors were requesting live traffic data on the pages 'trafikinfo' (traffic information) and 'vintertrafik' (winter traffic). The web traffic on the main DRD webpage has been unevenly spaced in time with 8-10.000 visitors on normal days and up to 250.000 visitors on busy days, which were invariably related to special weather or traffic conditions.

Previously, DRD conveyed online traffic information to road users via several Web-GIS applications, including applications dedicated to car drivers, lorry drivers, and cyclists. Furthermore, DRD displayed winter information on another separate Web-GIS application. The new traffic map consolidates all the existing applications and their geodata into one single Web-GIS application, which is reliable, highly scalable, and with a modernized interface, that caters PC, Tablet and Mobile users (Figure 1).

Figure 1. The Danish Road Directorate currently informs the public via several Web-GIS application (left side of figure). The new map (right side of figure) consolidates the existing applications and geodata into one single map-view, which can display all the information in a new highly reliable system with a modernized interface that caters PC, Tablet and Mobile users.



The new Web-GIS solution (right side of figure 1) is a thick client, using an array of front-end technologies, most notably Google Maps JavaScript API, AngularJS, and Bootstrap. The map functionality is based mainly on Google Maps JavaScript API, which gives the users a well-known, fast and responsive user interface. The users are served live traffic data that are continuously updated, but at various intervals, ranging from once every 5 seconds and up to daily. In the following, we present the geodata with a focus on how the cloud is used to handle and serve the data.

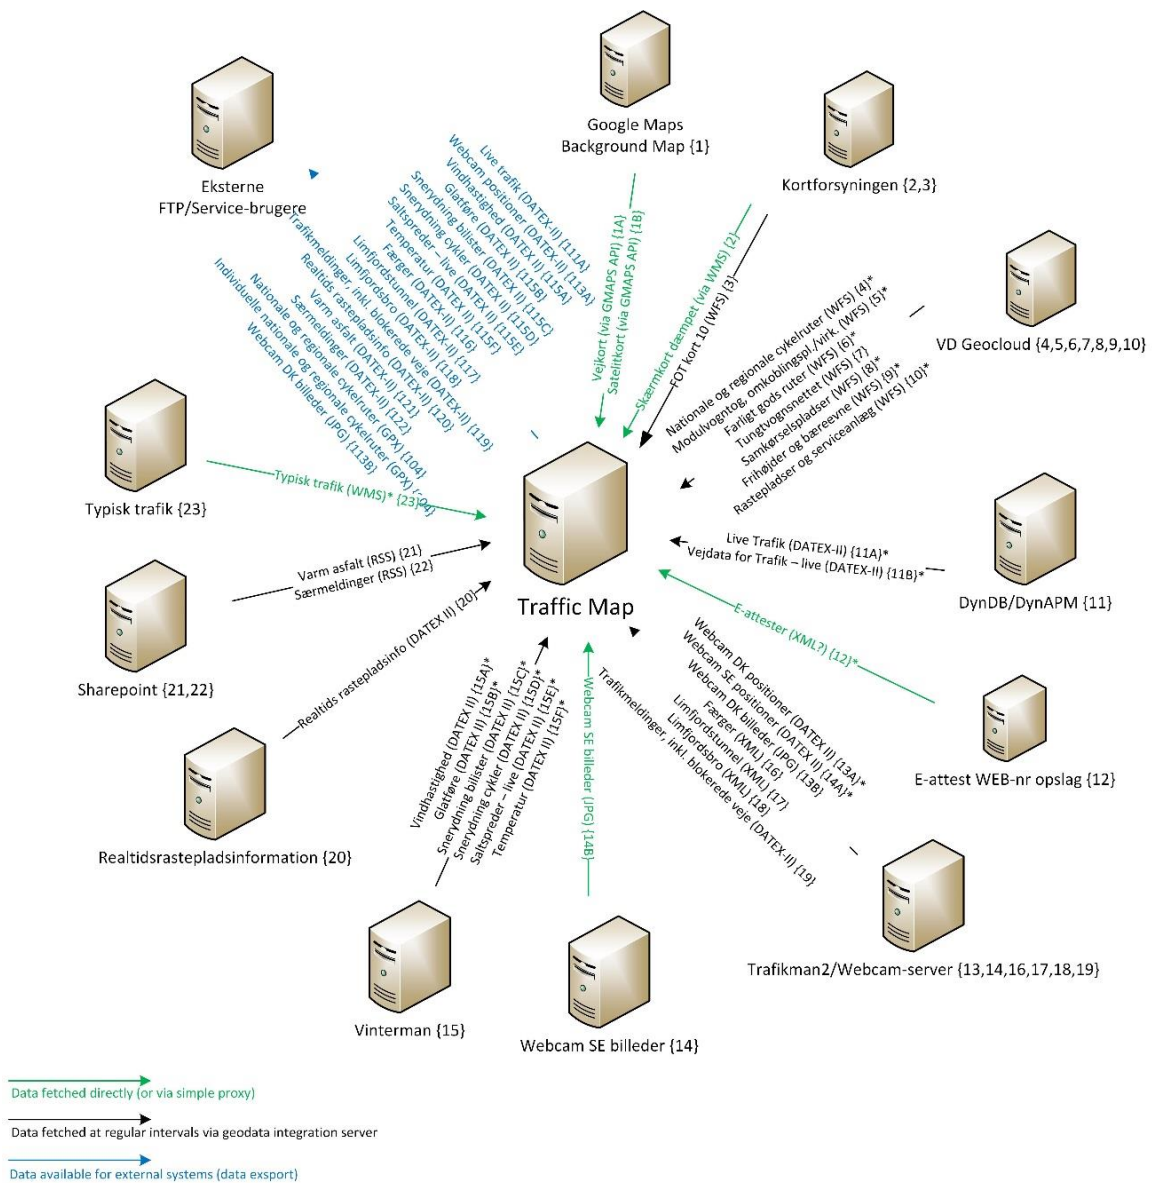
2.1 The Geodata

The new online traffic map is based on geodata retrieved from more than 20 web services hosted at different DRD web-servers (see Figure 2). The “raw” DRD data are served in various formats, most importantly DATEX II, which is a XML standard developed for information exchange between traffic management centres (DATEX II, 2015). In addition, some DRD data are served as Web Feature Service (WFS), as well as proprietary and nonstandard XML-formats. Some of the geodata are more or less static, being updated only once or twice a day, while other data are more dynamic and updated more frequently. Examples of large dynamic datasets are live traffic data for the entire Danish major road

network, which are updated once every minute, and webcam images that are updated every 5 seconds.

A central requirement for the new online traffic map is that the DRD geodata are cached by the solution, so that the DRD web-services and web-servers (Figure 2, next page) will not experience the expected high demand/load from the potentially up to several hundred-thousand users a day. The data are therefore retrieved from DRD web-services at predefined intervals and pushed into the cloud, wherefrom the end-users interacts with the geodata. On a yearly basis, hundreds of terabytes of geodata will be fetched from the DRD servers and pushed into the cloud.

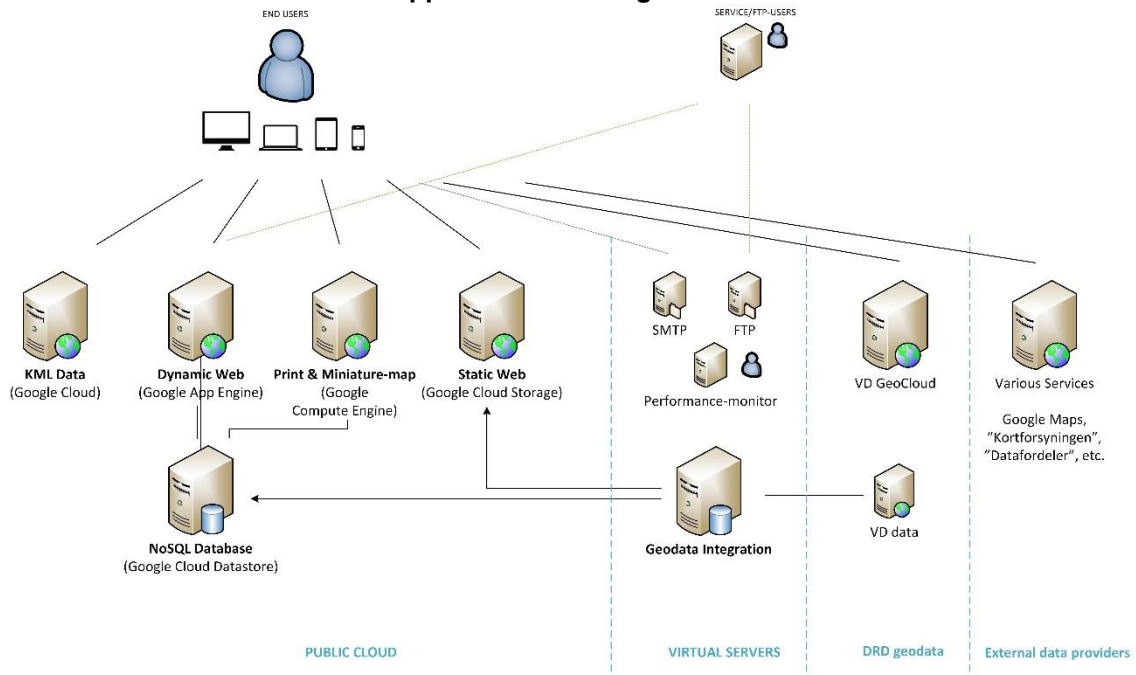
Figure 2. The new traffic map system in the centre of the figure is fetching geodata from a long list of more than 20 web services hosted at different Danish Road Directorate (DRD) web-servers. The original geodata are served in various formats, and fetched at intervals ranging from once every 5 seconds to once a day. Green arrows indicate geodata fetched directly (or via a proxy). Black arrows are geodata that are fetched and handled by the geodata integration server (see Figure 3). Blue arrows represent data made available to external systems.



2.2 Technical architecture

The overall technical architecture of the new Web-GIS solution is illustrated in figure 3. We will here shortly discuss the main architectural components, ignoring front-end technologies and backend details outside scope of present paper. The right part of figure 3, with the blue subtitles “DRD geodata” and “External data providers” represents the original “raw” geodata that the system builds upon (See also figure 2).

Figure 3. The technical architecture for the new online traffic map. The vertical hatched lines separates the architecture into four main parts: (i) External data providers; (ii) Danish Road Directorate geodata (see also Figure 1); (iii) Virtual Servers in Denmark, including the Geodata Integration application that represents the central back-end data engine; And (iv) the Public Cloud wherefrom the end-users interacts with the web application and its geodata.



The mid-part of figure 3 marked “Virtual Servers” are virtual servers hosted in Denmark. The most important element being the Geodata Integration application, which fetches geodata from DRD servers with the various formats and time-intervals discussed above (see paragraph 2.1). The principal tasks for the Geodata Integration application is to:

- Perform calculations on the “raw” DRD geodata, including for example the complex calculations carried out to define live traffic flow for the entire Danish major road network.
- Split and fuse the “raw” DRD geodata into data types that are logic for the traffic map end-user, for example road works, blocked roads, traffic events, etc. Moreover, the attribute values are translated into human-readable text strings.
- Convert the original geodata formats into Google Maps formats, which means GeoJSON for geodata that are points and KML for lines, i.e. the road network.

The Geodata Integration application is placed outside the cloud because its workload, although complex and heavy, is relatively constant, and so does not need the scalability and elasticity that is crucial in order to serve the geodata to the extremely varying amount of end-users. The most important reason that the Geodata Integration application is placed outside the cloud is to avoid the cloud vendor lock-in discussed further below (see paragraph 3.2).

The left hand side of figure 3 represents the cloud component of the system, consisting of several Google Cloud entities (see paragraph 2.3).

2.3 Cloud Architecture

Cloud computing is a general concept that comes in many shapes and formats, the main three models being (Rackspace, 2013):

- Infrastructure as a Service (IaaS) where the user has control and responsibility for the operating system and the application platform stack.
- Platform as a Service (PaaS) where the user gets a “fixed” framework she can build upon, for example to develop or customize an application.
- Software as a Service (SaaS) are software applications hosted by a vendor or service provider and made available to customers over the Internet.

Functionality is added as you go through the list from IaaS to PaaS to SaaS, but at the cost of freedom to design and customize the service. Within each of the three main cloud-computing models, a multitude of variations exists, as well as a long list of platform providers supplying both private and public cloud solutions. The online traffic map is placed in the Public Cloud of Google, utilizing different services of both PaaS and IaaS types. The goal of the chosen cloud architecture is to ensure both performance, scalability and elasticity, but at the same time also ensuring the required functionality, for example allowing end-users to interact with the application and subscribe to email

updates on the various geodata. The solution includes the following five main cloud-entities (see also left-hand side of figure 3):

- *Static Web* is hosted on Google Cloud Storage, which is a RESTful online storage web service for storing and accessing files. The Static Web Cloud component is used to store static files for the Web-GIS application e.g. HTML, CSS, and JavaScript files. Static Web is also the location where geodata in the form of GeoJSON- and KML files is continually uploaded from the Geodata Integration application (see figure 2). Google Cloud Storage is a Cloud Service of type IaaS, and it is highly elastic ensuring automatic scaling to suit the web traffic load on the Web-GIS site at any given point in time.
- *Print and Miniature-map* is hosted on Google Compute Engine, which is another IaaS service. Google Compute Engine is similar to virtual machines offering great freedom but, on the other, hand less functionality (Google Compute Engine, 2016). Most importantly, it does not offer automatic scaling. Google Compute Engine is used in the project to run PhantomJS (Friesel, 2014) as a webserver, allowing a user to generate PDFs and print images of exactly what she sees on her screen. The reason that the Print and Miniature-map is not placed on Google App Engine is that the solution uses Java features not available on the Google App Engine white-list (Google Cloud Platform, 2016).
- *Dynamic Web* is hosted on Google App Engine (GAE), which is a service of the type PaaS. Dynamic Web is a Java-based web application that handles dynamic aspects related to subscriptions and print queue handling. Dynamic Web integrates with Static Web, handling all the backend functionality that cannot be handled by Static Web. For the present project, the most important difference between GAE and Google Compute Engine is the elasticity and automatic scaling offered by GAE.
- *NoSQL database* is hosted on Google Cloud Datastore, which is a PaaS service for storing non-relational data. The NoSQL database is used by Dynamic Web to store backend-information concerning things such as the print-queue, subscriptions and system settings.
- *KML Data* is a cloud entity (figure 3) that is granted the system via the Google Maps JavaScript API. Keyhole Markup Language (KML) is an Open Geospatial Consortium XML-based standard used for describing two- and three-dimensional space. The national road network geodata are stored in KML format in Static Web, and based on disclosed Google cloud-services, the rather large sets of linear geodata is generalised and served in raster-format, so that end-users do not need to download more data than necessary for a given area and/or zoom-level. The

cloud functionality related to KML data is fundamental for the performance and responsiveness of the Web-GIS map.

Altogether, the five cloud entities ensure a highly scalable solution, with an excellent performance. The combination of the above listed cloud services collectively secures both the necessary functionality while still having the cloud computing advantage of almost infinite scalability. At the time of writing Google App Engine, serves hundred billion requests per day (Beyer et al., 2016), allowing a scalability that far exceeds the most optimistic web traffic expectations for the Web-GIS solution. Also considered in the cloud architecture is the overall cost. In order to minimise costs, the Web-GIS application is hosted to the greatest possible extent on the Static Web (Google Cloud Storage), which is cheaper than hosting and serving geodata from Dynamic Web (Google App Engine).

3. Discussions

Advantages and critical issues related to cloud computing has been thoroughly discussed and reviewed in a long list of studies, e.g. Armbrust et al. (2010); Avram (2014); and references therein. The issues are not fundamentally different for geodata and Web-GIS applications as compared to any other type of data and application. Nevertheless, each project is different, and the choice of hosting DRD geodata and the Web-GIS application in the cloud was based on lengthy considerations and discussions, which may serve as an inspiration and/or caution for interested readers. Here we shortly discuss the cloud-related issues, which was considered during the design and implementation of the new online traffic map.

3.1 Legal and Security issues

Major issues in cloud computing are legal and trust aspects (Ellegaard, 2011). In the case of the new traffic map, these issues are less relevant, as all the included DRD geodata are freely available and without the legal-restrictions attached to sensitive personal information. The DRD geodata also does not contain infrastructure information with significance for national security. Hence, although legal and security issues are highly relevant for the general usage of cloud computing, they were not considered an issue in the current project.

3.2 Vendor lock-in

Including public cloud services in the system design, the project automatically buys into the specific protocols, standards and tools of the cloud vendor, making future migration costly and difficult. Vendor lock-in is therefore an important and relevant issue to consider, and it is relevant for example in relation to Dynamic Web hosted on the Google App Engine (GAE). Although, being Java-based, the Dynamic Web is tied to specific GAE

components including the Google Data Store, which will make it difficult to port this part of the solution to another public cloud vendor should the desire arise.

Vendor lock-in was a major deliberation for the present project. Consequently, it was decided that backend geodata computations should be placed outside the public cloud (see Figure 3) also because the complex and heavy backend functionality relates to a near constant input of DRD geodata, and therefore does not need to implement scalability. The looser coupling to the cloud vendor that results from the placing Geodata Integration on a dedicated server was considered worth the higher hosting expenses as well as losing some of the other cloud computing benefits discussed below in paragraphs 3.3-3.6.

3.3 Scalability and Elasticity

Scalability and elasticity is one of the central motivations for the public cloud. Historical observations of the current DRD website usage demonstrate a highly variable load ranging from a few thousand daily visitors and up to a several hundred thousand visitors on busy days. The cloud architecture was designed with a strong focus on automatic scalability (see paragraph 2.3). As an interesting side-note, it is worth pointing out the potential problem related to elasticity and automatic scalability. The near infinite scalability means that the site can sustain malicious and massive web traffic designed to overwhelm and break down the site, which could result in a large bill considering the pay-per-use delivery model. The system design therefore includes features to counter potential malicious web traffic.

3.4 Performance and reliability

Performance and reliability is a main advantage of cloud computing. All major cloud providers have a history of very high up-times that would be near impossible to match on dedicated servers. The responsibility of system monitoring and fail-over implementations are excellently placed in the hands of cloud providers as they have vast and pooled resources to monitor the systems.

3.5 Administration

Another great advantage of the cloud is the excellent management and maintenance capabilities offered by professional cloud providers. All the cloud services in the present project include a simple web-console allowing the administrator to log on and easily carry out typical administration tasks, including monitoring and backing up.

3.6 Environmental issues

Cloud computing intelligently shares resources thereby lowering not only the cost, but also the energy efficiency and ultimately the Carbon footprint. Globally, cloud computing is considered an important factor to significantly cut CO₂ emissions from the IT industries (Carbon Disclosure Project Study, 2011). Although, it is difficult to

estimate the CO₂ reduction resulting from cloud hosting the new DRD Web-GIS solution compared to having it hosted on dedicated servers, it is nevertheless an important issue. Especially considering that, the Danish state has a strong focus on energy- and climate related goals and obligations, including the reduction of CO₂ emissions.

4. Conclusions

The first version of the new DRD Web-GIS solution was released on December 8, 2015 at <http://trafikkort.vejdirektoratet.dk/>. During its first five month in production, the new solution has received overwhelmingly positive feedback, and served live traffic information to more than 1.5 million online visitors (e.g. a threefold increase in online traffic information dissemination, as compared to the previously existing DRD online solutions). Altogether 46 percent of the users are mobile and tablets users, and the solution has hitherto had no downtime during its existence. The new DRD Web-GIS solution was awarded the 2016 national geodata prize (Brugstedet, 2016).

The new DRD Web-GIS solution consolidates several existing applications and combines them into one single map-view, which can display a long list of continuously updated geodata in a new highly reliable system with a modernized interface (Figure 1). Using public cloud services, the geodata are served with low response-time and through a Web-GIS interface that scales in an automated manner.

To the best of our knowledge, the map represents the first major Danish state-funded Web-GIS application hosted in the cloud. As such, the project may serve as an inspiration for other public and government agencies. This paper argues that based on a careful system design, the advantages of using public cloud services strongly outweigh the negatives. Note however, that a major part of the new DRD solution is placed on dedicated servers to avoid vendor lock-in, which is a typical and important negative of cloud computing. Combining the service oriented architecture of DRD web services, dedicated virtual servers, and altogether five different cloud services, the new DRD Web-GIS solution may represent the best of all possible worlds.

5. References

Armbrust, M., Fox, A., Griffith, R., Joseph, A.D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A. Stoica, I. and M. Zaharia (2010). A View of Cloud Computing, Communications of the ACM, Volume 53(4), P. 50-58.

- Avram, M.G. (2014). Advantages and Challenges of Adopting Cloud Computing from an Enterprise Perspective. *Procedia Technology*, Volume 12, 2014, Pages 529–534.
- Baurens, B. (2014). *InGeoCloudS - Inspired GEOdata Cloud Services*. EuroGeoSurveys Newsletter January 2014. P. 14-15.
- Beyer, B., Jones, C., Petoff, J. and Murphy, N.R. (2016). *Site Reliability Engineering, How Google runs Production systems*. O'Reilly
- Brugstedet (2016). Geodataprisen 2016, at <http://brugstedet.dk/> [accessed 6 June 2016].
- CartoDB (2015).
- Chou, D. (2015). Cloud computing: A value creation model, *Computer Standards & Interfaces* Volume 38, P. 72–77
- DATEX II (2015). DATEX II Data definitions, at <http://www.datex2.eu/> [accessed 25 October 2015].
- Dodge, M., Kitchin, R. and C. Perkins (2011). *The Map Reader: Theories of Mapping Practice and Cartographic Representation*. Wiley.
- Ellegaard, N.C. (2011). *The Legal Challenges of Cloud Computing*, Plesner International In-house Counsel Journal, Nordic Edition 2001, Volume 19, P. 19-28.
- Friesel, R. (2014). *PhantomJS Cookbook*, Packtpub.
- Google Compute Engine (2016). <https://cloud.google.com/compute/> [accessed 6 June 2016].
- Google Cloud Platform (2016). <https://cloud.google.com/appengine/docs/java/jrewhitelist> [accessed 6 June 2016].
- MapCentia (2015). Geospatial cloud solutions and services, at <http://www.mapcentia.com/> [accessed 25 October 2015].
- Moura, J. and Hutchinton, D. (2016). Review and analysis of networking challenges in cloud computing, *Journal of Network and Computer Applications*, Volume 60, P. 113–129
- Ohio Department of Administrative Services (2011). *Ohio's Location Based Response System*. White Paper September 15.
- United Nations (2011). Economic and environmental questions: cartography Global geospatial information management Report of the Secretary-General, E/2011/L.53
- Peterson, M.P. (2014). *Mapping in the Cloud*. Guilford Press.

- Pross, K.B. and H. Caumont (2014). *14-028r1 Testbed 10 Performance of OGC® Services in the Cloud: The WMS, WMTS, and WPS cases*, Open Geospatial Consortium, at <http://www.opengis.net/doc/ER/testbed10/cloud-performance>
- Rackspace Support (2013). *Understanding the Cloud Computing Stack: SaaS, PaaS, IaaS*, Rackspace Support Network, at http://www.rackspace.com/knowledge_center/whitepaper/understanding-the-cloud-computing-stack-saas-paas-iaas [accessed 10 September 2015].
- World Wide Web Consortium (2015). *First Web Page*, at <http://www.w3.org/History/19921103-hypertext/hypertext/WWW/TheProject.html> [accessed 25 October 2015]