SEED CRM and Data Visualization Needs Assessment

May 2018

Prepared by Neil Thomson, Executive Director, Open Green Building Society (OGBS) and Ryan McCuaig, Technical Lead, OGBS.

This work has been commissioned by the **City of Richmond**

Richmond

Executive Summary

The objective of this paper is to recommend the most appropriate Customer Relationship Management (CRM) and data visualization tools for the City of Richmond to use in administering its energy benchmarking program.

Although program reporting is done on an annual basis, management of a benchmarking program requires ongoing attention throughout the compliance year. Activities include extensive work to communicate important dates, collect accurate energy use data, maintain up-to-date information for building contacts, report on program performance, and enforce penalties for non-compliant participants. OGBS is developing CRM and data visualization tools for Richmond that are intended to make these workflows as efficient as possible.

Both the CRM and the data visualization tools must be built on the right database. A flexible and appropriately-structured database will enable administrators to scale their programs and achieve their goals with minimal cost and time investment. The Standardized Energy Efficiency Database Platform (SEED) has been designed specifically to support energy benchmarking programs, and is recognized as the best in class tool to support Richmond's end-to-end technology solution.

Based on the key CRM requirements captured during our analysis, we are confident that we can develop those features within a custom development through working with software firm, Regenerative Applications. We can also meet the data visualization requirements by preparing a map-based view for public disclosure of benchmarking data by adapting an existing tool used by several American cities, and providing an adapter between it and SEED.

For other municipalities considering adopting similar tools to support their energy benchmarking programs, important considerations in the selection process are an ability to meet the requirements listed in this document, an appropriate budget for the upfront development, ongoing budget for software licensing, maintenance and support, and the organization's internal IT capabilities to handle the work.

Introduction

Canadian municipalities are moving in tandem with the Federal government to enact energy benchmarking legislation for commercial buildings larger than 50,000 square feet. Most municipalities already report publicly on municipally owned buildings; however, they operate voluntary disclosure programs for private sector buildings. Both municipal and private sector building owners will be required to report on their energy use once the new legislation is in place.

This prospective legislation means that voluntary programs managing a small number of properties on Excel spreadsheets will become sophisticated programs requiring more comprehensive tools. Program managers will be required to communicate extensively with program participants, record interactions, and report on compliance status and program performance. These activities demand sophisticated technology solutions to ensure accuracy, and save the municipality time and money in operational overhead.

The SEED Platform has been identified as the best-in-class database for supporting an energy benchmarking program. SEED is an open-source database and web application that provides public agencies and private organizations with a standardized but flexible enterprise data platform. This platform helps users manage portfolio scale building performance data securely and cost-effectively, using data from the particular mix of sources that a municipality might have available to it. The original SEED Platform was developed by the US National Renewable Energy Laboratory (NREL) with funding from the US Department of Energy (DOE). Work began in 2014, and NREL still has a full-time team dedicated to updating and improving the tool.

On June 1st, 2017, OGBS started working to deliver a 'Canadianized' version of the SEED Platform. The main objective was to make the full scope of version 2.3 functionality available to Canadian users. This includes retrofitting the database so the system can store and process metric units, adding French localization of the user interface, and updating the property data configuration to account for the Canadian reality of postcode formats and tax lot to building number relationships. This work was delivered on December 14, 2017 and is referred to as Phase 1.

The next phase of the Canadian work, and the subject of this paper is referred to as Phase 2. Phase 2 involves identifying complementary technologies to SEED that will enable benchmarking program administrators to manage program administration in a low-touch and cost-effective way. Although SEED is effective as an energy benchmarking database, it is notably light on display capabilities and features to support specific benchmarking workflows. With this in mind, CRM and data visualization technologies have been identified as complementary tools to support the end-to-end technology solution. Phase 2 work began in October 2017 and is due for completion at the end of June 2018.

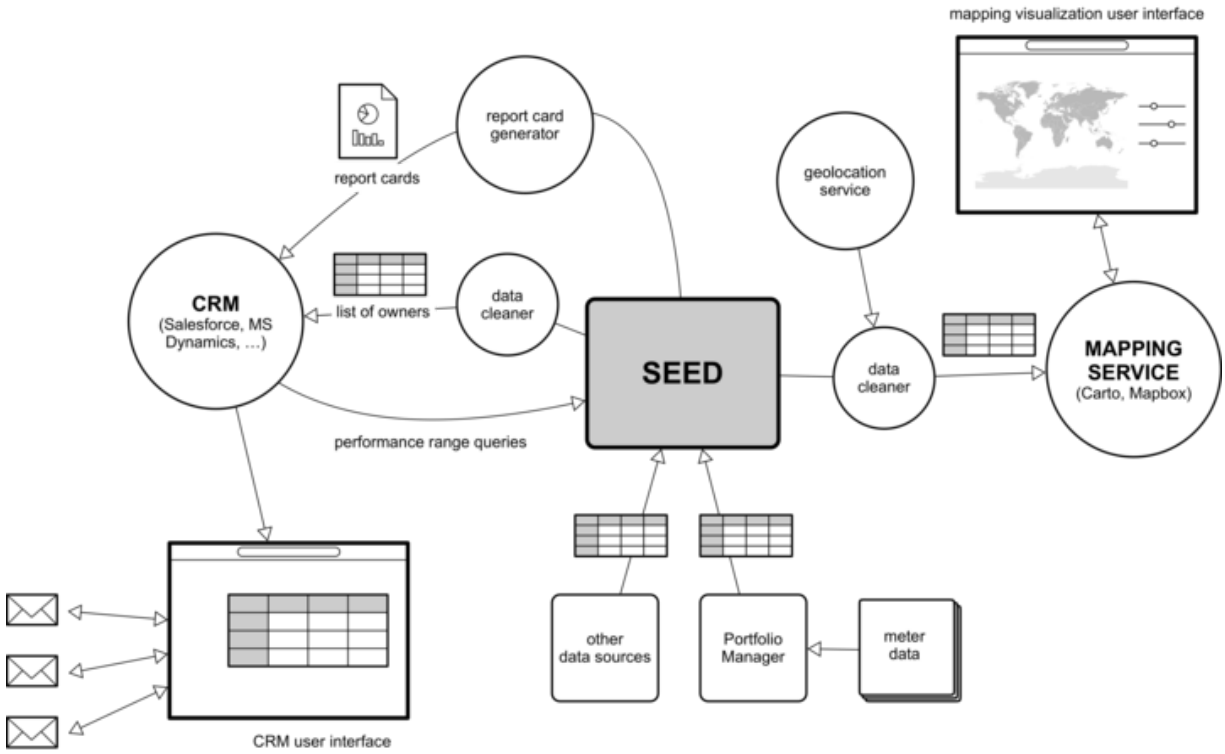

Phase 2 system architecture

Methodology

The first step in identifying the most appropriate CRM and data visualization tools is to understand in depth what it means to administer an energy benchmarking program. To achieve this, we conducted extensive primary and secondary research. The primary research included interviews with municipalities in both Canada and the US, and the secondary research was predominantly web-based.

The Canada interviews focussed on achieving a better understanding of the activities that municipalities will undertake to manage their mandatory benchmarking programs, and their level of awareness of CRM and data visualization tools. The people who kindly gave of their time for this work include:

- Lisa Dockman - City of Edmonton
- Brendan McEwan - City of Richmond
- Levi Higgs - City of Richmond
- Munim Ahmer - City of Toronto
- Micah Lang - City of Vancouver
- Monica Samuda - District of North Vancouver

Municipalities in the US are a few years ahead of their Canadian counterparts from a policy standpoint. Many have already started administering mandatory energy benchmarking programs. These interviews focussed more on understanding the value of the tools already in use to support program workflows, why they were chosen, and how they could be improved upon. The people who kindly gave of their time for this work include:

- Ben Silverman - City of Boston
- John Bolduc and Bronwyn Cooke - City of Cambridge
- Katrina Managan - City of Denver
- Luke Hollenkamp and Kelly Muellman - City of Minneapolis
- Rebecca Baker - City of Seattle
- Marshall Duer-Balkind and Katie Bergfeld - Washington D.C.

We structure this report as a presentation of the functional requirements identified during the research, and how they fit with the proposed CRM and data visualization technologies. Finally, we make a recommendation of the most appropriate tools for the City of Richmond to use for managing their program.

The energy benchmarking process

The central operational objectives of an energy benchmarking program are as follows:

1. Manage relationships with participants relative to compliance dates
2. Ensure accuracy of the data reported
3. Report on compliance and program performance
4. Take corrective action both internally and externally

Appendix A shows a typical energy benchmarking process and accompanying data flows in diagrammatic form.

Although reporting is done on an annual basis, management of a compliance program requires ongoing attention. The work involved in collecting accurate energy use data and maintaining up-to-date information for building contacts, reporting on compliance and program performance, in addition to enforcing penalties for non-compliant participants, is extensive.

CRM considerations

A CRM helps users track and manage their relationships with multiple ‘customers’. The goal for the energy benchmarking use case is to effectively communicate with all contacts responsible for compliance and, where necessary, assist with the activities required to meet their obligation.

A city like Toronto might have 5,000 buildings larger than 50,000 square feet reporting into their program. At a maximum, each building could have a new contact each year and potentially change ownership each year. The administrator will need a robust CRM to help manage all those relationships. Features like mass email and workflow automation will make the administrators life considerably easier.

Key requirements

Based on our research we've identified the following key functional requirements of the CRM. Note that the terms 'User', or 'Users' refer to an administrating organisation's collective user rights, without distinguishing between the different user profiles that an organization might set-up:

1.1	Users have the ability to generate individual building energy scorecards directly from the CRM
1.2	User can easily share the building energy scorecard with the building contact directly from the CRM
1.3	CRM interface is mobile friendly for benchmarking functions that lend themselves to on-the-go use
1.4	CRM interface can interpret data chronologically i.e. over compliance years
1.5	CRM must allow accurate recording of previous communications with the building contact for a particular property
1.6	CRM must allow for one to many relationships between building contacts and properties, and properties to building owners
1.7	CRM must display all the buildings contained in SEED for a specific municipal user account and recorded for each compliance year
1.8	CRM must clearly show which properties are in compliance and which ones aren't
1.9	SEED data quality errors are reported in the CRM
1.10	Building ID and parcel ID (tax lot ID) are connected with a specific building owner record for specific time periods to take into account ownership transfers
1.11	Buildings can be filtered in the CRM and easily identified by category i.e. public or private
1.12	Buildings can be filtered in the CRM and easily identified by type i.e. hospital, school, grocery store etc.
1.13	User should be able to send email notices to multiple building contacts in one action,

	much like Mail merge
1.14	CRM shows energy use intensity (EUI) and greenhouse gas emissions for each building
1.15	EUI and greenhouse gas emissions data can be viewed for different compliance years

Potential solutions

The visual display requirements listed above have geographic display dependencies, which isn't a feature typically associated with a CRM. Depending on the commercial mapping tile database chosen for integration with the CRM, this will likely have an additional cost for the organization in terms of development setup, licensing and ongoing maintenance.

Based on our research, the following CRMs should be top of mind for consideration by program administrators and their internal IT. The sections after provide information about each solution and a view on integration potential with SEED.

Computronix POSSE

Computronix was started in 1979 with a focus on providing "robust government workflow solutions" to state, provincial, and local governments. Their solution, known as POSSE, has been designed for processes related to regulatory compliance (permits, licensing, inspection, enforcement, etc.).[1]

POSSE is browser based and designed to automate a number of workflows and business rules. It is primarily built to be a land management software technology. A number of government agencies across North America have contracted Computronix to build custom solutions for their needs, including the Province of BC, City of Vancouver, City of Surrey and City of Edmonton in Canada, and Metro Denver in the USA.

Unlike Salesforce and Dynamics, POSSE has not been built with the intention of being highly customizable for integration with third party databases. POSSE is more a classic database than a CRM, to the extent that it isn't possible to run POSSE on top of SEED. POSSE would be more of a replacement for SEED than an integration option. The company has also informed us that for internal reasons they couldn't plug into an open-source solution.

Microsoft Dynamics

Dynamics offers a comprehensive and highly customizable CRM tool. The user interface is uncluttered and easy to use. Unlike Salesforce, which permits only one role per user, Dynamics

[1] Retrieved 15 March 2018 from https://www.computronix.com/

allows for multiple roles. Dynamics can be used as a convenient cloud based solution, and is also available for download if required. Users can even adopt a Hybrid Cloud Strategy, mixing and matching for optimal performance and security.

Dynamics is optimized for compatibility with Microsoft products. This means there's a risk that performance could be sub-optimal when used with other systems, such as SEED. Microsoft's intention is to encourage users who have a preference for Dynamics to use Access as the database for their CRM.

Dynamics is designed for larger organizations with complex customer relationship management requirements. This results in an often expensive and steep learning curve for users from small and medium sized organizations.[2] The upfront development work to customize Dynamics to meet the requirements listed in this document will be extensive. Ongoing operating costs of using Dynamics under the required Dynamics 365 Plan are $255.33 CAD per user per month, paid annually. A budget allocation for ongoing maintenance of the code should also be factored in.

Salesforce

Salesforce is a cloud computing SaaS provider based in San Francisco. The company was founded in March 1990 by a former Oracle executive and is considered by many to be the best CRM product on the market. Salesforce has 11 different cloud products, which are distinguishable by their focus on functionality for many variations that exist within the sales role.

Salesforce contains a highly customizable dashboard and application programming interface (API), and is accessible from anywhere through The Cloud. Consequently, there is no need to purchase software and hardware systems to get started. Salesforce provides continuous updates and maintenance for as long as an organization pays its monthly fee. The company also has a robust third party app marketplace and offers a 30 day free trial.

Salesforce has, however, been known to suffer from its own success. As the most widely used CRM in the world, there are reports that tech support can be slow to respond. Thus, many organizations prefer to work with third-party companies that provide support on their behalf. Also, when engineers are performing maintenance, some users have reported frustration due to times when their instance cannot be used.[3]

Upfront development work to customize Salesforce to meet the requirements listed in this document will be extensive. The ongoing operating costs of using Salesforce under the required Lightening Enterprise Plan are $150 USD per user per month, paid annually. It should be noted

---

[2] Rashid. F. 2015 August 13. Microsoft Dynamics CRM Online Professional. Retrieved from https://www.pcmag.com/article2/0,2817,2487416,00.asp

[3] Marvin. R and Krasnoff. B and McLaughlin. M. 2017 November 15. Salesforce Sales Cloud Lightning Professional. Retrieved from https://www.pcmag.com/article2/0,2817,2364726,00.asp

that Salesforce does have a government solution, specifically built with cities in mind. A budget allocation for ongoing maintenance of the code should also be factored in.

Custom development by Regenerative Applications (RA)

OGBS frequently works with Vancouver based software company, RA on development projects. The Phase 1 Canadianization of SEED completed in December 2017 was contracted to RA. Through this work RA is familiar with the technological elements of SEED and has established relationships with the SEED working group, both in Canada and the US. Phase 1 was successfully delivered on time and budget. This gives RA a proven track record of successful delivery.

In this scenario RA would develop a custom web based platform for municipalities to use in lieu of committing to full scale CRM customization. The RA platform would integrate directly with SEED and enable municipalities to easily manage compliance tracking, performance reporting and building contact information. The setup cost would be considerably lower than for a CRM customization, which may be appealing to smaller municipalities such as Richmond.

While the cost of setup will be lower than for a full CRM customization, some meaningful setup will still be required. RA will need to work to integrate the platform with SEED, setup hosting and the user accounts and settings for internal use. RA will also need to ensure that the platform remains updated as technology around it advances. It's expected that this cost will be covered with an ongoing user subscription payment.

Integration with SEED

To customize Salesforce to support an energy benchmarking program, OGBS will need to provide a data connection between the CRM and SEED, and develop an interface within the CRM (using the programming environments provided by the CRM vendor) to support the specific workflows Richmond wants to support. On a regular interval, the CRM will pull data from SEED using its API. Custom code within the CRM will respond to incoming data; for example, shortening the list of non-reporting owners with whom to follow up as reports come in, or alerting that a particular building is ±15% off year-over-year for follow-up action.

Note that SEED in all cases remains the 'source of truth' for building data, and the CRM is the 'source of truth' for the contact data. We expect most of the data flow to be one direction from SEED to the CRM. Our research suggests, however, that being able to make small adjustments and corrections to building attribute data while talking with owners is important. To that end, OGBS will provide a mechanism for building data corrections to flow back into SEED.

Data visualization considerations

Data visualization capabilities are a high-value extension for energy benchmarking program administrators. The core component is geographic information system (GIS) mapping. This enables the energy performance of individual buildings to be viewed on a map overlay and in accordance with data contained in SEED. The view can be filtered by criteria such as energy use intensity, GHG emissions and Energy Star Score.

Key requirements

Based on our research we've identified the following key functional requirements of data visualization for this Richmond work. The terms 'User', or 'Users' refer to an administrating organisation's collective user rights, without distinguishing between the different user profiles that an organization might set-up:

2.1	Organizations can show their data on a data visualization interface with GIS capabilities
2.2	Data visualization platform is mobile friendly
2.3	Users can easily identify the energy performance for a specific building address within the energy benchmarking program from the data visualization interface
2.4	Data visualization must display all the buildings contained in SEED that are subject to public disclosure
2.5	Data visualization allows for filtering by square footage
2.6	A municipality is able to decide whether the web-based data visualization display is public, or kept private to internal users only
2.7	Buildings can be filtered in the data visualization interface and easily identified by category i.e. public or private
2.8	Buildings can be filtered in the data visualization interface and easily identified by type i.e. hospital, school, grocery store etc.
2.9	Data visualization display can be filtered by energy use intensity and greenhouse gas emissions
2.10	Data visualization display can be filtered by Energy Star Score
2.11	Only data that doesn't flag any quality issues in SEED will pull through to the data visualization interface

Potential solutions

We've identified three data visualization tools that have the technical capabilities to satisfy all the functional requirements listed above:

CartoDB or Mapbox

In this solution SEED exports building energy data and building locations to a CartoDB or Mapbox geographical data source for accurate cartography. OGBS sets up a separate public website with a custom user interface.

We start from the public disclosure user interface code that is in use in several American municipalities, namely Seattle, Chicago, Washington DC, and others, which has been open-sourced by the City Energy Project in the United States. The custom front-end has sliders and filters that allow a viewer to zero in on a given year's data and interesting properties. These filters are translated into queries to the mapping backend. The mapping back-end is responsible for producing and delivering map tiles to exactly suit the current query, using the SEED data it hosts. One advantage of this server-side generation is that it keeps the visualization usable on lower-powered mobile devices.

Considerable upfront costs should be budgeted for to customize this Mapbox and CartoDB data visualization option. For publicly-accessible data sets, CartoDB currently waives all subscription and operating costs. This appears to be a major reason that many American disclosure sites have settled on it. CartoDB or Mapbox can also be used for privately-viewable geographical visualizations within the CRM. For this use the cost model shifts to a base subscription of several thousand dollars a year, with additional more minor fees based on usage. Licensing costs to Mapbox for using this option range from between $0 to $6,000 if an organization chooses public disclosure only, and $6,000 to $12,000 for both public disclosure and internal mapping for city staff.

Leaflet.js with OpenStreetMap or Richmond GIS tiles

Leaflet is a leading open-source JavaScript library for mobile-friendly interactive maps. It works efficiently across all major desktop and mobile platforms, and can be extended with a number of plugins.[4] Leaflet is another potential solution to build or retrofit the City Energy Project user interface to use Leaflet.js instead of CartoDB.

In the Leaflet.js scheme, OGBS generates the necessary map tiles (about half a million) annually from Richmond's existing GIS system. If that is infeasible for technical or copyright reasons, then the alternative is to use OpenStreetMap. The major differences between this and the CartoDB or Mapbox solution are higher development cost, higher cost of hosting tiles, and

[4] Retrieved 15 March 2018 from http://leafletjs.com/

that the tiles complete with map pins are not generated on the spot. Any filtered building data and map pins would instead be layered on in the user's browser. This is more taxing than the CartoDB approach of building the tiles on the server ahead of their use.

We see no reason to consider this option further for the public disclosure visualization. Because it is relatively technically similar, we bring it up as a backup plan just in case CartoDB or Mapbox's free cost tier is discontinued. This solution would also be suitable for privately-viewable geographic visualization if Richmond considers the ongoing cost of CartoDB or Mapbox's private hosting too high.

Tableau

Tableau is widely regarded as the number one off-the-shelf tool for creating data visualizations. The software's purpose is to make it cost-effective to engage in open-ended play to discover relationships in data through visualization with relatively little help from IT.[5] It makes it a simple drag-and-drop to, for example, change the value on a y-axis, plug in a new data source, or tweak colours to emphasize or de-emphasize interesting relationships and show them to small groups. By contrast, the other data visualization tools mentioned here tend to have a high-enough cost to implement a visualization that they are better suited to documenting relationships already known to be interesting, and to publish them widely.

Tableau could be suited for public disclosure of benchmarking data. They trade a low up-front development cost for an ongoing subscription cost (roughly $500 USD per user per year) for visualizations meant to be shared privately with small groups. Tableau, however, exempts visualizations that are meant for wide public consumption and provides free hosting for these. One remaining issue with it is that while there is a variety of canned visualizations available, the flexibility of each visualization will always be more limited relative to what can be achieved with skillfully-built custom code.

For energy benchmarking purposes, the task is not as much aimed at finding relationships as publishing a known set of data relationships; that is, the required visualization is known ahead of time and doesn't depend on the data set. Were the existing open-source City Energy Project code not available for use, we would strongly recommend using Tableau for the proof of concept. It may still have a place with any privately-viewable visualizations that the CRM can't provide.

Integration with SEED

The integration between SEED and the public disclosure visualization will be relatively similar to the integration with the CRM: OGBS will set up a data pipeline where, on an interval, SEED pushes publicly-disclosable geotagged data to the mapping backend. The front-end user

[5] Baker. P. 2018 January 30. Tableau Desktop. Retrieved from https://www.pcmag.com/article2/0,2817,2491943,00.asp

interface will be set up once and simply respond to new data entering the mapping backend. We recommend that SEED, as the 'source of truth' about building data, be where the necessary latitude and longitude data be stored. (It can be added as a straightforward import of building id + latitude + longitude rows as a CSV). The fact of a particular building being 'disclosable' will be implemented using SEED's labels: once a label such as 'Disclose-2018' is applied to a building within SEED, it will be pushed to the mapping backend on the next interval.

Other considerations

While CRM and data visualization tools mentioned here will do a great job of supporting most of the functional requirements of an energy benchmarking program, there are some gaps. The most notable gap is providing help centre support to program participants. Building contacts will have questions that relate to their compliance status, reporting deadlines, entering and submitting data using Energy Star Portfolio Manager, penalties for non-compliance etc. These calls and emails can feasibly be received internally by the municipality, however, there may be an appetite to outsource these activities to a third party service provider.

Recommendations

The City of Richmond already uses Salesforce in other departments of the city. We initially thought that this might allow us to fast track the Salesforce customization option; however, this is not the case. We discovered that approval for one department does not accelerate software approval for another department. A full Salesforce customization is also likely too much functionality for the 600 or so building that Richmond has to manage. Thus, the custom RA solution will be the best CRM option for Richmond.

The City Energy Project data visualization tool supported by CartoDB has already been shown to more than capably meet the visualization requirements listed in this paper. The fact that the technology is already being used by other large municipalities in the US and is performing admirably, makes us confident that it will also be a good choice for Richmond.

Selecting the most appropriate CRM and data visualization tools is the focus of this paper. Outstanding questions for Richmond include, How will they host their instance of SEED? Who will provide ongoing maintenance and support for SEED and their CRM and the data visualization solutions? Is support required in providing a help desk to program participants, or will this be handled by the city? We look forward to having these discussions with Richmond and helping them successfully implement their energy benchmarking program.

Conclusion

For other municipalities thinking about the implementation of their energy benchmarking programs, we hope this document has provided some clarity on the important elements to consider. Budget is obviously central to the decision making process, as is the number of buildings participating in the program. A smaller municipality with 150 buildings to manage might decided to save the cost of adopting complex technologies and use Excel instead. Other may decide that they only require the CRM functionality and not the data visualization functionality, or visa versa. For those municipalities with a larger number of buildings to manage and starting to get to grips with the operational complexity of the work at hand, this is a larger conversation. We invite all interested municipalities, regardless of size, to contact us about how OGBS might be able to help define the most appropriate solutions for your needs.

Appendix

Appendix A - The energy benchmarking workflow