

AUTOMATING MAPPING PRODUCTION FOR THE ENTERPRISE: FROM CONTRACT TO DELIVERY

R. Uebbing, C. Xie, B. Beshah, J. Welter

North West Geomatics Ltd., Suite 212, 5438 - 11th Street NE, Calgary, Alberta, T2E 7E9, Canada –
{robert.uebbing | chenglin.xie | belai.beshah | john.welter}@nwgeo.com

Commission IV, WG IV/5

KEY WORDS: Processing, Monitoring, Automation, Decision Support, Metadata, Spatial Infrastructures

ABSTRACT:

The ever increasing volume and quality of geospatial data has created new challenges for mapping companies. Due to increased image resolution, fusion of different data sources and more frequent data update requirements, mapping production is forced to streamline the work flow to meet client deadlines. But the data volume alone is not the only barrier for an efficient production work flow. Processing geospatial information traditionally uses domain and vendor specific applications that do not interface with each other, often leading to data duplication and therefore creating sources for error. Also, it creates isolation between different departments within a mapping company resulting in additional communication barriers.

North West Geomatics has designed and implemented a data centric enterprise solution for the flight acquisition and production work flow to combat the above challenges. A central data repository containing not only geospatial data in the strictest sense such as images, vector layers and 3D point clouds, but also other information such as product specifications, client requirements, flight acquisition data, production resource usage and much more has been deployed at the company. As there is only one instance of the database shared throughout the whole organization it allows all employees, given they have been granted the appropriate permission, to view the current status of any project with a graphical and table based interface through its life cycle from sales, through flight acquisition, production and product delivery. Not only can users track progress and status of various work flow steps, but the system also allows users and applications to actively schedule or start specific production steps such as data ingestion and triangulation with many other steps (ortho-rectification, mosaicing, accounting, etc.) in the planning stages. While the complete system is exposed to the users through a web interface and therefore allowing outside customers to also view their data, much of the design and development was focused on work flow automation, scalability and security. Ideally, users will interact with the system to retrieve a specific project status and summaries while the work flow processes are triggered automatically by modeling their dependencies. The enterprise system is built using open source technologies (PostGIS, Hibernate, OpenLayers, GWT and others) and adheres to OGC web services for data delivery (WMS/WFS/WCS) to third party applications.

INTRODUCTION

In recent years the management of geospatial data, imagery and point clouds, has become increasingly challenging for mapping companies due to higher resolution sensor technologies, new methods in data processing (GEHRKE et al, 2011), higher demand for faster delivery, but also the use of different software packages by various vendors. The inhomogeneous environment requires frequent re-formatting and import/export of data along the processing chain during the life cycle of the information from acquisition through to deliverable products.

In addition to the complexity of the data processing of the actual sensor data, the role of communication between departments within a mapping company plays a crucial role, more so due to the risen expectations of clients for faster turnaround times.

North West Geomatics (NWG), a mapping company located in Calgary, Canada, owns multiple ADS pushbroom sensors and one ALS LiDAR sensor. In addition to being a traditional mapping company with orthophoto and digital elevation model (DEM)

production NWG also provides an online image distribution service through VALTUS Imagery Services. Until recently, these two parts of the company have been treated as separate entities. Given the nature of the content (spatially referenced imagery) a central repository containing imagery and metadata provides an efficient and seamless integration between production and delivery.

But an enterprise catalogue goes beyond solving the technological gaps between the two parts of the company. By assigning resource load, budgeting information, quality control results to projects that act as containers for image data, it also make the catalogue an invaluable tool for planning, accounting and decision making.

By adding different completion states for images and projects during the production phase, stakeholders can monitor production progress. This information can also be utilized to automate work flow steps and further increase overall efficiency in production by reducing the amount of manual and interactive processing steps such as starting jobs for ortho production. In turn, this frees up

human resources to perform work flow steps that cannot be easily automated by machines such as visual quality control.

In the following the paper we will outline the stakeholders in a mapping production environment, describe various work flow steps in production, discuss the design and implementation of an enterprise catalogue at NWG and finally list the experiences the stakeholders had over the last year. While the enterprise catalogue discussed in this paper is a unique development and tailored specifically for NWG's needs and requirements, many of the challenges that were addressed can be found at other mapping companies.

STAKEHOLDERS

Throughout the life cycle of a mapping project there are various groups, or stakeholders, that have a vested interest the acquired and processed data. The following stakeholders are identified as:

Planning

Planning is responsible for the generation of a flight plan including spatial extent, flight line layout, distribution of ground control points and resource allocation.

Acquisition

Flight crews consisting of pilots and sensor operators execute the flight plans and are responsible for sending the acquired image data to the main office for quality control and processing.

Quality Control

Quality control (QC) makes assessments on the quality of imagery and metadata throughout the production work flow. While some of the QC is automated it is still in large parts a process requiring human interaction. The very first QC on incoming flight data is time crucial as it may call costly re-flights due to environmental impact such as clouds or strong turbulence and due to technical failures while the flight crew is still on site. Various software packages, especially for viewing and measuring, are utilized.

Production

Production generates deliverable image products such as ortho images, stereo viewable images, 3D point clouds and other derived products. The production group is also responsible for all intermediate processing steps such as geo-referencing and triangulation. The main software package used for creating image deliverables is the XPro software suite, although also other packages are used for auxiliary tasks.

Online Delivery

Provides access to clients and customers to view and downloading image products through a web browser and/or ArcGIS. Generic WMS/WCS delivery is also possible.

Accounting

Accounting is responsible for customer invoicing and tracking resource usage.

Management

Develops new business opportunities and maintains existing ones. Also, management is responsible for project budgeting, project life cycle monitoring and resource management. The challenge for any large mapping company like ours is how to facilitate better communication between these different stakeholders and reduce the inefficiency that results from working with multiple projects at the same time (status tracking, context switching, data handling etc.).

WORK FLOW

After having committed to a new contract, management and accounting will have done the resource allocation and the project goes into the technical planning stage. Here the flight plan is created taking all the contract specifications into account. This plan is very detailed with all the flight lines laid out for the flight crew to follow. At location the flight crew will execute the flight plan and send the data, even if only partially flown, back to the main office. QC assesses the imagery immediately at reception and checks for any type of issue such as clouds, smoke, turbulence, sensor failure, etc. Re-flights may then be called based on severity of the found issues. Given the cost of maintaining a flight crew at the acquisition site the timely detection of re-flights is crucial.

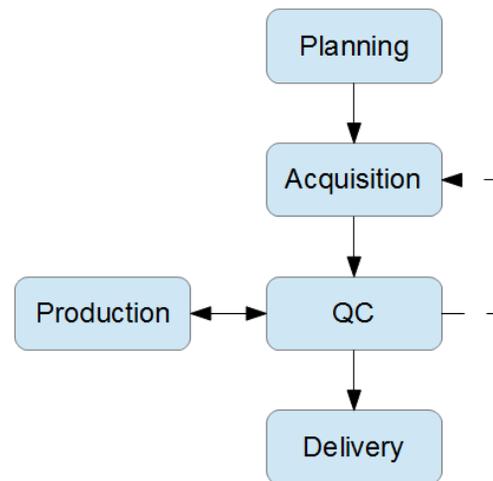


Figure 1: High level work flow diagram

Once the acquired imagery passes initial QC production starts on the actual processing of the data which involves geo-referencing, triangulation, ortho-rectification, mosaicing and generation of meta data that accompanies the deliverable imagery. Optional products may include stereo viewable image pairs and digital surface models. Most of the above processing steps are preceded by additional QC steps to ensure geometric and radiometric consistency of the final products. Finally, once the final deliverable products have met the original specifications the imagery is either sent by traditional means, i.e. stored on physical mass media storage, or made accessible through the VALTUS web store where clients can download the data or stream the imagery online into their software package of choice.

DESIGN OF AN ENTERPRISE SYSTEM

The enterprise mapping system developed at NWG was designed in the context of the preceding paragraphs on stakeholders and production work flow. But there were also other requirements that impacted the design of the system such as:

- Scalable to many simultaneous users
- Granular access rights for various roles
- Extensible for custom data
- Data management (copy, edit and move datasets)
- Task and process management
- Provide lineage across projects
- Support metadata standards (ISO and FDGC)
- Provide easy to use APIs
- Stable platform to build other products
- Server and storage scalability
- Support multiple OS platforms (Windows and Linux)

The above considerations led the design of an enterprise system based on a client/server architecture. At the core of the system is an application server that consists of multiple engines and services that manage the image data stored on a file server and metadata persisted in a spatial database. A web server layer sitting on top of the application server allows users to access functionality and data through a standard web browser. The system is also designed to allow other clients such as XPro or other applications to communicate directly with the application server using a remote messaging interface.

DATA MODEL

By design the data model for the enterprise mapping system is project centric, i.e. that nearly each entity modeled in the system resides in the context of a project. In addition to being project centric each entity is also derived from an abstract base object that contains primitives such as timestamps and geospatial extents, among many others, as its attributes. In practice, data such as images, flight lines, trajectories, acquisition blocks (and many more) have knowledge of their geometry and have temporal information in addition to the knowledge of their parent project.

Not only are the standard objects that are directly related to data processing modeled (images, sensor, calibration, control points, trajectories, etc.), but also auxiliary data such as aircraft id, aircraft hours, pilot's name, sensor operator's name, contract requirements, etc. are modeled. This allows the system to retrieve information based on a complex query. For example, searching for the number of hours pilot X has on flown project Y would retrieve a meaningful result. This may not be significant information from a pure processing perspective, but from an enterprise standpoint this may be relevant data for resource management. Generally, an attempt is made to capture any relevant data that contains value for the company in regards to a project, especially if this can be performed automatically without or with only little human intervention.

IMPLEMENTATION

The NWG next generation architecture was designed from the ground up to support the integrated vision from acquisition all the way to web service delivery. It was built using the best of open source components and open standards. A composite model of following the Java Platform Enterprise Edition (Java EE) standard that is the industry standard for enterprise Java computing with the injection of popular open source frameworks when required has resulted in an extensible platform. A high level architecture of the system is shown in Figure 2. The major components that make up this architecture are:

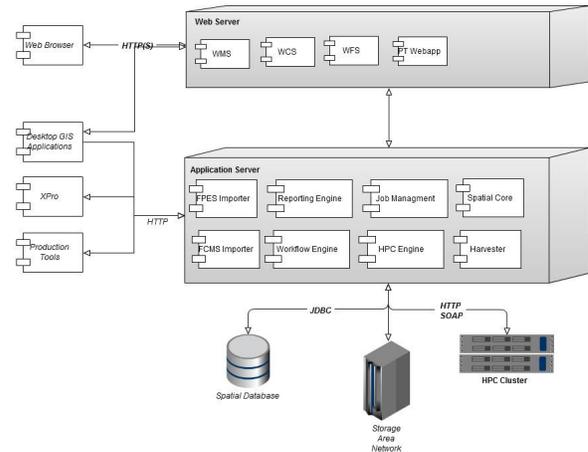


Figure 2: System Architecture

Harvester

The catalogue is able to harvest different kinds of resources (image coverages, flight export files etc.) using a simple plugin infrastructure. Multiple harvesters were developed using the plugin setup like FPESImporter that harvests the export from the Leica Flight Planning and Evaluation Software and FCMSImporter that imports the executed flights from the Leica Flight & Sensor Control Management System. The different harvesters in return use metadata harvesters registered in the system to provide metadata from diverse sources for each resource being harvested into the catalogue.

Work Flow Engine

The term Business Process Management (BPM) typically refers to business discipline or function that uses business practices, techniques and methods to create and improve business processes which are sequence of actions performed by humans and systems to achieve a business goal. Some of the advantages of using the BPM Systems are automation of the company business processes, declarative expression of the coordination between human tasks and automated systems and easier gathering of performance metrics.

The open source JBoss jBPM engine was selected as the work flow engine for the NWG architecture. jBPM is an extensible and flexible process engine that can run as a standalone server or embedded in any Java application. JBoss jBPM delivers the capability of developing new automated business processes and work flows with industry-standard orchestration using Business Process Model and Notation (BPMN V2.0). BPMN is a standard specification created by the Object Management Group (OMG) that defines a standard language to model business process including execution semantics that can be understood by both business analysts and technical people. The usage of jBPM has been instrumental in modelling our work flow from the start to finish including human steps and providing e-mail alerts for the different wait state in-between process.

Job Scheduling

There are a lot of job schedulers out there that have different capabilities. We selected Quartz the enterprise Job scheduler to be our scheduler of choice since it provides a very extensive sets of functionality and is a mature platform to build our applications on. Having a robust job scheduler integrated in the architecture has allowed us to run multiple jobs from harvesting daily flight logs to sending nightly what remains to fly reports efficiently in very customizable schedule.

Monitoring Services

The deployment stack consists of tools that provide a reliable means of detecting and monitoring faults in the production servers. Even though most production servers have redundancy, automating the response and alerting the operation group whenever there are interruptions in service and/or a resource reaches a critical state is a very important part of providing a reliable system. For these reasons we have developed monitoring tools based on open source that are able to diagnose problems and provide a self-healing environment if possible. For example, some of the capabilities of the alert and monitoring system are watch dogs for WMS servers using the GetCapabilities and GetMap responses.

High Performance Computing

The Condor® Project is our selected high performance computing platform to perform computationally intensive jobs on a grid of dedicated clusters. A module to interact with the Condor cluster for submission and monitoring of jobs within the application server was developed. This module allows us to increase the turnaround time of the business process and maximize the usage of the HPC cluster resource since the system is now able to submit jobs to the cluster as soon as the precondition for that step in the work flow engine are finished by either humans operators or another application.

Web Services

Standards enable systems to be built that easily discover and seamlessly combine spatial information from different sources and share it among many users and applications (Beshah et al., 2008). At NWG we have built an Enterprise Spatial Platform that is designed to securely manage and deliver geospatial data to users operating on rich client desktop systems and web clients. We

expose the different processed geospatial data and their status throughout the workflow using a number of OGC web services. Currently the system supports the OGC Web Mapping Service (WMS), Web Feature Service (WFS) and Web Coverage Service(WCS).

In addition, other core open source technologies like PostGIS, Hibernate, OpenLayers, GWT, GeoTools and others were integrated at different layers of the architecture. The selection and integration of these different frameworks and libraries to create a single coherent geospatial platform that is firmly based on both enterprise and high performance computing was the biggest challenge we faced.

REAL WORLD USAGE

NWG deployed the first version of the enterprise mapping system, coined Production Tracker, in 2010 and has since ingested every flight into the system as illustrated in Figure 3.

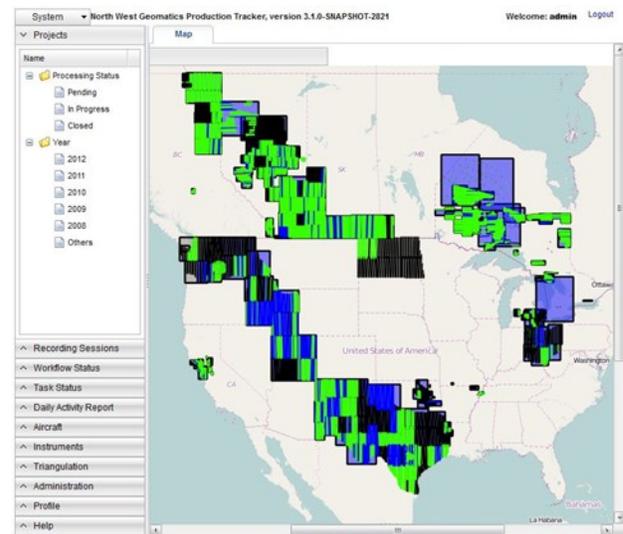


Figure 3: Overview of acquisition blocks in Production Tracker

Since the initial deployment the Production Tracker has been constantly updated and as of April 2012 the latest version exposes functionality that fulfils the requirements from Planning through Ingest and initial QC.

Once the flight crew has recorded imagery an email with an attached flight execution export that contains planned and recorded flight lines is sent to the office where by it is automatically ingested into the Production Tracker. The office can immediately check line coverage and consistency without having access to the full dataset. An example is given in Figures 4 & 5.

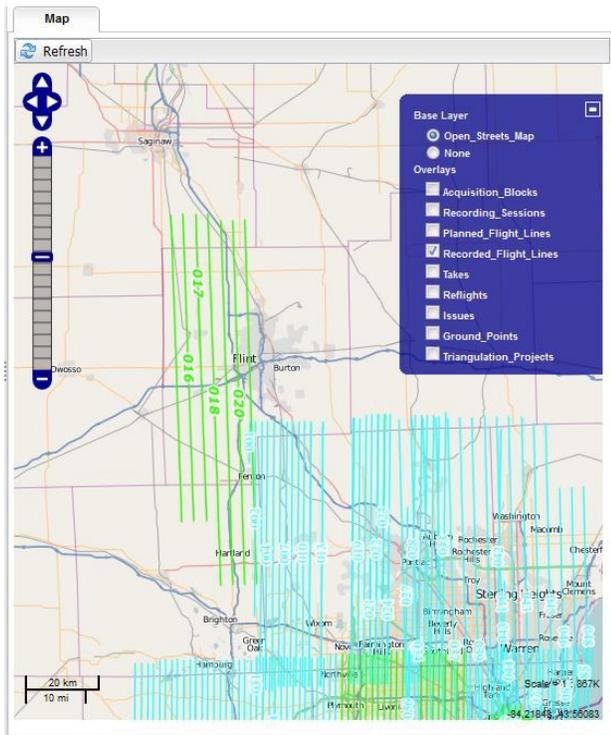


Figure 4: Planned flight lines

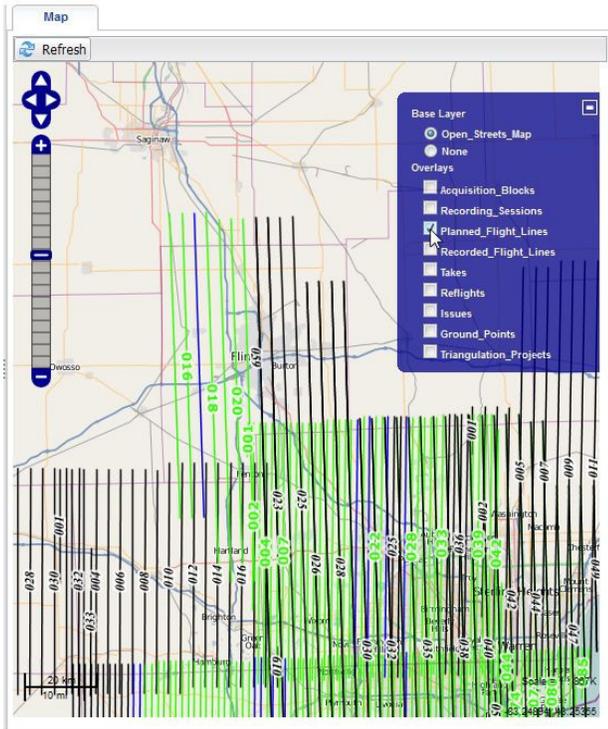


Figure 5: Recorded flight lines

Once the recorded imagery is received in the office the data is copied to the company network file storage and geo-referenced by an XPro application. The XPro Ingest package is fully connected to the Production Tracker allowing the actual footprints of the

imagery to now be viewable in the Production Tracker immediately— see Figure 6.

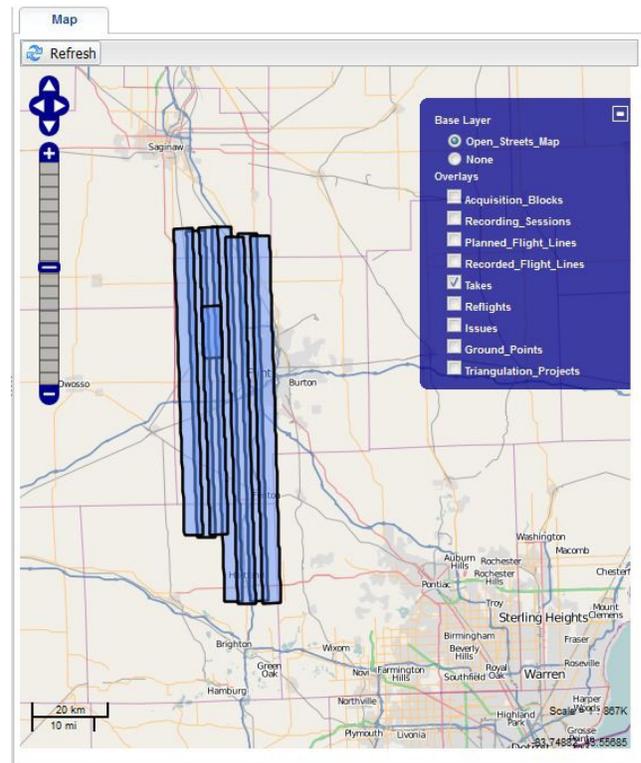


Figure 6: Footprints of ingested and geo-referenced flight lines

At this point all the imagery goes through a quality control process where any issues are recorded in the XPro QC image viewer which has been extended to communicate directly with the Production Tracker catalogue. Issues are traced with a feature drawing utility and persisted directly in the catalogue. A user with management rights then calls the re-flights and re-flight lists are sent automatically back to the flight crew.

The ability to view the actually flown flight lines before receiving the full data set together with the re-flight call functionality has helped optimize the flight order and therefore made the daily planning while the acquisition team is on location more efficient.

CONCLUSIONS

The enterprise mapping system as introduced in this paper has been a large step forward for the overall company productivity at NWG. The ability for stakeholders to view the progress of individual projects while being flown has helped identify potential problems earlier in the life cycle of a project and therefore reduced cost. The necessity for copying and moving data of the storage devices has also been greatly reduced as all the data is now stored in a central repository.

While only the work flow steps through initial QC have been fully implemented for production, the benefits have become obvious. The development of the remaining work flow with a tight coupling with XPro applications is underway. Initial plans to tightly integrate the Production Tracker with the online delivery system for other value added products like a point cloud that

contains calibrated color information (RGB and near infrared) for each point that is derived using a dense image matching method such as XPro SGM are underway. Additional work still needs to be done to integrate more with Business Intelligence (BI) products that provide a range of capabilities like reporting, data analysis, dashboards and data mining to improve the support for accounting and project management.

Phot. & Rem. Sens., Beijing, China, Vol. XXXVII, Part B1, pp. 1099 -1104.

Gehrke, S., K. Morin, M. Downey, N. Bohrer, and T. Fuchs, 2010. Semi-Global Matching: An Alternative to LiDAR for DSM Generation? *Int. Arch. Phot. & Rem. Sens.*, Calgary, AB, Vol. 38, Part B1.

Gehrke, S., R. Uebbing, M. Downey, and K. Morin, 2011. Creating and Using Very High Density Point Clouds Derived from ADS Imagery. *Proc. ASPRS Annual Conference*, Milwaukee, WI.

REFERENCES

Beshah, B., Martinez M., Stojic M., 2000, ISO/OGC Standards Impact On Next Generation Enterprise Spatial Platform, *Int. Arch.*