THE CITY OF CALGARY
INFORMATION TECHNOLOGY

The City of Calgary: Enterprise GIS Migration

Prepared for:
The 2008 ESRI International User Conference; Aug 4-8, 2008.
San Diego Convention Centre
San Diego, California

Author: Steven Waters

Abstract

The initial ESRI GIS environment at the City of Calgary was deployed in 1999. The goal of the GIS was to facilitate publishing of spatial information that originated in other systems and expand as data maintenance applications were migrated to the technology.

Fast forward 10 years, the City's vision of enterprise computing goals have evolved to encompass:

- Web centric applications;
- Supporting a mobile work force;
- Demonstrating the value of integration with new initiatives;
- Manage the lifecycle of the IT infrastructure.

The City chose to work with ESRI – System Design Team and ESRI Canada Professional Services to review our system architecture and validate our plans to support our corporate goals with modern ESRI technologies. The presentation reviews: where we started, existing corporate standards and our new target environment

City of Calgary Information

The City of Calgary is a rapidly growing community; the 2008 census measured 1,042,296 people residing within the City, a 13% increase since 2003. Calgary is well-known as a destination for winter sports and ecotourism with a number of major mountain resorts near the city and metropolitan area. Economic activity in Calgary is mostly centered on the petroleum industry; however, finance, transportation, agriculture, tourism, and high-tech industries also contribute to the city's fast economic growth.

City of Calgary Enterprise GIS

The Enterprise GIS system at The City of Calgary has grown and evolved significantly since the original deployment of ESRI technology in 1999. At the present time, the system has over 3000 users across the 30 lines of business that the City provides services. The vast majority of these users would not fit the traditional GIS power user who maintains data and prepares maps, but are data consumers that use GIS to provide valuable context and accurate information within various spatially enabled business applications.

The City of Calgary Enterprise GIS technology is managed by the GIS Centre within Information Technology on behalf of all City of Calgary GIS Users. The GIS Centre's role is to provide core technology, best practices and business consulting to City of Calgary business units.

Initial Implementation

The initial ESRI data serving technology put into place at The City of Calgary was an SDE 3 data repository. The SDE 3 environment was primarily used to share spatial data that had been created and maintained in other mapping systems (MicroStation) that did not have the scalability to serve information to numerous users within an enterprise (Figure 1). As little data was maintained in the ESRI GIS system there was a heavy reliance on data conversion methods and scheduled batch loading of data to keep the SDE repository reasonably current.

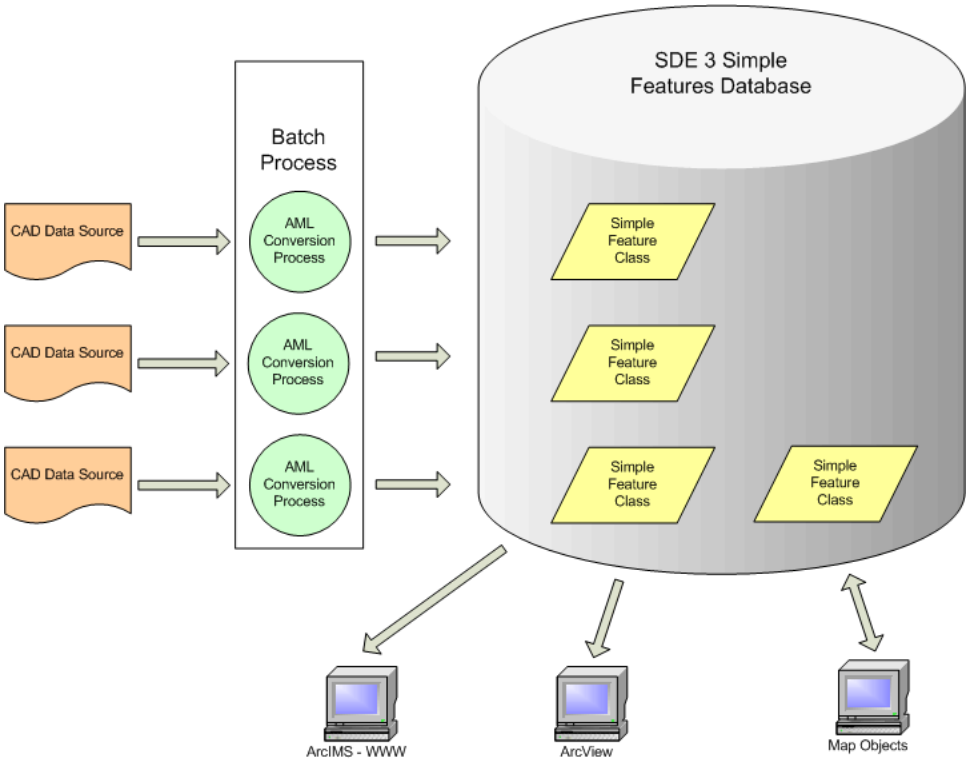

Figure 1: SDE 3 Architecture

Direct data maintenance was limited to a few focused applications built using MapObjects software. Several of these applications are still in use but are being transitioned to newer technology as opportunities are presented. This ESRI Simple Features architecture is still in production but been upgraded to the latest release of ESRI technology.

Advantages of the Simple Feature based architecture include:

- High performance and scalability;
- User access to continuous, non-tiled, layers of information.

Disadvantages of the Simple Feature based architecture include:

- Reliance on replication created latency issues;
- Input data systems collected geometry very well, but little or no attribution existed for most features;
- The features being published are not tied closely to the maintenance workflow;
- Few of the inherent relationships between the different features can be exposed;

The first major upgrade of the ESRI technology was from SDE 3 to ArcSDE 8.1.2. This was primarily a change in the underlying software technology that would position the Corporation to manage spatial relationships and data using the capabilities of the Geodatabase. In practice, the processes used to maintain the published geographic information within the Corporation were identical to the initial SDE 3 environment, yielding identical benefits.

The Corporation has currently deployed version 9.2 of the ArcSDE technology to support the maintenance and distribution of GIS data. Within in the current data architecture exist two primary data stores: PD1SPA for the storage of ArcSDE simple features, and PD2SPA for the maintenance of spatial information with enhanced Geodatabase properties (Figure 2).

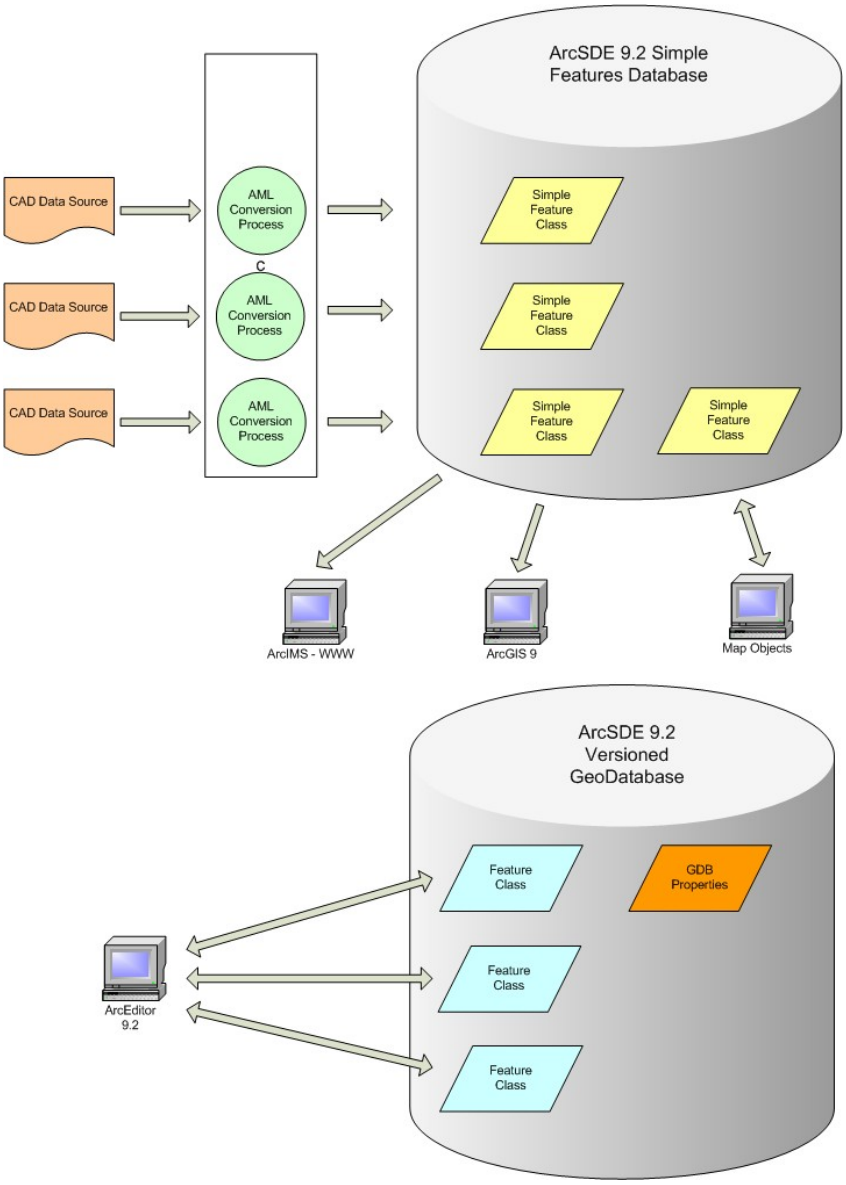

Figure 2: ArcSDE 9.2 Configuration

There are three ArcSDE instances currently in use:

1. SP1SPA - simple features staging and loading (Microstation maintained data is loaded and converted to this SDE server).
2. PD1SPA - simple features publishing.
3. PD2SPA - versioned Geodatabase editing (ArcMap).

The ArcSDE technology is implemented at Version 9.2 SP3 on a Linux RedHat 4, Oracle 10.2.0.3 database server. At this time, all applications connect through an application server process hosted directly on our server.

Client software deployed at the City of Calgary is extensive and is summarized in table 1. Deployment of desktop software is on various devices including, desktop connected LAN, Citrix remote desktop, disconnected laptop's. For applications that require mobile access that have extensive customizations and integrations with other systems, connecting laptops wirelessly to Citrix remote desktops has demonstrated itself to be useful in pilot exercises.

Desktop Software	Version	Deployments
ArcGIS – ArcView	9.2	1500
ArcGIS – ArcEditor	9.2	300
ArcGIS – ArcInfo	9.2	86
Map Objects	2.4	1600
ArcView	3.2	Decommissioned July 2008

Table 1 – Desktop GIS Technology Deployments

Internet server software deployed at the City of Calgary is in transition from a mature ArcIMS base, with the aim to transition to ArcGIS server as opportunities allow and summarized in table 2. Existing services are for both internal and external (citizen) use.

Server Software	Version	Services
ArcIMS	9.2	8
ArcGIS Server	9.2	2 in Development

Table 2 – Web GIS Services

At any given time during City of Calgary standard working hours, we typically observe 550 users connected to PD1SPA our simple features publishing environment and 70 users connected to PD2SPA our Geodatabase maintenance environment. GIS technology has clearly demonstrated it's value to The City of Calgary.

Drivers for System Renewal

Several sets of drivers for system renewal are present today. The groups influencing these drivers include: Information Technology Architecture requirements, Information Technology support, Business Unit operational requirements.

System Renewal from IT Architecture Perspective

Information Technology staff at The City of Calgary have been challenged to meet a number of high level goals to position itself so that it can ensure that we can continue to effectively deliver services to our internal business clients and citizens. These objectives include:

- Implement Microsoft .NET as the standard implementation platform for applications.
- Begin developing a standard set of web services for use by the public, vendors and internal application development
- Establish plan to integrate existing application portfolio into web-centric application architecture
- Develop GIS applications for the mobile environment.
- Demonstrate value of City integration technologies
- Transform and improve IT infrastructure
- Effectively manage the lifecycle replacement (LCR) of existing IT infrastructure.

System Renewal from System Support Perspective

The ESRI Simple Feature model does not support many of our business processes and integration requirements. As such many of our legacy systems are being renewed to take advantage of newer capabilities provided by ArcGIS Server and the Geodatabase.

A requirement of our Permitting and Addressing system is to store community name information with each record. Historically this information has been looked up on a map and then manually entered. Using a GIS web service we are now able to have the Community polygon layer be the definitive source of this information. The Community polygons are now being used in a proof of concept, that utilizes a service to discover the value for any point feature using web services and spatial processing.

The existing model of maintaining the data outside of the GIS, converting the data to a staging environment prior to publication creates numerous problems. These problems include: latency, long processing times and periods of time when data is not available and reduce our ability to deliver required services in a near 24*7 scenario. Examples relating to our address and ownership parcels include:

ADDRESS – a point layer contains approximately 850,000 spatial features across several types of civic addresses (parcel, suite, building, etc.). Typically there are 250 features changed on any given day, however, given the current maintenance practices and replications strategies, the layer must be rebuilt from scratch each day.

OWNERSHIP PARCEL – a polygon layer for ownership parcels within the City of Calgary. Currently 315,000 parcels are published in production, typically 100 parcels change on any given day, but the City rebuilds the entire layer each evening. The time taken to do full replication is increasing and makes the entire system more brittle.

Conversion, Staging, Loading and QA process introduce serious latency up to 48 hrs depending on source system until publication in GIS (PD1SPA). This builds redundancy into the system, as the City needs to check for conversion success, loading success to staging SDE, replication success between staging (SP1SPA) and published (PD1SPA). The process reduces the City's ability to deliver near-real time information within corporation and to citizens. A prominent example of critical data that experiences this latency is found within our Ownership Parcel base.

OWNERSHIP_PARCEL: Daily GIS Update Process

Day 1:

Cartographer checks out section from Oracle/MicroStation application

Checks in changes at the end of the day

Database extract via FME and features created and loaded to SP1SPA for staging

Day 2:

QA reports against SP1SPA

If no problems, automatic refresh

Day 3:

By 7:00 am (ideally sooner) data is refreshed to PD1SPA / PD2SPA. The refresh is based on SDEExport / SDEImport.

A project to migrate the maintenance of the Ownership Parcel base to the ArcGIS platform is underway and expected to deliver a base mapping system in 2010.

System Renewal from a Business Perspective

Today, significantly more data is being maintained natively in the GIS environment. This has created an additional challenge as there is no publication environment to support the broad consumption of the rich data and expose the relationships that the data contains. The growth in data maintenance with the GIS platform has grown from less than 10% to greater than 40% today (Figure 3).

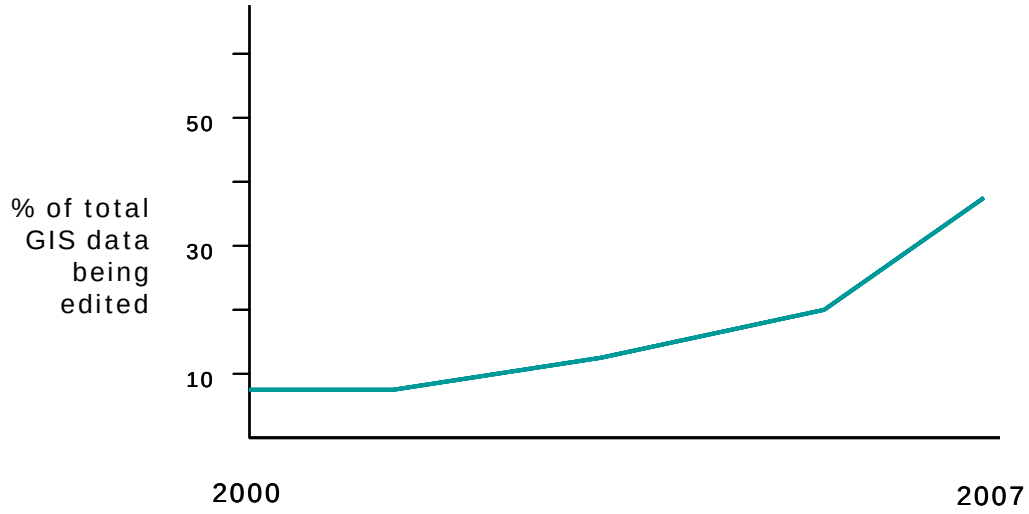

Figure 3 – Growth of GIS Data Maintenance

It is estimated that more than 75% of the spatial data at the City of Calgary will be maintained within the GIS by 2010 or 2011. This further necessitates the requirement for a Publication Geodatabase that will not have the service disruption associated with a lengthy GIS data refresh process.

A simplification of workflow has been key to gaining momentum in maintenance of GIS data in the ArcSDE environment. The City of Calgary has developed a generic versioning framework and set of tools to simplify the end user experience with versioned editing. This tool can be deployed on it's own as an extension to ArcGIS, or incorporated with any of the City's other GIS extensions. The tool and framework provide a persistent surrogate default version for users to work from. A custom "start edit" button automatically creates a child version and switches the users workspace to their personal workspace. A custom "stop edit" button automatically reconciles and posts the users changes to the surrogate default version. While working with the tool, all of the standard ArcMap functionality remains available to the user. A common approach to versioned editing also simplifies the administration of the Geodatabase.

While internal GIS analysts and operational staff may not require access to GIS data at all hours of the data, extended delivery of GIS data and applications are coming from specific groups including:

Business drivers for extended GIS support:
The City of Calgary customer contact centre is operated 24 hours per day and has mapping components.
Utility operations require timely access to spatial data in the field during off hours (when all emergencies seem to happen).
Citizen access to portals that offer web mapping including: Property, assessment and amenities outside of City business hours.

Evolution of an Enterprise System

The City of Calgary's GIS implementation, as-is, is best described as meeting our past requirements. As a rapidly growing Municipality we are being asked to deliver more types of services than we originally envisioned. The majority of these services can no longer be supported using our primary model of supporting Simple Features as a primary model and Geodatabase as a secondary model.

Upgrading our technical environment is also becoming increasingly difficult. With the current technology, our data server upgrade can only be completed once all of the applications have been migrated to the higher release. This requires an upgrade to the entire environment which includes:

- 60 ArcGIS extensions, 5 Map Objects applications, 8 ArcIMS Services;
- Nearly 1900 ArcGIS desktop deployments;
- Nearly 1600 Map Objects deployments;
- Integration with numerous other systems;
- Package and Deploy Issues.

For the majority of our applications, the upgrade process does not add any new business functionality to the end user. However the focused application or extension must be re-compiled and re-packaged and tested for the user; this requires significant effort given our processes.

A large number of deployments of stable Map Objects software were predicating the data model used in the publication environment. These mature applications are to be renewed as opportunities allow, due to the integration with other technologies, migration dates are also dependant on other vendor technology.

Creation of a Data Architecture to Support Multiple Requirements

Our solution to supporting what has effectively become a dual environment is to formally create a dual environment. We have built a formal architecture to support Geodatabase transactional editing and publication in addition to the environment for simple features.

The creation of the dual environment is the only practical way that we can meet the challenges for GIS services that Business, IT Architecture and System Support have placed while still maintaining continuity with our mature applications. An overview of our GIS Data architecture is presented in Figure 4.

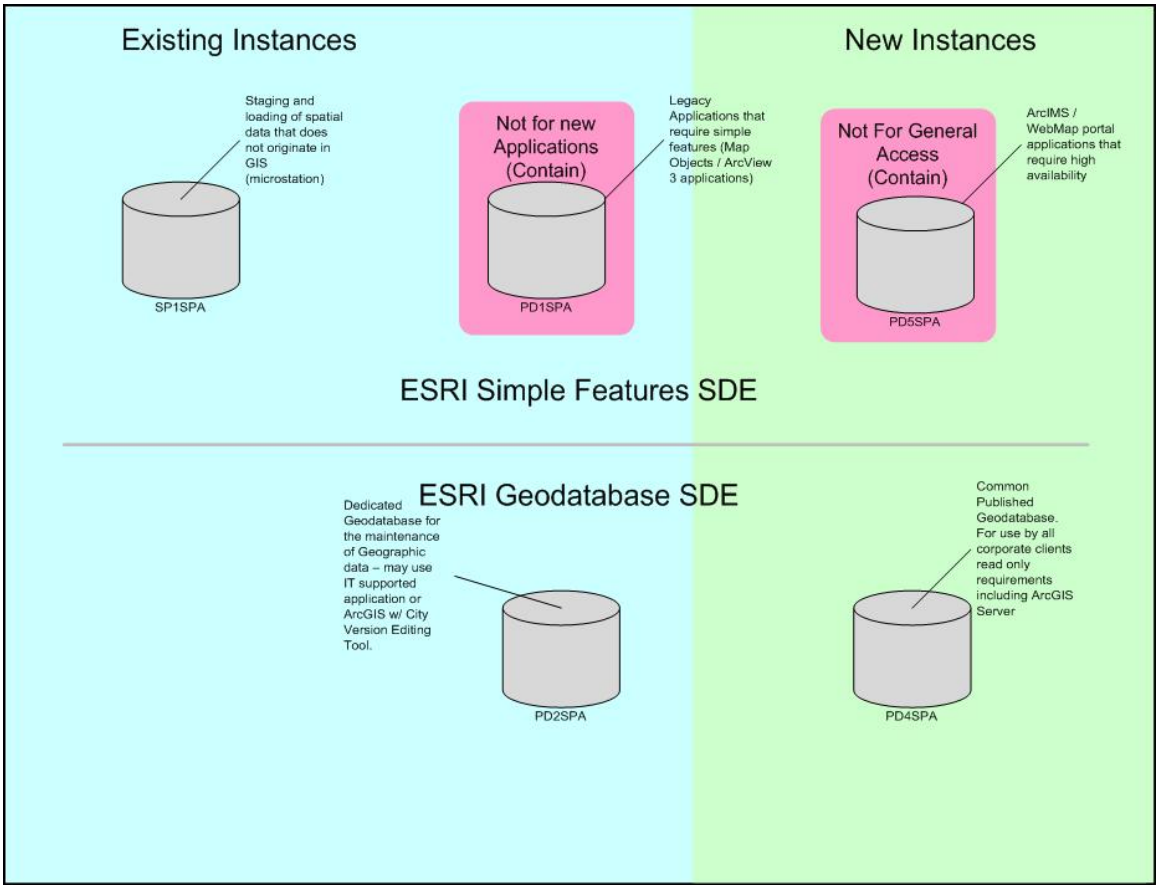

Figure 4 – Diagram of GIS Data Architecture

Each element in the data architecture has a specific function that is summarized in Table 3.

Instance	Status	Purpose
SP1SPA	Existing	Stage and load simple features from CAD systems
PD1SPA	Existing	Publish simple features
PD5SPA	New	Publish simple features for web applications only. Allows us greater scalability options for certain annual events that generate high usage.
PD2SPA	Existing	Maintain Geodatabase features via versioning
PD4SPA	New	Publish Geodatabase features utilizing Geodatabase replication and versioning

Table 3 – Elements of Data Architecture

Splitting our environment has created additional work from a maintenance perspective and has required careful planning to ensure that this is not the long term, normal course of operations. It is expected that that the architecture will shrink as the technologies for supporting the Simple Features architecture are retired.

In an effort to realize the goal of transitioning the environment the following high level tasks have been completed:

- Cost estimates for transitioning Geodatabase capable application to the new environment;
- Definition of publication subject areas and sources of data;
- Confirming the use of Geodatabase replication as viable technology for supporting the 1-way publication of data.

Our target GIS Data architecture for 2011 is dependant on the decommissioning of several technologies including: Map Objects and ArcIMS (figure 5). To be successful we must also have the support of vendor partners at The City of Calgary. The majority of our Map Objects deployments are used to support mapping functions within Posse our Permitting and Licensing work flow tool. The Posse application will need to support a map viewer other than Map Objects.

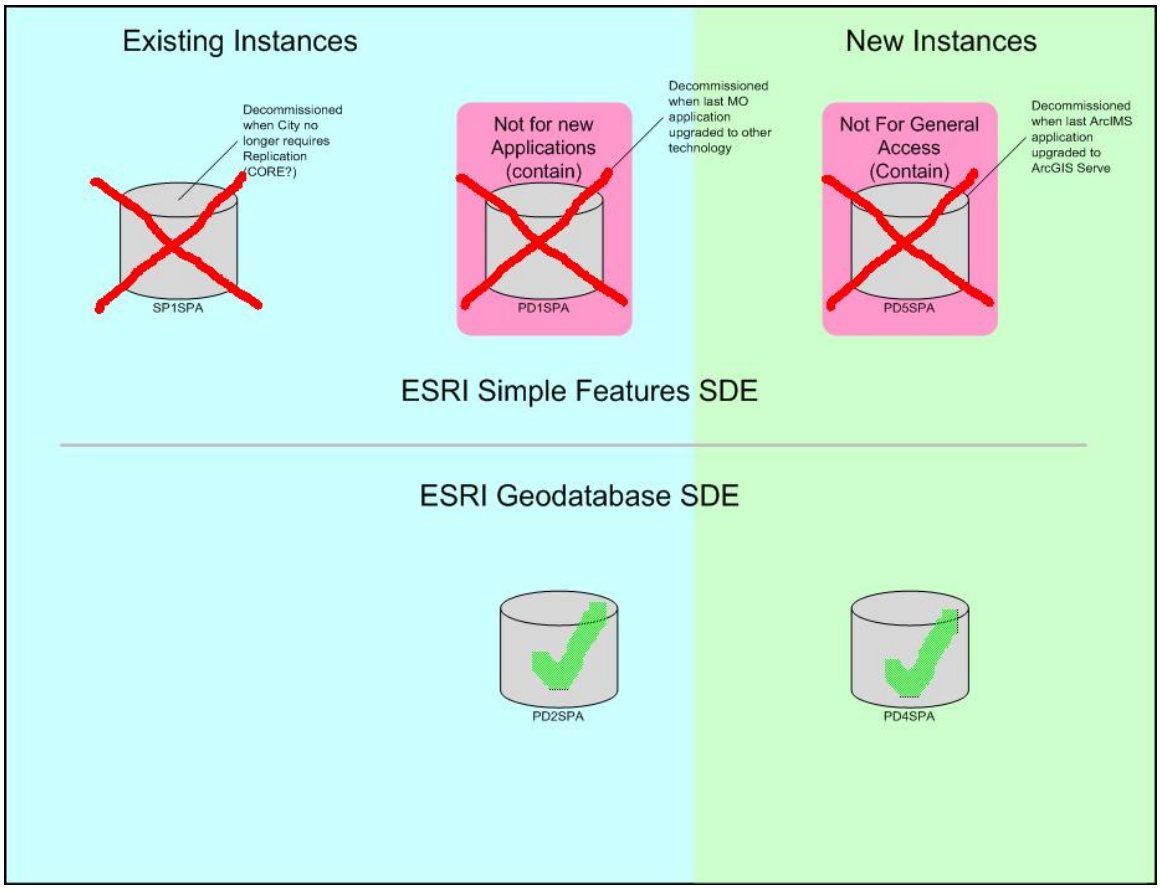

Figure 5 – Target GIS Data Architecture 2011.

Use of ESRI Systems Integration / ESRI Canada Professional Services

The City of Calgary recently engaged ESRI Systems Integration and ESRI Canada Professional Services to conduct an on-site GIS Systems Architecture review, analysis and capacity planning exercise that would work with our 2008 to 2011 timeframe for system renewal, utilizing ESRI's System Design Methodology. The exercise vetted our approach to transition and creation of a new data architecture to support upcoming initiatives and validate that the computing environment would support the required workflows.

The approach allowed us to model workflows that would support new GIS initiatives that would include:

- Citrix deployment to field based applications;
- ArcGIS Server Mobile Application Framework;
- Expanded use of Citrix for deployment of GIS technology to casual users;
- ArcGIS Server for traditional GIS / Geoprocessing;
- ArcGIS desktop editing workflows;
- ArcGIS desktop analysis workflows.

Over the next three years the growth in GIS use at The City of Calgary is expected to be moderate in terms of total new users. However, given the changes in workflows to maintenance based activities in the Geodatabase, the impact on our data servers will likely be greater than first anticipated. Estimates based on the calibration of custom workflows to existing server loads are presented in figure 6, these figures demonstrate the change in map displays per minute (see ESRI System Design Methodology), based on 2008 and observations and 2011 estimated workflows. It is estimated that our server loads placed on our GIS data maintenance environment will grow by over 300% in the time period. Using this information we were able to specify any new hardware requirements that had emerged based on the loads that would be placed on the system.

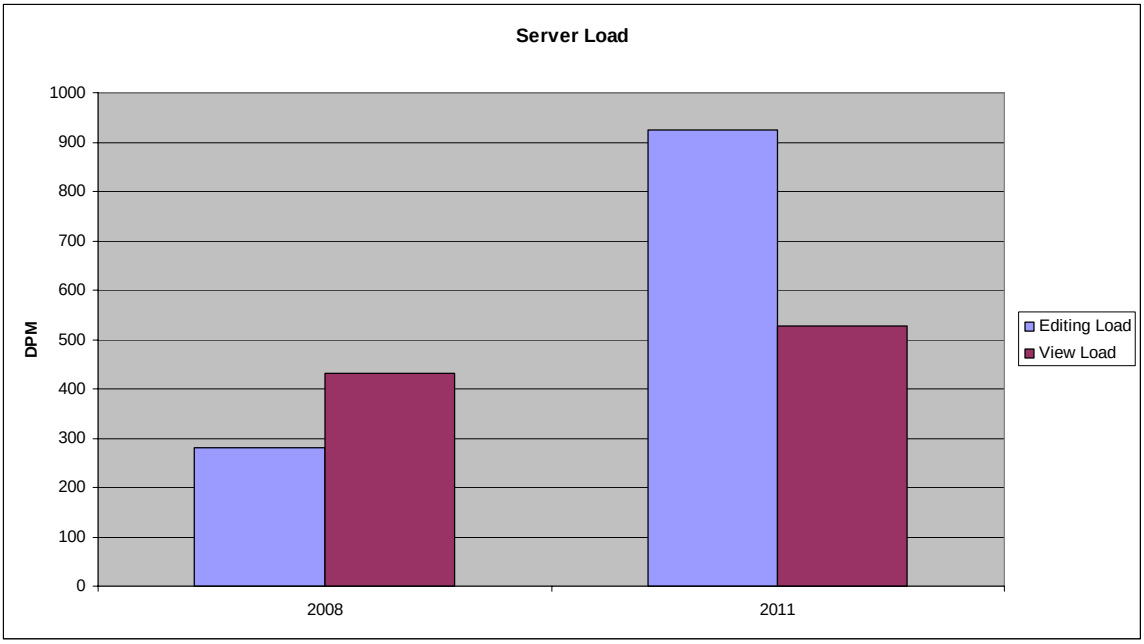

Figure 6 – 2008 to 2011 GIS Server Loads

General observations from the system review and analysis of current work flows include:

- The City has a lot of users working with light GIS workflows that are integrated with other application workflows.
- A lot of people using "a little bit of" GIS translates to moderate server loads, lower than would be expected based on traditional power user workflows and observed number of GIS users.
- The number of Asset Management applications on the horizon is likely to significantly grow the editing loads on the GIS data servers as more field based work is performed.

A Model of the 2008 and 2011 GIS environment was prepared utilizing the ESRI Capacity Planning tool along with our custom workflows. Utilizing this tool and the workflows we have been able to better measure the types of loads that our different workflows would be placing on our infrastructure and can sensibly evaluated different deployment options and their impacts to effectively manage the environment. This approach can help us evaluate if it is more appropriate for certain groups to utilize Citrix desktop rather than ArcGIS desktop in specific settings.

It is our intent to maintain the model and benchmark our observed performance with new workflows as additional GIS applications and maintenance activities are deployed within the organization. To ensure that the model is reflective of current use, touch points with the City's Application Development Methodology will gather appropriate information from any proposed system for incorporation to the capacity planning tool.

Additional benefits to the workshop and validation of the GIS Data Architecture put in place was a confirmation that the upgrade process will be simplified in a near release. ESRI has indicated that we will be able to deploy new database connectors to Geodatabase capable desktop applications (ArcGIS not Map Objects) that will facilitate

connecting an older generation client to a newer generation data server. This will have significant positive impact on the effort that is spent on upgrading our environment and allow for the City to be more nimble in terms of releasing GIS technology.

Conclusions

The City of Calgary will continue to grow its use of Enterprise GIS technology to support business and citizens. In response to a variety of drivers: business, IT infrastructure and system support, the City has developed a realistic plan to re-invent a popular GIS implementation that will take advantage of new vendor provided capabilities and deliver the requested functionality. The engagement of ESRI technical teams added significant value to our plans through the delivery of a documented model that will allow us to more easily gauge the impact of system changes prior to deployment through the 2008 to 2011 time period.

Author Information

Steven Waters
City of Calgary
Calgary, Alberta, Canada

Information Technology, GIS Centre
swaters@calgary.ca