

ActiveBatch®

Customer IT Automation Success Story

Graymont Services Turns to ActiveBatch to Automate the Mining of Data



Company: Graymont Services

Industry: Mining

Customer Site: Richmond, British Columbia, Canada

Brief Company Overview:

Incorporated in 1948, Graymont is a family owned and controlled Canadian private corporation headquartered in Richmond, British Columbia (a Vancouver suburb). Graymont is the second largest producer of lime in North America, with close to 20 locations serving key market regions across Canada and the United States. The Company also operates a construction materials business, Graymont Materials, which provides construction stone, sand and gravel, asphalt products and ready-mix concrete for infrastructure and general construction needs from four locations in upstate New York and southern Québec.





SUCCESS STORY HIGHLIGHTS

- Reduce Batch Runtimes by 55%
- Improve Processing Success to Over 95%
- End-to-End Automation of the Data Warehousing and IBM Cognos BI Processes via a Single Solution
- Centralize Task Scheduler, SQL Server Agent, Database Scripts and SQL Stored Procedures to Make Better Use of Hardware and Improve IT Service Levels to the Business

The “Elemental” Approach to IT Automation: Discovering Limitations

For Paul Epp, IS Manager and Cognos Application Manager at Graymont Services, taking an “elemental” approach to automating the organization’s data warehousing and BI processes was costing him both time and sleep. Graymont, a leading provider of lime and stone products in North America, relies on the timely and accurate delivery of Cognos reports to keep financing and accounting in order, but the end-to-end automation of the underlying data warehousing processes and scheduling of Cognos reports was unreliable and problematic.

Graymont’s previous IT automation strategy consisted of a collection of scripts and batch scheduling solutions. “We had Cron, we had Task Scheduler, we had SQL Server kicking off DB scripts and SQL Stored Procedures, we had it all,” Epp says. This collection of “point” scheduling tools was responsible for the end-to-end automation of ETL processes that upload budgeting, accounting and financial information into a data warehouse to be reported on via IBM’s Cognos Business Intelligence & Financial Performance Management solution.

This system presented a series of limitations. There was no centralized monitoring or alerting mechanism. Many times a failed job wouldn’t be discovered until the next day when a business user noted a report wasn’t delivered or the data was out of date. There was no ability to automatically restart/rerun a failed job and no ability to run jobs in parallel. Finally, Epp couldn’t build conditional dependencies and constraints between jobs. When a job did fail all other downstream jobs within the workflow would fail regardless if they were dependent on that job.

As a result, batch processing success rates hovered around 30%. These deficiencies were only magnified during Graymont’s budget cycle every autumn. “During that three week period I’d get paged almost every night at 1 or 2 AM and I’d have to login and restart a failed job,” Epp says.

Graymont Services went to market in search of a job scheduling and workload automation solution to unify its multiple scheduling tools within a single enterprise solution. The IT organization considered half-dozen vendors in addition to ActiveBatch. In the end, ease of use proved a critical factor. “For many of the legacy scheduling solutions the learning curve and complexity was too high; the workflow development and admin tools were too complex,” Epp says. “During the POC process, the fact that I was able to use the ActiveBatch documentation to start building workflows without any assistance spoke volumes in terms of the ease of use and ease of implementation.

“The amount of code we have to maintain has been reduced to virtually nothing, which is great because something as simple as adding an alert in the script could fail an entire overnight workflow due to one parsing error.”

- Paul Epp, IS Manager and Cognos Application Manager,
Graymont Services



The ActiveBatch System View

The ActiveBatch System View, which provides a graphical interface to build workflows and manage dependencies/constraints between jobs via drag and click commands, is a favorite of Epp. The System View provides the ability to build workflows that branch into multiple jobs that run in parallel, with dependencies and success/failure triggers between them, before concluding with a single "child job." "A lot of the other tools we evaluated made it very difficult to run multiple jobs and then bring the workflow back together into a single job," Epp says. "The System View makes it simplistic by allowing me to drag the steps into order and then click and drag between two jobs to establish a constraint or dependency."

"The System View makes it easier by allowing me to drag the steps into order and then click and drag between two jobs to establish a constraint or dependency."

- Paul Epp, IS Manager and Cognos Application Manager, Graymont Services

The screenshot displays the ActiveBatch Admin V9 interface. On the left is the Object/View Navigation tree, showing a hierarchy from ActiveBatch Network down to specific jobs like 'End of Day Run'. The main window shows a 'System View' of a workflow with 10 nodes. The workflow consists of the following steps:

- End On Day -- Start (Green)
- Pre Process Store Records (Blue)
- FTP Job Step_Copy1 (Yellow)
- Delete File -- File Operation (Yellow)
- Open a Web Site (Blue)
- Copy File -- File Operations (Yellow)
- Validation (Yellow)
- Report (Yellow)
- Complete Process (Green)
- Validation Review (Yellow)

Arrows indicate dependencies: 'End On Day -- Start' connects to 'Pre Process Store Records' and 'FTP Job Step_Copy1'. 'Pre Process Store Records' connects to 'Delete File -- File Operation'. 'Delete File -- File Operation' connects to 'Validation'. 'Validation' connects to 'Validation Review'. 'Open a Web Site' connects to 'Copy File -- File Operations'. 'Copy File -- File Operations' connects to 'Complete Process'. 'Report' connects to 'Complete Process'. There is also a dashed dependency from 'Pre Process Store Records' to 'Open a Web Site'.

At the bottom, the 'Instances' table shows the execution history for the selected job:

Name	ID	State	Execution Time	Duration	Exit Code	Exit Code Des...	Queue	Tags
End of Day Run	16704	Succeeded (1x)	7/5/2011 3:12:26 PM	1.78 min (-506.52 min)	0			
End of Day Run	16051	Succeeded (1x)	3/28/2011 2:00:53 PM	15923.23 min (+15597...	0		MainQueue	\$SLA
End of Day Run	16784	Succeeded	7/28/2011 4:25:56 PM	2.08 min (-252.97 min)	0		MainQueue	\$SLA
End of Day Run	55552	Succeeded	1/15/2014 12:34:54 PM	0.57 min (-15755.98 min)	0		MainQueue	\$SLA
End of Day Run	16734	Failed (1x)	7/22/2011 9:14:04 AM	9065.33 min (+8810.28 ...	-2144862...	Incomplete run	MainQueue	\$SLA
End of Day Run	16543	Failed (1x)	5/13/2011 11:11:12 AM	2.25 min (-1013.85 min)	-2144862...	Incomplete run	MainQueue	\$SLA
End of Day Run	16704	Succeeded (1x)	7/5/2011 3:12:26 PM	1.78 min (-506.52 min)	0		MainQueue	\$SLA
End of Day Run	16051	Succeeded (1x)	3/28/2011 2:00:53 PM	15923.23 min (+15597...	0		MainQueue	\$SLA

ActiveBatch's Queue Management

ActiveBatch's Queue Management capabilities were also important. Prior to ActiveBatch, Graymont would execute jobs sequentially on a single server. Using platform-specific scheduling tools such as Cron and SQL Server Agent meant running jobs in parallel across multiple servers was impossible. ActiveBatch's Queue Management capabilities allows users to assign jobs across multiple servers, either manually by creating an Execution Queue and assigning it to a server or by leveraging a Generic Queue to allow ActiveBatch to dynamically select where the job is run based on different scheduling algorithms. "We weren't making good use of the hardware we had available," Epp says. "We used ActiveBatch's Queue system to build a plan that covered all data dependencies and the Generic Queue handles how many jobs would run at once and which servers they run on." As a result, jobs now run in parallel, the IT organization is making better use of its hardware and batch runtimes have been cut by 55%.

Jobs now run in parallel and Graymont is making better use of its hardware and batch runtimes have been cut by 55%.

Cognos BI: The Last Piece of the Automation Puzzle

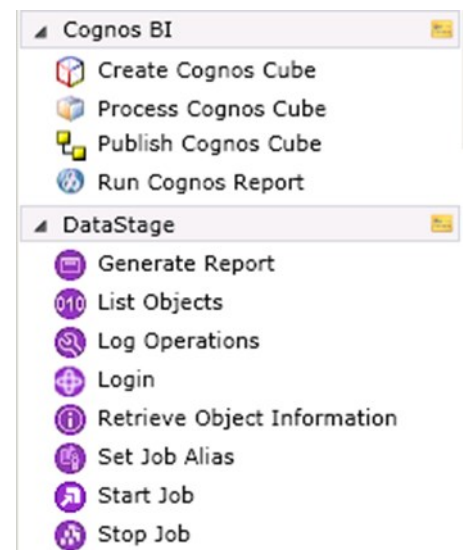
The batch processes and database workflows automated within ActiveBatch serve as the foundation on which Cognos pulls data and creates business reports for end users, including sales and operational reports, HR reports and others. Because the scheduling of Cognos BI jobs was dependent on the successful completion of upstream jobs executed within ActiveBatch, the next logical step was to bring the scheduling of Cognos within ActiveBatch. "It was the last piece of the data warehousing/BI automation puzzle."

Relying on the Cognos batch scheduler to schedule and execute Cognos reports presented a series of limitations. Its scheduler was limited to basic date/time functionality, forcing Epp to "pad" several hours between the completion of the ActiveBatch jobs and the scheduled start of Cognos tasks. Additionally, Cognos has no monitoring or alerting capabilities, which meant one of Epp's team members spent several hours each month manually checking to ensure the Cognos jobs had completed successfully. If a job didn't run successfully, often times the problem wouldn't be discovered until a business user noticed a report wasn't delivered or the data was outdated.

To automate Cognos from within ActiveBatch, Graymont Services licensed the ActiveBatch Extension for Cognos, which adds a series of Cognos production-ready Job Steps to the

Integrated Jobs Library to allow IT organizations to automate the generation and distribution of Cognos reports. This includes Job Steps for creating, processing and publishing a Cognos cube, a *Run Cognos Report* Job Step and the *Run Cognos Job*, which Epp leverages to run all of his preexisting Cognos jobs from within ActiveBatch. From within the *Run Cognos Report*, Epp executes a series of Cognos jobs that contain one or multiple reports and uses ActiveBatch variables to designate which report (HR, sales, etc.) the job should process.

The ActiveBatch Integrated Jobs Library Job Steps for IBM Cognos BI



Learn more at [ActiveBatch.com](https://www.activebatch.com) 

Achieving Success with an “Architectural” Approach to IT Automation

Due to the improved scheduling reliability of ActiveBatch and the reduced reliance on scripts, the IT organization has seen its batch processing success rate improve from 30% to over 95%, and the nightly failures of batch processes Epp would have to handle during the end-of-year budget cycle has been reduced from nearly every night to just once a month. When something does go wrong, ActiveBatch’s Instances View is leveraged to drill into individual jobs, view the log file, quickly diagnosis and fix the problem, while the Gantt and Operations View are popular choices to provide an enterprise-wide view of jobs running across the organization to identify processing bottlenecks.

“Consolidating the scheduling of scripts, database tasks and application processes such as Cognos into a centralized solution has proven invaluable,” Epps says. “It’s improved productivity within the IT organization, allows us to make better use of our hardware and flawlessly executes our batch processes, which makes the business happy and lets me sleep better at night.”

