HOW CONTROL-M HELPS BMC'S IT DEPARTMENT DELIVER DATA DEMOCRATIZATION

How do you give access to the enterprise data warehouse to line of business users plus put powerful, advanced analytics and data warehouse tools in their hands so they can create their own visualizations, reports and dashboards – without interfering with the thousands of workflows that keep the business running on a day-to-day basis? See how we safely opened up BMC's enterprise data warehouse (featuring Snowflake, Netezza, MicroStrategy, Oracle, Business Objects, Tableau, Alteryx and other technologies, including AI and machine learning), as part of a data democratization program that now enables business users across the company to use self service to create new data tools that are making the company more agile and responsive.

To be a modern competitive business, you have to be able to garner insights from data, and quickly make decisions on those insights. BMC knows this well, and we realized that for us to become more flexible and responsive, our enterprise data warehouse (EDW) – an asset critical to decision making and the delivery of essential business services – needed to do the same. Specifically, we needed to make our enterprise data warehouse easy to use and accessible to business users (the ultimate stakeholders) throughout our company. Despite prior investments in our EDW, business users found it hard to access and use data. To keep up with growing business demands, we knew it was crucial to develop a strategy based on self-service principals like mobile banking and ATMs. The more we could democratize data, the more value we could gain.

Self-service helped turn our enterprise data warehouse from what was considered an inflexible

vault to more of an ATM machine that many employees now use to conveniently access data, analytics and new services that are making BMC more responsive. By automating processes for accessing, merging, and analyzing various data sets, applying analytics and enabling output to MicroStrategy, Tableau, Alteryx, or other solutions, we've made our data more available and more valuable because employees can use it in new ways. Now, each week an average of 1,003 employees use the EDW and its more than three dozen source systems to run their own analytics and create reports, visualizations and dashboards. Overall, 54 percent of our employees have taken advantage of the self-service capabilities we created to access the EDW. Helped by the new information and analytics, they've reported 40 to 50 days saved per year in sales operations, more proactive customer service, countless hours saved during the time-sensitive quarterly close process. Automation has freed up lots of weekend time because employees no longer have to prep data for Monday reports or check to make sure reports and other jobs are running.

In this blog I'll share details of how and why we developed a data democratization strategy, some of the changes we made to our EDW ecosystem to make it happen, and how BMC's own Control-M application workflow orchestration platform became the key enabler for delivering powerful new self-service analytics and workflow management capabilities to thousands of business users around the world.

Setting the strategy

Like many companies, BMC is on the journey to becoming an Autonomous Digital Enterprise. One of the core pillars of this strategy is to become a data-driven business. To maximize the value of the vast data at our fingertips, we knew we needed to make more data sources and analytics capabilities available to our business users in real time. Our line-of-business users are eager to find new ways to use enterprise data to help the company. Many of them have the technical skills to do innovative, sophisticated work with our data, but had to rely on IT for access to it. Without significant changes to our IT systems, role, and culture, we simply couldn't keep up with the ever-growing volume of requests from our business users for new dashboards, visualizations, reports, and access to new data sources. Business users were frustrated. Some viewed IT as an obstacle to innovation, rather than an enabler. In response, we undertook a major a program of data democratization – making centralized data directly available to distributed business users and giving them business intelligence tools to get the information and insights they wanted anytime, anywhere.

In our case, that required major changes to the architecture and access points to our enterprise data warehouse and to how it interacted with our ERP and the other source systems that feed it. Along the way we identified the need for an overarching strategy and IT cultural change at the highest level. To achieve the promise of democratized data access, our strategy had to be supported with a comprehensive automation and orchestration strategy across our information systems. If we could automate operations and build in controls, we could offer new self-service capabilities at the user level so our employees in various roles could find new ways to take advantage of real-time access to enterprise data and analytics. When it became clear that automation, orchestration and self-service were the keys to becoming an Autonomous Digital Enterprise, it also became clear that we should make Control-M, which was already being used extensively within our IT infrastructure (as well as thousands of companies globally) – the conduit to carry out enterprise strategy and the end-user level.

Navigating the environment

Our EDW isn't really a singular item, it's a collection of assets. Netezza has been the foundation, but we've introduced Snowflake to supplement and potentially replace it. We've identified more than 35 source systems, including our Oracle ERP, Salesforce, CallidusCloud, Hyperion, OneStream, plus many function-specific applications and numerous scripts. We use six different ETL tools to bring data into the EDW environment. Once it's there, it is reorganized into multiple data lakes, other pools, folders, and now, user-managed databases (UMDBs) that support specific needs for finance, marketing, sales operations, customer support, and other business functions.

When business users want a new dashboard or other service request, the workflow for collecting, updating, and processing the data often cuts across several sources, including the hundreds of applications used throughout the company. As volume grew, it became harder to orchestrate these processes so the workflows did not interfere with others while still using the most up-to-date data. Giving thousands of business users self-service access to the EDW and its dozens of source systems would make it harder than ever to coordinate data refreshes, ETL operations, dashboard updates, and other workload scheduling and execution. It also became even more important to prevent errors and hang-ups, because a single failed workflow execution now had the potential to affect many, many others.

We had learned in the past that when analytics and other data operations scaled, data quality problems often surfaced. For example, at higher scale the ETL process will take longer to complete. Without planning, that could cause a dependent job (say a dashboard update) to run before the ETL process was done refreshing the data. The result would either be a job failure, or execution with out-of-date data, which represents a data quality problem. We were about to scale our operations in a big way through self service, so we had to proactively address data quality.

We knew improving the quality of data going into our data warehouse and improving the orchestration of data warehouse-dependent jobs with systems throughout the enterprise, were what we needed to satisfy our volume and speed pressure. When we dug deeper, we determined it was critical to have an abstraction layer that can sit above applications to manage ETL, workflows and custom scripts.

Control-M gave us the flexible guardrails we needed. It gives us a single interface to orchestrate all the file transfers, ETL tasks, and other handoffs between the enterprise data warehouse and our source systems, and provides proactive alerts about potential workflow failures that will affect others. It helps us ensure self-service and other jobs run on the most up-to-date data, and even has some features that automate the data quality checks.

New users, new approach

The new ecosystem needed to accommodate everyone from executives and end users to analysts and data scientists. It had to allow users to blend trusted enterprise data with their own internal and external data sources.

But that's not all that had to change. We recognized that we as an IT department needed to change. An ecosystem that focuses more on enabling citizen developers and analysts than on delivering IT-oriented solutions requires a shift in mindset. To become innovation enablers, the new engagement model had to become "guide and recommend," rather than "guard and control."

While we needed to share control, we were still responsible for information security and system reliability. That's an uncomfortable position that required more system and cultural changes to navigate. We knew there was a lot of great data in the enterprise data warehouse that users could benefit from. We were confident in the data warehouse but were concerned about the security implications of opening it to users. Control-M helped give us the confidence we needed because it puts controls in place and automatically enforces data access policies and workload execution hierarchy while giving users self-service tools.

The new environment

To open the EDW and analytics resources to a wide user base, we created a multi-layer data ecosystem that is both centralized and distributed at the same time. Control-M manages workflow execution across the EDW and all the systems it interacts with, provides the user interface, and applies task automation and controls. The new ecosystem includes a trusted enterprise data warehouse and a metadata layer, enhanced with departmental self-service data preparation, visualization, and storage. Four frameworks make up the ecosystem:

Enterprise Data Warehouse – This IT-developed warehouse layer serves as the foundation for most analyses. It includes a corporate information factory (CIF) that stores historical fact and dimension data in third normal form (3NF), and a reporting-friendly dimensional view layer. Most of the data loads run once or twice per day, with some exceptions, and are developed following industry best practices for ETL and data architecture. The EDW is the single source of truth of data from all major BMC source systems.

Enterprise Data Hub – The MicroStrategy schema layer serves as the enterprise metadata for BMC, and it provides a user-friendly abstraction layer that points to warehouse tables and columns. Without coding, users can drag and drop pre-defined attributes, metrics, filters, and prompts into their reports. Because the schema, joins, and logic are already defined by IT, answers are consistent and accurate. Users that have already invested time and money into their own preferred visualization tools are welcome to keep using them, but they need to use the Enterprise Data Hub to source data when it exists. For these populations we support REST API-based connectors and data mart writeback capabilities that enable moving certified MicroStrategy data into departmental databases and BI tools, like SQL Server, Tableau, and Power BI.

Self-Service Analytics (SSA) and Data Visualization – Using MicroStrategy, all BMC employees can create their own reports and dashboards, which can be sourced from either the enterprise schema, bring-your-own-data sources, or a blended combination of the two. Each enterprise subject area is stewarded by a handful of business SMEs that serve as content owners. These analysts create pre-defined analytics content for the end-users in their areas, organize and secure the content of shared folders, and partner with IT to support their constituents' analytics needs. Several business teams have created data-driven mobile applications for their end-user populations to consume on-the-go.

Self-Service Data Preparation and User-Managed Databases (UMDBs) – Many of BMC's business users have advanced skillsets. The needs of such teams usually lean toward added control over data, whether for ad-hoc analysis or downstream reporting. Using Alteryx, these users typically query data from the EDW that they have been granted access to see, perform joins, cleansing, grouping and derived attribution, and capture data at different time intervals, including snapshots. This data can be consumed by a visualization tool of choice (primarily MicroStrategy and Tableau, and sometimes both). Most of the same teams that participate in the self-service data prep offering

need a place to store the data they have prepared. For this, we provide a user-managed database framework. Co-located on the high-speed Netezza platform alongside the EDW, UMDBs typically include custom views directly on EDW entities, and serve as a pre-viz landing zone for data.

These user-focused frameworks get IT out of the way as much as possible, allowing us to break down the usual barriers to accessing, onboarding, creating, deploying, and maintaining citizen-developed BI applications. Control-M is the key enabler through its easy-to-use self-service interface for creating workflows, checking their status, and automating many of the handoffs and tasks that need to happen for jobs to execute. For example, Control-M automates ETL and file transfer operations through the six different third-party tools BMC uses for moving data in and out of the EDW, executes event-based vs. scheduled workflow execution, and more. Control-M frees IT operations from having to handle all these tasks while also managing SOX compliance and internal data access policies.

Turning strategy into action, with automation

Here's an example of how data democratization helps us navigate the environment and makes BMC a more responsive business. Our finance department wanted to improve forecasting, and specifically to do a better job predicting what customers would do when their software licensing contracts would expire. Its idea was to create a customer health dashboard that presented relevant customer satisfaction scores and other metrics for each customer, then apply machine learning to predict whether the customer would renew.

The value to having this insight was clear, but the path to getting it was not. The dashboard and ML analysis would require access to the latest customer service and support data, historical contract information, and more. Once finance identified the metrics it wanted, Control-M automated the file transfer and other data extraction from multiple source systems, and managed the export to Tableau, which creates the dashboards. Analytics are used to calculate a customer health score, and a machine learning model predicts the likelihood that the customer will renew its contract.

Control-M also automatically sends updates and alerts to account managers when customer contracts are coming up for renewal. The finance team did much of this on their own, with Control-M enabling the self-service aspects.

"Before Control-M this was a nightmare," says Robert Hanley, BMC's director of finance. "We had to have an analyst dedicated to updating our machine learning models. Now the update process is zero percent of the analyst's job. The process is completely automated, and we can easily get an up-to-date customer health score in real-time."

In future blog posts we plan to cover that program in depth and share others too. Our business users have proven to be very creative and have developed many resources, including data marts, mobile apps, automated reporting, and more. In one case, all the data needed to run customer support operations, including the executive dashboards that guide decision making, is now refreshed and ready by 4 a.m. instead of 8 a.m. That improvement has improved our customer responsiveness while making it easier for hundreds of employees in the organization to focus on additional high-value projects. In another, our sales operations team used MicroStrategy, Tableau, Alteryx, multiple databases and other systems to create a workflow that combines previously hard-to-access data to improve customer support, simplify tasks for hundreds of business users, and give executives more strategic insight.

Our data democratization efforts never end but we are proud of our success so far, which was validated when we won TDWI's Best Practices Award for Emerging Technologies and Methods – an award that recognizes tech users, not tech vendors.

While our business users that benefit from data democratization are on a broad spectrum of technical skill and have many diverse use cases, there's a common thing they keep telling us: They are getting their weekends back thanks to Control-M. We will share some detailed use case examples in future posts. Until then, click here to learn more about Control-M capabilities and how other customers are using it to transform.