New Jersey Legislature

Office of LEGISLATIVE SERVICES

OFFICE OF THE STATE AUDITOR

Office of Information Technology
Enterprise Data Warehouse

October 2, 2017 to August 31, 2018

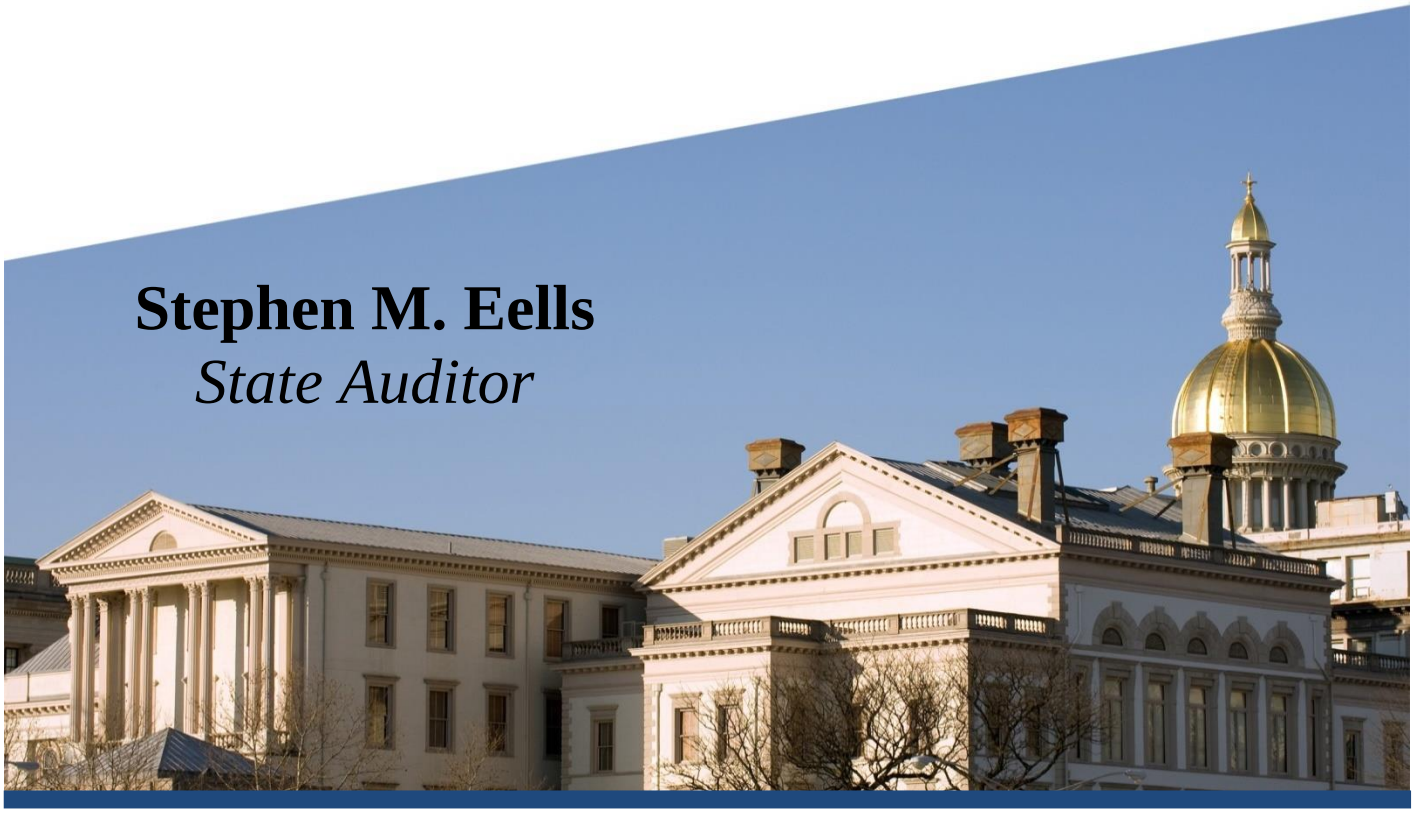

Stephen M. Eells
State Auditor

New Jersey State Legislature

OFFICE OF LEGISLATIVE SERVICES
OFFICE OF THE STATE AUDITOR
125 SOUTH WARREN STREET
PO BOX 067
TRENTON NJ 08625-0067

PERI A. HOROWITZ
Executive Director
(609) 847-3901

OFFICE OF THE STATE AUDITOR
(609) 847-3470
FAX (609) 633-0834

STEPHEN M. EELLS
State Auditor

DAVID J. KASCHAK
Assistant State Auditor

THOMAS TROUTMAN
Assistant State Auditor

The Honorable Philip D. Murphy
Governor of New Jersey

The Honorable Stephen M. Sweeney
President of the Senate

The Honorable Craig J. Coughlin
Speaker of the General Assembly

Ms. Peri A. Horowitz
Executive Director
Office of Legislative Services

Enclosed is our report on the audit of the Office of Information Technology, Enterprise Data Warehouse for the period of October 2, 2017 to August 31, 2018. If you would like a personal briefing, please call me at (609) 847-3470.

Stephen M. Eells
State Auditor
April 2, 2019

Table of Contents

Scope 1

Objective 1

Methodology 1

Conclusions 2

Background 2

Findings and Recommendations

Items Reported Under Separate Cover 4

Data Integrity 4

Auditee Response 10

Scope

We have completed an audit of the Office of Information Technology (OIT), Enterprise Data Warehouse (EDW) for the period of October 2, 2017 to August 31, 2018. The scope of the audit included the data integrity, security, and accessibility of data in the Office of Information Technology Enterprise Data Warehouse. Specifically, our audit focused on data from the New Jersey Comprehensive Financial System (NJCFS), Payroll, and Personnel Management Information System (PMIS), collectively known as the New Jersey Administrative Warehouse System (NJAWS); the electronic Cost Accounting and Timekeeping System (eCATS) static data; and the EDW administrative sequences, which include all NJAWS data validation reports. Excluded from this audit were the data warehousing activities of other agencies, boards, commissions, and authorities of the executive branch, as well as the judicial and legislative branches.

Objective

The objective of our audit was to determine whether adequate controls are in place to ensure the confidentiality, integrity, and reliability of the data contained in the areas of the OIT Enterprise Data Warehouse included in our scope.

This audit was conducted pursuant to the State Auditor's responsibilities as set forth in Article VII, Section I, Paragraph 6 of the State Constitution and Title 52 of the New Jersey Statutes.

Methodology

Our audit was conducted in accordance with *Government Auditing Standards*, issued by the Comptroller General of the United States. Those standards require that we plan and perform the audit to obtain sufficient, appropriate evidence to provide a reasonable basis for our findings and conclusions based on our audit objectives. We believe that the evidence obtained provides a reasonable basis for our findings and conclusions based on our audit objective. Additional guidance for the conduct of the audit was provided by the standards promulgated by the Center for Internet Security, as well as reviews of industry standards and best practices.

In preparation for our testing, we studied legislation; executive branch policies and procedures; and industry standards and best practices for data warehousing. Provisions we considered significant were documented, and compliance was verified by interviews of key personnel, observations, and access to the OIT EDW's information technology assets. Data integrity controls, as well as controls in place to protect the data from unauthorized or improper access, modification, and deletion, were tested as considered necessary.

A non-statistical sampling approach was used. Our samples were designed to provide conclusions on our audit objectives as well as internal controls and compliance. Sample items were judgmentally selected for testing.

Conclusions

We found that controls existed to protect the confidentiality, integrity, and reliability of the data in the areas of the OIT Enterprise Data Warehouse included in the scope of the audit. In the area of data integrity, our testing did not identify a specific instance where we determined that the data in the Office of Information Technology (OIT) EDW were not accurate, complete and timely; however, the results of our testing indicate there is a risk that an error affecting the accuracy, completeness, and timeliness of the data would not be detected before impacting the data available to end-users. We also noted areas for improvement that are being communicated only to management in a confidential management letter.

Background

The Oracle Corporation defines a data warehouse as "a database designed to enable business intelligence activities: it exists to help users understand and enhance their organization's performance." To achieve this business intelligence, a data warehouse brings together data from disparate sources, normalizes it, and structures it in a way that it can be used for comprehensive analysis and reporting. To perform this function, the data must be of high quality. This depends on the usage of the data, as well as the data itself. The most common attributes of data quality are accuracy, timeliness, relevance, completeness, understandability, comparability, and reliability. The Office of Information Technology (OIT) Enterprise Data Warehouse (EDW) has the responsibility for those attributes only as they apply to the data which is contained within the warehouse. Issues with data quality in the source systems are the responsibility of those systems' owners. The OIT EDW should reflect – completely, accurately, and timely – the data in the corresponding source system, and the OIT is responsible for the institution of adequate controls to ensure this, as well as ensure the data is protected from unauthorized or improper access, modification, or deletion.

To provide a basic understanding of the process used by the OIT to move data from a source system to data accessible for reporting and analysis in the EDW, an overview of the automated process is presented. First, data is extracted from the source system and sent to a secure file transfer server. It is then collected by DataStage, the extract, transform, and load (ETL) software utilized by the state. DataStage ETL jobs are grouped and ordered into sequences which start with the extracted source system file, transform and move the data through the staging area, and end with the data in the warehouse. This data is read-only and is not directly accessible by any end-user. These read-only views of the data are created and structured in ways that make it easier for end-users. End-users access the data through a Business Objects interface, which is business intelligence software that provides structure and context to the data so users can more easily understand and utilize it.

There are various controls necessary to ensure the integrity of the data during the process and protect it from unauthorized or improper access, modification, and deletion. In the ETL process, the OIT EDW staff asserted that the following controls are in place to ensure data integrity is maintained.

- A manual review of all DataStage sequences is performed daily to look for jobs and sequences that may not have run successfully, and to correct any issues.

- The DataStage sequences are programmed to send email notifications to appropriate staff when steps in the sequence succeed, fail, or run with warnings.

- DataStage contains a sequence that was developed by the OIT EDW staff and reports on all jobs that did not run or ran with warnings. This is run every Monday through Saturday and is sent to appropriate staff.

- The ETL Load report is generated Monday through Friday and extracts data from an EDW table that records the last load date of data in various EDW tables.

- DataStage runs a series of data validation queries which extract the last load dates and calculate control totals from selected EDW tables. This information is either reconciled internally within the report, or compared to other data sources to verify integrity. Review of the results of the queries and the subsequent comparisons are manual processes performed by OIT EDW staff.

- For some data, reports are run by the OIT EDW staff against the source systems, and the information is reconciled to the aforementioned data validation reports to ensure source-to-target integrity. These reports require manual review by OIT EDW staff.

- Data owners are ultimately responsible for ensuring that the data in the final EDW target destination is an accurate, complete, and current representation of the source system. Data owners achieve this by performing reconciliations between the EDW data and the various source systems.

Items Reported Under Separate Cover

Our audit disclosed reportable conditions deemed confidential in nature. These conditions were communicated in a confidential management letter provided to agency management only. The findings and recommendations contained in the management letter are subject to the Office of the State Auditor's compliance process as required by N.J.S.A. 52:24-4.

Data Integrity

Controls in the data integrity process designed to ensure accurate, complete, and current data in the Enterprise Data Warehouse need improvement.

Articles published by ISACA and the International Journal of Computer Science and Information Technology cite research that reinforces the critical role that data integrity plays in the creation and maintenance of a warehousing solution that is useful to an organization. If there is a lack of data integrity in the warehouse, decision makers receiving the reports and analysis cannot trust the results. For the warehouse to provide useful reporting and analysis, it must accurately and completely reflect the most recent data in the source system. In order to maintain data integrity throughout, each step of the extract, transfer, and load process should have data integrity checks in place to ensure that source system data is the same as the data in the final Enterprise Data Warehouse (EDW) destination.

We tested all of the data integrity controls discussed in the background section of this report, and found issues in multiple areas. Although our testing did not identify a specific instance where we determined that the data in the Office of Information Technology (OIT) EDW were not accurate, complete, and timely, the results of our testing indicate there is a risk that an error affecting the accuracy, completeness, and timeliness of the data would not be detected before impacting the data available to end-users. This conclusion is based on the following three areas of concern:

Lack of Regular and Complete Data Reconciliations

We tested 15 days of processing in the New Jersey Administrative Warehouse System (NJAWS) and the electronic Cost Accounting and Timekeeping System (eCATS) sequences, including manually reviewing sequence results, obtaining all relevant reports, and recreating data reconciliation procedures described to us by the OIT EDW staff. Based on this review, we identified 2,557 data integrity verifications (checks) that could have been performed during our review period. Of these, 1,265 (49 percent) are date matches between a combination of the last run date obtained from our manual review of DataStage sequence success, the ETL Load report last load date for EDW tables, the last load date from the applicable data validation report, and the report date from the applicable source system reports. These checks are designed to validate that the warehouse contains current data. Additional integrity checks include matching data totals between EDW data tables, as well as to data in the source systems. Our review found that 748

(29 percent) of all the possible data integrity checks either could not be performed because of a lack of necessary data, errors during the manual reconciliation process, or other reasons. The 748 affected data integrity checks are broken down as follows:

- 532 of the 748 affected integrity checks could not be performed because necessary information to complete the check was not available. This included 207 where the last load date of a table populated by a DataStage sequence could not be verified because the table was not listed on the ETL Load report; 121 where the last load date of a table populated by a DataStage sequence could not be compared to the ETL Load report because the ETL Load report was not run, either by design or failure; 194 where a date, count, or amount could not be reconciled because the required data validation report did not run when expected; and 10 where date, count, or amount checks could not be run because a required source system report did not run. Although discrepancies in the reconciliation of control totals were a minor issue in our testing, it should be noted that the total could be higher if the missing information that prevented the 532 checks from being run was available.

- 175 of the 748 affected integrity checks involved a discrepancy in the reconciliation between two pieces of information. This included 100 where the manual review last run date for the DataStage sequence and the ETL Load report last load date for the corresponding EDW table did not match; 59 where the ETL Load report last load date and data validation report last load date fields for a particular EDW table did not match; 14 where the manual review last run date did not match the data validation report last load date; and 2 where the ETL Load report last load date did not match the date of the source system report.

- The remaining 41 affected integrity checks included 3 where an internal reconciliation performed by a data validation query did not reconcile; 27 where the data validation report and the source system report control totals did not reconcile; and 11 where the DataStage sequence either did not run when scheduled (with no recovery), or the sequence ran when it was not scheduled to.

At least one of the data integrity check issues occurred in 13 of the 15 days tested. The only days with no affected checks were the two Sundays included in our testing period; however, it should be noted that only one DataStage sequence is scheduled to run on Sunday. The rest of the data integrity checks performed on Sunday only verify that the other DataStage sequences did not run.

Lack of Source System Reconciliations

Part of the data integrity process is a final reconciliation between the source system data and the corresponding data in its final destination in the EDW to ensure that the data warehouse completely and accurately reflects the current data in the source system. We reviewed the source system-to-EDW-final-destination reconciliation process for the NJAWS and eCATS data and found the following:

- For the NJCFS source system, there appears to be no source-to-final-destination reconciliation taking place. The OIT EDW staff runs reports against the NJAWS data daily

and sends them to the Department of the Treasury, Office of Management and Budget (OMB), who are the data owners. However, these reports only compare one NJAWS table against another (usually general ledger to subsidiary ledger), and replace the previous version of the report daily. No report history is kept unless there is a discrepancy. The OMB runs system assurance reports within the NJCFS to test data integrity within the source system, and monitors for the success or failure of the mainframe job which extracts and exports the initial raw data from the NJCFS to the EDW for processing. There is no direct reconciliation between the NJCFS source data and the NJAWS EDW destination data.

- For the Payroll source system, we identified four reports that are run by the OIT EDW staff against the Payroll source system, the results of which are manually compared to data validation reports produced by DataStage. As was discussed previously, there were days when one or more of these reports did not run, thereby preventing the reconciliation from being performed. The data owner does perform a weekly source-to-destination reconciliation on a single control total and communicates the results to the OIT EDW staff.

- For the PMIS source system, we identified two source system reports that are run by the OIT EDW staff and manually compared to EDW data validation reports. As with the other source system reports, there were days when one or both of these reports did not run, thereby preventing the reconciliation from being performed. Through discussions with the data owners, we learned that they do not reconcile source system data to EDW final destination data.

- For the electronic eCATS source system, we found from discussions with the eCATS staff that they perform all of their data validation checks on the live source system, which is also accessible through the EDW. The warehoused version of the data, which is used when the live source system is unavailable, is not validated against the live source system after the data load is completed. The only time the warehoused version would be addressed is when the data extract from the live source system used to populate the warehoused version is changed. At that point, the eCATS staff would look at the data extract file only to determine if the changes were successful. No reconciliation is done by the EDW staff on any eCATS data, live or warehoused.

It is the assertion of the OIT EDW staff that the responsibility for the source-to-final-destination reconciliation lies with the agency owning the data. However, OIT EDW staff are running reports and performing some source-to-final-destination reconciliations for the Payroll and PMIS systems.

Incorrect, Missing, or Undelivered DataStage Notifications

The OIT populates the EDW through the use of DataStage sequences that are programmed to notify relevant OIT EDW staff of processing results. In addition, DataStage is programmed to generate data validation and job status reports and send them to OIT EDW staff. We found the following issues with those notifications and reports:

- A sequence runs daily which checks if the previous day was a holiday. If so, it sets a trigger that other sequences may use to determine if source data will be present. Without source data to process, running the sequence will produce no change in the destination data. Our 15-day testing period included the day after a holiday, and we noted that 11 of the sequences that could utilize that trigger were not programmed to check for it. This should be mitigated by the fact that the source data file would not be received. Each of these sequences does have a notification programmed to let appropriate personnel know the source data file did not arrive; however, this notification was not received.

- There were eight instances during our test period where the success or failure notice in a sequence was sent only to a non-existent recipient. Seven of these instances were notifications of the success or failure of a single step in a sequence, and one was the success notification for the entire sequence. Although our manual review verified that all of the sequences did successfully run, sending notification emails only to non-existent users negates the purpose of the integrity control.

- Six sequences contained a step that ran with warnings on all of the days it was scheduled to run, but no notification was sent because the sequence is not programmed to do so. This is a practice also used for other sequences. Our manual review, however, verified that the six sequences did run successfully.

- Three sequences did not run on any of the days they were scheduled to because the source data file was not received. Although there is no expectation that the file will be received on any particular day, there is no notification sent concerning whether the file is received.

- There were two sequences that run Monday through Friday, but are only expected to have source data to process once a week. There is no success notification at the end of the sequence, a practice used for other sequences.

- One sequence runs on Wednesdays and Thursdays because source data is only expected to be present one of those two days each week. The sequence does not send a notification if the data does not arrive on those days, nor does it have a notification if the data is received and processed successfully, though notifications for similar events are sent for other sequences.

- DataStage has a sequence which is designed to report daily all jobs that failed to run, ran with warnings, or were still running at the time the report was produced. This alerts OIT EDW staff to potential issues with sequence processing. We found that these reports had multiple issues which rendered them ineffective as a tool for sequence monitoring. The sequence that generates the reports failed to run for three consecutive days during our testing period without being corrected and reset. In addition, our analysis of the information contained in the reports for the nine days they were successfully generated disclosed that the data was identical for all days, and no report contained data with a load date later than "3/20/2017", though we documented sequence failures and warnings during our testing period, which was after that date.

There seem to be multiple reasons why these issues exist. Although there are currently many individual components of data integrity checking being performed, there is no single comprehensive process for completing and documenting data integrity checking in the EDW. Some of the tools in place are not properly configured or utilized. The completion and results of the manual review of sequence success performed by the OIT EDW staff is not documented. The ETL Load report, which is used to determine the last load date of data into the EDW, does not contain all of the tables that are populated by the NJAWS and eCATS sequences, and is not run on all days when processing takes place. When it was run during our testing period, we found instances where our manual review of the last sequence run date did not match the ETL Load report last load date. Lastly, data validation queries exist for many tables in the OIT EDW, but not for all tables populated by the sequences.

With regard to checking the original source data to the final EDW warehouse data, the OIT EDW has not required data owners to perform, and provide evidence of, successful reconciliations. In the past, there were additional source system reports developed by the OIT EDW staff for additional source-to-target reconciliations in the NJAWS universe. The OIT EDW staff stated that these reconciliations were taken over by the data owners; however, we found that this had not occurred. The data owners do reconcile the source system data internally within their system for integrity purposes, and some of the NJAWS warehouse data is reconciled internally as well, but the two data sets are not reconciled to each other. In the eCATS, the warehoused data is a straight migration from the live source system, without any data transformation. This may have created the false impression that reconciliation is unnecessary.

Finally, although success, warning, and failure notifications were built into many DataStage sequences, they do not exist for all sequences. Also, the lack of report generation and incorrect data we found in the report of jobs that did not run, or ran with warnings, shows that this is not an effective tool, and the OIT EDW staff stated that these reports are not used.

The lack of effective data integrity controls in the data warehouse could negate its primary goal of providing a complete, accurate, and current set of data to be used for analysis and reporting that improves the organization. The various issues we noted with the process could contribute to errors in the extract, transfer, and load process, including data values that do not accurately reflect the contents of the source system, data missing from the EDW that is contained in the source system, duplicate or extraneous data loaded into the EDW, or the existence of non-current data. Any of these situations would cause errors in the reporting and analysis generated by users of the EDW, thereby rendering them inaccurate.

Recommendation

We recommend the OIT assess and improve the data integrity process to ensure that all data extracted, transformed, and loaded from source systems to the data warehouse is complete, accurate, and timely. This solution should, at a minimum, address the following:

- All tables populated by the NJAWS and eCATS sequences are reported to OIT EDW staff for verification of current load date on a daily basis.

- Information needed for reconciliations and data integrity checks is produced for processed data when necessary.

- Necessary control totals are being produced for tables in the EDW and validated to other relevant EDW tables, the source system, or both.

- Data integrity validation is being performed daily and the results are documented, including resolution of any discrepancies.

- Data owners utilizing the EDW are required to develop adequate source-to-target reconciliation procedures and communicate the results to the OIT EDW staff when they are performed.

- Necessary notifications and reports in the extract, transform, and load process are configured properly, and new ETL sequences are required to have proper notifications.

PHILIP D. MURPHY
Governor

SHEILA Y. OLIVER
Lt. Governor

Office of Information Technology
P.O. Box 212
Trenton, New Jersey 08625-0212

CHRISTOPHER J. REIN
Chief Technology Officer

March 29, 2019

Mr. Stephen M. Eells – New Jersey State Auditor
Office of Legislative Services
125 South Warren Street
Trenton, NJ 08625-0067

RE: Office of Information Technology's Data Warehousing Audit Report

Dear Mr. Eells,

Please accept this letter from the New Jersey Office of Information Technology ("NJOIT") in response to the report sent by the Office of Legislative Services ("OLS") on March 8, 2019, addressing the audit report for NJ Office of Information Technology, Enterprise Data Warehouse ("EDW") covering October 2, 2017 to August 21, 2018.

As Chief Technology Officer, I realize that assessments and audits from neutral third-parties are extremely valuable to help identify areas of strengths or those in need of improvement. From the depth of the auditors' questions, and the data source analyses they performed, it is clear that they are quite thorough and knowledgeable in their tradecraft. I value the information in this audit; and as provided below our Agency has, and will continue to, put controls in place to strengthen these areas.

Data Warehousing Background

The New Jersey Office of Information Technology, NJOIT, provides and maintains the information technology infrastructure of the Executive Branch of State Government, including ancillary departments and agencies, and coordinates and conducts all information technology infrastructure operations in the Executive Branch of State Government. The EDW is the result of multiple agency source data integrations, including (but not limited to):

- New Jersey Administration Warehouse Solution (NJAWS) – from New Jersey's Office of Management and Budget (OMB), Treasury (Payroll), and Civil Service Commission (CSC).
- New Jersey electronic Cost Accounting and Timesheet System (eCATS)
- New Jersey Enterprise Analysis System for Early Learning (NJ-EASEL) – Integration of source information from the Departments of Education, Health, Human Services, & Children and Families.
- New Jersey Department of Transportation Data Warehouse (NJDOT-TransInfo & EIS)

The Auditor found that controls do exist to protect the confidentiality, integrity and reliability of the information and data in the OIT EDW; and further, that no specific data was found to be inaccurate; However, OIT agrees with the Auditor's finding that process weaknesses exist and continual process improvement and quality controls are needed to ensure that the EDW contains complete, accurate and current data.

There are many challenges to planning, creating, and maintaining an EDW environment. As the report points out, this technology design exists to help users better understand and enhance New Jersey state agency performance.

By using sound practices, enterprise-scale tools, and incremental improvements, the NJ EDW is creating statewide efficiencies. It is providing reusable data, eliminating redundant interfaces and reduced maintenance support cost.

OLS Recommendations and OIT Acknowledgment

In its *Conclusions (p.2)*, the OLS audit team identified inconsistent data warehousing notifications, stating that its testing indicated *"a risk that an error affecting the accuracy, completeness and timeliness of the data would not be detected before impacting the data available to end users."* OIT recognizes this area as one of significant concern, and one that requires both tools (software) and process/policy changes to mitigate. The current data integrity validation (and exception) process is composed of three components: email notifications, automated alert generation, and manual job log checking. OIT has budgeted for and planned software upgrades that include features that enable improved monitoring and notification to address these concerns.

- **Data Integrity Recommendation: Improvement in the controls in the data integrity process designed to ensure accurate, complete and current data in the EDW, assessment and improvement of the data integrity process to ensure that all data extracted, transformed, and loaded from source systems to the data warehouse is complete, accurate and timely.** The audit team also recommends that notifications and reports in extract, transform, and load process be configured properly, and that new ETL sequences are required to have proper notifications.

  OIT Response: As we migrate to the latest version of InfoSphere DataStage toolset, OIT plans to leverage the Operations Console, a new feature within the IBM Information Server Suite, to monitor our jobs, job activity, system resources, and workload management queues for each of our InfoSphere Information Server engines. The Operations Console provides engine-wide information about job runs, system resources, workload management queues, and engine status. Timely and granular notifications will be sent to data owners at the source, and this information is also used to control or halt downstream processes. The operation console provides a complete and centralized view of all ETL processes that failed to run, making the need for notifications for each individual process obsolete.

- **Lack of Source System Reconciliations - Recommendation: Data owners utilizing the EDW should be required to develop adequate source-to-target reconciliation procedures and to communicate the results to the OIT EDW staff when they are performed.** The audit team noted that tables that are populated by the NJAWS and eCATS sequences are reported to OIT EDW staff for verification of current load data on a daily or agency-specified basis; information needed for reconciliations and data integrity checks is produced for processing data when necessary; necessary control totals are being produced for tables in the EDW and validated to other relevant EDW tables, the source system, or both.

  OIT Response: Currently, part of the interim solution includes OIT continuously adding more tables with record counts and refresh dates to our centralized report table, giving us quick insight to those tables that have become stale. This is a long running process, which requires some care to ensure it completes successfully. In some instances, we ensure that if there are several tables in a sequence of loads we report the first and last loads. This alerts us to the fact that if the last table is loaded we can assume they all were loaded and if the first table failed to load then none of the tables loaded. Today we have over 80% of the NJAW objects tracked in the centralized report table. We plan to add the remaining tables as well as address the gap in eCATS Data Mart tracking. OIT also has reached out to the Data Owners and Technology Units, such as Office of Management and Budget (CFS), Civil Service Commission (PMIS), and Treasury (Payroll) for additional information about the data sent in the form of flat files to allow us to complete data reconciliation and integrity validation.

- **Data Integrity Validation - Recommendation**: **Data integrity validation should be performed daily and the results should be documented, including resolution of any discrepancies.**

  OIT Response: Data integrity validation can be automated using data testing software that identifies bad data, tags it with any required resolutions, and provides a comprehensive view of the data's soundness. It is used for high-volume data testing during the development cycle by assisting in pinpointing issues with source data (aka data profiling). The automated process will be developed using an OIT-evaluated software package specifically designed for validation of data warehouse loads. We are presently early in the procurement cycle.

Actions Taken and Tasks Planned

Over the course of the audit, and since the audit period ended, OIT has made changes in policy, practice, and configuration - to improve the data integrity and verification processes by taking the following steps:

- Working with agency contacts, OIT replaced the manual system checks with an interim solution: we now review automated processes that raise exceptions to both developer and business user attention based on data quality and process metrics.
- We have obtained a short-term procurement of a valuated collaborative data testing software that identifies bad data and provides a comprehensive and interconnected view of the data's soundness. Both of these measures have improved OIT's ability to verify the integrity of data.
- OIT undertook an assessment of Quality Assurance software (as referenced above).

The following platform Upgrades have been performed:

- Upgrade OIT Enterprise Business Intelligence (EBI) server environment – completed January 2018
- Migration of OIT Enterprise Business Intelligence (EBI and EDW) production data to the latest version of DB-Hosting platform (Exadata) - completed in July 2018
- Upgrade OIT Enterprise Business Intelligence (EBI) production software toolset to the latest version – completed November 2018

The following projects are scheduled, regarding this EDW infrastructure:

- Procurement, Testing, & Configuration of process validation & dashboard software - 2Q-4Q 2019
- Upgrade of OIT EDW Enterprise server environment – planned for 2Q-3Q 2019
- Migration of OIT EDW Enterprise software toolset – planned for 3Q 2019 – 1Q 2020

Conclusion

The EDW continues to support projects, data requests, and analytics within an environment that is hosted in a secure manner. It provides a repository for both historical and integrated data, and over one hundred thousand business intelligence reports are run by State agency staff based on information within the EDW. It is an essential tool for our State's operations, analysis and other management purposes. To this end, we acknowledge the opportunity for process improvement and greater quality assurance with our data; OIT accepts these audit recommendations as valuable input to continuous improvement efforts which the citizens of New Jersey deserve.

Respectfully,

Christopher Rein
Chief Technology Officer
State of New Jersey