

**TECHNOLOGY AND INNOVATION
DEPARTMENT
BACKUP AND RECOVERY REVIEW
AUDIT 14-08
SEPTEMBER 23, 2014**



CITY OF TAMPA

Bob Buckhorn, Mayor

Internal Audit Department

Christine Glover, Internal Audit Director

September 23, 2014

Honorable Bob Buckhorn
Mayor, City of Tampa
1 City Hall Plaza
Tampa, Florida

RE: Backup and Recovery Review, Audit 14-08

Dear Mayor Buckhorn:

Attached is the Internal Audit Department's report on Technology and Innovation (T&I) Department's Backup and Recovery processes.

We thank the management and staff of T&I for their cooperation and assistance during this audit.

Sincerely,

/s/ Christine Glover

Christine Glover
Internal Audit Director

cc: Dennis Rogero, Chief of Staff
Sonya Little, Chief Financial Officer
Russell Hauptert, Chief Technology Officer
Eric Hayden, Infrastructure Services Manager

**TECHNOLOGY & INNOVATION DEPARTMENT
BACKUP AND RECOVERY REVIEW
AUDIT 14-08
SEPTEMBER 23, 2014**

/s/ Stephen Mhere

Auditor

/s/ Christine Glover

Audit Director

**TECHNOLOGY & INNOVATION DEPARTMENT
BACKUP AND RECOVERY REVIEW
AUDIT 14-08**

BACKGROUND

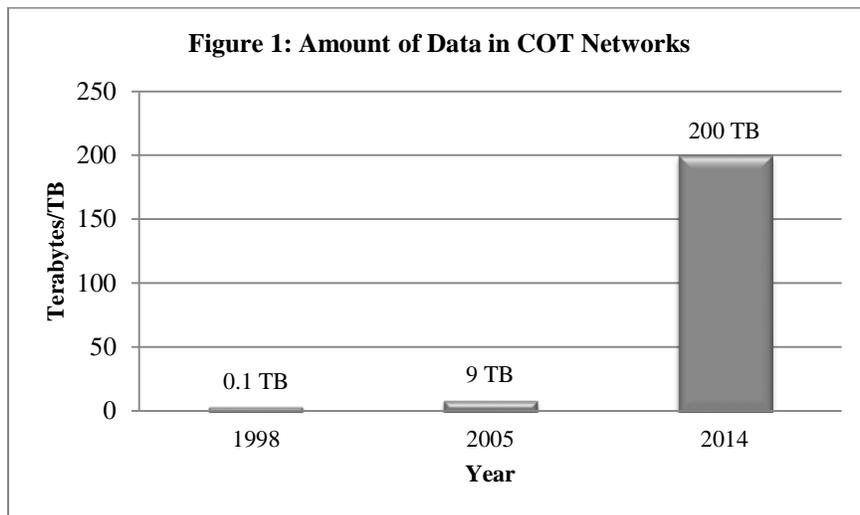
The City of Tampa (COT)'s most critical operations are based on electronic data generated, processed, and maintained in the various components of its computer networks. Service interruption due to operational data loss is one of the challenges the City faces on a daily basis. To mitigate the risk of service interruption, the City, through its Technology and Innovation (T&I) Department, has implemented a Data Backup and Recovery program. Data backup is a process of replicating and retaining copies of electronic data files, operating systems, and applications software that are used in daily operations. Data recovery or restore is a process of making backup copies available for use when the originals are inaccessible as might result from data corruption, system failure, or inadvertent deletions, etc. Data backup/restore is a critical component of the City's larger Emergency Disaster Recovery Plan.

T&I performs electronic backup and restore operations spanning 100 physical servers, 261 virtual machines, as well as the mainframe system. The City maintains a primary data center and a secondary facility that runs other critical City operations. The two sites act as reciprocal backup centers: data generated at the primary site is backed up through the network and stored in servers at the secondary facility and vice versa. The primary site also runs operations on a mainframe system, which is backed up to tape.

An independent contractor performs tape vaulting operations on behalf of the City. The contractor transports mainframe tapes from the data center to a vault located offsite. Other contractor responsibilities include daily rotation and courier services for tapes, as well as tape delivery round-the-clock whenever required, including during weekends and holidays. In Fiscal Year 2013 the City spent \$11,635 for these services.

STATISTICS

Figure 1 below shows the rapid growth in the amount of data in city networks since the late 1990s.



Source: City of Tampa's Technology and Innovation Department Internal Records

Potentially, all this data has to be backed up, thus requiring the City to implement innovative technology and cost strategies to effectively perform its functions.

STATEMENT OF OBJECTIVES

This audit was conducted in accordance with the Internal Audit Department's FY 2014 Audit Agenda. Its objectives were as follows:

1. To determine if retention policies T&I implemented in the City's backup/restore systems are consistent with business and/or compliance needs.
2. To evaluate the backup/restore testing performed and planned to be performed in FY 2013 and 2014 as well as the testing protocols used or planned to be used.

STATEMENT OF SCOPE

The audit covers fiscal years FY 2013 and FY 2014, and includes only T&I activities relevant to backup and recovery processes, including policy documentation that might have been originated prior to those years but were operational during that period. The audit does not include T&I's emergency disaster recovery planning, of which backup and restore are a part.

STATEMENT OF METHODOLOGY

The audit involved reviewing T&I's backup and recovery processes and documentation. Policies and practices were evaluated against the Information Technology Infrastructure Library (ITIL), which T&I has adopted as its operational standard, as well as Common Objectives for Information and Related Technology (COBIT 5) and industry best practices. Evaluation of performance was also measured against legal standards, particularly Florida statutes.

Information used in this audit was gathered from T&I employees and business owners through interviews and survey questionnaires. Fieldwork also included physical inspections of the primary data center, the secondary backup facility, and the tape vaulting center. Employees from the tape vaulting contractor were also interviewed and the tape transportation van inspected. Data backup and system or file restoration records were reviewed. Two tests were performed, one to evaluate the procedures to restore a file on the mainframe, and the other to assess the effectiveness to recover a Linux-based web server.

STATEMENT OF AUDITING STANDARDS

We conducted this performance audit in accordance with generally accepted government auditing standards. 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 objectives.

AUDIT CONCLUSIONS

Based upon the test work performed and the audit findings noted below, we conclude that:

1. Retention policies T&I implemented in COT's backup/restore systems are consistent with business and/or compliance needs.
2. Backup/restore processes do not encrypt data, there is no policy document currently in place

that defines a comprehensive restore testing protocol, and backup/restore processes do not include data integrity verification or backup checksum.

NOTEWORTHY ACCOMPLISHMENTS

T&I has adapted to a rapidly changing environment of rising demand for server, storage, and backup requirements. The department has implemented leading industry technology by virtualizing both the primary and secondary data centers. Virtualization of the server environment has seen virtual machine inventory grow from zero in 1998, to 20 in 2005, and then to 261 in 2014¹. Because of that, physical server inventory grew at a much slower rate, growing from 30 to 120 and down again to 100, respectively. Also, T&I has migrated about one third of backup data to disk (versus tape), which has improved efficiency in restore performance. Data restore activities have had a high success rate over the years. T&I has not had a tape or system problem that prevented data from being restored.

¹ Virtual servers are generally regarded to be better than traditional physical servers. Some of the benefits of virtual servers include lower energy costs, less data center floor space, easier system provisioning and deployment, increased uptime, improved disaster recovery, easier isolation of applications, etc.

While the findings discussed below may not, individually or in the aggregate, significantly impair T&I operations, they do present risks that can be more effectively controlled.

ENCRYPTION OF BACKED UP DATA

CONDITION: COT collects and maintains personal information in its networks. Operational data in City networks include employee and citizens' names along with their social security, driver's license, bank account and credit card numbers, as well as employee health information. To protect against the loss of data, T&I implemented a backup and recovery policy.

T&I backs up physical and virtual servers at the primary data center across networks to servers at COT's backup site and vice versa. Mainframe data volumes are backed up to tapes after which a private contractor transports them to an offsite storage vault. Neither disk-to-disk nor mainframe backup process includes encryption to protect backed up data.

CRITERIA: Section 5.2.3.1 of the ITIL version 3 provides that organizational data, including backup data, be protected. Also, Section 501.171 of the Florida Information Protection Act of 2014 requires organizations to secure personal information.

CAUSE: T&I does not believe backed up data encryption is critical because there are safeguards in place to protect data from unauthorized access. These include using file and device level access restrictions, physical security such as key-card entry into facilities housing backup tapes and disk drives, and using the Fast Dump Restore (FDR) program to back up tapes which makes them unreadable without that program. Also, the tape vaulting service contractor uses secure protocols to transport tapes to the vault.

EFFECT: All computer networks, as well as password-protected devices, are susceptible to hacker exploitation. As such, network-based disk-to-disk backup processes can be hacked, and tape-based backed up data can be stolen. The absence of encryption in both situations leaves data vulnerable to unauthorized access, corruption, or theft. Unauthorized access compromises the confidentiality of personal information in COT networks, corruption of data results in loss of its integrity, and stolen data may result in loss of its availability for backup/restore purposes.

Furthermore, if security of unencrypted personal data belonging to more than 500 people is compromised, COT will be required, under the breach notification clause of the Florida statute, to notify both the State and each individual whose information may have been breached. Also, if breach involves credit card information for more than 1,000 people, COT would be required to notify all consumer reporting agencies that compile and maintain files on consumers nationwide.

RECOMMENDATION 1: T&I should formulate and implement strategies to use encryption in both network-based backup procedures as well as mainframe backup tapes to provide data protection against hacking exploitation.

MANAGEMENT'S RESPONSE: Thank you for presenting your finding. Management agrees that if sensitive data on any electronic medium is stolen and accessed, it is possible for data to be exploited. At current, T&I does not own the necessary software encryption modules. To address this vulnerability specifically for our backup media, as stated in the finding, several safeguards were put into effect including requiring our vaulting vendor lock media containers before leaving our data center prior to transporting to their vault site. The same applied for returning media to our data centers.

T&I will purchase the additional software to encrypt the mainframe backup jobs that contain sensitive data. For the non-mainframe disk-based backups and tape backups of servers, T&I is in the process of changing to a new licensing option for our open-systems backup software. The new licensing option includes an encryption module allowing for real time data encryption during future backups. This is currently funded within our Server Virtualization project.

Both encryption products will be acquired, installed and placed into operation by January 1st, 2015.

TARGET IMPLEMENTATION DATE: January 1, 2015.

BACKUP/RESTORE TESTING

CONDITION: While T&I continually performs backup/restore operations for data owners, there is no policy document currently in place defining a comprehensive backup/restore testing protocol. Records show the only restore testing guidelines for FY 2014 are provided in the Continuity of Operations and Emergency Preparedness Disaster Recovery Plan (COOP-DRP). However, the COOP-DRP evaluation is very wide in scope and provides for conducting tests only twice per year.

T&I does not perform restore tests on a regular basis, and does not test all conceivable recovery scenarios. One of the COOP-DRP tests conducted in June 2014 was a disaster simulation involving the recovery of the Exchange email system. Records show this test was done, as were recoveries of two Windows servers, one Windows XP machine, and one Linux server. During the year, four database restore tests were also performed on the mainframe. However, there are no records of tests involving various other possible conditions in the restore environment.

Also, T&I's backup/restore processes do not include data integrity verification or backup checksum.

CRITERIA: According to ITIL, backup strategies should be properly tested. The standard provides that protocols should include test restores as well as data integrity and verification tests. COBIT 5 also emphasizes data verification, stating that procedures should be in place to ensure that usability of back-ups is regularly verified.

CAUSE: The audit period coincided with T&I transitioning its physical servers into a new virtual server environment, implementing a storage area network, and archiving Exchange email. This on-going project has delayed the documentation of a comprehensive backup/restore testing policy.

Between FY 2013 and FY 2014, T&I performed more than 20 actual restores on the mainframe, three on Linux servers, and about 15 per month on Windows servers, all at the request of data owners, in addition to the COOP-DRP restore activities. As a result, regular testing is not done because procedures for actual data recovery have proved successful.

According to T&I, checksum is not included in backup/recovery procedures because the CommVault software used does not have an option to configure or enable the tool. Also, T&I has not experienced problems with the reliability of the software, therefore lowering the perceived risk exposure resulting from the lack of data verification and checksum testing.

EFFECT: Without a formally documented policy, much needed consistency in restore testing is difficult to achieve. Documentation is also important for training and preservation of effective restore procedures.

Regular testing of various restore scenarios provides the best preparation for future responses to data losses. Testing is particularly important in an environment with many variables - different user applications software and hardware systems, disparate operating systems, different types of backup media, etc - each of which may require a unique data recovery approach. When all

possible restore scenarios are not regularly tested, an opportunity to gain information potentially beneficial for future recoveries is lost.

Sometimes backed up data may be unrestorable because it is not physically intact on destination media or is unreadable. Data may also be corrupt, making it valueless even if restored. These risks are significant without data verification and backup checksum, potentially undermining COT's ability to run effective data restore operations.

RECOMMENDATIONS 2: T&I should formulate and document a systematic policy that provides guidelines on the performance of restore testing. The policy should include a guideline of the frequency of restore tests, who performs them, as well as how restores requested by data owners are to be incorporated in the testing regime if they are to be used in lieu of actual testing. Rather than wait for the ongoing technology projects to complete before documenting such a policy, it may be better to write one that reflects how to accomplish restore tests in the current environment and, if necessary, update it at an appropriate time.

T&I should also design a testing protocol where all possible recovery scenarios are tested regularly as well as each time there is a change in the operational environment (e.g., change of user applications, operating systems, backup media, patch updates). All restore tests performed under the testing protocol should be completely documented and system logs preserved.

T&I should incorporate backup checksums in their backup/restore procedures where available.

MANAGEMENT'S RESPONSE: Thank you for presenting your finding. In response, we agree our existing recovery procedures for each server technology is included within our disaster recovery plan documentation. This plan document does not include a testing program for proving server recovery as a scheduled recurring activity.

To remedy this condition, T&I will immediately begin reviewing backup and recovery policies and testing procedures from organizations using similar technology to aide in developing a formal recovery policy and testing procedures. As a future ongoing activity, we will practice and prove testing recovery across "classifications" of servers versus individual servers in order to sweep through all essential servers every 24 months. Our goal is to have a formal recovery testing program in place by January 1st, 2015.

We also recognize the value in data validation algorithms. Unfortunately, checksum data validation is not a feature we can turn on in all of our hardware. T&I will research our solutions to determine if there are equivalent technologies available in our hardware. Data validation is essential, most of our replication hardware has been engineered to use some form of this capability to provide exact data duplication.

TARGET IMPLEMENTATION DATE: January 1, 2015.

PROXIMITY OF BACKUP LOCATION TO PRIMARY DATA CENTER

CONDITION: T&I has adopted ITIL as its information technology operations standard. ITIL and the City's Records Management Manual (which addresses off-site storage of archival records and backed up data) provide that backup facilities be in remote locations or far enough away so that they are not affected by a disaster that affects the primary site. Although these standards are not explicit in terms of the mileage separation between the backup and primary sites, industry professionals advise that data centers in the Gulf Coast should locate backup sites at least 100 miles away to mitigate the regional threat of hurricanes.

COT's backup/recovery site and the tape vaulting location are only four miles and three miles away from the primary data center, respectively. Also, the tape storage site is less than half a mile from McKay Bay and only about 45 feet above sea level.

CRITERIA: ITIL standards and COT's Records Management Manual, as described above, provided criteria for this evaluation. Opinion from information technology industry professionals was used to provide expert interpretation of the standards.

CAUSE: The Communications Center was chosen as the backup site because T&I believes that a catastrophic event in the immediate area of the primary location would not adversely affect operations at the Center. This opinion is based on (i) T&I's belief that the distance from the data center to the site is significant and (ii) the fact that the site is rated to withstand Category 5 hurricane wind speeds. Also, T&I does not consider the vault vulnerable because it is built to withstand Category 5 hurricanes.

EFFECT: The close proximity of the data center to the backup site and tape vaulting location means that a regional threat like a hurricane may affect all three locations, rendering backup and recovery operations ineffective or impossible. Also, the proximity of the vault to McKay Bay means that even if it can withstand Category 5 hurricanes, tapes may be inaccessible when needed if the area is flooded.

RECOMMENDATION 3: The City should consider (i) moving the backup site to a location far enough away from Tampa so that a widespread threat that affects the data center may not also affect the backup site, and (ii) moving the vault further away from the primary location and to a higher elevation so that backup/restore efforts are not hindered by the same risk elements and by limited access to tapes that may be caused by flooding.

MANAGEMENT'S RESPONSE: Thank you for presenting your finding. T&I management agrees it is beneficial having remote off-site vaulting storage facilities for protecting recovery media.

We are under a city contract to the vendor for media retention services including electronic tape vaulting. In a recent discussion with our vendor, they report having another more secure facility 40 miles from our data center, over 140 feet above sea level (highest land elevation in central Florida) and rated to withstand category 5 winds. This would certainly be a desirable facility.

We will alter our approach to incorporate both locations to rotate off-site media storage. By

having off-site media rotated between two hardened and secure facilities more than 30 miles apart and 40 miles from our data center, this option appears to mitigate any conceivable risk compared to using a single secure storage facility.

To put this into practice, T&I prepares daily backup tape rotation reports for the vendor and will need to amend our backup rotation schedules to work with two vaults instead of one. The scheduling work will be completed by October 1st, 2014.

TARGET IMPLEMENTATION DATE: October 1, 2014.