



*The Return Review Program Enhances the
Identification of Fraud; However, System
Security Needs Improvement*

July 2, 2015

Reference Number: 2015-20-060

This report has cleared the Treasury Inspector General for Tax Administration disclosure review process and information determined to be restricted from public release has been redacted from this document.

Phone Number / 202-622-6500

E-mail Address / TIGTACommunications@tigta.treas.gov

Website / <http://www.treasury.gov/tigta>



HIGHLIGHTS

THE RETURN REVIEW PROGRAM ENHANCES THE IDENTIFICATION OF FRAUD; HOWEVER, SYSTEM SECURITY NEEDS IMPROVEMENT

Highlights

Final Report issued on July 2, 2015

Highlights of Reference Number: 2015-20-060 to the Internal Revenue Service Chief Technology Officer.

IMPACT ON TAXPAYERS

During Fiscal Year 2013, there were almost 146 million individual income tax returns filed. Individual income tax withholding and tax payments totaled more than \$1.5 trillion, and almost \$312.8 billion in refunds were issued. Undetected tax refund fraud, including identity theft, has a significant impact on tax administration. It has the potential to erode taxpayer confidence in our Nation's tax system and results in significant unintended Federal expenditures.

WHY TIGTA DID THE AUDIT

Tax fraud is a major challenge for the IRS. In February 2009, the IRS chartered the initiation of a new program called the Return Review Program (RRP). The IRS plans to replace the Electronic Fraud Detection System with the RRP. Development of the RRP entered a strategic pause in January 2014 to allow the IRS time to evaluate the performance and design of the parallel processing database and to revisit strategic business fraud detection goals. Our overall objective was to determine if the RRP effectively meets requirements and identifies fraudulent tax returns.

WHAT TIGTA FOUND

The RRP models flagged potential identity theft fraud not detected by the Electronic Fraud Detection System models. The IRS initiated a pilot of the RRP Identity Theft Model. Processing only 32 days (one day per week for 32 weeks) over the duration of the pilot, the RRP identified 51,946 returns as potential identity

theft cases. The IRS confirmed that 41,311 of the 51,946 returns were identity theft. Of the confirmed identity theft cases, the IRS determined that 10,348 cases (25 percent) totaling \$43 million in refunds were not detected by the Electronic Fraud Detection System or the Dependent Database. In addition, IRS tests showed that eight million returns a day can be loaded to the RRP database as required. For example, over a one-week period, the RRP consistently loaded between seven million and nine million returns a day.

However, the IRS classified the RRP as a Level 3 system (an information resource instead of a major system). By classifying the RRP as a Level 3 Federal Information Security Management Act system, RRP-specific security issues may not be effectively addressed. In addition, identified security vulnerabilities were not remediated. For example, the October 2014 network scans identified two RRP servers that were still vulnerable to the Heartbleed bug six months after the vulnerability was announced.

WHAT TIGTA RECOMMENDED

TIGTA recommended that the Chief Technology Officer: 1) ensure that IRS personnel completing the Federal Information Security Management Act system classifications are familiar with the Act's requirements; 2) ensure that the validation of system classification and reclassification is discussed, reviewed, and documented during the biweekly Cybersecurity management meeting; and 3) ensure that all critical and high-risk RRP vulnerabilities are resolved.

In their response to the report, IRS officials agreed with all three recommendations. The IRS plans to brief personnel on the Federal Information Security Management Act requirements for each level of classification; enhance its current process for the validation of system classification and reclassification as discussed, reviewed, and documented during the biweekly Cybersecurity management meeting; and focus on resolving the critical vulnerabilities in production and then the lower environments.



TREASURY INSPECTOR GENERAL
FOR TAX ADMINISTRATION

DEPARTMENT OF THE TREASURY
WASHINGTON, D.C. 20220

July 2, 2015

MEMORANDUM FOR CHIEF TECHNOLOGY OFFICER

FROM: Michael E. McKenney
Deputy Inspector General for Audit

SUBJECT: Final Audit Report – The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement (Audit # 201420017)

This report presents the results of our review to determine if the Return Review Program effectively meets requirements and identifies fraudulent tax returns. This review is included in the Treasury Inspector General for Tax Administration's Fiscal Year 2015 Annual Audit Plan and addresses the major management challenges of Modernization, Tax Compliance Initiatives, Fraudulent Claims and Improper Payments, and Achieving Program Efficiencies and Cost Savings.

Management's complete response to the draft report is included as Appendix V.

Copies of this report are also being sent to the IRS managers affected by the report recommendations. If you have any questions, please contact me or Alan R. Duncan, Assistant Inspector General for Audit (Security and Information Technology Services).



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

Table of Contents

Background	Page 1
Results of Review	Page 3
Return Review Program Models Identified Additional Fraud Not Detected by the Electronic Fraud Detection System Models.....	Page 4
Tests Showed That Eight Million Returns per Day Can Be Loaded to the Return Review Program Database As Required.....	Page 6
The Return Review Program Was Incorrectly Classified As a Level 3 System	Page 6
<u>Recommendation 1:</u>	Page 7
<u>Recommendation 2:</u>	Page 8
Identified Security Vulnerabilities Are Not Remediated.....	Page 8
<u>Recommendation 3:</u>	Page 9
Appendices	
Appendix I – Detailed Objective, Scope, and Methodology	Page 10
Appendix II – Major Contributors to This Report	Page 11
Appendix III – Report Distribution List	Page 12
Appendix IV – Cross-Industry Standard Process for Data Mining	Page 13
Appendix V – Management’s Response to the Draft Report	Page 14



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

Abbreviations

CRISP-DM	Cross-Industry Standard Process for Data Mining
EFDS	Electronic Fraud Detection System
FISMA	Federal Information Security Management Act
IRS	Internal Revenue Service
RRP	Return Review Program
VRIP	Vulnerability Remediation Implementation Process



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

Background

The Internal Revenue Service (IRS) implemented the Electronic Fraud Detection System (EFDS) in 1994 to identify questionable and/or potentially fraudulent tax returns.¹ The IRS determined that numerous inefficiencies and operational challenges render the EFDS too risky to maintain, upgrade, or operate long term. The IRS reports that the long-term limitations of the EFDS include its inability to keep pace with increasing levels of fraud or to serve the organization's evolving compliance needs.² The IRS plans to replace the EFDS with the Return Review Program (RRP), an automated system that will enhance the IRS's capabilities to prevent, detect, and resolve criminal and civil tax noncompliance.

The IRS determined that numerous inefficiencies and operational challenges render the EFDS too risky to maintain, upgrade, or operate long term. The IRS plans to replace the EFDS with the RRP.

In October 2006, the IRS established the Pre-Refund Program Office to develop an enterprise vision and strategy for IRS prerefund activities. The Pre-Refund Program Office gathered business requirements and developed fraud patterns and scenarios. These business requirements, fraud patterns, and scenarios were modeled, and the resulting selection criteria were deployed through the EFDS and the Dependent Database.³ Once deployed, the selection criteria were frequently updated. For example, the Pre-Refund Program looked for characteristics of identity theft returns that could be encoded in "rules" – software codes written inside the system – that flagged returns suspected of being identity theft fraud, which began the process of requesting further authentication from the filer before disbursing the refund.

In February 2009, the IRS Commissioner approved a program charter authorizing formation of the RRP Office under joint leadership provided by the Wage and Investment Division and Criminal Investigation. The Wage and Investment Division is responsible for RRP requirements development, risk management, governance, project management, and deployment support. Criminal Investigation is responsible for supporting the RRP by identifying and developing schemes to refer and support high-impact criminal tax and related financial investigations.

¹ Treasury Inspector General for Tax Administration, Ref. No. 2013-40-083, *Income and Withholding Verification Processes Are Resulting in the Issuance of Potentially Fraudulent Tax Refunds* (Aug. 2013).

² Treasury Inspector General for Tax Administration, Ref. No. 2013-20-063, *Improvements Are Needed to Ensure Successful Development and System Integration for the Return Review Program* (Jul. 2013).

³ The Dependent Database is a risk-based audit selection tool used by the IRS to identify tax returns for audit. The Dependent Database is made up of a collection of information databases that include birth certificate information and court documents used to establish a relationship and residency between the taxpayer and the qualifying children claimed on the tax return.



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

A successful RRP system is critical to the IRS mission because it will be the key automated component of the IRS's prerefund initiative. The RRP system will implement the IRS's new business model for a coordinated criminal and civil tax noncompliance approach to prevent, detect, and resolve prerefund tax fraud. Based on fraud detected by the EFDS and supplemented by manual detection methods, the IRS estimates that prerefund tax fraud is more than \$19.2 billion per fiscal year. During Fiscal Year 2013, there were almost 146 million individual income tax returns filed. Individual income tax withholding and tax payments totaled more than \$1.5 trillion and almost \$312.8 billion in refunds were issued. Undetected tax refund fraud, including identity theft, has a significant impact on tax administration, has the potential to erode taxpayer confidence in our Nation's tax system, and results in significant unintended Federal expenditures.

The RRP project is following the Waterfall system development methodology⁴ with the implementation of RRP functionality via four Transition States. Transition State 1 includes the release of a new relational database and supporting business intelligence tools. However, RRP development entered a strategic pause in January 2014 to allow the IRS time to evaluate the performance and design of the Transition State 1 parallel processing database and to revisit strategic business fraud detection goals. To exit the strategic pause, the IRS developed a restart plan that was approved by the Executive Steering Committee in January 2015. In addition, the Change Control Board approved the RRP Transition State 1 Milestone 5 exit for March 2015. Currently, the RRP project team is developing a release strategy that will define the long-term plan for delivering new system functionality in subsequent transition states.

This review was performed at the RRP project office in New Carrollton, Maryland, during the period August 2014 through March 2015. 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 objective. We believe that the evidence obtained provides a reasonable basis for our findings and conclusions based on our audit objective. Detailed information on our audit objective, scope, and methodology is presented in Appendix I. Major contributors to the report are listed in Appendix II.

⁴ The Waterfall model describes a development method that is linear and sequential. Once a phase of development is completed, the development proceeds to the next phase, and there is no turning back.



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

Results of Review

The IRS uses the Cross-Industry Standard Process for Data Mining (CRISP-DM) methodology to develop each model in the RRP system. The CRISP-DM is a cyclical process that entails running the models, presenting the business with projected case volumes based on the initial parameters, receiving feedback from the business on changes to improve model performance, implementing business rule modifications based on business feedback, and rerunning the models. The iterative nature of the CRISP-DM, along with the flexibility to change RRP business rules, enables the IRS to refine the accuracy of known fraud models as well as identify new fraud schemes during the filing season. Appendix IV provides an overview of the CRISP-DM methodology.

The RRP system is comprised of three major components:

1. **Detection** – This part of the system incorporates several existing models as well as new models. By using algorithms and business rules, the system detects errors on the tax return as the return is filed and routes the return to the correct treatment stream, thereby allowing the taxpayer to receive one notice with all the issues that must be resolved before the refund is released. The system also detects returns with potential fraud characteristics and routes those returns to the treatment stream, which allows Criminal Investigation to associate/link and analyze groups of returns to identify schemes for potential criminal prosecution.
2. **Resolution** – This part of the system contains existing treatment streams as well as new treatment streams. Returns are routed systemically to a treatment stream and opened into that treatment stream's inventory. In addition, initial contact letters are sent to the taxpayer.
3. **Prevention** – This part of the system allows the results of the resolution to be sent and updated into the detection models systemically. Both outreach and education inventory can be selected through the system to allow for early intervention to stop the noncompliance before the next filing season. It also allows for the analysis of additional fraud not identified by the detection models.

The RRP met the requirements of Transition State 1. Specifically, the RRP deployed a relational database and met the business requirements. The IRS planned to deliver a workflow management system (*i.e.*, case management) after Transition State 1. Without a workflow management system, the resolution and prevention components of the RRP cannot deliver a fully integrated and unified cross-functional system that will enhance IRS capabilities to detect, resolve, and prevent criminal and civil tax noncompliance.



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

Return Review Program Models Identified Additional Fraud Not Detected by the Electronic Fraud Detection System Models

The RRP Project Charter states, “The key objective of the RRP is to deliver an integrated and unified system that is cross-functional and will enhance IRS capabilities to detect, resolve, and prevent criminal and civil tax noncompliance.” Congress enacted the Improper Payments Elimination and Recovery Act of 2010⁵ into law in July 2010 with a goal to reduce wasteful, improper payments by \$50 billion.

The IRS implemented a controlled launch of RRP Transition State 1 in March 2014 to identify how many more returns the RRP would have selected into inventory versus the EFDS. During the controlled launch, the RRP models ran parallel with the EFDS in the production environment. The RRP tolerances and thresholds were also set similar to the EFDS. However, the EFDS remained the system of record for working suspected fraudulent returns.

Through July 2014, the RRP identified approximately one million potentially fraudulent returns. Almost 350,000 of those potentially fraudulent returns were not detected by the EFDS. Figure 1 provides a breakdown of the confirmed fraudulent tax returns identified.

Figure 1: RRP Confirmed Fraudulent Tax Returns Identified (March 2014 – July 2014)

Tax Fraud Identified by System	Number of Confirmed Fraudulent Tax Returns	Refund Amount (in Millions)
Detected by RRP Models and Detected by EFDS Models	668,470	\$9,154
Detected by RRP Models; Not Detected by EFDS Models	220,508	\$1,001
Detected by RRP Linked Return Analysis; Not Detected by EFDS Models	128,490	\$470

Source: IRS RRP Predictive Analytics Performance Report Detection Summary.

One reason the RRP detected more fraud is that the EFDS focuses on income, withholding, and prior year fraud examples whereas the RRP uses data from a broader number of sources. Using the analytics capability in the RRP, the IRS can create predictive fraud and noncompliance detection models that will seek out subtle data patterns to determine reliability of return data, including the filer’s identity. The RRP generates a scorecard for questionable returns, evaluating consistency and dependability.

⁵ Pub. L. No. 111-204, 124 Stat. 2224.



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

In contrast, the EFDS system processing minimally uses predictive analytics. Cross-functional collaboration is difficult, and returns with multiple issues are often partially worked. With the RRP, the IRS is able to respond to multiple issues of noncompliance on a single return. The RRP incorporates linked return analysis, a tool that reveals patterns and relationships in masses of return data. Linked return analysis allows the RRP to identify clusters of returns that share characteristics indicative of schemes and other tax fraud or noncompliance.

For the 2014 Filing Season, the EFDS employed 15 models. In comparison, the RRP enables the IRS to employ 34 models in production. Additionally, the RRP generates 15 scores for each return to identify potential fraud, whereas the EFDS generates only one score per return.

Specifically, in reviewing the results of potential identity theft, we determined that the RRP Identity Theft Model was more effective than the EFDS model. Approximately one month after the RRP controlled launch, the IRS initiated a pilot of the RRP Identity Theft Model. The pilot took information from the RRP once a week, from April through November 2014, and fed it into the EFDS reporting database in an effort to identify additional potential fraud previously undetected by the EFDS and the Dependent Database. The IRS also worked the new RRP potential identity theft returns along with the potential identity theft returns identified by the EFDS and the Dependent Database.

Figure 2 illustrates the number of identity theft cases identified by the RRP versus the EFDS and the Dependent Database.

Figure 2: Returns Identified As Identity Theft by the RRP Versus the EFDS and the Dependent Database

Identity Theft Identified by System	Number of Confirmed Identity Theft Tax Returns
Detected by RRP Identity Theft Models and by the EFDS and the Dependent Database Identity Theft Models	16,321
Detected by RRP Identity Theft Models. The EFDS Models Identified As Suspicious but Not Identity Theft	14,642
Detected by RRP Identity Theft Models and Not by the EFDS or the Dependent Database Identity Theft Models	10,348
Total Number of Confirmed Identity Theft Returns From the Pilot	41,311

Source: IRS Identity Theft Selections for Taxpayer Protection Program Processing as of November 23, 2014.

Processing just 32 days (one day per week for 32 weeks) over the duration of the pilot, the RRP identified 51,946 returns as potential identity theft. The IRS confirmed that 41,311 of those returns were identity theft, of which 10,348 (25 percent) were not detected by either the EFDS or the Dependent Database. The remaining 10,635 of the 51,946 potential identity theft returns



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

were determined to be either not identity theft or were still being evaluated. The IRS determined that the incremental revenue protected by the RRP Identity Theft Model totaled \$43 million. Based on the success of the 2014 RRP Identity Theft pilot, the IRS received approval for the 2015 Filing Season to expand the RRP Identity Theft pilot to run daily instead of weekly. In addition, the IRS will use the RRP Identity Theft Model, as opposed to the EFDS Identity Theft Model, as the basis for generating all identity theft leads.

Tests Showed That Eight Million Returns per Day Can Be Loaded to the Return Review Program Database As Required

The Business System Architecture Report documents the solution architecture for RRP Transition State 1 in support of the IRS's Compliance Domain, Enforcement Division. The Business System Architecture Report specifies the key driving capability statements, functional requirements, and non-functional requirements that affect the RRP's architecture scope. Among the key driving requirements, the system shall have the capability to load eight million individual tax returns for a peak day in the first processing year.

For the 2014 Filing Season, the IRS recorded only one day when the individual tax return volume was sufficient to test the RRP's ability to load eight million returns in a day. To simulate additional days of peak volumes, the IRS processed the catch-up data, *i.e.*, the individual tax returns from the start of the 2014 Filing Season. Over a one-week period, the test showed that the RRP consistently loaded between seven million and nine million "catch-up" returns a day to the RRP database. The IRS is also planning to conduct additional performance testing after the start of the 2015 Filing Season.

The Return Review Program Was Incorrectly Classified As a Level 3 System

To ensure compliance with Federal information technology security standards, the IRS developed standard operating procedures⁶ that describe the processes and governance used by the IRS. Authorizing officials and the security administration team use the document for assistance with system classification and the management of information systems and resources. This document also helps define the IRS's Federal Information Security Management Act of 2002⁷ (FISMA) reportable inventory of information systems (General Support Systems and Major Applications), collectively known as "Major Information Systems," as well as Information Resources as defined by the Office of Management and Budget.

⁶ Federal Information Security Management Act Master Inventory Standard Operating Procedures (Version 2.0, dated June 4, 2014).

⁷ Pub. L. No. 107-347, Title III, 116 Stat. 2899, 2946-2961 (2002) (codified as amended in 44 U.S.C. §§ 3541-3549).



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

Major information systems are defined as those having any of the following characteristics:

- 1) Any system with a Federal Information Processing Standards 199 categorization level of High.
- 2) Any system called out in a major Security Capital Planning and Investment Control investment.
- 3) Any system that is comprised of (or contains) a Treasury-designated Critical Infrastructure asset.
- 4) Any system defined as “major” by bureau policy, bureau management, or the business owner due to the level of attention required for security.

The IRS identified several conditions during its security assessment of the RRP that culminated in the issuance of a conditional Authorization to Operate in December 2013. The IRS informed us that the Conditional Authorization to Operate was no longer valid because the RRP had been reclassified from a Level 1 (major) system to a Level 3 (non-FISMA information resource) system in March 2014. We noted that the questionnaire used for the reclassification stated that the RRP was only a batch program and therefore should be classified as an information resource. However, our analysis of the document identified several questions that were answered incorrectly that culminated in the lower classification level. For example, the first error, and the building block for the incorrect classification of the RRP, occurred when the IRS stated that the RRP system is not listed on a Capital Planning and Investment Control major investment report. We determined that the RRP system was listed as a major system within Exhibit 53A – Agency Information Technology Investment Portfolio 2015.

The IRS agreed with our assessment that the RRP should be a Level 1 FISMA system. In December 2014, the IRS changed the classification and issued a new Authorization to Operate to reflect the RRP as a Level 1 major system.

The security posture of major systems subject to FISMA reporting receives a higher level of scrutiny. By classifying the RRP as a Level 3 FISMA system, the RRP would be tested and reported as part of a larger General Support System instead of being tested and reported separately. This would increase the risk that RRP-specific security issues would not be effectively addressed.

Recommendations

Recommendation 1: The Chief Technology Officer should ensure that IRS personnel completing FISMA system classifications are familiar with the FISMA requirements for each level of classification.

Management’s Response: The IRS agreed with this recommendation. Throughout the 2016 FISMA reporting period, Information Technology Cybersecurity will brief IRS



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

personnel completing the FISMA system classifications to ensure that they understand the FISMA requirements for each level of classification.

Recommendation 2: The Chief Technology Officer should ensure that the validation of system classification and reclassification is discussed, reviewed, and documented during the biweekly Cybersecurity management meeting.

Management's Response: The IRS agreed with this recommendation. Information Technology Cybersecurity will enhance its current process for the validation of system classification and reclassification as discussed, reviewed, and documented during the biweekly Cybersecurity management meeting.

Identified Security Vulnerabilities Are Not Remediated

The Internal Revenue Manual requires the IRS to reduce its exposure to new vulnerabilities by requiring security patches to be installed within a set amount of time based on the criticality of the vulnerability. Specifically, it states that the IRS should implement patches for critical vulnerabilities within 72 hours, while patches for high vulnerabilities should be implemented within five business days.⁸

The IRS is running three different vulnerability scans. These scans are: 1) noncredentialed network scans of all systems on the internal network, 2) credentialed policy compliance scans of all servers and network devices, and 3) database vulnerability scans. All three scans found issues on the RRP servers currently running. The IRS had not applied critical patches, within the required time frames, to servers and databases supporting the RRP system. For example, the most recent scans found that:

- 1) The October 2014 network scans identified two RRP servers that were still vulnerable to the Heartbleed bug six months after the vulnerability was announced.
- 2) The policy compliance scans rated four of 131 systems at less than 80 percent compliant with required security settings; an additional 12 servers failed a high-risk check.
- 3) The database scans identified 322 critical failed tests.

We informed the IRS in November 2014 and December 2014 of the vulnerabilities noted during our analysis of the IRS scan results. The IRS stated that it remediated the Heartbleed bug vulnerability from the servers by removing a software package from those servers and renewing the certificates. The IRS also provided a network scan from December 8, 2014, for those four servers that shows the Heartbleed bug has been fully remediated.

The Government Accountability Office reported a similar issue in April 2014 and stated, “the IRS had not applied critical patches within required time frames to servers supporting multiple

⁸ Internal Revenue Manual 10.8.50, *Information Technology Security, Servicewide Security Patch Management*.



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

systems we reviewed, including the authorization, procurement, and e-mail systems. By not installing critical patches in a timely fashion, IRS increases the risk that known vulnerabilities in its systems may be exploited.”⁹

In addition, in April 2014, the IRS implemented a pilot, the Vulnerability Remediation Implementation Process (VRIP), involving 12 systems (not including the RRP) to manage identified vulnerabilities. The VRIP is a framework to report, assign, and facilitate remediation of database vulnerabilities detected by the compliance and vulnerability scanning tools. The VRIP is a vulnerability-centric (versus a server-centric or tool-centric) approach to database vulnerability remediation. This approach not only minimizes manpower requirements but also standardizes and expedites the remediation of security weaknesses. However, during the pilot, the IRS is working the vulnerabilities on only the 12 systems in the pilot, and the RRP is not one of those systems. This pilot is not complete; therefore, we could not validate the effectiveness of the VRIP pilot.

Patch management is an important element in mitigating the risks associated with known vulnerabilities. When vulnerabilities are discovered, the vendor may release an update to mitigate the risk. If the software update is not applied in a timely manner, an attacker may exploit a vulnerability not yet mitigated, enabling unauthorized access to an information system. Untimely patch management may also enable users to have access to greater privileges than authorized. The IRS is leaving itself open to potential security breaches from a lack of action to resolve known security issues.

Recommendation

Recommendation 3: The Chief Technology Officer should ensure that all critical and high-risk RRP vulnerabilities are resolved regardless of whether the system is part of the VRIP pilot. Once all of the highest risk issues are resolved, the IRS should work on resolving the remaining issues.

Management’s Response: The IRS agreed with this recommendation. The IRS continues to improve its remediation process to ensure that vulnerabilities are resolved. To date, the IRS has remediated all but 52 of the previously identified 322 vulnerabilities. Based on continued monitoring, the IRS will focus on resolving critical vulnerabilities in production and then the lower environments. The RRP will be added to the VRIP pilot, which will provide the IRS with the framework to report, assign, and facilitate remediation of database vulnerabilities detected by the compliance and vulnerability scanning tools.

⁹ Government Accountability Office, GAO-14-405, *IRS Needs to Address Control Weaknesses That Place Financial and Taxpayer Data at Risk* p.13 (Apr. 2014).



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

Appendix I

Detailed Objective, Scope, and Methodology

Our overall objective was to determine if the RRP effectively meets requirements and identifies fraudulent tax returns. To accomplish our objective, we:

- I. Evaluated the RRP technology and determined if it could support the RRP requirements.
 - A. Determined if the central processing unit and memory would support tax processing and the data models.
 - B. Determined if the technology could support running the RRP daily.
- II. Determined if the RRP was accurately identifying fraudulent tax returns and if the data models were capable of identifying new fraud schemes.
 - A. Evaluated the data models and traced the business requirements to the models.
 - B. Evaluated if the RRP processes could support new fraud schemes identified during the current filing season.
 - C. Determined if there was a feedback process to reduce the RRP's false positive identification.
- III. Determined if the RRP system was authorized to operate beyond October 2014.
 - A. Determined if the RRP met the conditions stated within the conditional Authorization to Operate granted in December 2013.
 - B. Determined if the vulnerability scanning required by the IRS and the National Institute of Standards and Technology was performed and vulnerabilities were mitigated or corrected.

Internal controls methodology

Internal controls relate to management's plans, methods, and procedures used to meet their mission, goals, and objectives. Internal controls include the processes and procedures for planning, organizing, directing, and controlling program operations. They include the systems for measuring, reporting, and monitoring program performance. We determined the following internal controls were relevant to our audit objective: the IRS's Internal Revenue Manual and the National Institute of Standards and Technology requirements that require the IRS to reduce vulnerabilities to systems by scanning its networks and computers to identify vulnerabilities, assessing the criticality of each identified vulnerability, and installing patches in a timely manner. We evaluated these controls by interviewing management and reviewing IRS documentation containing the results of various vulnerability scans performed.



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

Appendix II

Major Contributors to This Report

Alan R. Duncan, Assistant Inspector General for Audit (Security and Information Technology Services)
Danny Verneuille, Director
Kevin Liu, Audit Manager
Mike Mohrman, Lead Auditor
Joan Bonomi, Senior Auditor
Larry Reimer, Senior Auditor



*The Return Review Program Enhances the Identification
of Fraud; However, System Security Needs Improvement*

Appendix III

Report Distribution List

Commissioner C
Office of the Commissioner – Attn: Chief of Staff C
Deputy Commissioner for Operations Support OS
Deputy Commissioner for Services and Enforcement SE
Commissioner, Wage and Investment Division SE:W
Deputy Chief Information Officer for Operations OS:CTO
Associate Chief Information Officer, Applications Development OS:CTO:AD
Associate Chief Information Officer, Enterprise Operations OS:CTO:EO
Associate Chief Information Officer, Enterprise-Program Management Office OS:CTO:EPMO
Associate Chief Information Officer, Enterprise Services OS:CTO:ES
Director, Customer Account Services, Wage and Investment Division SE:W:CAS
Director, Office of Audit Coordination OS:PPAC:AC
Director, Office of Program Evaluation and Risk Analysis RAS:O
Chief Counsel CC
National Taxpayer Advocate TA
Office of Internal Control OS:CFO:CPIC:IC
Audit Liaison: Director, Risk Management Division OS:CTO:SP:RM



Appendix IV

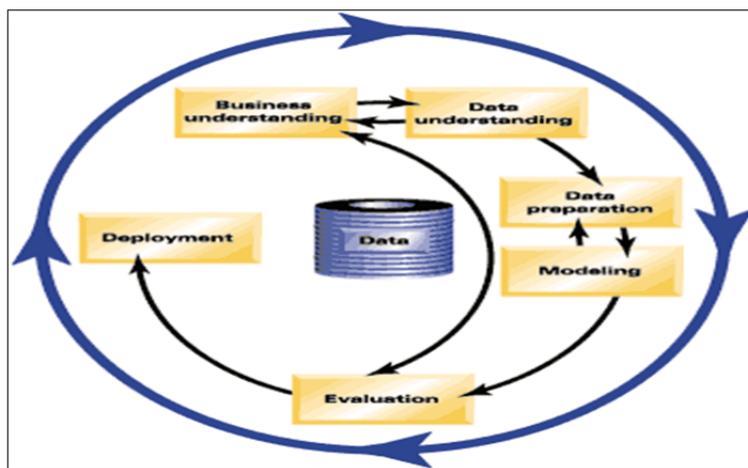
Cross-Industry Standard Process for Data Mining

The CRISP-DM is a proven, well-structured process for predictive modeling. This life cycle consists of six phases:

- **Business Understanding** focuses on understanding the project objectives and requirements.
- **Data Understanding** enables the IRS to enhance understanding of relevant data sources and identify any data quality problems.
- **Data Preparation** covers transformation, integration, and cleaning activities needed to prepare the data for modeling.
- **Modeling** applies various analytic techniques to develop predictive models.
- **Evaluation** focuses on thoroughly assessing, refining, and validating the model. In this phase, the IRS measures performance using quantitative measures such as the false positive rate and false negative rate.
- **Deployment** involves moving new models into production to score returns. This phase includes rigorous testing to validate that the models are correctly deployed.

The following figure is a graphical depiction of the CRISP-DM that shows the interconnection and iterative characteristics of the CRISP-DM methodology:

CRISP-DM Process Diagram



Source: IRS Return Review Program Fraud Analytics Overview.



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

Appendix V

Management's Response to the Draft Report



CHIEF TECHNOLOGY OFFICER

DEPARTMENT OF THE TREASURY
INTERNAL REVENUE SERVICE
WASHINGTON, D.C. 20224

JUN 10 2015

MEMORANDUM FOR DEPUTY INSPECTOR GENERAL FOR AUDIT

FROM:

Terence V. Milholland
Chief Technology Officer

SUBJECT:

Draft Audit Report – The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement, Audit (#201420017)

Thank you for the opportunity to review your draft audit report and to discuss the earlier draft report observations with the audit team. I appreciate the positive feedback received from the TIGTA team regarding the Return Review Program (RRP) and agree that a successful implementation of RRP is critical to the IRS mission. This is particularly true since our recent success in using massively parallel processing technology has broadened our vision for RRP to an enterprise platform, which will be utilized across the agency for identifying anomalies in different business processes.

We appreciate that the IRS and TIGTA teams were very much engaged and worked in partnership in producing this report. The TIGTA team provided valuable feedback during the process, such as information on a Heartbleed security vulnerability, which the IRS was able to immediately remediate prior to the close of the audit fieldwork.

We remain committed to managing the security risks in our IT infrastructure as required by the Federal Information Security Management Act, National Institute of Standards and Technology guidance, and other appropriate standards. We continue to actively monitor the IT environment and improve processes to ensure vulnerabilities are effectively prioritized and remediated.

In conclusion, we agree with TIGTA's recommendations, and we are committed to improving the overall security posture for RRP and our other IT systems, albeit tempered by limitations in people resources and organizational capacity associated with budget cuts over the past few years. Attached is our Corrective Action plan which outlines our steps for mitigation.

We value your continued support and the assistance and guidance your team provides. If you have any questions, please contact me at (240) 613-9373 or Karen Mayr at (202) 368-8396.



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement (201420017)

RECOMMENDATION #1: The Chief Technology Officer should ensure IRS personnel completing FISMA system classifications are familiar with the FISMA requirements for each level of classification

CORRECTIVE ACTION #1: The IRS agrees with this recommendation. Throughout the FISMA 2016 reporting period, IT Cybersecurity will brief IRS personnel completing the FISMA system classifications to ensure they understand the FISMA requirements for each level of classification.

IMPLEMENTATION DATE: June 15, 2016

RESPONSIBLE OFFICIAL: ACIO, IT Cybersecurity

CORRECTIVE ACTION MONITORING PLAN. We enter accepted Corrective Actions into the Joint Audit Management Enterprise System (JAMES) and monitor them on a monthly basis until completion.

RECOMMENDATION #2 The Chief Technology Officer should ensure the validation of system classification and reclassification are discussed, reviewed, and documented during the bi-weekly Cybersecurity management meeting.

CORRECTIVE ACTION #2: The IRS agrees with this recommendation. IT Cybersecurity will enhance its current process of validation of system classification and reclassification as discussed, reviewed, and documented during the bi-weekly Cybersecurity management meeting.

IMPLEMENTATION DATE: August 15, 2015

RESPONSIBLE OFFICIAL: ACIO, IT Cybersecurity

CORRECTIVE ACTION MONITORING PLAN. We enter accepted Corrective Actions into the Joint Audit Management Enterprise System (JAMES) and monitor them on a monthly basis until completion.



The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement

The Return Review Program Enhances the Identification of Fraud; However, System Security Needs Improvement (201420017)

Recommendation #3 The Chief Technology Officer should ensure all critical and high risk RRP vulnerabilities are resolved, regardless of whether the system is part of the VRIP pilot or not. Once all of the highest risk issues are resolved, the IRS should work on resolving the remaining issues

CORRECTIVE ACTION #3: The IRS agrees with this recommendation. The IRS continues to improve its remediation process to ensure vulnerabilities are resolved. To date, the IRS has remediated all but 52 of the previously identified 322 vulnerabilities. Based on our continued monitoring, we will focus on resolving our critical vulnerabilities in production and then our lower environments. We will add RRP to the VRIP pilot, which will provide us with the framework to report, assign and facilitate remediation of database vulnerabilities detected by the compliance and vulnerability scanning tools.

IMPLEMENTATION DATE: November 15, 2015

RESPONSIBLE OFFICIAL: Associate CIO, Enterprise Program Management Office

CORRECTIVE ACTION MONITORING PLAN. We enter accepted Corrective Actions into the Joint Audit Management Enterprise System (JAMES) and monitor them on a monthly basis until completion.