The Census Bureau Must Correct Fundamental Cloud Security Deficiencies in Order to Better Safeguard the 2020 Decennial Census

FINAL REPORT NO. OIG-19-015-A
JUNE 19, 2019

U.S. Department of Commerce
Office of Inspector General
Office of Audit and Evaluation

UNITED STATES DEPARTMENT OF COMMERCE
Office of Inspector General
Washington, D.C. 20230

June 19, 2019

MEMORANDUM FOR: Dr. Steven Dillingham
Director
U.S. Census Bureau

FROM: Frederick J. Meny, Jr.
Assistant Inspector General for Audit and Evaluation

SUBJECT: *The Census Bureau Must Correct Fundamental Cloud Security Deficiencies in Order to Better Safeguard the 2020 Decennial Census*
Final Report No. OIG-19-015-A

Attached for your review is our final report on the audit of the U.S. Census Bureau's (the Bureau's) cloud-based systems supporting the decennial census. Our objective was to determine the effectiveness of security processes and controls for select cloud-based information technology systems supporting the 2020 Census.

We found the following:

I. Unsecured GovCloud root user keys caused severe risks to 2020 Census cloud environments.

II. Unimplemented security baselines that document system settings and configurations left critical systems vulnerable.

III. Basic security practices were not fully implemented to protect Title 13 data hosted in the cloud.

On May 23, 2019, we received the Bureau's response to the draft report's findings and recommendations, which we include within the report as appendix C. The Bureau concurred with all eight report recommendations, and noted actions it has taken and will take to address them.

Pursuant to Department Administrative Order 213-5, please submit to us an action plan that addresses the recommendations in this report within 60 calendar days. This final report will be posted on OIG's website pursuant to sections 4 and 8M of the Inspector General Act of 1978, as amended (5 U.S.C. App., §§ 4 & 8M).

We appreciate the cooperation and courtesies extended to us by your staff during our audit. If you have any questions or concerns about this report, please contact me at (202) 482-1931 or Clark Morsbach, Director for Audit and Evaluation, at (202) 482-5509.

Attachment

cc: Dr. Ron Jarmin, Deputy Director and Chief Operating Officer, Census Bureau
Terryne Murphy, Acting Chief Information Officer, Department of Commerce
Kevin B. Smith, Chief Information Officer, Census Bureau
Jeffery Jackson, Acting Chief Information Security Officer, Census Bureau
Colleen Holzbach, Program Manager for Oversight Engagement, Census Bureau
Jean McKenzie, IT Security Audit Liaison, Census Bureau
Corey J. Kane, Audit Liaison, Census Bureau
Kemi A. Williams, Program Analyst for Oversight Engagement, Census Bureau
Joselyn Bingham, Audit Liaison, Office of the Chief Information Officer
Maria Dumas, IT Security Audit Action Officer, Office of the Chief Information Officer
MaryAnn Mausser, Audit Liaison, Office of the Secretary

Report in Brief

June 19, 2019

Background

The U.S. Census Bureau (the Bureau) is responsible for conducting a decennial census to ensure an accurate count of the U.S. population. During the 2020 decennial census (the 2020 Census), the Bureau will use the Internet to collect sensitive data of U.S. individuals and businesses protected under U.S. Code Title 13. These protected Title 13 data include personally identifiable information, such as names, addresses, dates of birth, and telephone numbers. The far-reaching consequences of altered, lost, or stolen Title 13 data emphasize the necessity to safeguard the Bureau information technology (IT) systems that will support the 2020 Census.

Every aspect of the 2020 Census related to Title 13 data collection and storage will rely upon commercial cloud services for its primary means, and will therefore require unique security precautions.

As part of the preparation for the 2020 Census, the Bureau conducted the 2018 End-to-End (E2E) Test to assess and validate the 2020 Census operations, procedures, systems, and infrastructure. To execute an effective test of the IT systems that will support the 2020 Census, Title 13 data were collected and stored within the Bureau's 2020 Census cloud environments.

Why We Did This Review

The objective of this audit was to determine the effectiveness of security processes and controls for select cloud-based IT systems supporting the 2020 Census.

U.S. CENSUS BUREAU

The Census Bureau Must Correct Fundamental Cloud Security Deficiencies in Order to Better Safeguard the 2020 Decennial Census

OIG-19-015-A

WHAT WE FOUND

We found that the Bureau's cloud-based IT systems—which will support the 2020 Census—contained fundamental security deficiencies that violated federal standards and U.S. Department of Commerce policies. Many of these deficiencies indicate that the Bureau was behind schedule and rushed to deploy its systems to support the 2018 E2E Test and the 2020 Census. Specifically, we found that (1) unsecured GovCloud root user keys caused severe risks to 2020 Census cloud environments; (2) unimplemented security baselines that document system settings and configurations left critical systems vulnerable; and (3) basic security practices were not fully implemented to protect Title 13 data hosted in the cloud.

Throughout this audit, we worked closely with Bureau system administrators, security staff, and senior leadership so that the security issues we identified could be addressed. This coordination allowed the Bureau to remediate some of these issues before the conclusion of our audit. However, these findings demonstrate that the Bureau did not securely use commercial cloud services to host its cloud environments during 2020 Census preparations, which placed the sensitive Title 13 data collected by the Bureau during the 2018 E2E Test at increased risk of potential misuse or loss. Our recommendations, if fully implemented, will help the Bureau manage its cloud environments in a more secure manner.

WHAT WE RECOMMEND

We recommend that the Chief Information Officer of the U.S. Census Bureau do the following:

1. Manage the GovCloud root user account according to federal and Departmental requirements. This must include a standardized, documented process to disable the use of all GovCloud root user accounts during the environment creation process for any new GovCloud environments.
2. Assess all Amazon Web Services user accounts in accordance with National Institute of Standards and Technology (NIST) account management requirements and conduct periodic reviews as part of Office of Information Security assessments.
3. Reassess, implement, and continuously monitor security baselines within all cloud environments.
4. Perform technical assessments to validate implementation of security baselines as part of the Bureau's cloud systems' initial and ongoing assessments.
5. Track all Title 13 data that are stored and processed in Bureau cloud environments. This must include coordination between cloud administrators, operational staff, and Office of Information Security personnel.
6. Expedite the implementation of the backup solution in progress and ensure it is operating in accordance with NIST guidance.
7. Formally document and ensure the implementation of controls compensating for lack of disaster recovery planning or engage in disaster recovery planning if the Bureau is unable to meet its obligation to compensate for the lack of disaster recovery planning.
8. Develop and approve an exit strategy for all Bureau cloud systems, including details for completely and securely removing data from the cloud service provider.

Contents

Cover: Herbert C. Hoover Building main entrance at 14th Street Northwest in Washington, DC. Completed in 1932, the building is named after the former Secretary of Commerce and 31st President of the United States.

Introduction

The U.S. Census Bureau (the Bureau) is responsible for conducting a decennial census as mandated by the United States Constitution[1] to ensure an accurate count of the U.S. population. Data collected during a decennial census are used to apportion the number of seats each state will have in the U.S. House of Representatives, define congressional districts, and distribute billions of dollars in federal funds for things like highways, hospitals, and schools. During the 2020 decennial census (the 2020 Census), the Bureau will use the Internet to collect sensitive data of U.S. individuals and businesses protected under U.S. Code Title 13. These protected Title 13 data include personally identifiable information (PII), such as names, addresses (including GPS coordinates), dates of birth, and telephone numbers. The far-reaching consequences of altered, lost, or stolen Title 13 data emphasize the necessity to safeguard the Bureau information technology (IT) systems that will support the 2020 Census.

Figure 1. Bureau AWS Cloud Environments

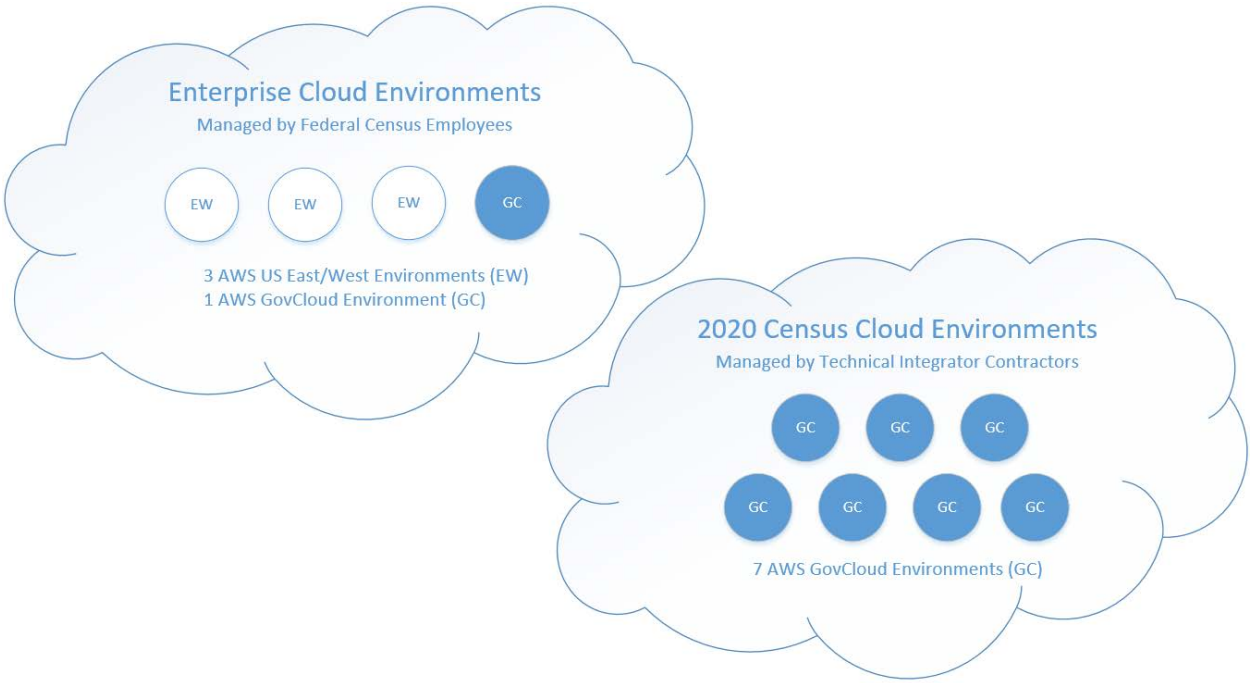

Source: Figure created by OIG based upon the Bureau's AWS cloud environments assessed during our audit.

Every aspect of the 2020 Census related to Title 13 data collection and storage will rely upon commercial cloud services for its primary means, and will therefore require unique security precautions. Commercial cloud services are services provided by a company that rents its data center computing resources to other entities for hosting web-based environments and services. The Bureau is using Amazon Web Services (AWS) to host its enterprise and 2020 Census environments, which are composed of digital storage, servers, and databases. The Bureau's enterprise cloud environments support its ongoing operations, while the 2020 Census cloud

[1] U.S. Const. art. I, § 2.

environments were created to prepare for and host the 2020 Census. Bureau employees manage the enterprise cloud environments, while technical integrator (TI) contractors manage the 2020 Census cloud environments. Most of these environments (8 of 11) are hosted within the AWS GovCloud Region (see figure 1). The GovCloud Region supports a higher level of security, is built and operated within the United States, and is only available to U.S. entities such as government, contractor, and educational organizations.[2]

As part of the preparation for the 2020 Census, the Bureau conducted the 2018 End-to-End (E2E) Test to assess and validate the 2020 Census operations, procedures, systems, and infrastructure. To execute an effective test of the IT systems that will support the 2020 Census, Title 13 data were collected and stored within the Bureau's 2020 Census cloud environments.

Proper implementation of commercial cloud services can provide highly flexible and efficient computing resources. However, the shared responsibility for securing systems hosted in commercial cloud environments is clearly defined between the cloud service provider and customer. As the customer, the Bureau is responsible for maintaining the confidentiality and integrity of Title 13 data in the cloud. Unfortunately—and according to the U.S. Government Accountability Office[3]—the Bureau has experienced delays and has fallen behind schedule in implementing its IT systems. These delays have required the Bureau to operate under a compressed schedule due to the immutable deadline of the 2020 Census.

[2] All GovCloud environments referenced in this report were hosted within the GovCloud West region. On November 12, 2018, AWS announced the availability of GovCloud East, an additional GovCloud region. Before November 12, 2018, GovCloud West was the only GovCloud region available.

[3] U.S. Government Accountability Office, August 2018. *Continued Management Attention Needed to Address Challenges and Risks with Developing, Testing, and Securing IT Systems*, GAO-18-655. Washington, DC: GAO.

Objective, Findings, and Recommendations

The objective of this audit was to determine the effectiveness of security processes and controls for select cloud-based IT systems supporting the 2020 Census. Our review primarily focused on the Bureau's IT systems hosted by AWS cloud services. We briefed senior Bureau personnel, such as the chief information officer, chief technology officer, and acting chief information security officer regarding our initial findings on November 28, 2018. See appendix A for further details regarding our objective, scope, and methodology.

We found that the Bureau's cloud-based IT systems—which will support the 2020 Census—contained fundamental security deficiencies that violated federal standards and U.S. Department of Commerce policies. Many of these deficiencies indicate that the Bureau was behind schedule and rushed to deploy its systems to support the 2018 E2E Test and the 2020 Census. Specifically, we found that

1. unsecured GovCloud root user keys caused severe risks to 2020 Census cloud environments,
2. unimplemented security baselines that document system settings and configurations left critical systems vulnerable, and
3. basic security practices were not fully implemented to protect Title 13 data hosted in the cloud.

Throughout this audit, we worked closely with Bureau system administrators, security staff, and senior leadership so that the security issues we identified could be addressed. This coordination allowed the Bureau to remediate some of these issues before the conclusion of our audit. However, these findings demonstrate that the Bureau did not securely use commercial cloud services to host its cloud environments during 2020 Census preparations, which placed the sensitive Title 13 data collected by the Bureau during the 2018 E2E Test at increased risk of potential misuse or loss. Our recommendations, if fully implemented, will help the Bureau manage its cloud environments in a more secure manner.

I. Unsecured GovCloud Root User Keys Caused Severe Risks to 2020 Census Cloud Environments

The Bureau did not secure the most privileged user account, the root user account, in each of its AWS GovCloud environments, several of which contained Title 13 data collected during the 2018 E2E Test. AWS GovCloud environments contain a root user account which has unlimited privileges and cannot be restricted within its environment. For instance, a root user account can modify other GovCloud user accounts, remove data from the environment, and delete any virtual server from the environment. The GovCloud root user account accesses the GovCloud environment infrastructure by issuing API (Application Programming Interface) calls using an identification key and secret key. These keys can be used to access the cloud environment via the Internet from anywhere in the world. We found that access keys for the Bureau's GovCloud root users had been lost for a prolonged period, root user accounts had not been secured, and root user account capabilities were not understood.

A. The Bureau lost GovCloud root user keys for a prolonged period

When an AWS GovCloud environment is initialized, a root user account is automatically created for the environment. AWS, industry best practices, and the Bureau's security documentation all recommend that GovCloud root user keys be disabled after initial environment setup. However, we found that the Bureau did not disable the GovCloud root user keys, which were then lost for all eight of the Bureau's GovCloud environments. The oldest of these keys were created more than 2 years prior to our discovery of this issue in April 2018. Evidence collected from the system indicates that, shortly after the GovCloud environments were created, all the keys were used[4]—suggesting that someone had access to the root keys. While lost, the Bureau did not have control of the root user accounts and could not disable access to them.

The root user keys for the seven 2020 Census cloud environments—which will be responsible for hosting the 2020 Census—were lost by the commercial cloud reseller.[5] The root user keys for the one enterprise GovCloud environment were lost by an employee who had left the Bureau before our audit (see figure 1). The Bureau's Office of Information Security (OIS)—which is responsible for the security of the cloud environments—did not identify this issue as part of its required security assessments before the environments were put into operation.

After identifying this issue as part of our assessment, we provided the results to OIS on April 25, 2018, and discussed the results with Bureau staff on May 8, 2018. We briefed the Bureau's chief information officer on May 23, 2018, and issued a memorandum to Bureau leadership on June 4, 2018, informing them of this issue and recommending immediate action to secure all root users keys for its cloud environments (see appendix B). While the Bureau never officially responded to this memorandum, we determined it did eventually disable all GovCloud root user keys. Specifically, the Bureau relied upon technical intervention by AWS to reset the Bureau's root user keys. This process—that is, to rotate and disable the GovCloud root user keys—took the Bureau's enterprise cloud administrators 42 days (from May 2 to June 13, 2018) and the TI contractor cloud administrators 12 days (from May 23 to June 4, 2018). This involved coordination between the Bureau, the commercial cloud reseller, and AWS. The 6-week timeframe it took to secure its root user accounts demonstrates the Bureau's inability to have stopped a potential attacker with stolen root keys from modifying or destroying all cloud system resources hosted in its GovCloud environments.

[4] We observed that all GovCloud root keys had been rotated, which disables the current keys and creates new ones. Rotating GovCloud root keys is only possible by using the existing root keys.

[5] A commercial cloud reseller allows government agencies to procure cloud services from commercial cloud providers that do not have existing contract vehicles with the government. (See https://www.fedramp.gov/cloud/ [accessed December 17, 2018].)

B. *The Bureau's GovCloud root user accounts were not secured which left them vulnerable to misuse*

Multi-factor authentication (MFA) is a method of authentication that requires the use of two or more pieces of evidence[6] before a user is allowed access to a system. MFA helps protect a user's account from an attacker who has compromised the account's credentials, like a username and password. The Bureau's GovCloud root users' identification key and secret key, which function much like a username and password, were not being protected by MFA. Not having MFA configured for the most privileged user accounts of the Bureau's GovCloud environments left the keys vulnerable to theft and misuse.

We contacted AWS to discuss this issue, and found that MFA is not supported for the GovCloud root user account in an attempt to discourage the use of the account. This limitation is unique to the GovCloud root user account; MFA is supported for all other AWS user accounts. Federal Standards require[7] privileged user accounts[8] to be configured to use MFA. However, since MFA could not be implemented on the root user account the Bureau should have disabled its GovCloud root user accounts to adhere to federal and Departmental requirements. Since the root user account is created by default when a GovCloud environment is established, the Bureau must ensure it is disabled after initial GovCloud environment setup.

C. *Bureau personnel did not understand the capabilities of GovCloud root user accounts*

Both OIS and TI cloud administrators were unaware that the API keys had been lost for the seven 2020 GovCloud environments (see figure 1). However, prior to our audit in May 2017, enterprise cloud administrators realized that they did not possess the ability to modify or otherwise manage their GovCloud root user account because the API keys had been lost. Accordingly, the Bureau of the Census Computer Incident Response Team (BOC CIRT) opened a case, and reported the incident to the US-CERT (United States Computer Emergency Readiness Team). AWS Support was also contacted to assist the enterprise cloud administrators to disable the lost API keys. Unfortunately, in July 2017, AWS Support mistakenly led the cloud administrators to believe that the Bureau was not responsible for managing the GovCloud root user account. Consequently, the BOC CIRT closed its case based upon this incorrect guidance. The US-CERT ended its review after determining the incident was a violation of Census' own policies, and therefore outside of its purview. Nevertheless, the Bureau is

[6] MFA is a security enhancement that allows users to present two or more pieces of evidence—their credentials—when logging in to an account. Their credentials fall into any of these three categories: (1) something they know (like a password or PIN), (2) something they have (like a smart card), or (3) something they are (like a fingerprint). Credentials must come from two different categories to enhance security—thus, entering two different passwords would not be considered multi-factor. (See https://www.nist.gov/itl/tig/back-basics-multi-factor-authentication [accessed December 17, 2018].)

[7] National Institute of Standards and Technology, April 2013. *Security and Privacy Controls for Federal Information Systems and Organizations*, NIST SP 800-53, Rev. 4. Gaithersburg, MD: NIST, F-90–F-91.

[8] A privileged user is one that is authorized (and therefore, trusted) to perform security-relevant functions that ordinary users are not authorized to perform.

responsible for managing the root user account for its AWS GovCloud environments. The Bureau may have avoided acting upon incorrect guidance if it had sufficiently understood the capabilities and its responsibilities for the GovCloud root user account.

During our audit, the same Bureau cloud administrators who attempted to troubleshoot this issue in 2017 contacted AWS Support to disable their GovCloud root user's API keys. However, in May 2018, AWS Support representatives again provided inaccurate information, this time by mistakenly assuring the Bureau cloud administrators that there was no GovCloud root user account associated with their environment. Ultimately, the Bureau cloud administrators convinced AWS Support that the GovCloud root user was a functioning account and that it can perform all functions within the GovCloud environment without limitation. AWS eventually assisted the Bureau to reset its GovCloud root user keys, which were in turn disabled.

Failure to secure the GovCloud root user keys exposed preparations for the 2020 Census to severe security risks. If we had not identified this issue, it is plausible that the GovCloud root users would have remained in an insecure state up to and during the 2020 Census. This is due to the Bureau not identifying the issue during security assessments, inherent limitations in the AWS infrastructure, and security decisions based on inaccurate information. However, OIS and the Bureau's cloud administrators should have sufficiently understood their responsibilities for the GovCloud root user account, which were clearly defined in both AWS's and the Bureau's own security documentation.

While the lost root user keys may not have been easily guessed by an attacker due to their length and complexity, the uncertainty of who may have had the keys resulted in severe risk to Bureau operations and the Title 13 data collected as part of the 2018 E2E Test. Fortunately, we did not find evidence of the lost root keys being used maliciously. However, the Bureau could not know if they had been stolen or sold and, having lost the root user keys, would have been powerless to stop an attacker from causing irreparable harm to the cloud environments. Therefore, we conclude that the Bureau exposed the 2020 Census preparations to potentially catastrophic risk by not securing the root user accounts.

Recommendations

We recommend that the Chief Information Officer of the U.S. Census Bureau do the following:

1. Manage the GovCloud root user account according to federal and Departmental requirements. This must include a standardized, documented process to disable the use of all GovCloud root user accounts during the environment creation process for any new GovCloud environments.
2. Assess all AWS user accounts in accordance with National Institute of Standards and Technology (NIST) account management requirements and conduct periodic reviews as part of OIS assessments.

II. Unimplemented Security Baselines That Document System Settings and Configurations Left Critical Systems Vulnerable

One fundamental and required method of securing any IT system is by creating and adhering to a security baseline, which documents the agreed upon system settings and configurations. Based on industry best practices, the Center for Internet Security (CIS) created a benchmark for AWS environments that customers can implement to strengthen the security posture of their cloud-based systems. The Bureau modeled its AWS cloud security baselines after the CIS benchmark, which prescribes specific AWS environment configuration conditions in areas such as identity and access management, system logging, and system monitoring.

Properly implemented security baselines serve as the backbone of security and protection of an IT system. However, we found that the Bureau did not securely configure its cloud environments before putting them into production, and OIS did not effectively oversee the implementation of cloud security baselines.

A. The Bureau did not securely configure its cloud environments before putting them into production

We assessed the Bureau's 11 AWS cloud environments (8 hosted in GovCloud and 3 in US East/West[9]) against the Bureau's defined baselines, and found that none had fully implemented its baseline security configurations. However, these critical environments were put into production despite being insecurely configured, some of which then collected and stored Title 13 data during the 2018 E2E Test. The following are significant examples of baseline settings not implemented:

Unused credentials for 90 days or greater were not disabled. We found a large number of inactive users of the Bureau's AWS infrastructure that should have been disabled. The CIS benchmark includes the condition to disable user credentials unused in the last 90 days. The Bureau elected an even stricter requirement for its cloud environment baselines, reducing the time a user account can be left inactive to 30 days before it should be disabled. Unfortunately, the Bureau's operational and auditing safeguards failed to identify and disable inactive users by even the less restrictive condition of the CIS benchmark. We found 10 out of 52 users (19 percent) in the enterprise cloud environments and 141 out of 527 (27 percent) in the 2020 Census environments had not been used in more than 90 days. The practice of disabling inactive users, when thoroughly implemented, helps prevent the compromise and misuse of rarely used user credentials by bad actors.

MFA was not enabled for several users. In addition to the GovCloud root user accounts identified in finding I, we found nine user accounts in two of the Bureau's AWS cloud environments that had not been configured to use MFA. After providing our results to the Bureau, we found that some of these user accounts were originally for testing purposes and should have been disabled after the testing period

[9] AWS US East/West is a cloud region independent from GovCloud and is available for worldwide use by the public.

was finished. Not having MFA configured leaves the affected user accounts vulnerable to password attacks (e.g., brute force[10] and/or key-logger[11]). The occurrence of just one unsecured account increases the risk of unauthorized access.

Alarms for infrastructure configuration changes were largely not utilized. Alarms are used to alert system administrators of selected authorized or unauthorized changes to an IT environment. When alarms are in place and operating as intended, security professionals can attempt to identify, eliminate, or rollback unauthorized changes in a timely manner. However, the Bureau did not consistently implement alarms within its cloud environments, ignoring many conditions of the Bureau security baselines. We found that enterprise cloud environments only utilized 15 of 56 (27 percent) alarms while 2020 Census environments only utilized 23 of 56 (41 percent) alarms. By not having these alarms in place, changes to the system that could cause extensive harm to the environment may go unnoticed.

B. OIS did not effectively oversee the implementation of cloud security baselines

The Bureau's OIS is responsible for creating and distributing the cloud security baselines for both the enterprise and 2020 Census environments according to federal standards.[12] It is also responsible for performing assessments to determine if the baselines have been fully implemented. We found that the system administrators from the Bureau's enterprise and 2020 Census cloud environments had significantly different security baselines to protect Title 13 data. Regardless of enterprise and 2020 Census administrators using different baseline versions, neither had fully implemented the baseline it had received. These inconsistently implemented baselines led to a weakened and vulnerable security posture of the environments. As a result, the Bureau was in violation of federal standards to implement its security baselines.

Assessments performed by OIS should have identified that baselines had not been implemented for the Bureau's cloud systems. However, we found that either no assessments were performed or the assessments when performed were insufficient. For example, in February 2017, OIS assessors only requested a copy of the baseline document to determine whether one existed. We found no indication that the assessors took actions to verify whether the AWS cloud environments actually adhered to the security baselines. These insufficient assessments were particularly damaging since the assessors indicated that the cloud environment baselines had been implemented. Ultimately, this could cause unimplemented baselines to be largely ignored for at least another year until the baselines would be reassessed as part of its reassessment schedule. OIS' inattention to the fundamental security requirements for critical 2020 Census systems caused security vulnerabilities to persist with no awareness of the weaknesses.

[10] A method of accessing a device or user account through attempting multiple combinations of numeric/alphanumeric passwords.

[11] A program designed to record which keys are pressed on a computer keyboard used to obtain passwords or encryption keys and thus bypass other security measures.

[12] See NIST SP 800-53, Rev. 4, p. F-70.

Recommendations

We recommend that the Chief Information Officer of the U.S. Census Bureau do the following:

3. Reassess, implement, and continuously monitor security baselines within all cloud environments.
4. Perform technical assessments to validate implementation of security baselines as part of the Bureau's cloud systems' initial and ongoing assessments.

III. Basic Security Practices Were Not Fully Implemented to Protect Title 13 Data Hosted in the Cloud

During the 2018 E2E Test, the Bureau collected and stored Title 13 data in its 2020 Census cloud environments. We assessed how the Bureau managed the Title 13 data in these cloud environments and found that many basic security practices had not been sufficiently implemented. Specifically, we found that the Bureau (a) did not sufficiently track Title 13 data in its cloud environments; (b) did not complete disaster recovery planning; (c) had not fully implemented its intended data backup program; and (d) had not created an exit strategy for the data stored in the cloud. These conditions indicate that the Bureau was not equipped to ensure the security of the sensitive data collected and stored in its cloud environments, which increased the risk to the 2020 Census preparations.

A. The Bureau did not sufficiently track the location of Title 13 data

In our attempts to determine whether Title 13 data hosted in the 2020 Census cloud environments were being backed up, 2020 Census cloud administrators and OIS staff provided us incomplete and contradictory information about the data's location within the cloud environments. Ultimately, the cloud administrators indicated that they were uncertain which 2020 Census virtual servers and databases contained Title 13 data. One method of tracking where data reside within a system is through data tagging, which is used to digitally mark objects that contain sensitive information. We found that because the Bureau was not utilizing data tagging or a similar practice, it could not identify which servers hosted in the cloud contained Title 13 data.

By not having a detailed record of where sensitive data were located, the Bureau could not effectively prioritize the security of its data. Further, it would be inappropriate to treat all Bureau data stored in a cloud environment as Title 13 data because this will result in the Bureau's limited resources being used to protect nonsensitive data. By not implementing a practice to track its sensitive data, the Bureau made it more difficult to effectively safeguard, backup, and migrate sensitive data within its cloud environments. Further, by not sufficiently tracking its sensitive data, the Bureau could incur greater levels of risk to the Title 13 data provided by Census respondents and effectively lessens its ability to protect the data.

B. *The Bureau lacked disaster recovery options to safeguard against data loss*

We found that the 2020 Census cloud environments did not have disaster recovery options capable of restoring data lost in the event of a large-scale attack or disaster. The GovCloud West Region—where the 2020 Census cloud environments were hosted during our audit—does not currently provide disaster recovery options since it is composed of a single geographic region. NIST states, "[i]f an organization relies on a cloud service for data storage and processing, it must be prepared to carry on mission critical operations without the use of the service for periods when the cloud experiences a serious outage."[13] Based on this guidance, the Bureau should consider implementing a disaster recovery solution that exists separate from commercial cloud services. For example, this solution could utilize a federally owned data center to provide disaster recovery capabilities.

We also found that the current disaster recovery plan for the 2020 Census cloud environments had not been completed or approved by Bureau IT management. The Bureau is required to conduct a functional test of its disaster recovery capabilities that includes an element of system recovery. However, we found that there had been no functional test of these capabilities for the 2020 Census cloud environments. For example, no comprehensive testing had been conducted to recover from data backups. The Bureau purportedly conducted a tabletop exercise (i.e., a verbal walkthrough with key stakeholders that does not actually test recovery capabilities) in April 2017, but was unable to provide an account of who participated in the test. Therefore, we were unable to validate the effectiveness of the exercise. Based upon these observations, we determined that the Bureau did not adequately prepare for disaster recovery for its 2020 Census cloud environments.

In October 2018, we asked the 2020 Census disaster recovery team about their confidence in a full recovery in the event of a major incident. While the team was not confident in the existing disaster recovery planning, they did express confidence in the duplication of data across multiple availability zones[14] to compensate for the lack of completed disaster recovery plans. Subsequently, on November 8, 2018, the Bureau's associate director for Decennial Census Programs and the chief information officer decided not to engage in traditional disaster recovery for its 2020 Census IT systems.[15] This decision was made, in part, based on the AWS cloud infrastructure providing high

[13] NIST, December 2011. *Guidelines on Security and Privacy in Public Cloud Computing*, NIST SP 800-144. Gaithersburg, MD: NIST, 32.

[14] Availability zones provide AWS cloud infrastructure customers with functionality to duplicate virtual servers and data across multiple data centers. While this functionality can provide higher availability to the data when configured correctly, it does not provide disaster recovery options.

[15] The Bureau is authorized to make this decision under NIST guidance if the Bureau (1) selects compensating controls for the lack of traditional disaster recovery planning; (2) provides "a supporting rationale for how compensating controls provide equivalent security capabilities . . . and why the baseline security controls could not be employed"; and (3) "assess[es] and accept[s] the risk associated with implementing compensating controls . . . ." See NIST SP 800-53, Rev. 4, p. 36.

availability for the Bureau's data by using multiple availability zones (i.e., data centers).[16] However, we found that 25 percent (26 of 103) of the Bureau's cloud databases with an operational requirement to use multiple availability zones had not been configured to do so. By not using multiple availability zones, the data stored within these databases were less tolerant to unexpected data loss, and therefore did not provide sufficient compensation for not engaging in traditional disaster recovery. We conclude that without disaster recovery capabilities or correctly configured AWS cloud environments to compensate for a lack of traditional disaster recovery controls, the Bureau would be unable to carry on mission critical operations in the event of a major disruption or outage.

C. Data backups were not sufficiently implemented or tested before Title 13 data were collected as part of the 2018 E2E Test

The Bureau had not fully implemented—and was not able to test—its backup capabilities as part of the 2018 E2E Test. These capabilities include creating backup copies of Title 13 data collected during the 2018 E2E Test and 2020 Census. We found that although the backup solution was available in January 2018, TI cloud administrators did not begin to implement the backup solution until July 2018 because the Bureau had not granted the backup solution an authorization to operate. This lack of implementation illustrates our overall conclusion that the Bureau was behind schedule to have fundamental security functions in-place before collecting Title 13 data. We followed up with the Bureau in October 2018 and found that it had still not configured 50 percent of the 2020 cloud environment servers to use the backup solution.

During the 2018 E2E Test, the Bureau relied on AWS snapshots within the cloud environment to back up its data. Snapshots take a point-in-time copy of the host system and data stored in the cloud, then take point-in-time copies of data that have been added or updated since the previous snapshot. Within the 2020 cloud environments, these snapshots occurred every 24 hours. However, as of October 2018, the Bureau had not determined whether this snapshot interval was sufficient. Furthermore, we found the 2020 cloud environment disaster recovery plan being developed defined the recovery point objective (RPO) of Census respondent data (e.g., Internet self-response) as 10 minutes. The RPO is the point in time, prior to a disruption or outage, to which data must be recovered (see figure 2). In the event of a disaster, the Bureau would be unable to meet a 10-minute RPO due to snapshots only having been taken every 24 hours. For instance, a system failure after a snapshot is taken could result in up to 24 hours of response data being lost (see figure 3). This could lead to an increase in cost and time for the Bureau to complete the 2020 Census.

[16] The Bureau made the decision not to engage in traditional disaster recovery as we were concluding our fieldwork, which limited our ability to analyze the compensating controls or to evaluate whether the decision should have been reconsidered.

Figure 2. Definition Diagram of RPO

Source: Figure created by OIG based upon industry accepted disaster recovery concepts.

Figure 3. State of the Bureau's 24-hour Snapshot Schedule

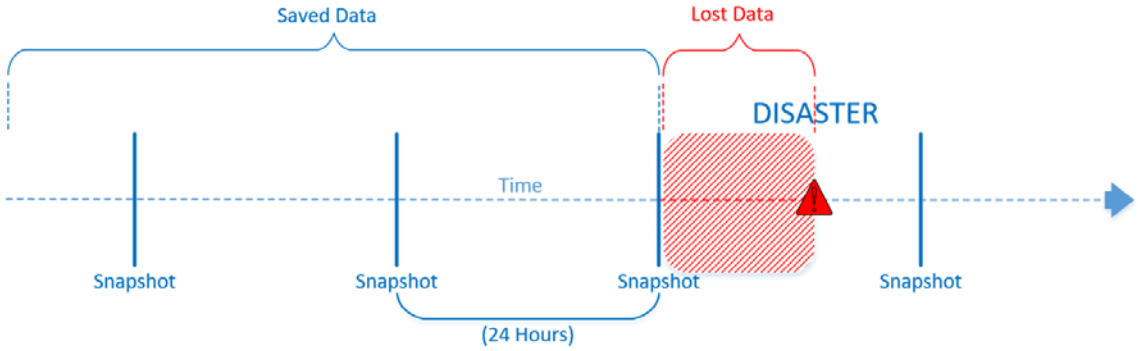

Source: Figure created by OIG based upon the Bureau's current snapshotting process.

Until a thorough test of its backup capabilities is conducted, the Bureau will have no assurance that the backup mechanisms intended for the 2020 Census will be effective. Since the backup solution was not sufficiently implemented as part of the 2018 E2E Test, and snapshots had not been configured to achieve the Bureau's anticipated data recovery requirements, we conclude that recovery of sensitive data would not have been achievable following a major incident.

D. The Bureau did not develop an exit strategy for its cloud environments

We found that the Bureau did not have an exit strategy in place for data stored in its 2020 Census cloud environments. Removing data from the cloud requires unique considerations because of the risk of data remnants left within the cloud environment. This is partly because of the customer's inability to retain or destroy the physical storage drives used within the cloud infrastructure. Unauthorized individuals who later have access to the cloud storage could potentially retrieve some of the residual data long after the Bureau has left the cloud provider. While the Bureau had a general vision for Title 13 data collected and produced from the 2018 E2E Test and the 2020 Census, there were no finalized plans for migrating this data off the commercial cloud provider. This illustrates that the Bureau was behind schedule and did not fully consider the security needs of the data collected during the 2018 E2E Test and the 2020 Census.

According to NIST,[17] establishing an exit strategy is an important part of the planning process before cloud system deployment. NIST states that an exit strategy should

[17] See NIST SP 800-144, p. 43.

include plans for termination of a commercial cloud service. For instance, termination includes the expiration of the service agreement or poor performance. Organizations using commercial cloud services should formulate a plan to remove data from the cloud through secure, reliable, and efficient means, and in a timely manner. By not being prepared with an exit strategy for its cloud environments, the Bureau risks exposing Title 13 data collected as part of the 2018 E2E Test and 2020 Census.

Recommendations

We recommend that the Chief Information Officer of the U.S. Census Bureau do the following:

5. Track all Title 13 data that are stored and processed in Bureau cloud environments. This must include coordination between cloud administrators, operational staff, and OIS personnel.
6. Expedite the implementation of the backup solution in progress and ensure it is operating in accordance with NIST guidance.
7. Formally document and ensure the implementation of controls compensating for lack of disaster recovery planning or engage in disaster recovery planning if the Bureau is unable to meet its obligation to compensate for the lack of disaster recovery planning.
8. Develop and approve an exit strategy for all Bureau cloud systems, including details for completely and securely removing data from the cloud service provider.

Summary of Agency Response and OIG Comments

On May 23, 2019, we received the Bureau's response to the draft report's findings and recommendations. In response to our draft report, the Bureau concurred with all eight recommendations and described both completed and planned actions to address each recommendation.

We have included the Bureau's formal response as appendix C of this report.

Appendix A: Objective, Scope, and Methodology

Our audit objective was to determine the effectiveness of security processes and controls for select cloud-based IT systems supporting the 2020 Census.

We reviewed internal security controls significant within the context of our audit objective and employed a comprehensive methodology to evaluate the security posture of the Bureau's AWS cloud environments. This included review and assessment of two Bureau systems listed in table A-1.

Table A-1. Census Systems Selected for Review

System Name	Brief Description
CEN08 TI—AWS GovCloud PaaS[a]	This system was hosted in the AWS cloud and will facilitate the 2020 Census. All seven environments were hosted in the GovCloud Region. Contracted TIs managed this system.
CEN09—Enterprise Cloud Services	This system was hosted in the AWS cloud and includes four environments that facilitate Census enterprise workloads. Three of these environments were hosted in the AWS US East/West Regions and one in the GovCloud Region. Federal Census Bureau employees managed this system.

Source: Table created by OIG based upon the Bureau's system descriptions.

[a] PaaS – Platform as a Service

We reviewed the implementation status of fundamental security controls defined in NIST Special Publication 800-53, Rev. 4, including account management, access control, configuration management, identity and authentication, and contingency planning.

To do so, we

- reviewed system-related artifacts, including policy and procedures, planning documents, and security control documentation;
- interviewed Bureau officials, including system owners, IT security and operations staff, and management;
- assessed the Bureau's AWS infrastructure configuration using specialized assessment tools and manual techniques; and
- validated the results of vulnerability scanning of Bureau AWS environments' configuration.

We reviewed the Bureau's compliance with the following applicable internal controls, provisions of law, and mandatory guidance:

- The Federal Information Security Modernization Act of 2014

- U.S. Department of Commerce Information Technology Security Program Policy, Version 3.2, and applicable Commerce Information Technology Requirements (CITR):
  - CITR-015, *Contingency Plan Testing*
  - CITR-016, *Vulnerability Scanning and Patch Management*
  - CITR-019, *Risk Management Framework (RMF)*
  - CITR-021, *Password Management*
  - CITR-024, *FedRAMP Applicability*
- NIST Special Publications:
  - 800-37, Rev. 1, *Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach*
  - 800-53, Rev. 4, *Security and Privacy Controls for Federal Information Systems and Organizations*
  - 800-53A, Rev. 4, *Assessing Security and Privacy Controls in Federal Information Systems and Organizations: Building Effective Assessment Plans*
  - 800-144, *Guidelines on Security and Privacy in Public Cloud Computing*
  - 800-145, *The NIST Definition of Cloud Computing*

We also used CIS industry best practice security benchmark as criteria for testing cloud-based infrastructure.

We collected computer-generated data directly from the Bureau's AWS cloud environments. We verified this data by interviewing appropriate Bureau officials and provided them the data to eliminate the possibility of false positive results. We determined that the data were sufficiently reliable for the purposes of this report. We conducted our review from March 2018 through November 2018 under the authority of the Inspector General Act of 1978, as amended (5 U.S.C. App.), and Department Organization Order 10-13, April 26, 2013. We performed our fieldwork at Department of Commerce headquarters in Washington, DC, as well as Bureau headquarters in Suitland, Maryland. 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.

Appendix B: 2020 Decennial Cloud-Based Systems Root Accounts Memorandum

UNITED STATES DEPARTMENT OF COMMERCE
Office of Inspector General
Washington, D.C. 20230

June 4, 2018

MEMORANDUM FOR: Ron Jarmin
Performing the Nonexclusive Functions and Duties
of the Director
U.S. Census Bureau

FROM: Frederick J. Meny, Jr.
Assistant Inspector General for Audit and Evaluation (Acting)

SUBJECT: IT Security Concerns with the 2020 Decennial Cloud-Based Systems Root Accounts

As part of our fiscal year 2018 audit of the Census Bureau's cloud-based systems, we are evaluating the cloud-based systems that are the primary mechanisms supporting the 2020 Census. These systems provide resources in the processing, storing, and transmitting of Title 13 data. Our objective is to determine the effectiveness of security processes and controls for select cloud-based IT systems supporting the 2020 decennial census. This memorandum provides you with our preliminary observations concerning root accounts that warrant your immediate attention. We have already discussed this matter with your chief information officer on May 23, 2018.

We found that the root accounts for the Census Bureau's Amazon Web Services (AWS) GovCloud environments (the oldest of which was created more than 2 years ago) have not been secured for a prolonged period. The Census Bureau does not have control of the root accounts and cannot disable access to them.

The root accounts are the AWS environment's most privileged accounts with unlimited power to perform any administrative action and read, modify, or delete information. Therefore, securing them is of paramount importance to Census Bureau operations. Left unchanged, this issue presents significant and undue risks to Census Bureau systems currently facilitating the 2018 End-to-End Test, a critical preparation for the 2020 decennial census. If these vulnerable accounts are exploited, the credibility of and participation in the 2020 decennial census could be significantly impacted.

Therefore, we recommend that the Census Bureau take immediate action to secure all root accounts for its cloud environments. Further details of our review and additional recommendations will follow in our forthcoming audit report.

The information contained in this memorandum may adversely affect information security if disclosed. Accordingly, the public release of this memorandum and information herein is prohibited by 44 U.S.C. § 3555(f), and this memorandum may be exempt from release in response to requests under 5 U.S.C. § 552 (the Freedom of Information Act). The Census Bureau should take appropriate steps to ensure the protection of the information in this memorandum.

If you have any questions regarding this matter, please contact me at (202) 482-1931 or Clark Morsbach, Director for Audit and Evaluation, at (202) 482-5509.

cc: Rod Turk, Acting Chief Information Officer
Kevin Smith, Chief Information Officer, Census Bureau
Colleen Holzbach, Program Manager for Oversight Engagement, Census Bureau
Jean McKenzie, Audit Liaison, Census Bureau
Corey J. Kane, Program Analyst, Census Bureau
Maria Dumas, Audit Liaison, Office of the Chief Information Officer

Appendix C: Agency Response

UNITED STATES DEPARTMENT OF COMMERCE
Economics and Statistics Administration
U.S. Census Bureau
Office of the Director
Washington, DC 20233-0001

May 23, 2019

MEMORANDUM FOR Frederick J. Meny, Jr.
Assistant Inspector General for Audit
and Evaluation (Acting)
Office of Inspector General

FROM: Steven D. Dillingham
Director
U.S. Census Bureau

SUBJECT: Response to *OIG Report: The Census Bureau Must Correct Fundamental Cloud Security Deficiencies in Order to Better Safeguard the 2020 Decennial Census*

This memo serves as a response to *OIG Report: The Census Bureau Must Correct Fundamental Cloud Security Deficiencies in Order to Better Safeguard the 2020 Decennial Census*.

The U.S. Census Bureau appreciates the review by the Office of Inspector General (OIG) of the cloud instances used during the 2018 End-to-End test. We acknowledge there are additional steps the Census Bureau can take to improve monitoring of the security posture of our cloud-based computing environment as we move out of development and toward production operations for the 2020 Census. As the Census Bureau and the OIG both concluded, neither found any evidence of the lost root keys being used maliciously. Further, no systems or data held and processed by the Census Bureau on behalf of the public have been identified as lost or compromised as a result of the issue in the recommendations in the OIG's report.

The Census Bureau understands that the specific recommendations put forward by the OIG will serve to enhance the overall security posture of the 2020 cloud environment and increase confidence on the part of the public in the Census Bureau's mission. As such, we will implement a significant amount of the recommendations in the report and continue to partner with the OIG in preparing for the Census Bureau's 2020 Census efforts.

OIG Recommendations/Responses

Recommendation #1 – Manage the GovCloud root user account according to Federal and Departmental requirements. This must include a standardized, documented process to disable the use of all GovCloud root user accounts during the environment creation process for any new GovCloud environments.

United States Census Bureau | United States Census 2020

census.gov

Response – *The Census Bureau agrees with this recommendation.*

In April 2018, the OIG made the Census Bureau aware of the insecure root user keys. Upon notification, the enterprise cloud administrators and the Technical Integrator cloud administrators began working to secure the root user keys, and the keys were secured within two months.

The Census Bureau has implemented a documented, standardized process to disable root user keys upon the establishment of any new GovCloud environments.

Recommendation #2 – Assess all AWS user accounts in accordance with National Institute of Standards and Technology (NIST) account management requirements and conduct periodic reviews as part of OIS assessments.

Response – *The Census Bureau agrees with this recommendation.*

In December 2018, the Census Bureau began using the tool Prowler to run scans against the GovCloud environments to assess account management. Prowler is now used to run quarterly continuous monitoring scans on all cloud instances according to the Census Bureau benchmark, which is based on the Center for Internet Security. The results of these scans are analyzed and deviations are documented in the form of accepted baseline configurations.

Recommendation #3 – Reassess, implement, and continuously monitor security baselines within all cloud environments.

Response – *The Census Bureau agrees with this recommendation.*

Security baselines are reassessed annually through the Census Information Security Continuous Monitoring (ISCM) process. Additionally, as mentioned in the response to Recommendation #2, the tool Prowler has been implemented to improve our ability to monitor continuously the security baselines within all cloud environments. The Census Bureau will update the policy for assessments to establish the frequency in which these scans will be completed.

Recommendation #4 –Perform technical assessments to validate implementation of security baselines as part of the Bureau's cloud systems' initial and ongoing assessments.

Response – *The Census Bureau agrees with this recommendation.*

The Census Bureau has improved its ability to perform technical assessments with the tool Prowler, which conducts assessments to validate the implementation of security baselines

as part of the initial and ongoing assessments. As mentioned in the response to Recommendation #3, the Census Bureau will update the policy for assessments and will also require conducting the scan during the initial Authority to Operate assessment.

Recommendation #5 – Track all Title 13 data that is stored and processed in Bureau cloud environments. This must include coordination between cloud administrators, operational staff, and OIS personnel.

Response – *The Census Bureau agrees with this recommendation.*

For the 2020 Census, the Census Bureau protects data at the highest watermark and does not differentiate the security requirements of data. We are protecting and securing all data in the entire cloud environment at the Title 13 level. In the future as technology evolves to allow greater granularity of data tracking, the Census Bureau will incorporate additional enhanced visibility into the storage and processing of already protected Title 13 data. This will include coordination among cloud administrators, operations staff, and Office of Information Security personnel, and others as needed.

Recommendation #6 – Expedite the implementation of the backup solution in progress and ensure it is followed in accordance with NIST guidance.

Response – *The Census Bureau agrees with this recommendation.*

In order to reduce risk, the Census Bureau transitioned to the enterprise tool to backup the 2020 cloud instances in July 2018. This backup approach was implemented for all servers on a rolling basis and was completed at the end of October 2018.

Recommendation #7 – Formally document and ensure the implementation of controls compensating for lack of disaster recovery planning.

Response – *The Census Bureau agrees with this recommendation.*

The Census Bureau established a Memorandum that outlines the decision not to engage in a traditional Disaster Recovery posture, which was acknowledged and approved by both the Chief Information Officer (CIO) and the Associate Director for Decennial Census Programs. This memo outlined the rationale and compensating controls for the decision. To formally document this decision, as recommended, the Census Bureau will also establish a risk acceptance memorandum for this decision.

Recommendation #8 – Develop and approve an exit strategy for all Bureau cloud systems, which includes details for completely and securely removing data from the cloud service provider.

Response – *The Census Bureau agrees with this recommendation.*

The Census Bureau will develop an exit strategy for all GovCloud systems.

If you have any questions regarding this matter, please contact me at 301-763-2135 or Kevin Smith, CIO, at 301-763-2117.

cc: Terryne Murphy, Acting Chief Information Officer, Department of Commerce
Jeffrey Jackson, Acting Chief Information Security Officer, Census Bureau
Colleen Holzbach, Program Manager for Oversight Engagement, Census Bureau
Jean McKenzie, IT Security Audit Liaison, Census Bureau

18CENS182314