Smithsonian Institution

**Office of the Inspector General**

Date June 8, 2005

To Sheila P. Burke, Deputy Secretary and Chief Operating Officer
Dennis R. Shaw, Chief Information Officer

cc Cristián Samper, Director, National Museum of Natural History

Debra S. Ritt

From Debra S. Ritt, Inspector General

Subject Management Advisory Report (05-01)
Web4 Server Failure

On February 9, 2005, a computer server operated by the Office of the Chief Information Officer (OCIO) failed, resulting in over 30 Institution websites going off-line for a period of two to six days. The server, Web4, was a legacy server that hosted Smithsonian Institution public websites as well as intranet websites and applications used by the National Museum of Natural History (NMNH), the Archives of American Art, the National Museum of the American Indian, and many others.

We conducted a review to determine the causes of the server failure and its impact on the OCIO customer community. OCIO also conducted a root cause analysis of the Web4 server failure and issued its report on March 18, 2005. Our review confirms the root causes identified in OCIO's report and provides additional information about the nature of the crash, its impact on server users, and the reasonableness of preventative measures planned by OCIO to mitigate future service interruptions and data losses.

We found the Web4 server crashed when two of its hard drives failed in succession. While the exact cause of the crash could not be determined, the age of the system and its heavy use may have been contributing factors. A mission-critical database and application used by NMNH's Research Training Program had not been backed up and were irretrievably lost, resulting in the cancellation of its 2005 program. OCIO had backed up other databases on the Web4 server but, according to OCIO staff, it was experiencing delays in backing up Web4 and other legacy servers it manages.

OCIO's root cause analysis report outlined a number of measures that it planned to take to prevent or mitigate future service disruptions and data losses on the legacy servers. However, we found that many of these measures have not been implemented as promised because of resource constraints. For example, OCIO indicated it would install an extra drive on certain servers to allow automatic rebuilding of drives when more than one fails, but OCIO staff informed us that such a measure is labor intensive and that resource constraints have limited its support of the legacy servers beyond routine maintenance and operations. For this reason, OCIO has encouraged its customer community to accelerate the upgrading of their publicly accessible websites so that they can be moved to the newer servers. The customer community reports that it lacks the resources to do so. Because resource constraints will make legacy servers a necessary component of OCIO operations for the near term, we recommend that OCIO work with its customer community to identify a strategy for timely backups of legacy servers; inform customers how website expansions affect backups; clarify when or if it will add spare drives to servers under its control; clarify how it will test the restore process; and revise the timeframe for testing the recovery plan for one of the clustered webservers. Further details on our findings are provided below.

## Nature and Impact of the Server Crash

Our review disclosed that the Web4 server crashed when two of its four hard drives failed in succession. The server was an older model that hosted some of the Institution's public websites and applications that were not in compliance with the new standards developed by OCIO.[1] The Web4 server was configured with built-in redundancy so that if one of the four hard drives failed, the data would migrate and be shared with the remaining drives until the failed drive could be replaced.[2] If a second drive failed, however, the system could not automatically rebuild itself and the data could not be shared between the two remaining drives.

When the server failed, an audio alarm sounded, alerting OCIO that there had been a serious malfunction in the secured room that housed the server. The OCIO staff member who examined the Web4 server in response to the alarm saw that the hard drive failure indicator light was on.[3] He went to his office to retrieve a replacement hard drive and on his return to the secured room saw that a second drive had failed.

The Web4 server was sent to a private company to attempt a recovery of the data. The company noted that in attempting to rebuild the system, OCIO staff had overwritten data on one of the remaining drives, inadvertently wiping out whatever data had been recorded. Although the exact cause of the drive failures could not be determined[4], the age of the server and its heavy use may have been contributing factors. The server was eight months beyond its three-year warranty period. OCIO officials told us that because they operate on a four-year replacement cycle, they accept a planned risk on expired warranties. We noted that 8 of the 31 servers managed by OCIO's Web Server Division

[1] These standards include Smithsonian Directive 920, Life Cycle Management, and OCIO's Technical Reference Model, IT-940-01.

[2] Web4 is a RAID5 server. RAID stands for "redundant array of independent disks." It is a way of storing the same data in different places on multiple hard drives.

[3] OCIO staff in the Operations Center would check the indicator lights on all of the servers in the morning. The alarm indicating a failure sounded before the Operations Center staff conducted its check.

[4] The engineers at the company could not determine the cause of the failure.

are out of warranty. Had the customer community whose operations were hosted on the Web4 server been aware that the server was out of warranty, they might have been better prepared to address the risk of the server failure with OCIO.

When the server failed, over 30 websites[5] were temporarily lost, for two to six days.[6] These sites included public web pages and services, such as the Archives of American Art and the History Wired site of the National Museum of American History, and Institution intranet sites, such as the National Museum of the American Indian's intranet and the Woodrow Wilson International Center's site.

OCIO was able to restore all of the downed websites by February 15, 2005 from backup copies maintained by OCIO staff, except for the Smithsonian Online Academic Appointments (SOLAA) database.[7] SOLAA, which provides mission-critical support to NMNH's Research Training Program, had not been backed up for approximately 18 months. Program staff ultimately located a December 2004 copy of the database from a contractor, but it was insufficient to allow the program to proceed. As a result, NMNH cancelled the 2005 Research Training Program and the planned fundraising efforts to celebrate the program's 25th anniversary. NMNH officials told us the 2006 program may also be in jeopardy if the SOLAA application is not recreated by the start of the program year.

Given that the Web4 and other legacy servers like it may have reliability problems, it is critical that adequate backups be made so that data is not irretrievably lost when a server fails. However, OCIO, which is responsible for backing up the servers over which it exercises control, mistakenly had not backed up the SOLAA database. OCIO officials also stated that they lacked the resources to make the backups without temporarily removing the system from production. To make the backups, data on the Web4 server would have to be sent to OCIO's server, which resides behind a firewall,[8] where the data is copied and retransmitted back to the Web4 server. This would have required that the SOLAA site be taken down for a short period of time. We found that OCIO performs this type of backup for other systems it operates. While there is a temporary disruption of service, OCIO posts notices to announce the backup schedule to alert users that the system will be unavailable.

Moreover, on this and other legacy servers, customers were allowed to expand their individual websites with new data and features, which resulted in more data needing to be

[5] The larger sites that were hosted on the Web4 server included Affiliations, SOLAA, Archives of American Art, HistoryWired, Smithsonian Institution Libraries, Smithsonian Press's Smithsonian Legacies, and the Woodrow Wilson International Center.

[6] All sites but one were returned to service by February 15, 2005.

[7] Data from the National Museum of the American Indian was also irretrievably lost, but it was not deemed mission-critical.

[8] A firewall is a system designed to prevent unauthorized access to or from a private network such as an intranet.

backed up. OCIO committed to its customer community on its website that it would be responsible for maintaining the equipment and backing up data under its control. However, OCIO staff told us that they had been experiencing problems meeting the backup schedules requested by its customer community. These schedules largely fell within non-production hours (8:00 p.m. to 8:00 a.m.).

Preventative Measures Planned by OCIO

In its root cause analysis report, OCIO outlined a number of measures it plans to take to prevent or otherwise mitigate future service disruptions and data losses on legacy servers. Since issuing the report, OCIO has taken many positive steps to implement these measures. For example, OCIO adopted new policies in March 2005 requiring that: (1) no data be excluded from standard backups of OCIO-maintained servers; (2) any failed hard drives removed from servers be kept pristine until data recovery has occurred; and (3) when a data loss occurs, all of the server's hard drives be sent for recovery within one day.

Moreover, OCIO is making a strong effort to communicate with and involve the customer community in preventing and mitigating such problems in the future. OCIO is now posting weekly backup status reports on the Institution's intranet (Prism) and has updated its "frequently asked questions" page on Prism to include recommendations on how to protect certain production systems. OCIO has also posted on Prism a contact list and other key information for servers that it maintains and will record and track future prevention actions through completion.

However, a number of other measures described in its report have not been or cannot be implemented, primarily because of resource constraints. For example, the report stated that OCIO will:

- Reconfigure all RAID5 servers under its control by April 1, 2005, to add a fifth drive as a "hot spare" to allow automatic rebuilding of drives in the event of two drives failing. OCIO officials told us that they presently do not have the resources to accomplish this because it would require a rebuild of each server down to the operating system level. Moreover, such a task would direct their limited resources to the legacy servers rather than to the newer technologies and the newer systems they are developing.

- Test the restore process[9] at the request of owners of systems and applications, and then have owners test and verify the accuracy of the restoration. OCIO staff noted that such a process would require exactly mirroring one system onto another system, and OCIO lacks the legacy hardware needed to perform the restoration --

[9] To restore is to copy backup files from secondary storage to hard disk to return data to its original condition if data has been damaged or to copy or move data to a new location.

hardware that is becoming obsolete and, therefore, is not worth purchasing. Moreover, members of the customer community we interviewed were not sure that if they made such a request it would be acted upon, or whether OCIO could meet the demand if several customers simultaneously requested such service. The customer community expressed these concerns at the draft stage of OCIO's report, but the final report does not address them.

- Test the recovery plan[10] for one of the clustered web servers in April 2005. As of May 9, 2005, OCIO had not done so, although the customer community told us that such an exercise was important to regain their confidence.

[10] A recovery plan consists of the precautions taken so that the effects of a disaster (e.g. loss of computers and data) will be minimized, and the organization will be able to either maintain or quickly resume mission-critical functions.

## Conclusions and Recommendations

A substantial number of legacy servers managed by OCIO are operating on expired warranties and are likely to face the same vulnerabilities as the Web4 server. OCIO has accepted server failure as an operating risk because resource constraints have limited its support of the legacy servers beyond routine maintenance and operations.

Given that additional failures of the legacy servers are likely, it is imperative that OCIO perform timely backups of customer data and applications so that they can be restored in the event of a server failure. While OCIO has identified preventative measures it will take to mitigate future service disruptions and data losses, it has not addressed how it will overcome current delays in performing timely backups of customer data on servers that it manages. It will also need to address with customers how expansion of the customers' individual websites will affect OCIO's ability to meet data backup requirements and whether controls should be imposed on such expansions.

Further, other measures aimed at providing additional drives and testing the restore and recovery process may not be implemented as promised or may transfer responsibility to the customers for services that OCIO should provide. Because legacy servers will remain an essential component of OCIO operations for the foreseeable future, OCIO will need to provide its customer community assurances that any lost data or applications on the legacy servers can be adequately recovered. To provide these assurances, we recommend that OCIO:

1. Develop a plan, in coordination with its customers that describes how it will ensure that timely backups on OCIO-maintained servers are performed.

2. Inform customers on how further website expansions on the legacy servers will affect scheduled backups and what controls should be exercised over such expansions.

3. Clarify whether RAID5 servers under its control will be reconfigured to add spare drives and, if so, develop a timeframe for completing such actions.

4. Clarify how it would test the restore process for customer applications and systems given that it lacks the legacy hardware required for such tests.

5. Provide a revised timeframe for testing the recovery plan for the clustered web server it reported it would test in April 2005.

Management Comments and Office of Inspector General Response

We discussed this report with OCIO officials, and their written comments (attached to this report) have been incorporated, as appropriate. OCIO concurred with the report's findings, conclusions, and recommendations and identified corrective actions to prevent or otherwise mitigate future service disruptions and data losses on legacy servers. By June 30, 2005, OCIO will develop a plan for web server infrastructure operations that will address timely data backups and the reconfiguration of RAID5 web servers to use hot spares. At an upcoming monthly Webmasters meeting OCIO also will discuss its plans for an automated backup solution, the resource requirements associated with website expansions, and a test of the restore process for customer applications and systems. Finally, by January 2006, after changes to enhance the redundancy of the web server infrastructure have been completed, OCIO will test the recovery plan for the clustered web server which was originally to be tested in April 2005.

OCIO's proposed actions are responsive to our recommendations, and once implemented, should address the issues raised in this report.

## Management Response

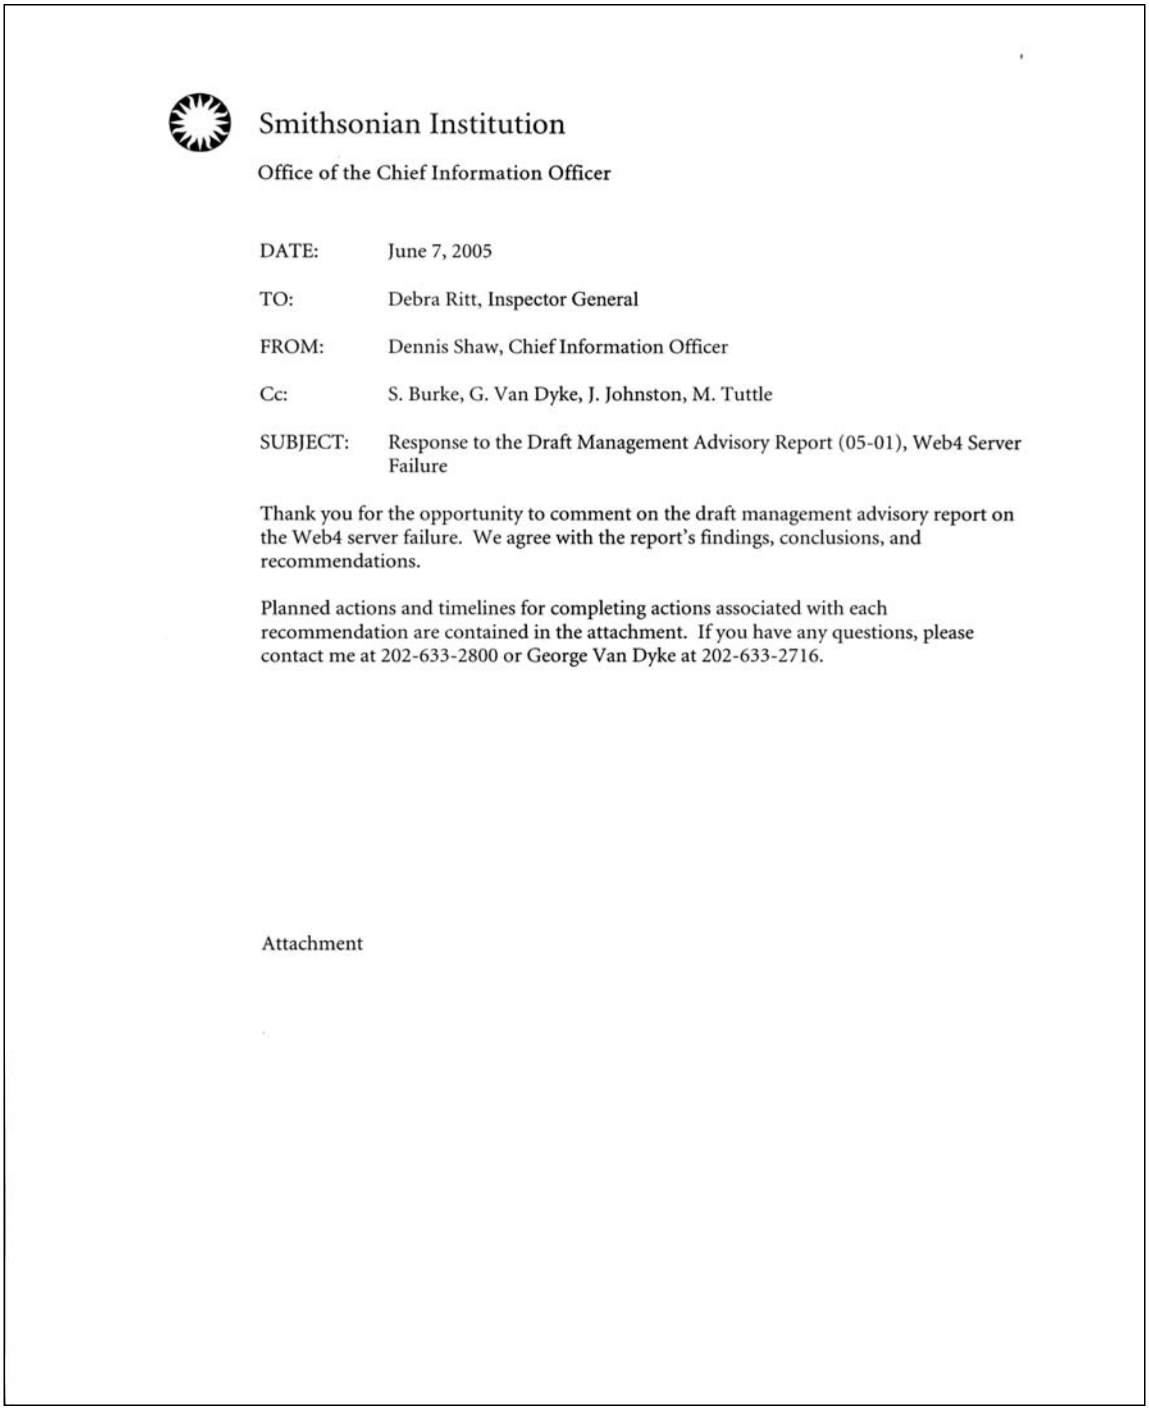

Smithsonian Institution

Office of the Chief Information Officer

DATE: June 7, 2005

TO: Debra Ritt, Inspector General

FROM: Dennis Shaw, Chief Information Officer

Cc: S. Burke, G. Van Dyke, J. Johnston, M. Tuttle

SUBJECT: Response to the Draft Management Advisory Report (05-01), Web4 Server Failure

Thank you for the opportunity to comment on the draft management advisory report on the Web4 server failure. We agree with the report's findings, conclusions, and recommendations.

Planned actions and timelines for completing actions associated with each recommendation are contained in the attachment. If you have any questions, please contact me at 202-633-2800 or George Van Dyke at 202-633-2716.

Attachment

Management Response (continued)

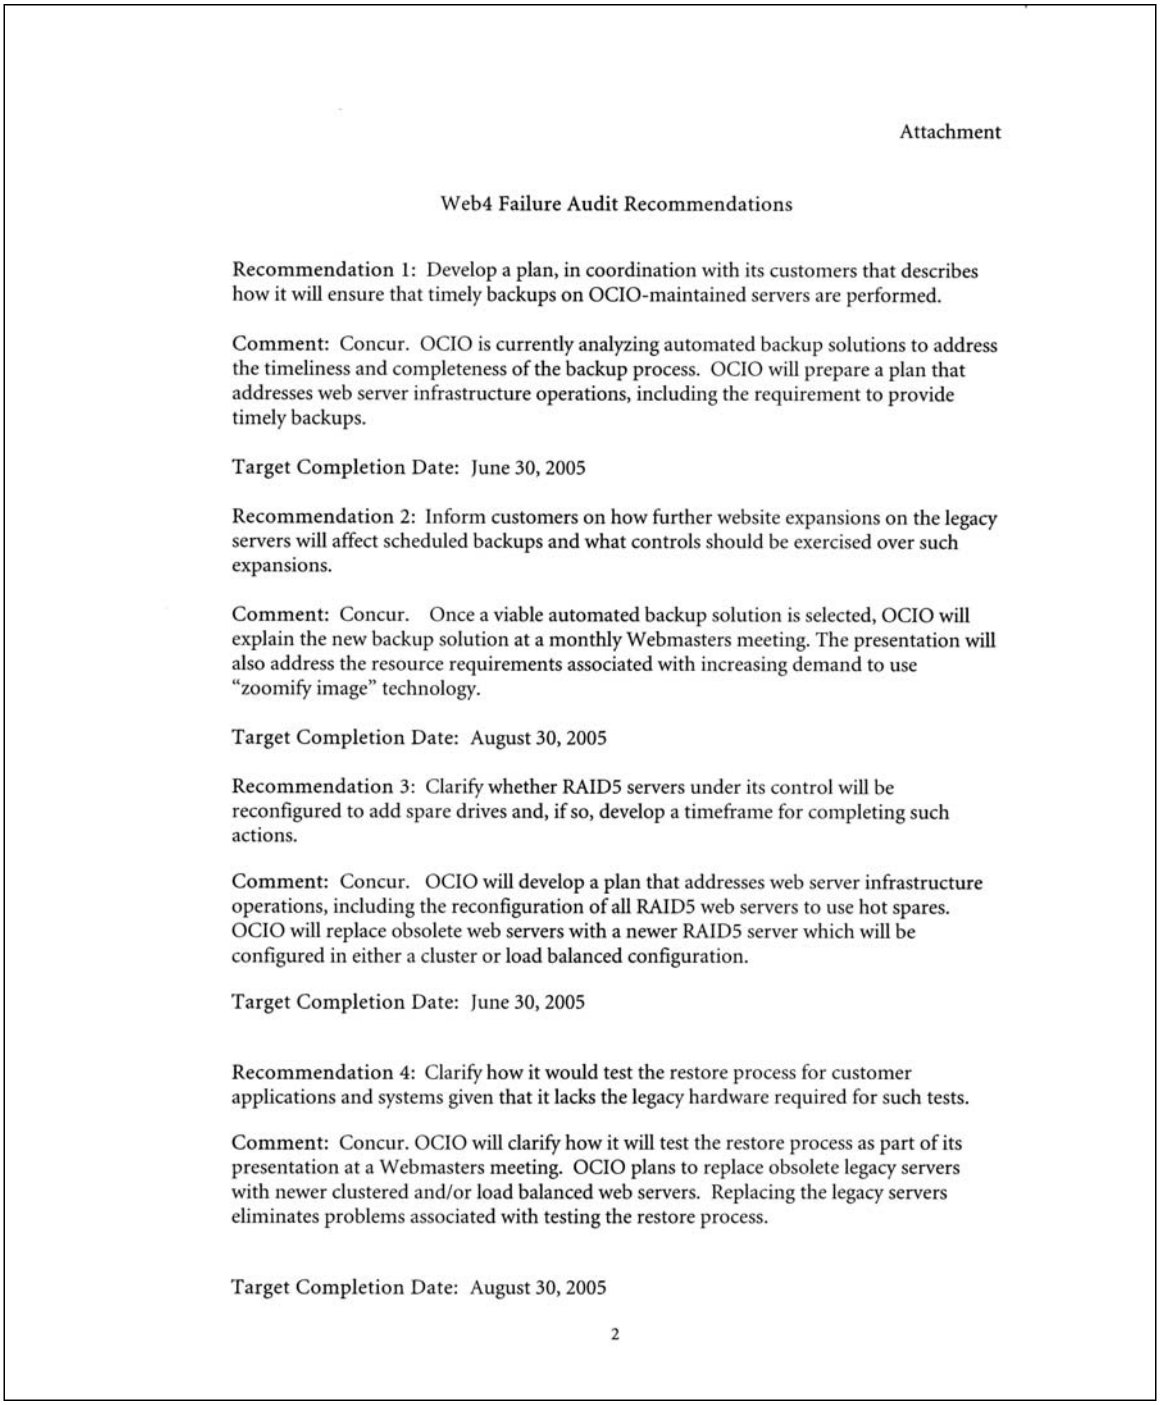

Attachment

## Web4 Failure Audit Recommendations

Recommendation 1: Develop a plan, in coordination with its customers that describes how it will ensure that timely backups on OCIO-maintained servers are performed.

Comment: Concur. OCIO is currently analyzing automated backup solutions to address the timeliness and completeness of the backup process. OCIO will prepare a plan that addresses web server infrastructure operations, including the requirement to provide timely backups.

Target Completion Date: June 30, 2005

Recommendation 2: Inform customers on how further website expansions on the legacy servers will affect scheduled backups and what controls should be exercised over such expansions.

Comment: Concur. Once a viable automated backup solution is selected, OCIO will explain the new backup solution at a monthly Webmasters meeting. The presentation will also address the resource requirements associated with increasing demand to use "zoomify image" technology.

Target Completion Date: August 30, 2005

Recommendation 3: Clarify whether RAID5 servers under its control will be reconfigured to add spare drives and, if so, develop a timeframe for completing such actions.

Comment: Concur. OCIO will develop a plan that addresses web server infrastructure operations, including the reconfiguration of all RAID5 web servers to use hot spares. OCIO will replace obsolete web servers with a newer RAID5 server which will be configured in either a cluster or load balanced configuration.

Target Completion Date: June 30, 2005

Recommendation 4: Clarify how it would test the restore process for customer applications and systems given that it lacks the legacy hardware required for such tests.

Comment: Concur. OCIO will clarify how it will test the restore process as part of its presentation at a Webmasters meeting. OCIO plans to replace obsolete legacy servers with newer clustered and/or load balanced web servers. Replacing the legacy servers eliminates problems associated with testing the restore process.

Target Completion Date: August 30, 2005

Management Response (continued)

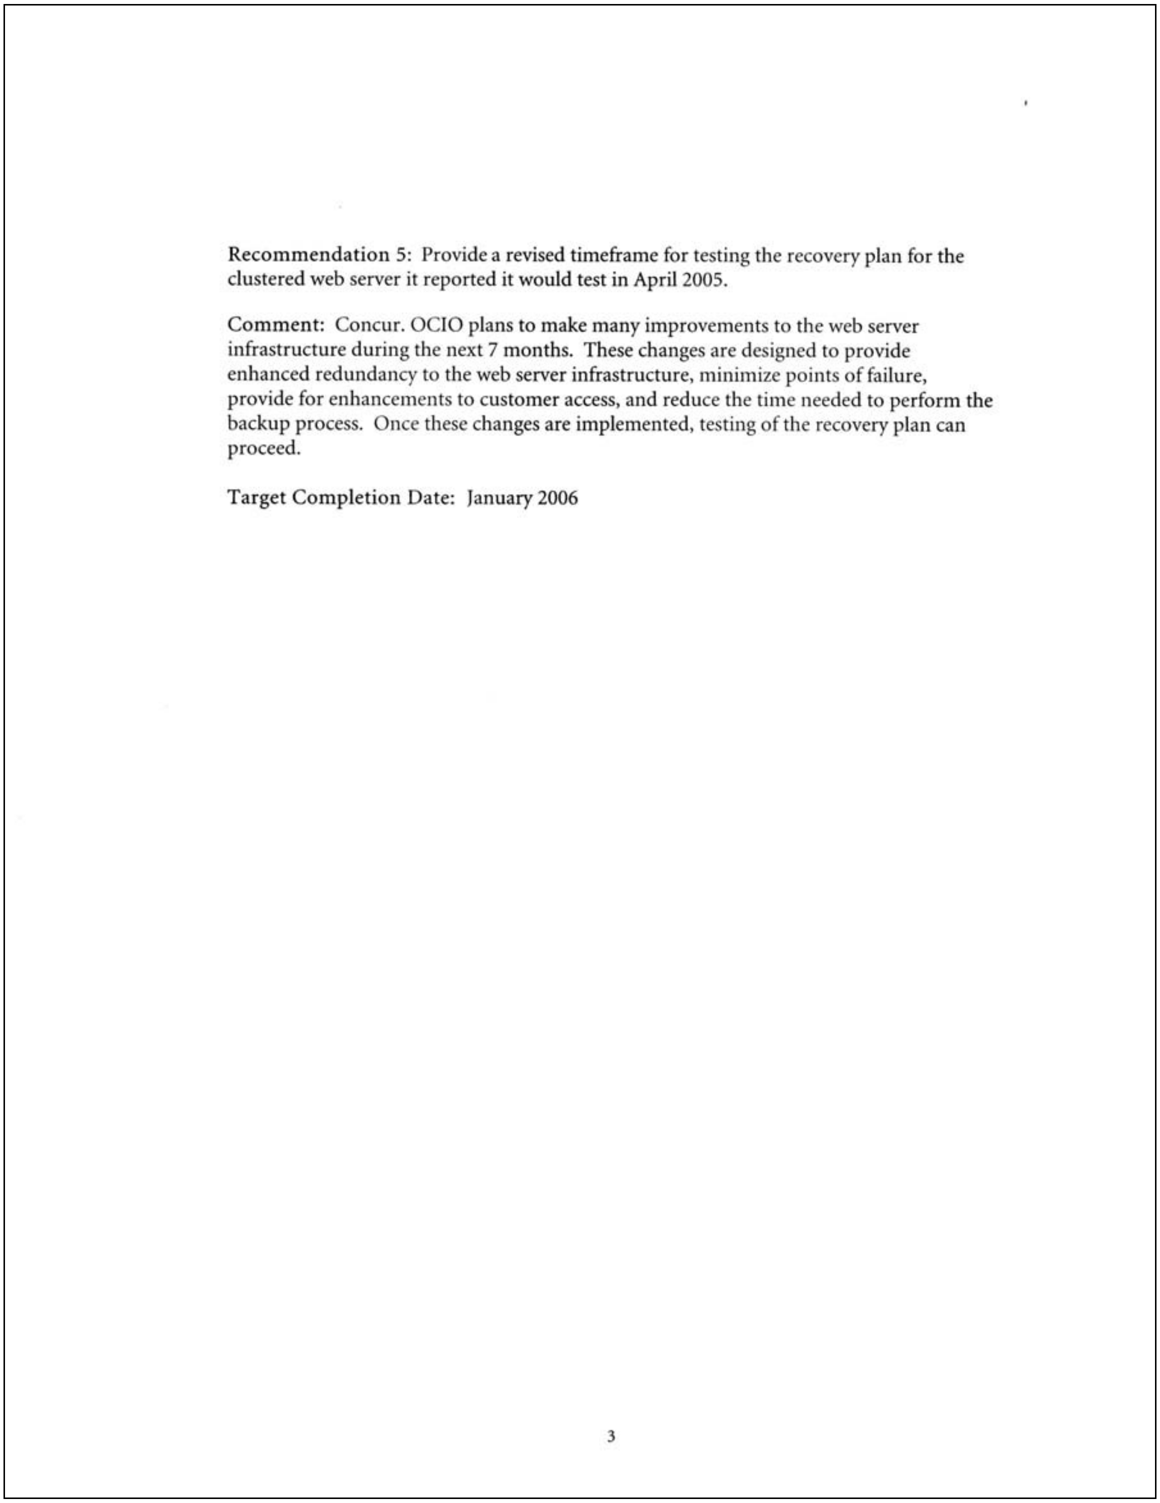

**Recommendation 5:** Provide a revised timeframe for testing the recovery plan for the clustered web server it reported it would test in April 2005.

**Comment:** Concur. OCIO plans to make many improvements to the web server infrastructure during the next 7 months. These changes are designed to provide enhanced redundancy to the web server infrastructure, minimize points of failure, provide for enhancements to customer access, and reduce the time needed to perform the backup process. Once these changes are implemented, testing of the recovery plan can proceed.

**Target Completion Date:** January 2006