

# Degraded Service Event

**Event Period:** 9/17/09 09:25 – 9/17/09 12:45  
**System(s) Affected:** Operations/Partner Test/Testbed  
**Product(s) Affected:** WIST/ECHO/Ingest/Website

## Executive Summary:

On Thursday September 17<sup>th</sup>, at 9:25am (EDT), ECHO experienced a power outage that affected operational ECHO services. This outage was the result of two separate combined events.

1. GSFC Facilities Management Department (FMD) had scheduled to perform a repair to a GSFC power system component. The ECHO team did not get notification of this planned outage and was not prepared for it.
2. In conjunction to this power outage, a component of the ECHO's redundant power system had failure on 9/11/09. Although a service appointment from the vendor had been scheduled for Wednesday 9/16, other administrative issues prevented the technician from accessing the site. The technician returned on Thursday 9/17 and was repairing the failed component at the time of the outage. Had this repair been performed on Wednesday, there would have been no ECHO outage.

The effect of the combined factors resulted in a power loss to a critical storage component for the Operational database.

The power outage lasted for approximately 30 minutes, and the associated service outage for approximately 3½ hours. After power was restored to ECHO components, Operational services remained unavailable while the ECHO system components were restored to a fully operational state, including necessary hardware and software integrity verifications. No data loss or integrity issues occurred.

## Detailed Summary:

Root causes for this outage are as follows

1. Communications breakdown between GSFC FMD and ECHO system administrators. There was insufficient ECHO subscription to power outage notifications and due to GSFC FMD POC personnel outages and recent NOMAD account changes, ECHO system administrators did not become aware of the planned outage.
2. Delay in repair of the ECHO redundant power system component. The technician did not have a valid badge to enter into center due to escort requirements and recent changes between GSFC and the vendor (Sun). As a result the repair work had to be delayed a full day.

Upon losing power to the RAID storage, the ECHO database correctly identified that its redolog storage was missing and immediately stopped performing database transactions. The ECHO Ingest, Kernel, and WIST components all retained connectivity, but were bound by the database's halted processing. When power was restored, the ECHO system administration and

database team worked to restore connectivity between the database and its redolog storage. This required a reboot of the database host and associated integrity verification.

**Timeline:**

- 09:25 (EST) – Power was lost to PDU6
- 09:25 (EST) – All ECHO systems with redundant power supplies notified ECHO sysadmins of power loss by email
- 09:37 (EST) – Nagios monitoring software alerts sysadmin and ops team of systems outage by email
- 09:57 (EST) – Email sent to echo-team mailing list informing of power outage
- 09:57 (EST) – Power restored to PDU6
- 09:57 (EST) – Systems notified sysadmins of power return via email. Single-powered systems on PDU6 return
- 12:30 (EST) – Database host boots cleanly and all zfs filesystems check out ok
- 12:45 (EST) – ECHO services are restored.

**Associated Tickets/NCRs:**

- ECHO\_SA\_TTs –
  1. 14000283 – ECHO Power Outage – Written to record the incident details.
  2. 14000275 - Failed power and cooling unit on Sun StorEdge 6120
  3. 14000285 – Adequate inclusion on B32 Tenant Contact List
- ECHO\_TTs – None
- ECHO\_NCRs – None

**Future Mitigation:**

As a part of ECHO's ongoing effort to improve availability and stability, the following actions have been identified to mitigate future outages:

1. Work with GSFC FMD for a clear understanding of how outages should be communicated. This will require a revision of who are the POCs on both sides and what email addresses or lists should be used for communications.
2. Components should be repaired quicker, commensurate with equipment impact (OPS, Development, Test) and support level (24x7) and ECHO must ensure that all service calls are properly supported by providing adequate access for service personnel to the center.
3. FMR alerts need to be fully reviewed for impact, and SA\_TT/NCR logged for visibility.