

# CERN Service Incident Report

## Description

CMSR database is used mostly by CMS computing. This database is access directly from Tier 1s and Tier 2 by Phedex agents.

The CMSR database was migrated to new hardware and new firewall rules were enabled. After the migration Phedex agents in CNAF and Wisconsin reported database errors related to transactions replay. The issue was traced to new hardware using MTU 9000 network settings (jumbo frames). As it was discovered MTU negotiation is not possible with CERN network and >1500MTU packets are dropped, resulted in Phedex clients going into hung state when executing long queries. Tier-1s not allowing jumbo frames were automatically using 1500 MTU packets passing through without issues, but bigger packets allowed by CNAF and Wisconsin were dropped at the gateway to CERN. A proxy using 1500MTU had to be put in place to work around this issue.

## Impact

Phedex data transfers where not working for 48 hours in CNAF and Wisconsin

## Timeline of the incident

Date	Action
21.03.2017 13:30	The HW migration of CMSR database has started
21.03.2017 15:21	The HW migration of CMSR database has been completed, new firewall deployed
21.03.2017 17:34	Phedex agents not responding in CNAF were reported to CMS computing operations
21.03.2017 17:51	DBA informed about the CNAF issue
21.03.2017 18:02	DBA investigation started
21.03.2017 19:35	A set workarounds for ORA-25408 issue proposed by the DBA
21.03.2017 20:13	Stuck connections to the database reported from CNAF
21.03.2017 20:20	Firewall stopped
21.03.2017 21:00	Connection details for CNAF confirmed
21.03.2017 21:30	Wisconsin agents report same issue
21.03.2017 21:30 – 22.03.2017 00:00	Ongoing investigation of the ORA-25408 issue both in CNAF and Wisconsin
22.03.2017 00:10	DBA requests tests of Phedex using a tunnel.
22.03.2017 13:50	CNAF reproduces the issue using the

	tunnel
22.03.2017 14:28	DBA Team granted access to CNAF Phedex hosts
22.03.2017 16:14	DNS misconfiguration confirmed
22.03.2017 17:02	DB connections and simple queries confirmed working for CNAF and Wisconsin both using SQL Plus and perl libraries
22.03.2017 19:20 – 23.03.2017 01:44	Phedex code confirmed working properly until a long query executed (than hangs)
22.03.2017 20:34	DNS resolution forced in Wisconsin
23.03.2017 09:41	Long query hang confirmed from CNAF
23.03.2017 09:58	Jumbo frame loss confirmed from CNAF along with 9000MTU network configuration
23.03.2017 10:08	Phedex in CNAF confirmed working after network changed to 1500 MTU
23.03.2017 10:30	CNAF hosts experience issues running 1500 MTU
23.03.2017 14:00	MTU negotiation with CERN network confirmed impossible by network experts @CERN
23.03.2017 14:55	Connection Manger proxy to connect to the CMSR database put in place
23.03.2017 15:03	CNAF agents working
23.03.2017 17:07	Wisconsin agents working

## Analysis

From the events our first suspicion of network not working as expected was covered up by the database connection working properly for short connections and the fact that small packets were transferred properly. Long part of the investigation was devoted to tracing issues related to perl connection pool or Oracle client bug.

CMSR remain the only migrated database with external network access.

The analysis was made additionally difficult by including teams working in different time zones.

## Follow up actions

A ticket has been open with network support to explain issues related to MTU negotiations.

Developments to connection manager proxy monitoring and migration of CMSR external DB clients to proxy has been given higher priority.

Tier 1 site will be warned about possible issues using 9000 MTU network for CERN connections.

## **Summary**

CMSR database used by Phedex agents was migrated to new hardware using MTU 9000. MTU negotiation is not possible with CERN network and >1500MTU packets are dropped, resulted in phedex clients going into hung state when executing long queries. A proxy using 1500MTU had to be put in place to work around this issue.