DCC Major Incident Summary Report

(Produced in accordance with Section H9 of the SEC)

Date of Incident	02/06/2020
DCC Incident Reference Number	INC000000595628
DCC Problem Reference Number	PBI000000120302
Service Impacted	All Service Requests targeting SMETS1 devices were failing. All pending and intended SMETS1 migrations were halted (a vast majority of migrations had completed with only circa 50 migrations left to complete).
Date/ Time Incident reported	02/06/2020 13:59 (Actual start of impact time)
Date &time incident resolved	02/06/2020 18:10 (Service restoration time)
Time taken to restore Service(s) (Hours)	4 hours 11 minutes
Resolution within SLA (Y/N) ***[SEC 9.14(b)]***	No

Nature of the Major Incident / Short Description

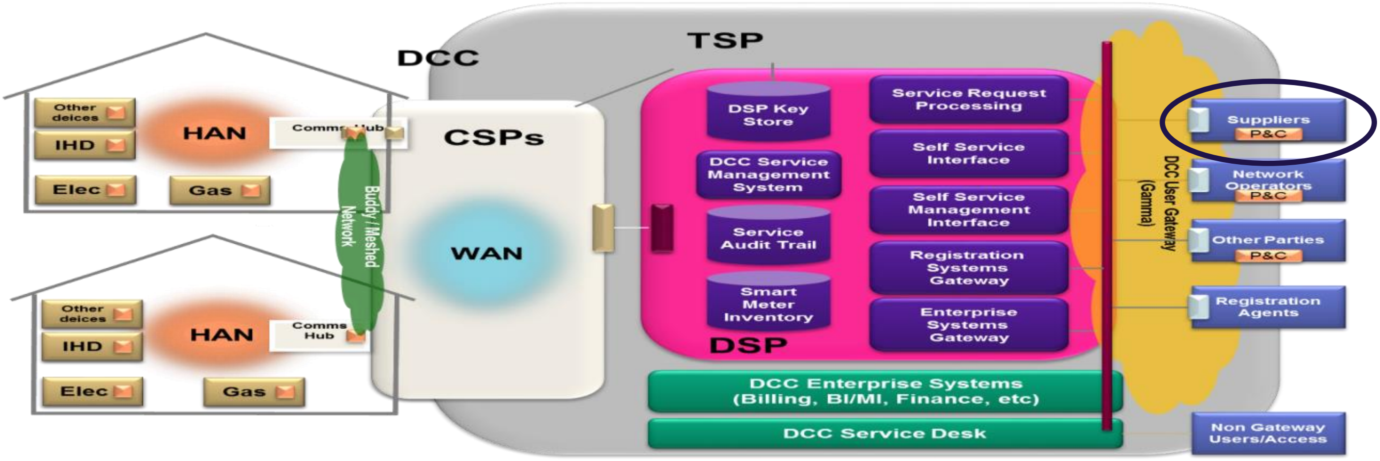

On 2 June 2020 at 14:55 the DCC Service Centre received an email from a single DCC User advising that they had not been able to communicate with any SMETS1 device since 14:00. The Service Centre engaged DCC Major Incident Management (MIM) to enquire if they were aware of any SMETS1 issues.

INC000000595628 was raised by the DCC Service Centre at 15:24, following a request by the DCC MIM to provide an incident ticket. At 16:08 this was reclassified as a Category 2.

A Technical Bridge was convened at 16:30 with DCC MIM, the Dual Control Organisation (DCO), SMETS1 Service Provider (S1SP) and Data Service Provider (DSP). On the call it was confirmed that 100% of traffic from S1SP and DSP was failing in the DCO space. This resulted in all SMETS1 Service Requests failing and migrations to stall. At 16:48 the incident was reclassified as Category 1.

DCO confirmed that the issue was a repeat of INC000000594100 (27/05/2020), whereby the primary database had fallen out of the cluster. Services failed over to the secondary database, however this was again in a read-only state. Both databases were restarted, along with all DCO system components, but this did not have a positive impact.

Further investigations identified a disk space issue on the database nodes. The primary database was extended first to rectify this issue and the same action was then completed on the other database nodes. All components were restarted however this did not resolve the issue. It was found that the application was reporting that the primary database node was in a read-only state, despite being configured in a read/write state. The decision was made to fail over services to the secondary database node.

Once the secondary database node had been configured correctly for read/write access service was restored and Service Requests could be seen to be completing as normal. A small number (circa 50) of migrations also completed, indicating that the migration service was also restored.

In order to ensure ongoing stability and resilience throughout the night, other database nodes in the cluster were configured as standby nodes and underwent full health-checks to ensure that they were suitably set-up should the need arise.

A follow up call was held 3 June at 08:30 and it was confirmed that no other issues had arisen, and the current database configuration was capable of maintaining service while root cause investigations continued into the original, affected database node.

Region / Location impacted

All SMETS1 devices across all regions were impacted.

Summary of impact / Likely future impact of the Major incident

No SMETS1 business as usual activity could be carried out. All Service Requests targeted at SMETS1 devices were failing.

Migration Control Centre (MCC) were unable to continue with any planned SMETS1 migration activity for the duration of the incident.

Once service was restored all pending migrations completed and Service Requests were completing as normal. No additional impact has been experienced as a result of this incident.

Resolving actions taken

Secondary and tertiary DCO database nodes were re-configured to ensure that they were correctly set up in read/write mode.

The secondary database node was configured as primary and the tertiary node set up as secondary. This effectively moves the initially affected primary node into a ‘non-used’ state, keeping the cluster correctly configured and ensuring full resilience.

Root Cause, if known

Root cause analysis is ongoing. However, there is strong evidence to suggest that the primary node initially ‘fell-off’ the cluster due to it running out of log file space. Log file space has been extended across all nodes while investigations continue into the cause of the log file increase.

Additional investigations are also ongoing into why application services reported primary and secondary database nodes in read-only state, when they were correctly configured in read/write mode. Initial investigations suggest that this may be due to a routing issue between the application and database services.

Table of linked incidents

Incident	Linked incident	Nature of link
INC000000595628	INC000000595460	Related
	INC000000595466	Related
	INC000000595711	Related
	INC000000595621	System Generated Alert