

What happened to cluster AZLCL01-CL on 1 October 2026

- **Cluster:** AZLCL01-CL (Azure Local, two servers: AZL-N01 and AZL-N02)
- **Date of the event:** 1 October 2026, 06:40 to 06:43 UTC
- **Based on:** the ClusterDown evidence package collected on 2 October 2026
- **Report date:** 2 October 2026

Management summary

What happened. On 1 October at 06:41 UTC, one of the two servers in the cluster (AZL-N01) briefly stopped responding. The other server (AZL-N02) took over all the work. After about 80 seconds AZL-N01 was back. When the data was collected the next day, both servers were part of the cluster again.

Was anything lost? No server crashed or restarted. The cluster did what it is designed to do: when one server stops answering, the other one carries on. For about one and a half minutes the cluster ran on a single server, and one storage area paused briefly before it continued on the other server. Whether any applications noticed this is not visible in the data we received.

Why it happened. At the moment the server stopped responding, a full memory snapshot of that server was being taken. This is a diagnostic action that someone or something has to request. While the snapshot is taken, the server is paused. Here it was paused for about 50 seconds; the cluster only waits 25 seconds before it treats a server as gone.

What we do not know yet. Who or what asked for that memory snapshot, and why. That is the real question, and the data we received does not answer it.

What we recommend.

1. Find out who or what requested the memory snapshot at 06:40 UTC on 1 October (your administrator can check this; see "What we need from you").
2. Agree that diagnostic snapshots are only taken on a server after it has been taken out of the cluster for maintenance.
3. Consider larger servers: each server has only 4 processors, which leaves little room for extra work like this.

Risk if nothing is done: low, as long as nobody takes such a snapshot on a running cluster server again. If it was triggered automatically, it can happen again, which is why point 1 matters.

For the administrator

In one paragraph

AZL-N01 did not crash. A live memory dump was being captured on it from 06:40:49 to about 06:41:41 UTC. During the capture the server was frozen, so it stopped sending its "I'm alive" signals (heartbeats) to AZL-N02. After 25 missed heartbeats AZL-N02 removed AZL-N01 from the cluster, at 06:41:41, and took over all roles. When AZL-N01 came back to life, it found that the cluster had moved on without it, so its own Cluster service stopped (06:42:12) and was started again. AZL-N01 rejoined at 06:43:02.

What happened, step by step

Time (UTC)	What happened
06:40:49	A live memory dump starts on AZL-N01. From this second on, AZL-N01 sends no more heartbeats.
06:41:00	AZL-N02 notices it is missing heartbeats from AZL-N01, on all three networks at once.
06:41:41	The dump capture finishes its last step. In the same second, AZL-N02 removes AZL-N01 from the cluster and starts taking over its work.
06:41:44 to 06:41:54	The shared storage volume "Infrastructure_1" pauses and moves to AZL-N02.
06:42:12	AZL-N01 wakes up, sees that it is no longer part of the cluster, and its Cluster service stops itself. Windows starts it again.
06:42:43	The Azure Stack HCI Orchestrator Service on AZL-N01 also stops unexpectedly.
06:43:02	AZL-N01 is back in the cluster.

Why we think this is the cause

- **The timing matches to the second.** The heartbeats stopped in the same second the dump started (06:40:49). AZL-N01 was removed in the same second the dump finished (06:41:41).
- **AZL-N01 was silent in between.** Apart from the dump's own messages, AZL-N01 logged nothing between 06:40:49 and 06:41:41.
- **AZL-N01 did not notice anything until afterwards.** Only at 06:41:55 did it report missing heartbeats, after it had "woken up". That is what a paused server looks like, not one with a broken network.

What we ruled out

Possible cause	Why we ruled it out
The server crashed	Both servers were last started on 30 September. No crash, no blue screen, no crash dump.
A network problem	All three networks failed at exactly the same moment, and no network adapter logged a problem.
A hardware fault	No hardware errors were logged on either server.
A storage problem	The storage messages only start after AZL-N01 was removed. They are a result, not the cause.

What you can check yourself

- **Was the dump requested by a person or a tool?** Look at what was running on AZL-N01 around 06:40 UTC on 1 October:
 - the PowerShell logs (PowerShell/Operational, events 4103 and 4104);
 - the Security log (event 4688, if process auditing is on);
 - any monitoring, backup or support tool that can take memory dumps.
- **Where is the dump file?** It was not in the usual place (C:\Windows\LiveKernelReports). A tool that saves it somewhere else, for example under C:\Temp, is a likely candidate.
- **Are the applications on the cluster fine?** Check whether any of them logged errors around 06:41 to 06:43 UTC.

What we need from you for a complete answer

1. Who or what started the memory dump at 06:40 UTC, and the location of the dump file.

2. What you or your team were doing on the cluster at that time.
3. If you have it: CPU and memory graphs of the two servers (and of the host they run on, since these servers are virtual machines) for 06:30 to 06:50 UTC.

What we recommend

1. **Before taking a memory dump on a cluster server, drain it first:** `Suspend-ClusterNode -Name <node> -Drain`, then take the dump, then `Resume-ClusterNode -Name <node>`. Then the cluster expects the server to be away and nothing fails over unexpectedly.
2. **Do not raise the cluster's heartbeat timeouts** to cover this. The cluster reacted correctly to a server that did not answer.
3. **Give the servers more room.** Each has 4 processors, and AZL-N01 also runs the Arc Resource Bridge virtual machine. A dump takes longer on a busy server.
4. **Turn on process auditing and PowerShell logging**, so that next time the question "who started this?" can be answered from the logs.
5. **Make the Health event log bigger.** On AZL-N02 it had already overwritten the minutes before the event.

Something else we noticed (not related)

Since 2 October, AZL-N01 from 00:43 UTC and AZL-N02 from 01:16 UTC, the servers log the same error every 30 seconds: "SDDC Management: Failed to get dedup jobs. Invalid namespace". It started 18 hours after this event and is not related to it, but it is worth looking at.

A technical version of this report, with all evidence references, is available separately (AZLCL01-CL_RCA_2026-10-01).