



CERNphone clients connectivity degraded | OTG0150847
CERNPhone service down | OTG0150866

Ihor Olkhovskyi

12/07/2024

Agenda

- Summary and impact of the incident
- Trigger and the root cause
- Detection
- Steps taken to diagnose, assess, and resolve
- Timeline and communication
- What went wrong and where we got lucky
- Followup actions
- Takeaways

Summary & Impact | OTG0150847

Summary:

- Some CERNphone clients, including mobiles, desktops, IP Phones have experienced difficulties to connect to the servers due to a memory issue in one of them.

Impact:

- In the time period 05/07/2024 20:47 - 06/07/2024 08:30 there were random difficulties to establish or receive call with CERNphone clients needing to register---so, primarily, CERNphone mobile clients. In particular, CERNphone mobile clients in the CCC were unable to make calls. TONE/CERNphone Call Centres were not affected to the main CCC number (72201) continued to work normally as did the Fire Brigade's SCR and the phone for the Security/Guards.

Root cause, Detection, Resolving

Root cause:

- Memory leak in the Frontend CERNphone component software (Kamailio)

Detection:

- Impossibility to make calls and visual loss of registration on all CERNphone endpoints, eg. mobile, desktop, IP phones

Assess and resolve:

- Checking log files and service restart

Timeline

- **05/07/2024 23:59**
TI noticed the problem, but persons listed in the contact list (Ihor Olkhovskyi, Rodrigo Sierra, Joao Fernandes (should not be there)) were not accessible
- **06/07/2024 08:25**
Ihor Olkhovskyi visually noticed issues on his personal device
- **06/07/2024 08:30**
Problem resolution, creation of OTG, etc

Luck and not so

What went wrong:

- Lack of monitoring of memory consumption. This issue could exist before the migration but went unnoticed because of a workaround regarding a Kamailio TLS problem (solved in the current version).

Where we got lucky:

- Critical services who are using call centers were not affected, but it's not luck as it was an initial design decision

Follow up actions:

- [DONE] Adding monitoring for specific log entries.
- [DONE] Workaround for automated service restart 2 times a week
- [IN PROGRESS] Investigation of a memory leak point (high suspicion on the python component of the service, as during migration the service was migrated from python2 to python3)

Takeaways

- Workaround for one issue can hide the others
- Add a stress testing in a development pipeline

Summary & Impact | OTG0150866

Summary:

- CERNphone service down

Impact:

- No possibility to make and receive calls with a CERNphone incl. mobile, desktop, IP phones. Call Centers (Fire Brigade, CCC TI, Security Gurads) were not affected.

Root cause, Detection, Resolving

Root cause:

- Server misconfiguration in LanDB

Detection:

- Happened during resolution of the other issue and was noticed immediately

Assess and resolve:

- Bound the “real” server IP to a hardware interface in LanDB

Timeline

- **07/07/2024 22:37**
Spectrum alarm indicating one of the CERNphone fronted server down (not the start of the incident)
- **08/07/2024 ~ 08:45**
TI informs TR team over the phone about one of Cp server is down
- **08/07/2024 ~ 09:50**
Rebooting affected server over IPMI
- **08/07/2024 ~ 10:00**
Loss of CERNphone services (actual incident start)
- **08/07/2024 ~ 10:08**
OTG creation
- **08/07/2024 ~ 10:30**
Identifying the core issue with NetCOM team
- **08/07/2024 ~ 10:37**
Service restoration

Detailed description

In CERNphone frontend service uses a so-called “floating IP”, that is automatically bound to a main or backup server. This IP needs to be declared as an interface in LanDB and attached to the server. But the server also has a “primary” IP, and a “floating” is secondary. In this case in LanDB “primary” IP was not bound to the hardware interface, so on DHCP renewal, the DHCP server returned a “floating” IP instead of “main” which caused IP conflict in the network segment.

Luck and not so

What went wrong:

- Interface misconfiguration in LanDB during server creation

Where we got lucky:

- It happened during the working hours, consequently the network operation team was quickly available to debug the issue together with the NetCom team.

Follow up actions:

- [IN PROGRESS] Check with the software team (CS-CT) if a check can be implemented in the network database to prevent having two (or more) fixed IP interfaces in the same network service without being bound to the corresponding hardware address.
- [IN PROGRESS] Check with the software team (CS-CT) if a check can be implemented in the network database default interface bound to hardware address automatically

Takeaway:

Check your servers config in LanDB

▼ CERNPHONE-FLOAT-FE.CERN.CH
0058-SA-SI1

Service 0058-SA-SI1
Domain **GPN**

IPv4 128.141.18.117
Subnet 255.255.255.192
Gateway 128.141.18.65

MAC Address -
Outlet Not connected to structured cabling

DNS Servers
IPv4 137.138.16.5, 137.138.17.5
IPv6 2001:1458:0201:1000:0000:0000:0000:0005,
2001:1458:0201:1100:0000:0000:0000:0005

Risky *

*means it's done on purpose

▼ CERNPHONE-M-FE.CERN.CH
0058-SA-SI1

Service 0058-SA-SI1
Domain **GPN**

IPv4 128.141.18.126
Subnet 255.255.255.192
Gateway 128.141.18.65

MAC Address 7C-C2-55-99-A6-E0
Outlet

DNS Servers

OK



home.cern