

EXAScaler Loss of ARP Network Settings for Multi-Rail on Reboot



Issue Summary

An EXAScaler server with the LNET multi-rail feature enabled may lose its Address Resolution Protocol (ARP) settings on reboot, leading to network timeouts.

ARP performs a required function in internet routing (that is, routing between devices on different networks) by mapping a dynamic 32-bit IP address (in IPv4) to a fixed 48-bit hardware address, also known as the Media Access Control (MAC) address, of a host. Multi-rail network interfaces add complexity to this task. When multi-rail is enabled, the same EXAScaler file system can serve two parallel streams of data over two ports on the same InfiniBand subnet. Although each port has its own IP address and MAC address, the EXAScaler file system may respond to an ARP request received on one port with an ARP reply sent from the other. This is known as the ARP flux problem, and it is resolved by the proper ARP configuration settings.

However, in EXAScaler versions 5.0.1 through 5.2.1, the necessary ARP configuration settings may be lost on server reboot. The ARP flux problem can then emerge and cause timeouts that impair system performance.

More information about the EXAScaler LNET multi-rail feature can be found in Section 13.9, "Multi-Rail LNet," in the *EXAScaler Installation & Administration Guide*.

Affected Products

EXAScaler versions 5.0.1, 5.1.0, 5.1.1, 5.2.0 and 5.2.1 are affected by this issue.

Resolution

This issue is resolved in EXAScaler versions 5.2.2 and later. Affected customers should contact DDN Support at their earliest convenience to obtain an upgrade plan.

Workaround

If you are experiencing ARP-related network timeouts but cannot upgrade right away, add the recommended ARP configuration settings to the EXAScaler configuration file as follows.

1. Back up the existing EXAScaler configuration file.

```
cp -vp /etc/ddn/exascalr.conf /etc/ddn/exascalr.conf.BAK
```

2. Use a text editor to edit the configuration file `/etc/ddn/exascalr.conf`. At the end of the `[sysctl_defaults]` section, append the following ARP configuration entries:

```
net.ipv4.conf.all.accept_local: 1
net.ipv4.conf.all.arp_announce: 2
```

```
net.ipv4.conf.all.arp_filter: 0
net.ipv4.conf.all.arp_ignore: 1
net.ipv4.conf.all.rp_filter: 0
net.ipv4.conf.default.accept_local: 1
net.ipv4.conf.default.arp_announce: 2
net.ipv4.conf.default.arp_filter: 0
net.ipv4.conf.default.arp_ignore: 1
net.ipv4.conf.default.rp_filter: 0
```

3. Distribute the edited configuration file to all nodes in the cluster.

```
sync-file /etc/ddn/exascalr.conf
```

4. Use the `es_install` command to update the `sysctl` configuration. The `clush` utility may be used to update all nodes in a cluster at once.

```
clush -a "es_install --yes --steps os"
```

Contacting DDN Technical Support

Please contact DDN Technical Support at any time if you have questions or need assistance. Support can be reached online, by email, or by phone as listed below.

Web

DDN Community Support Portal <https://community.ddn.com/login>
Portal Assistance webportal.support@ddn.com

Email

Support Email support@ddn.com

Telephone

DDN Support Worldwide Directory <https://www.ddn.com/support/global-services-overview/>

Bulletins & Notices

Support Bulletins <http://www.ddn.com/support/technical-support-bulletins>
End-of-Life Notices <http://www.ddn.com/support/end-of-life-notices>
Release Notes <https://community.ddn.com/login>
Subscription Requests support-tsb@ddn.com