

## Ceph - Bug #43417

Since the local loopback address is set to a virtual IP, OSD can't restart .

12/24/2019 03:33 AM - David Lee

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	Normal	<b>Spent time:</b>	0.00 hour
<b>Assignee:</b>			
<b>Category:</b>	common		
<b>Target version:</b>			
<b>Source:</b>	Development	<b>Affected Versions:</b>	v14.2.4
<b>Tags:</b>		<b>ceph-qa-suite:</b>	
<b>Backport:</b>	nautilus	<b>Pull request ID:</b>	32420
<b>Regression:</b>	No	<b>Crash signature (v1):</b>	
<b>Severity:</b>	3 - minor	<b>Crash signature (v2):</b>	
<b>Reviewed:</b>			

### Description

I set a local loopback ip on the same network segment as the cluster, like lo:0. The network configuration is as follows

```
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.211.42 netmask 255.255.255.0 broadcast 192.168.211.255
    inet6 fe80::ea61:1fff:fe16:e7b7 prefixlen 64 scopeid 0x20<link>
    ether e8:61:1f:16:e7:b7 txqueuelen 1000 (Ethernet)
    RX packets 2845662864 bytes 2454056684356 (2.2 TiB)
    RX errors 0 dropped 8338 overruns 0 frame 0
    TX packets 2573756982 bytes 2075412304729 (1.8 TiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
    device memory 0xc5800000-c5ffff
lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
    inet 127.0.0.1 netmask 255.0.0.0
    inet6 ::1 prefixlen 128 scopeid 0x10<host>
    loop txqueuelen 1 (Local Loopback)
    RX packets 61609662 bytes 155594418449 (144.9 GiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 61609662 bytes 155594418449 (144.9 GiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
lo:0: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
    inet 192.168.211.200 netmask 255.255.255.255
    loop txqueuelen 1 (Local Loopback)
```

The ceph configure is as follows:

```
[global]
mon_initial_members = 172e18e211e**, 172e18e211e**, 172e18e211e**
mon_host = 192.168.211.**, 192.168.211.**, 192.168.211.***
auth_cluster_required = cephx
auth_service_required = cephx
auth_client_required = cephx
public network = 192.168.211.0/24
```

The the problem is :

```
[root@172e18e211e42 ~]# systemctl status ceph-osd@8
● ceph-osd@8.service - Ceph object storage daemon osd.8
```

```
Loaded: loaded (/usr/lib/systemd/system/ceph-osd@.service; enabled; vendor preset: disabled)
Active: active (running) since Tue 2019-12-24 09:33:42 CST; 12s ago
Process: 1553818 ExecStartPre=/usr/lib/ceph/ceph-osd-prestart.sh --cluster ${CLUSTER} --id %i (code=exited, status=0/SUCCESS)
Main PID: 1553823 (ceph-osd)
CGroup: /system.slice/system-ceph\x2dosd.slice/ceph-osd@8.service
└─1553823 /usr/bin/ceph-osd -f --cluster ceph --id 8 --setuser ceph --setgroup ceph
```

```
Dec 24 09:33:42 172e18e211e42 systemd[1]: Starting Ceph object storage daemon osd.8...
Dec 24 09:33:42 172e18e211e42 systemd[1]: Started Ceph object storage daemon osd.8.
Dec 24 09:33:42 172e18e211e42 ceph-osd[1553823]: 2019-12-24 09:33:42.981 7f6304c02d80 -1 Falling back to public interface
Dec 24 09:33:45 172e18e211e42 ceph-osd[1553823]: 2019-12-24 09:33:45.820 7f6304c02d80 -1 osd.8 9095 log_to_monitors {default=true}
Dec 24 09:33:45 172e18e211e42 ceph-osd[1553823]: 2019-12-24 09:33:45.849 7f62f722a700 -1 osd.8 9095 set_numa_affinity unable to identify public interface 'lo:0' numa node: (2) No such file or directory
Dec 24 09:33:51 172e18e211e42 ceph-osd[1553823]: 2019-12-24 09:33:51.758 7f62f722a700 -1 osd.8 9131 set_numa_affinity unable to identify public interface 'lo:0' numa node: (2) No such file or directory
```

The the osd can't restart . And I pinpoint the problem at /src/common/ipaddr.cc ,

```
const struct ifaddrs *find_ipv4_in_subnet(const struct ifaddrs *addrs,
                                         const struct sockaddr_in *net,
                                         unsigned int prefix_len,
                                         int numa_node)
```

can't get the right ip.

#### Related issues:

Related to Ceph - Bug #48893: Ceph-osd refuses to bind on an IP on the local ...	Resolved
Related to Ceph - Bug #49938: daemons bind to loopback iface	Resolved
Copied to Ceph - Backport #49202: nautilus: Since the local loopback address ...	Resolved

#### History

##### #1 - 12/25/2019 09:53 AM - David Lee

PR:<https://github.com/ceph/ceph/pull/32420>

##### #2 - 12/26/2019 09:59 AM - Kefu Chai

- Status changed from New to Resolved

- Pull request ID set to 32420

##### #3 - 01/15/2021 02:40 PM - Kefu Chai

- Related to Bug #48893: Ceph-osd refuses to bind on an IP on the local loopback lo added

##### #4 - 02/07/2021 12:54 PM - Nathan Cutler

- Status changed from Resolved to Pending Backport

- Target version deleted (v14.2.5)

- Backport set to nautilus

##### #5 - 02/07/2021 12:55 PM - Nathan Cutler

- Copied to Backport #49202: nautilus: Since the local loopback address is set to a virtual IP, OSD can't restart . added

**#6 - 03/09/2021 08:35 AM - Nathan Cutler**

- *Status changed from Pending Backport to Resolved*

While running with --resolve-parent, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved" or "Rejected".

**#7 - 03/31/2021 06:54 AM - Nathan Cutler**

- *Related to Bug #49938: daemons bind to loopback iface added*

**#8 - 04/22/2021 02:36 PM - Kefu Chai**

- *Description updated*