

mgr - Bug #43313

telemetry module can crash on entity name with multiple '.' separators

12/14/2019 02:49 AM - Dan Mick

Status: Resolved	% Done: 0%
Priority: Urgent	
Assignee: Dan Mick	
Category:	
Target version:	
Source: Community (user)	Reviewed:
Tags:	Affected Versions:
Backport: nautilus,mimic	ceph-qa-suite:
Regression: No	Pull request ID: 33094
Severity: 3 - minor	Crash signature:
Description	
A user discovered this:	
<pre># ceph telemetry show Error EINVAL: Traceback (most recent call last): File "/usr/share/ceph/mgr/mgr_module.py", line 914, in _handle_command return self.handle_command(inbuf, cmd) File "/usr/share/ceph/mgr/telemetry/module.py", line 754, in handle_command channels=command.get('channels', None) File "/usr/share/ceph/mgr/telemetry/module.py", line 660, in compile_report report['crashes'] = self.gather_crashinfo() File "/usr/share/ceph/mgr/telemetry/module.py", line 337, in gather_crashinfo (etype, eid) = c.get('entity_name', '').split('.') ValueError: too many values to unpack</pre>	
The entity name in question was	
<pre>client.rgw.stage2-rgw02-object02.s2r6</pre>	
, which indeed won't split correctly. The fix is to add a max-split count of 1 so that only the first '.' is considered a separator.	
Related issues:	
Copied to mgr - Backport #44057: nautilus: telemetry module can crash on enti...	Resolved
Copied to mgr - Backport #44058: mimic: telemetry module can crash on entity ...	Rejected

History

#1 - 12/14/2019 02:49 AM - Dan Mick

- Project changed from Ceph to mgr

#2 - 01/14/2020 11:46 PM - Alexandre Marangone

Dan, I'm looking into this. I'm not certain on how you expect the crash to look on the server side.

It would be easy to specify a maxsplit of 1 but then the reported entity_name in the crash would be client.\$hash. Would it be more useful to get client.<client_type>.\$hash?

Also it looks like any radosgw crash would trigger this since the expectation is that there's only 1 dot in the name and the default naming convention for rgw is client.rgw.`hostname`-s`

#3 - 01/15/2020 04:15 AM - Dan Mick

with maxsplit 1:

```
>>> s='client.rgw.stage2-rgw02-object02.s2r6'  
>>> s.split('.', 1)  
['client', 'rgw.stage2-rgw02-object02.s2r6']
```

which is what you want, I think. As for the hash, the point of the hash is to obscure any customer identifiers like hostname or zone name or whatever is stuffed into the name.

#4 - 01/15/2020 04:49 AM - Alexandre Marangone

Gotcha. I thought it could have been important to identify client type (i.e. rgw) in the entity name

#5 - 02/06/2020 03:21 AM - Sage Weil

- Priority changed from Normal to Urgent
- Backport set to nautilus,mimic

#6 - 02/07/2020 04:53 AM - Dan Mick

- Status changed from In Progress to Fix Under Review
- Pull request ID set to 33094

#7 - 02/08/2020 02:08 PM - Sage Weil

- Status changed from Fix Under Review to Pending Backport

#8 - 02/10/2020 05:48 AM - Konstantin Shalygin

- Copied to Backport #44057: nautilus: telemetry module can crash on entity name with multiple '.' separators added

#9 - 02/10/2020 05:48 AM - Konstantin Shalygin

- Copied to Backport #44058: mimic: telemetry module can crash on entity name with multiple '.' separators added

#10 - 01/27/2021 08:03 PM - 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".