Contents

Introduction

This article describes the apparent false trigger of the ThreshDNSLookupFailure trap when an Service Redundancy Protocol (SRP) connection bounce occurs on an SRP standby node. Infrastructure Domain Name Service (DNS) is used on various nodes in the Long Term Evolution (LTE) network indirectly as part of the call setup process. On a Packet Data Network Gateway (PGW) it can be used to resolve any Fully Qualified Domain Names (FQDNs) returned in S6b authentication, as well as to resolve FQDNs specified as peers in the various Diameter endpoint configurations. If DNS timeouts (failures) occur on an active node processing calls, then this can negatively affect call setups depending on what components rely on the DNS functioning properly.

Problem

Starting in StarOS v15 there is a configurable threshold to measure infrastructure DNS failure rate. In the case where the PGW is implemented with Inter-Chassis Session Recovery (ICSR), there is the likelihood that if the SRP connection between both nodes goes down for whatever reason, and the ensuing Standby node goes into pending Active state (but not fully active because the other node remains fully SRP active assuming no other issues), then the associated DNS alarm/trap is triggered. This is because in pending active state, the node attempts to establish the various diameter connections for the various diameter interfaces in the ingress context in preparation of potentially becoming fully SRP active. If the configuration for ANY of the diameter connections is based on specifying peers in the endpoint configuration that are FQDNs instead of IP addresses, then those peers need to be resolved via DNS with A (IPv4) or AAAA (IPv6) queries. Since the node is in pending active state, such queries ALL FAIL because the responses to the requests will be routed to the active node (which will drop the responses), which results in 100% failure rate which in turn causes the alarm/trap to be triggered. While this is expected behavior in this scenario, the potential result is an opened customer ticket regarding the significance of the alarm.

Here is an example of such an alarm where Diameter Rf is configured with FQDNs and therefore requires DNS to resolve. Shown is an FQDN that needs to be resolved by DNS.

The SRP connection goes down for some reason (external to the pair of PGW nodes and the reason not important for the purposes of this example) for 7+ minutes, and the SNMP trap ThreshDNSLookupFailure triggers.

```
Tue Nov 25 08:43:42 2014 Internal trap notification 1037 (SRPConnDown)
vpn SRP ipaddr 10.211.220.100 rtmod 3 Tue Nov 25 08:43:42 2014 Internal trap notification 120
(SRPActive)
vpn SRP ipaddr 10.211.208.165 rtmod 3 Tue Nov 25 08:51:14 2014 Internal trap notification 1038
(SRPConnUp)
vpn SRP ipaddr 10.211.220.100 rtmod 3 Tue Nov 25 08:51:14 2014 Internal trap notification 121
(SRPStandby)
vpn SRP ipaddr 10.211.208.165 rtmod 9 Tue Nov 25 09:00:08 2014 Internal trap notification 480
```

```
(ThreshDnsLookupFailure)
context "XGWin" threshold 5% measured value 12%
```

Here is the alarm and associated log:

```
[local]XGW> show alarm outstanding verbose

Severity Object          Timestamp                         Alarm ID
-------- ----------      --------------------------------- ------------------
  Alarm Details
------------------------------------------------------------------------------
Minor    VPN XGWin       Tuesday November 25 09:00:0       3611583935317278720
<111:dns-lookup-failure> has reached or exceeded the configured threshold <5%>,
the measured value is <12%>. It is detected at <Context [XGWin]>.

2014-Nov-25+09:00:08.939 [alarmctrl 65201 info]
[5/0/6050 <evlogd:0> alarmctrl.c:192] [context: XGWin, contextID: 6] [software internal system
critical-info syslog] Alarm condition: id 321eec7445180000 (Minor):
<111:dns-lookup-failure> has reached
 or exceeded the configured threshold <5%>, the measured value is <12%>.
It is detected at <Context [XGWin]>.
```

Bulkstats confirms 100% failure for Primary and Secondary AAAA DNS queries attempting to resolve Diameter Rf peers:

%time%	%dns-central-aaaa-atmpts%	%dns-primary-ns-aaaa-atmpts%	%dns-primary-ns-aaaa-fails%	%dns-primary-ns-query-timeouts%	%dns-secondary-ns-aaaa-atmpts%	%dns-secondary-ns-aaaa-fails%	%dns-secondar query-timeouts%
08:32:00	16108	16098	10	10	10	0	0
08:34:00	16108	16098	10	10	10	0	0
08:36:00	16108	16098	10	10	10	0	0
08:38:00	16108	16098	10	10	10	0	0
08:40:00	16108	16098	10	10	10	0	0
08:42:00	16108	16098	10	10	10	0	0
08:44:00	16236	16162	74	74	74	64	64
08:46:00	16828	16466	362	362	362	352	352
08:48:00	17436	16770	666	666	666	656	656
08:50:00	18012	17058	954	954	954	944	944
08:52:00	18412	17250	1162	1162	1162	1152	1152
08:54:00	18412	17250	1162	1162	1162	1152	1152
08:56:00	18412	17250	1162	1162	1162	1152	1152

Solution

This trap/alarm can be ignored and cleared since the node is not truly SRP active and not handling any traffic. Note the failure rate in the example above is much lower than the expected 100% and bug CSCuu60841 has now fixed that issue in a future release so that it will always report 100%.

clear alarm outstanding

OR

To just clear that particular alarm:

clear alarm id <alarm id>

Another twist of this issue can occur on a newly SRP Standby chassis after an SRP switchover has taken place. The alarm should be ignored in that scenario also since the chassis is SRP Standby and DNS failures are therefore irrelevant.

Finally, it goes without saying that the cause for this alarm needs to be immediately investigated on a truly SRP active PGW, as subscriber or billing impact will likely occur depending on what types of FQDNs are attempting to be resolved.