

CRS-1 路由器FABRIC报错处理案例

目录

[硬件平台](#)
[软件版本](#)
[案例简介](#)
[故障排除思路](#)
[故障诊断步骤](#)
[经验总结](#)
[相关命令](#)

硬件平台

CRS-1 多框路由器

软件版本

以IOS XR 3.6.3 举例说明

案例简介

在CRS-1多框路由器的日常维护过程中，我们可能会看到设备日志中有这样的告警：

```
LC/1/5/CPU0:Dec 3 04:37:05 : fabricq_mgr[136]: %FABRIC-FABRICQ-3-PCL_PKT : Minor
error in
PCL of fabricq asic 0. PCL UC Partial Packet: CAOPCI: 0x70 (1/8, UC, LO)
LC/0/1/CPU0:Dec 3 04:37:05 : fabricq_mgr[136]: %FABRIC-FABRICQ-3-PCL_PKT : Minor
error in PCL of fabricq asic 1. PCL UC Partial Packet: CAOPCI: 0x74 (1/9, UC, LO)
LC/0/13/CPU0:Dec 3 04:37:05 : fabricq_mgr[136]: %FABRIC-FABRICQ-3-PCL_PKT : Minor
error in PCL of fabricq asic 1. PCL UC Partial Packet: CAOPCI: 0x70 (1/8, UC, LO)
```

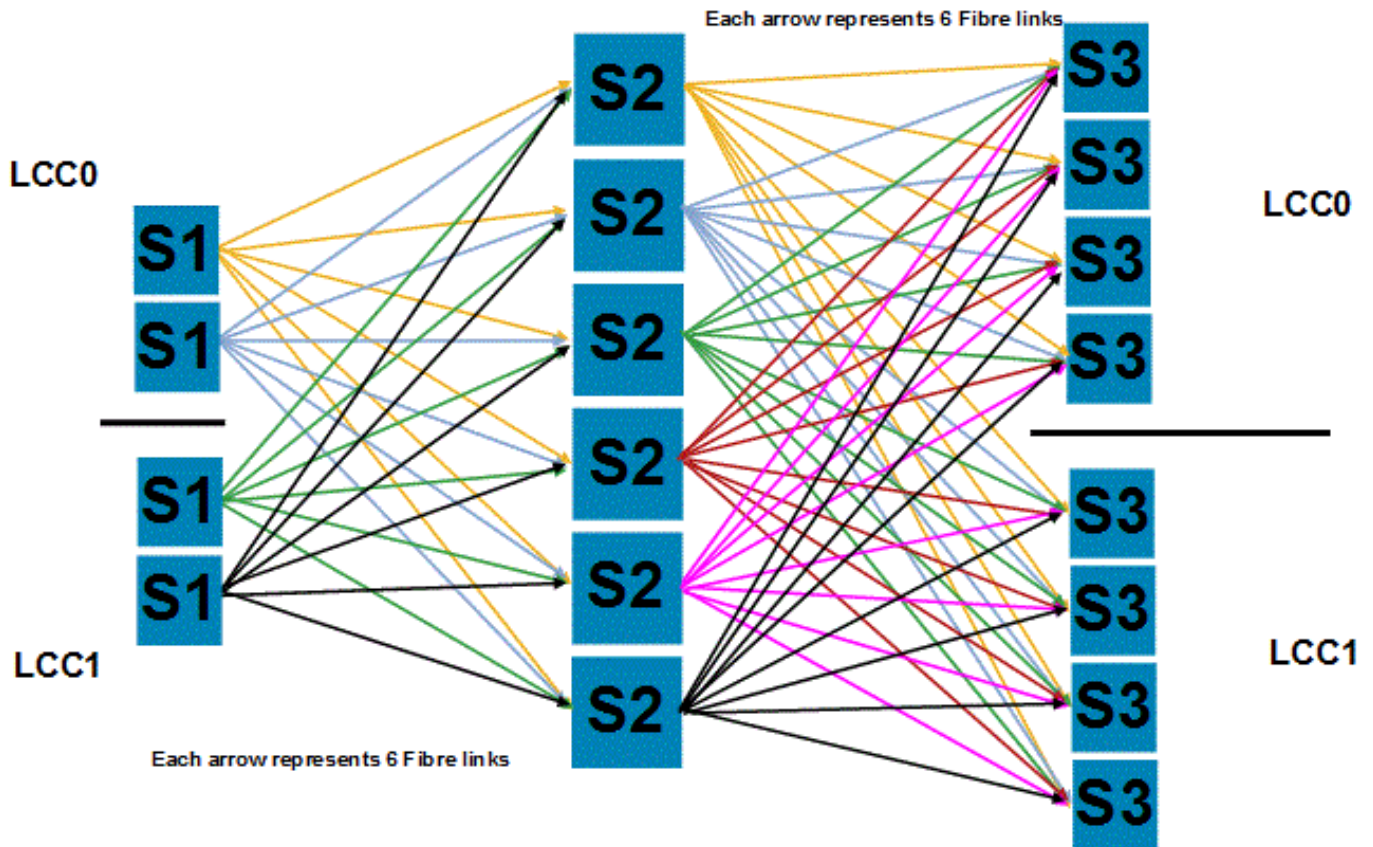
同时，还可能伴随有少量丢包的现象。

下面，我们将讨论一下这种情况的处理。

故障排除思路

首先，我们需要知道Fabric是怎么工作的。CRS-1路由器的包转发是由FABRIC来实现的，Fabric的包转发有3个阶段：S1, S2, S3。具体到多框的环境下，S1和S3是通过在LCC上的S13卡来实现的，S2是通过FCC上的S2卡来实现的。

Fabric connectivity (per plane)



S13与S2卡是通过fabric光缆相接的。每个fabric光缆接口包括六组共72根独立的小光纤（如下图所示）。如果在过程中不注意，有可能使光纤头进灰。或者因为安装时封口不严，在使用过程中导致微量着尘。这样，有可能对海量数据高速转发时出错带来隐患。这就导致了我们的开篇提出的问题。



回顾一下告警信息：

```
LC/1/5/CPU0:Dec 3 04:37:05 : fabricq_mgr[136]: %FABRIC-FABRICQ-3-PCL_PKT :  
Minor error in PCL of fabricq asic 0. PCL UC Partial Packet: CAOPCI: 0x70 (1/8, UC,  
L0)
```

这条告警告诉我们，从板卡(MSC)的fabricq asic 收到了错误。MSC的fabricq asic是与fabric card 相连的芯片，与S3芯片直接相接。这个错误可能是在S1，S2，和S3之中的任何一个阶段产生的，需要逐段排查。

故障诊断步骤

以下数据来自真实网络环境。为保护客户资料，隐去敏感信息，同时不影响故障排查示例。

一般说来，接收端对错误的探测更为敏感。我们常常从接收端查起。查看s1rx, s2rx, s3rx。在本例中，我们可以看到在s2rx的几条fabric link 探测到了错误。以下略去对s1rx, s3rx, 以及发送端的排查输出。

```
RP/0/RP0/CPU0:CRS(admin)#show controllers fabric link port s2rx all statistics |  
exclude 0.*0.*0
```

Total racks: 4

Rack 0:

SFE	Port	In	In	CE	UCE	PE
R/S/M/A/P		Data Cells	Idle Cells	Cells	Cells	Cells

Rack 1:

SFE	Port	In	In	CE	UCE	PE
R/S/M/A/P		Data Cells	Idle Cells	Cells	Cells	Cells

Rack F0:

SFE	Port	In	In	CE	UCE	PE
R/S/M/A/P		Data Cells	Idle Cells	Cells	Cells	Cells
F0/SM21/SP/2/34		98537181536	448554293397	273	12	0

Rack F1:

SFE	Port	In	In	CE	UCE	PE
R/S/M/A/P		Data Cells	Idle Cells	Cells	Cells	Cells
F1/SM12/SP/1/23		194837049246	22429631246318	216	22	0
F1/SM12/SP/3/23		177896462986	21951736335508	89	12	0
F1/SM12/SP/4/22		1214039534516	18732861653988	152	8	0

我们主要关注的是不可修复的错误 UCE, 这种错误有可能与物理问题相关。

Correctable Error (CE) – A cell with an error that was detected via the Forward Error Correction (FEC) code and is fixed.

Uncorrectable Error (UCE) – A cell with an error that was detected via the FEC code and was not able to be fixed.

我们看到报错的芯片在 F1/SM12/SP，由此我们可以进一步看看是什么方面的错误。我们知道S2板卡有6块芯片，于是我们一一检查：

```
RP/0/RP0/CPU0:CRS(admin)#show asic-errors s2 0 all location F1/SM12/SP
```

```

*****
*                               Single Bit Errors                               *
*****
*****
*                               Multiple Bit Errors                             *
*****
*****
*                               Parity Errors                                   *
*****
*****
*                               CRC Errors                                     *
*****
*****
*                               Generic Errors                                 *
*****

```

```

Name           : QRL_RS_THRSH_ERROR-GENERIC
Node Key       : 0x1050037
Thresh/period(s) : 2/172800      Alarm state: OFF
Error count    : 32
Last clearing  : Fri Nov 6 00:12:59 2009
Last N errors  : 32

```

First N errors.

@Time, Error-Data

```

-----
Nov 6 00:12:59.664: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 6 14:44:24.546: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 7 02:02:52.482: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 10 14:03:18.649: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 11 21:36:10.289: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Nov 12 16:26:09.211: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 14 20:51:17.168: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 14 21:27:53.209: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Nov 15 01:57:18.119: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 18 15:07:13.375: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS

```

error
count exceeded
Nov 19 22:26:28.606: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 21 08:22:27.709: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 22 04:46:49.269: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Nov 22 04:46:49.270: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 23 10:25:58.324: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Nov 23 10:42:26.323: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 23 20:39:52.038: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 23 22:04:14.612: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Nov 24 17:40:13.150: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 25 08:19:01.483: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Nov 25 10:56:10.571: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 26 00:33:31.008: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Nov 26 12:14:30.236: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 27 12:14:39.284: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 27 20:36:39.892: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Last N errors.
@Time, Error-Data

Nov 28 10:50:33.219: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded

```

Nov 28 13:33:45.543: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 28 20:01:04.387: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 28 22:53:11.078: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Nov 28 22:53:11.080: Link: s2rx/F1/SM12/SP/0/23, qrl: 5, link: 2, Uncorrectable RS
error
count exceeded
Nov 29 14:05:43.246: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
Nov 29 16:56:33.135: Link: s2rx/F1/SM12/SP/0/22, qrl: 4, link: 2, Uncorrectable RS
error
count exceeded
-----
Name          : SLOW_FLAP_ERR-GENERIC
Node Key      : 0x1050060
Thresh/period(s): 1/0   Alarm state: OFF
Error count   : 2
Last clearing : Sat Nov 28 22:54:27 2009
Last N errors : 2
-----
First N errors.
@Time, Error-Data
-----
Nov 28 22:54:27.963: s2rx/F1/SM12/SP/0/23 flaps slowly
Nov 29 16:57:38.111: s2rx/F1/SM12/SP/0/22 flaps slowly
-----
*****
*                   ASIC Reset Errors                   *
*****

```

以下省略了另外五块芯片的检查结果。

```

show asic-errors s2 1 all location F1/SM12/SP
show asic-errors s2 2 all location F1/SM12/SP
show asic-errors s2 3 all location F1/SM12/SP
show asic-errors s2 4 all location F1/SM12/SP
show asic-errors s2 5 all location F1/SM12/SP

```

从上面的例子，我们可以看到，错误类型为不可修复的RS错误。什么是RS错误呢？原来，Reed/Solomon (RS) 是一种编码方法，当编码进行时遇到问题，就会报RS错误。RS错误一般发生在系统的启动过程中；如果某一fabric link脏了，可能会使fabric 芯片收到噪声信号，也会产生RS错误。当信号被噪声污染，衰减到一定程度，就会报UCE (不可修复的错误)，因为信号无法被还原了。

小结我们排错的结果，我们可以看到以下四条link在一周之内翻转 (或者说flapping, up/down) 最频繁。

link s2rx/F1/SM12/SP/5/23: 3 次

link s2rx/F1/SM12/SP/0/22: 2 次
link s2rx/F1/SM12/SP/3/22: 3 次
link s2rx/F1/SM12/SP/4/23: 4 次

由于CRS-1的fabric 光缆每个接口有72条光纤，只有四条报噪声，我们可以考虑通过shutdown/no shutdown, 或者把这四条光纤admin down (管理down)的方式来作为临时解决方案。CRS-1的冗余性非常好，把这四条光纤shutdown一点都不会影响业务。等到有维护窗口的时候，我们再对这四条光纤所在的光缆进行清洁工作。

shutdown 的命令如下。

```
admin
config
(admin-config)#controller fabric link port s2rx/F1/SM12/SP/5/23 shutdown.
(admin-config)#commit
```

清洁的时候，请参照下图寻找光纤在光缆中的位置。

命令示例如下：

```
RP/0/RP0/CPU0:CRS(admin)#show controllers fabric link port s2rx F1/SM12/SP/5/23
detail
```

```
Flags: P - plane admin down,          p - plane oper down
       C - card admin down,           c - card oper down
       L - link port admin down,      l - linkport oper down
```

```

A - asic admin down,          a - asic oper down
B - bundle port admin Down,  b - bundle port oper down
I - bundle admin down,       i - bundle oper down
N - node admin down,         n - node down
o - other end of link down   d - data down
f - failed component downstream
m - plane multicast down,    s - link port permanently shutdown
t - no barrier input

```

```

Sfe Port          Admin Oper  Down   Sfe BP  Port BP  Other
R/S/M/A/P        State State  Flags   Role    Role    End
-----
F1/SM12/SP/5/23  UP     UP
1/SM6/SP/1/16

```

```

Connection Details for s2rx/F1/SM12/SP/5/23
-----

```

```

Type: Inter-chassis bundle
Near-end bundle port: bport/F1/SM12/5 ribbon 1 fiber 5
Far-end bundle port : bport/1/SM6/2  ribbon 4 fiber 5
HBMT pin name       : P7L3_5
Fabric group offset : (unknown)
Fabric group        : (unknown)

```

经验总结

由于CRS-1的Fabric 排错相对来说比较复杂，需要对CRS-1的FABRIC体系架构有一定的认识，对于shutdown光纤数量对系统的影响（内容较多，本文不予讨论）也要有正确的评估，本示例仅作为快速处理的参考。建议您碰到CRS-1 fabric 相关问题时，联系Cisco TAC来帮助您进行故障排查。

相关命令

```

(admin)#show controllers fabric link port s[x][r/t]x all statistics | exclude .* 0.* 0.* 0
(admin)#show asic-errors s[x] 0 all location [x/x/x]
(admin-config)#controller fabric link port [x/x/x/x/x/x] shutdown.
(admin-config)#commit
(admin)#show controllers fabric link port s[x][r/t]x [x/x/x/x/x] detail

```