

记一次线上服务器频繁宕机

简单点，说话的方式简单点

1、查看dmesg日志

```
grep -E "error|Error|ERROR|fail|Fail|FAIL" dmesg
[Hardware Error]: This system BIOS has enabled interrupt remapping
ERST: Error Record Serialization Table (ERST) support is initialized.
ACPI Error: No handler for Region [IPMI] (ffff88042a610300) [IPMI]
(20090903/evregion-319)
ACPI Error: Region IPMI(7) has no handler (20090903/exfldio-295)
ACPI Error (psparse-0537): Method parse/execution failed [_SB_.PMIO._GHL] (Node
ffff88082a7aeab0), AE_NOT_EXIST
ACPI Error (psparse-0537): Method parse/execution failed [_SB_.PMIO._PMC] (Node
ffff88082a7aeb00), AE_NOT_EXIST
```

2、查看messages日志

```
grep -E "error|Error|ERROR|fail|Fail|FAIL" messages
Mar 15 11:32:07 rsyslogd: UDP message reception disabled due to error logged in
last message.
Mar 15 11:32:07 kernel: [Hardware Error]: This system BIOS has enabled interrupt
remapping
Mar 15 11:32:07 kernel: ERST: Error Record Serialization Table (ERST) support is
initialized.
Mar 15 11:32:07 kernel: ACPI Error: No handler for Region [IPMI]
(ffff88042a610300) [IPMI] (20090903/evregion-319)
Mar 15 11:32:07 kernel: ACPI Error: Region IPMI(7) has no handler
(20090903/exfldio-295)
Mar 15 11:32:07 kernel: ACPI Error (psparse-0537): Method parse/execution failed
[_SB_.PMIO._GHL] (Node ffff88082a7aeab0), AE_NOT_EXIST
Mar 15 11:32:07 kernel: ACPI Error (psparse-0537): Method parse/execution failed
[_SB_.PMIO._PMC] (Node ffff88082a7aeb00), AE_NOT_EXIST
Mar 15 11:32:14 mcelog: failed to prefill DIMM database from DMI data
Mar 15 11:32:14 snmpd[20522]: /etc/snmp/snmpd.conf: line 61: Error: ERROR: This
output format has been deprecated - Please use the 'extend' directive instead
Mar 15 11:32:14 snmpd[20522]: /etc/snmp/snmpd.conf: line 62: Error: ERROR: This
output format has been deprecated - Please use the 'extend' directive instead
Mar 15 11:32:14 snmpd[20522]: /etc/snmp/snmpd.conf: line 63: Error: ERROR: This
output format has been deprecated - Please use the 'extend' directive instead
Mar 15 11:32:14 snmpd[20522]: /etc/snmp/snmpd.conf: line 64: Error: ERROR: This
```

```
output format has been deprecated - Please use the 'extend' directive instead
Mar 15 11:32:14 snmpd[20522]: /etc/snmp/snmpd.conf: line 65: Error: ERROR: This
output format has been deprecated - Please use the 'extend' directive instead
Mar 15 11:32:14 snmpd[20522]: /etc/snmp/snmpd.conf: line 66: Error: ERROR: This
output format has been deprecated - Please use the 'extend' directive instead
Mar 15 11:32:14 snmpd[20522]: /etc/snmp/snmpd.conf: line 67: Error: ERROR: This
output format has been deprecated - Please use the 'extend' directive instead
Mar 15 11:32:14 snmpd[20522]: /etc/snmp/snmpd.conf: line 68: Error: ERROR: This
output format has been deprecated - Please use the 'extend' directive instead
Mar 15 11:32:14 snmpd[20522]: net-snmp: 8 error(s) in config file(s)
Mar 15 11:32:37 kernel: oddjobd[26628]: segfault at 3700000000 ip
00000037975292b0 sp 00007fff5397ef28 error 4 in libc-2.12.so[3797400000+18a000]
Mar 15 11:32:37 oddjobd: oddjobd startup failed
```

3、查看kdump的log文件

```
grep -E "error|Error|ERROR|fail|Fail|FAIL" vmcore-dmesg.txt
```

```
<4>[Hardware Error]: This system BIOS has enabled interrupt remapping
<6>ERST: Error Record Serialization Table (ERST) support is initialized.
<4>ACPI Error: No handler for Region [IPMI] (ffff88042a610300) [IPMI]
(20090903/evregion-319)
<4>ACPI Error: Region IPMI(7) has no handler (20090903/exfldio-295)
<4>ACPI Error (psparse-0537): Method parse/execution failed [\_SB_.PMIO._GHL]
(Node ffff88082a7aeab0), AE_NOT_EXIST
<4>ACPI Error (psparse-0537): Method parse/execution failed [\_SB_.PMIO._PMC]
(Node ffff88082a7aeb00), AE_NOT_EXIST
<6>oddjobd[26656]: segfault at 3700000000 ip 00000037975292b0 sp
00007ffffb3fd00d8 error 4 in libc-2.12.so[3797400000+18a000]
<6>chntf[42911]: segfault at 10 ip 000000000045afb4 sp 00007f374386b940 error 4
in chntf[400000+12e000]
<6>dmagic93[26512] general protection ip:3797476065 sp:7f20c803b500 error:0 in
libc-2.12.so[3797400000+18a000]
```

分析:

由于BIOS中开启了中断重映射（这是个复杂的东西，也可以简单讲明白，我就不讲了。），在ERST（芯片集中的错误校验表）校验时发生错误，导致高级配置电源管理模块无法处理IPMI驱动请求，预存数据到内存发生错误，引发kernel上演了一出找到空指针的戏法。

解决方法:

在grub.conf的内核启动参数中添加 `intremap=off` 或者 `intremap=no_x2apic_optout`
`intremap={on,off,nosid,no_x2apic_optout}`

on (默认值) 开启中断重映射, BIOS中默认开启

off 关闭中断重映射

nosid 重映射时不对SID(Source ID)做检查

no_x2apic_optout 无视BIOS的设置，强制禁用x2APIC特性，主要用于解决某些对x2APIC支持有缺陷的BIOS导致的故障