

GIS - Support #13235

Server hardware failure (sdc hard drive)

18/11/2021 15:17 - Selvarani C

Status:	Resolved	Start date:	18/11/2021
Priority:	Normal	Due date:	
Assignee:	Philippe May	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Description Hi Philippe, The Gisaf is not Connecting. so can you please check it. I have attached the png for your reference. Thank you.			

History

- #1 - 18/11/2021 16:33 - Philippe May
- Project changed from Gisaf to GIS
 - Status changed from New to In Progress

Definitely not a problem with Gisaf, but with the server.

I could connect with VPN on gisdb.csr.av, but not with ssh to the dom0 (dream.csr.av). Coming to CSR office: the server console shown a lot of messages related to disk OK errors on sdc2.

Rebooted in safe mode: could not fsck /dev/sdc2 since it's the root FS. Continued the boot and the services and VMs started fine.

BUT: it's a sign that the sdc drive is having issues, and action needed very soon.

#2 - 18/11/2021 16:38 - Philippe May

Some information:

smartctl -a /dev/sdc

```
smartctl 7.2 2020-12-30 r5155 [x86_64-linux-5.10.0-9-amd64] (local build)
Copyright (C) 2002-20, Bruce Allen, Christian Franke, www.smartmontools.org

=== START OF INFORMATION SECTION ===
Model Family:      Seagate BarraCuda 3.5
Device Model:      ST2000DM006-2DM164
Serial Number:      Z8E0SXSW
LU WWN Device Id:  5 000c50 0a5b727e4
Firmware Version:  CC26
User Capacity:      2,000,398,934,016 bytes [2.00 TB]
Sector Sizes:       512 bytes logical, 4096 bytes physical
Rotation Rate:      7200 rpm
Form Factor:        3.5 inches
Device is:          In smartctl database [for details use: -P show]
ATA Version is:     ACS-2, ACS-3 T13/2161-D revision 3b
SATA Version is:    SATA 3.1, 6.0 Gb/s (current: 6.0 Gb/s)
```

Local Time is: Thu Nov 18 16:32:05 2021 IST
SMART support is: Available - device has SMART capability.
SMART support is: Enabled

=== START OF READ SMART DATA SECTION ===

SMART overall-health self-assessment test result: PASSED

General SMART Values:

Offline data collection status: (0x82) Offline data collection activity
was completed without error.

Auto Offline Data Collection: Enabled.

Self-test execution status: (0) The previous self-test routine completed
without error or no self-test has ever
been run.

Total time to complete Offline
data collection: (80) seconds.

Offline data collection

capabilities: (0x7b) SMART execute Offline immediate.

Auto Offline data collection on/off support.

Suspend Offline collection upon new
command.

Offline surface scan supported.

Self-test supported.

Conveyance Self-test supported.

Selective Self-test supported.

SMART capabilities: (0x0003) Saves SMART data before entering
power-saving mode.

Supports SMART auto save timer.

Error logging capability: (0x01) Error logging supported.

General Purpose Logging supported.

Short self-test routine

recommended polling time: (1) minutes.

Extended self-test routine

recommended polling time: (210) minutes.

Conveyance self-test routine

recommended polling time: (2) minutes.

SCT capabilities: (0x1085) SCT Status supported.

SMART Attributes Data Structure revision number: 10

Vendor Specific SMART Attributes with Thresholds:

ID#	ATTRIBUTE_NAME	FLAG	VALUE	WORST	THRESH	TYPE	UPDATED	WHEN_FAILED	RAW_VALUE
1	Raw_Read_Error_Rate	0x000f	058	057	006	Pre-fail	Always	-	95563133
3	Spin_Up_Time	0x0003	096	096	000	Pre-fail	Always	-	0
4	Start_Stop_Count	0x0032	100	100	020	Old_age	Always	-	465
5	Reallocated_Sector_Ct	0x0033	100	100	010	Pre-fail	Always	-	0
7	Seek_Error_Rate	0x000f	073	060	030	Pre-fail	Always	-	23428726
9	Power_On_Hours	0x0032	063	063	000	Old_age	Always	-	33124
10	Spin_Retry_Count	0x0013	100	100	097	Pre-fail	Always	-	0
12	Power_Cycle_Count	0x0032	100	100	020	Old_age	Always	-	469
183	Runtime_Bad_Block	0x0032	100	100	000	Old_age	Always	-	0
184	End-to-End_Error	0x0032	100	100	099	Old_age	Always	-	0
187	Reported_Uncorrect	0x0032	001	001	000	Old_age	Always	-	65535
188	Command_Timeout	0x0032	100	100	000	Old_age	Always	-	0 0 0
189	High_Fly_Writes	0x003a	100	100	000	Old_age	Always	-	0
190	Airflow_Temperature_Cel	0x0022	056	049	045	Old_age	Always	-	44 (Min/Max 44/46)
191	G-Sense_Error_Rate	0x0032	100	100	000	Old_age	Always	-	0
192	Power-Off_Retract_Count	0x0032	100	100	000	Old_age	Always	-	406
193	Load_Cycle_Count	0x0032	093	093	000	Old_age	Always	-	14277
194	Temperature_Celsius	0x0022	044	051	000	Old_age	Always	-	44 (0 24 0 0 0)
197	Current_Pending_Sector	0x0012	100	100	000	Old_age	Always	-	8
198	Offline_Uncorrectable	0x0010	100	100	000	Old_age	Offline	-	8
199	UDMA_CRC_Error_Count	0x003e	200	200	000	Old_age	Always	-	0
240	Head_Flying_Hours	0x0000	100	253	000	Old_age	Offline	-	32007h+16m+24.524s
241	Total_LBAs_Written	0x0000	100	253	000	Old_age	Offline	-	4757015136
242	Total_LBAs_Read	0x0000	100	253	000	Old_age	Offline	-	6203030490

SMART Error Log Version: 1

ATA Error Count: 2030 (device log contains only the most recent five errors)

CR = Command Register [HEX]

FR = Features Register [HEX]

SC = Sector Count Register [HEX]

SN = Sector Number Register [HEX]

CL = Cylinder Low Register [HEX]

CH = Cylinder High Register [HEX]

DH = Device/Head Register [HEX]

DC = Device Command Register [HEX]
ER = Error register [HEX]
ST = Status register [HEX]
Powered_Up_Time is measured from power on, and printed as
DDd+hh:mm:ss.sss where DD=days, hh=hours, mm=minutes,
SS=sec, and sss=millisec. It "wraps" after 49.710 days.

Error 2030 occurred at disk power-on lifetime: 33124 hours (1380 days + 4 hours)
When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:
ER ST SC SN CL CH DH
-- -- -- -- -- -- --
40 53 00 ff ff ff 0f Error: WP at LBA = 0xffffffff = 268435455

Commands leading to the command that caused the error were:
CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name
-- -- -- -- -- -- --
61 00 58 ff ff ff 4f 00 00:54:01.824 WRITE FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:53:59.625 READ FPDMA QUEUED
61 00 08 ff ff ff 4f 00 00:53:59.625 WRITE FPDMA QUEUED
61 00 08 ff ff ff 4f 00 00:53:59.625 WRITE FPDMA QUEUED
61 00 08 ff ff ff 4f 00 00:53:59.625 WRITE FPDMA QUEUED

Error 2029 occurred at disk power-on lifetime: 33124 hours (1380 days + 4 hours)
When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:
ER ST SC SN CL CH DH
-- -- -- -- -- -- --
40 53 00 ff ff ff 0f Error: WP at LBA = 0xffffffff = 268435455

Commands leading to the command that caused the error were:
CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name
-- -- -- -- -- -- --
61 00 08 ff ff ff 4f 00 00:53:57.213 WRITE FPDMA QUEUED
61 00 08 ff ff ff 4f 00 00:53:57.213 WRITE FPDMA QUEUED
61 00 08 ff ff ff 4f 00 00:53:57.213 WRITE FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:53:55.970 READ FPDMA QUEUED
ea 00 00 00 00 00 a0 00 00:53:55.962 FLUSH CACHE EXT

Error 2028 occurred at disk power-on lifetime: 33123 hours (1380 days + 3 hours)
When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:
ER ST SC SN CL CH DH
-- -- -- -- -- -- --
40 53 00 ff ff ff 0f Error: WP at LBA = 0xffffffff = 268435455

Commands leading to the command that caused the error were:
CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name
-- -- -- -- -- -- --
61 00 50 ff ff ff 4f 00 00:14:00.991 WRITE FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:14:00.991 READ FPDMA QUEUED
ef 10 02 00 00 00 a0 00 00:14:00.987 SET FEATURES [Enable SATA feature]
27 00 00 00 00 00 e0 00 00:14:00.987 READ NATIVE MAX ADDRESS EXT [OBS-ACS-3]
ec 00 00 00 00 00 a0 00 00:14:00.986 IDENTIFY DEVICE

Error 2027 occurred at disk power-on lifetime: 33123 hours (1380 days + 3 hours)
When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:
ER ST SC SN CL CH DH
-- -- -- -- -- -- --
40 53 00 ff ff ff 0f Error: WP at LBA = 0xffffffff = 268435455

Commands leading to the command that caused the error were:
CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name
-- -- -- -- -- -- --
61 00 50 ff ff ff 4f 00 00:13:59.636 WRITE FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:13:57.368 READ FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:13:57.367 READ FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:13:57.367 READ FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:13:57.367 READ FPDMA QUEUED

Error 2026 occurred at disk power-on lifetime: 33123 hours (1380 days + 3 hours)
When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:
ER ST SC SN CL CH DH
-- -- -- -- -- -- --
40 53 00 ff ff ff 0f Error: UNC at LBA = 0x0fffffff = 268435455

Commands leading to the command that caused the error were:
CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name
-- -- -- -- -- -- --
60 00 28 ff ff ff 4f 00 00:13:53.658 READ FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:13:53.657 READ FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:13:53.657 READ FPDMA QUEUED
60 00 08 ff ff ff 4f 00 00:13:53.656 READ FPDMA QUEUED
60 00 20 ff ff ff 4f 00 00:13:53.656 READ FPDMA QUEUED

SMART Self-test log structure revision number 1
Num Test_Description Status Remaining LifeTime(hours) LBA_of_first_error
1 Short offline Completed without error 00% 7346 -

SMART Selective self-test log data structure revision number 1
SPAN MIN_LBA MAX_LBA CURRENT_TEST_STATUS
1 0 0 Not_testing
2 0 0 Not_testing
3 0 0 Not_testing
4 0 0 Not_testing
5 0 0 Not_testing

Selective self-test flags (0x0):
After scanning selected spans, do NOT read-scan remainder of disk.
If Selective self-test is pending on power-up, resume after 0 minute delay.

lsblk

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
sda	8:0	0	1.8T	0	disk	
└sda1	8:1	0	512M	0	part	/boot/efi
└sda2	8:2	0	23.3G	0	part	
└sda3	8:3	0	9.3G	0	part	
└sda4	8:4	0	7.9G	0	part	
└sda5	8:5	0	1.9G	0	part	
└sda6	8:6	0	1.8T	0	part	
└└md0	9:0	0	1.8T	0	raid1	
└└└dream.csr-gisaf.csr.av--swap	253:0	0	8G	0	lvm	
└└└dream.csr-gisaf.csr.av--disk	253:1	0	1000G	0	lvm	
└└└dream.csr-infra.csr.av--swap	253:2	0	1G	0	lvm	
└└└dream.csr-infra.csr.av--disk	253:3	0	10G	0	lvm	
└└└dream.csr-samba.csr.av--swap	253:4	0	1G	0	lvm	
└└└dream.csr-samba.csr.av--disk	253:5	0	100G	0	lvm	
└└└dream.csr-gisaf2.csr.av--swap	253:6	0	1G	0	lvm	
└└└dream.csr-gisaf2.csr.av--disk	253:7	0	10G	0	lvm	
└└└dream.csr-gisdb.csr.av--swap	253:8	0	1G	0	lvm	
└└└dream.csr-gisdb.csr.av--disk	253:9	0	10G	0	lvm	
└└└dream.csr-jupyter.csr.av--swap	253:10	0	1G	0	lvm	
└└└dream.csr-jupyter.csr.av--disk	253:11	0	20G	0	lvm	
sdb	8:16	0	931.5G	0	disk	
└sdb1	8:17	0	512M	0	part	
└sdb2	8:18	0	23.3G	0	part	
└sdb3	8:19	0	9.3G	0	part	
└sdb4	8:20	0	7.9G	0	part	
└sdb5	8:21	0	1.9G	0	part	
sdc	8:32	0	1.8T	0	disk	
└sdc1	8:33	0	1.3T	0	part	/var/backups
└sdc2	8:34	0	23.4G	0	part	/
└sdc3	8:35	0	7.5G	0	part	[SWAP]
sdd	8:48	0	1.8T	0	disk	
└sdd1	8:49	0	512M	0	part	
└sdd2	8:50	0	23.3G	0	part	
└sdd3	8:51	0	9.3G	0	part	
└sdd4	8:52	0	7.9G	0	part	
└sdd5	8:53	0	1.9G	0	part	

└sdd6	8:54	0	1.8T	0	part
└┬md0	9:0	0	1.8T	0	raid1
│└dream.csr-gisaf.csr.av--swap	253:0	0	8G	0	lvm
│└dream.csr-gisaf.csr.av--disk	253:1	0	1000G	0	lvm
│└dream.csr-infra.csr.av--swap	253:2	0	1G	0	lvm
│└dream.csr-infra.csr.av--disk	253:3	0	10G	0	lvm
│└dream.csr-samba.csr.av--swap	253:4	0	1G	0	lvm
│└dream.csr-samba.csr.av--disk	253:5	0	100G	0	lvm
│└dream.csr-gisaf2.csr.av--swap	253:6	0	1G	0	lvm
│└dream.csr-gisaf2.csr.av--disk	253:7	0	10G	0	lvm
│└dream.csr-gisdb.csr.av--swap	253:8	0	1G	0	lvm
│└dream.csr-gisdb.csr.av--disk	253:9	0	10G	0	lvm
│└dream.csr-jupyter.csr.av--swap	253:10	0	1G	0	lvm
└┬dream.csr-jupyter.csr.av--disk	253:11	0	20G	0	lvm

lshw -c disk

```

*-disk:0
  description: ATA Disk
  product: ST2000DM006-2DM1
  physical id: 0
  bus info: scsi@0:0.0.0
  logical name: /dev/sda
  version: CC26
  serial: Z8E0QMHC
  size: 1863GiB (2TB)
  capabilities: gpt-1.00 partitioned partitioned:gpt
  configuration: ansiversion=5 guid=674b4b97-ac3c-4564-b4fe-e7f7be9b4fa9 logicalsectorsize=512 sectorsize
=4096
*-disk:1
  description: ATA Disk
  product: TOSHIBA DT01ACA1
  vendor: Toshiba
  physical id: 1
  bus info: scsi@1:0.0.0
  logical name: /dev/sdb
  version: A810
  serial: 77I739JMS
  size: 931GiB (1TB)
  capabilities: gpt-1.00 partitioned partitioned:gpt
  configuration: ansiversion=5 guid=59098f62-b43f-4f54-b510-3548c7c21fec logicalsectorsize=512 sectorsize
=4096
*-disk:2
  description: ATA Disk
  product: ST2000DM006-2DM1
  physical id: 2
  bus info: scsi@2:0.0.0
  logical name: /dev/sdc
  version: CC26
  serial: Z8E0SXSX
  size: 1863GiB (2TB)
  capabilities: gpt-1.00 partitioned partitioned:gpt
  configuration: ansiversion=5 guid=2883c10e-8661-48e9-af21-81073d972719 logicalsectorsize=512 sectorsize
=4096
*-disk:3
  description: ATA Disk
  product: ST2000DM006-2DM1
  physical id: 3
  bus info: scsi@3:0.0.0
  logical name: /dev/sdd
  version: CC26
  serial: Z4Z98DCQ
  size: 1863GiB (2TB)
  capabilities: gpt-1.00 partitioned partitioned:gpt
  configuration: ansiversion=5 guid=59098f62-b43f-4f54-b510-3548c7c21fec logicalsectorsize=512 sectorsize
=4096

```

#3 - 18/11/2021 16:40 - Philippe May

- Subject changed from *Gisaf is not connecting to Server hardware failure (sdc hard drive)*

#4 - 18/11/2021 18:49 - Philippe May

sdb is not used (was used initially during server install) and does not show any error.

So, i re-partitioned it with a sdb2 ext4 partition, and:

```
root@dream:~# dd if=/dev/sdc2 of=/dev/sdb2 bs=64K conv=noerror,sync
dd: error reading '/dev/sdc2': Input/output error
102468+1 records in
102469+0 records out
6715408384 bytes (6.7 GB, 6.3 GiB) copied, 77.6204 s, 86.5 MB/s
382991+1 records in
382992+0 records out
25099763712 bytes (25 GB, 23 GiB) copied, 261.595 s, 95.9 MB/s
```

Also, i copied the EFI:

```
root@dream:~# dd if=/dev/sda1 of=/dev/sdb1 bs=64K conv=noerror,sync
8192+0 records in
8192+0 records out
536870912 bytes (537 MB, 512 MiB) copied, 5.4858 s, 97.9 MB/s
```

Since there was 1 error reported by dd: complete fsck:

```
root@dream:~# fsck -pvcf /dev/sdb2
fsck from util-linux 2.36.1
/dev/sdb2: Updating bad block inode.
```

```
46166 inodes used (3.01%, out of 1534896)
185 non-contiguous files (0.4%)
76 non-contiguous directories (0.2%)
# of inodes with ind/dind/tind blocks: 0/0/0
Extent depth histogram: 41521/139/4
1009063 blocks used (16.47%, out of 6127616)
0 bad blocks
1 large file

36869 regular files
4699 directories
7 character device files
0 block device files
0 fifos
13 links
4580 symbolic links (4485 fast symbolic links)
2 sockets
-----
46170 files
```

Looks like all is OK.

Next: config to boot with root on /dev/sdb2

#5 - 19/11/2021 17:59 - Philippe May

- % Done changed from 0 to 100
- Status changed from In Progress to Resolved

Today the dom0 (dream.csr.av) was not accessible.

I came to CSR again this afternoon and rebooted the server, manually starting it with root set to sdb2 with grub options. So far so good: no error.

Try to make sure that the server's EFI config boots on sdb2:

- changed /etc/fstab to make sure that /dev/sda1 is used as /boot/efi (needed because of the confusion of UUID since i cloned yesterday with dd)
- mount manually /dev/sda1 as /boot/efi
- ran grub-install: hopefully the server will eventually reboot on sdb2.

So, in short, this new setup should prevent the need for replacement of hardware.

4 hard drives:

- sdc (size: 2TB) has shown a worrying error: 1 bad sector. The rest (backup space) seems alright.
- sdb: (1TB): used for the dom0, seems all OK.
- sda and sdd: RAID for all application, data and VMs: all OK.

Marking this ticket as resolved. Will follow up eventually if there's an issue when the server restarts.

Files			
Not Connected.png	68.4 KB	18/11/2021	Selvarani C