SYA: CSAA861E FAILURE DETECTED IN MVS COMMON STORAGE TRACKING

Case
TS008834715

Status
IBM is working

Product
z/OS

Description
Hello,
today at 10:44:34 system named SYA captured an SVC dump on MASTER AS. Over this, a message: CSAA861E FAILURE DETECTED IN MVS COMMON STORAGE TRACKING was issued.
I'm going to attach to the case DUMP, syslog, and EREP in that timeframe.
Please may you investigate this issue?
Many thanks,
kind regards
Pier Paolo Mattavelli

Product Area
z/OS zOS BCP General <5752OS390>

Product Version
2.3.0

System Down
No

Service Type
BreakFix

Severity
2 - Significant impact (any system is down)

Created
24 Mar 2022(147 days old)
3:52 PM

Attachments
SYA.D22083.RMFASD.TRS
SYA.D22083.SMF7X.TRS
SYA.D22083.EREPD08.TRS
SYA.D22083.EREPD08.TRS
SYA.D22083.EREPS08.TRS
SYA.D22083.DSQD.TRS
SYA.D22083.EREPDET.TRS
SYA.D22083.ALLOCAS.TRS
SYA.D22083.SYSLOG.TRS
SYA.D22083.EREPSUM.TRS
SYA.D22083.DUMPMSTR.TRS
IODF19.CSSCU.txt
IODF19.COMPARE.txt
mail.html

Customer
CEDACRI SPA ONLY CODE CTR OIO

Client reference number

IBM customer number
0461723

Geography
IT

Team

Pier Paolo Mattavelli **(case owner)**

Alessandro Studer

Rina Forni

Case history

---

4 Aug 2022

09:26 PM CEST

Jonathan.Avila (IBM)

Hello,
We are still currently reviewing when the PTFs will be available.
Please let us know if you have any questions or concerns.
Thank You,
Jonathan Avila,
DFSMS Support - Device Services
Action Taken: Follow up with customer
Action Planned: Await customer response
FUP: 8/11/2022

---

21 Jul 2022

09:05 AM CEST

v_gay (IBM)

Hi Jonathan,
Thanks a lot or opening these 2 new ideas!
"Add SORT command on panels to help in avoiding problems for apar OA63553 (https://ibm-z-hardware-and-operating-systems.ideas.ibm.com/ideas/ZOS-I-3372)"
"New cross reference function (https://ibm-z-hardware-and-operating-systems.ideas.ibm.com/ideas/ZOS-I-3373)"
I notify Jonathan.

We are tracking APAR status and we will inform you as soon as PTFs are available
Kind regards. Vincent (for Jonathan)

Action taken:

- acknowledged client comments
- noted link to both ideas
- notify Jonathan

Action plan: to Jona for APAR tracking
Notified: @Jonathan Avila
08:40 AM CEST
Mario.Dimitrov (IBM)
Hello Rina,
Thank you for the feedback.
I will notify Jonathan.
Regards,
Mario
action plan: notify Jonathan FUP 7/21/2022
Notified: @Jonathan Avila
08:40 AM CEST
Mario.Dimitrov (IBM) changed Status from ***Waiting for IBM*** to ***IBM is working***.

---

20 Jul 2022

11:44 PM CEST
Rina Forni (Customer)
Hello Jonathan,
we will wait for the ptf! Thanks
Just for your info, I opened two "ideas" against HCD:
"Add SORT command on panels to help in avoiding problems for apar OA63553"
"New cross reference function"
I don't know if they will be implemented but, as the program tells, they are just ideas!
ciao, Rina
11:44 PM CEST
Rina Forni (Customer) changed Status from ***Awaiting your feedback*** to ***Waiting for IBM***.
10:54 PM CEST
Jonathan.Avila (IBM)

Hello Pier and Rina,
We understand your last update and we are moving forward with an APAR. The apar is opened at OA63553 and will be fixing it at the current release. With respect to HCD panels not providing a SORT command, It definitely sounds like a good idea. I would suggest that you open an RFE (Now called Ideas) which can be accessed at the following link:
https://www.ibm.com/support/pages/node/6438917 (https://www.ibm.com/support/pages/node/6438917)
Let us know if you have any additional questions, otherwise we will let you know when the PTFs are available.
Thank You,
Jonathan Avila,
DFSMS Support - Device Services
Action Taken: Inform customer of apar
Action Planned: Await customer response
FUP: 7/27/2022

10:53 PM CEST

Jonathan.Avila (IBM) changed Status from ***IBM is working*** to ***Awaiting your feedback***.

10:52 AM CEST

v_gay (IBM)
Hi Rina,
Thanks for your comments.
I notify Jonathan who will update you during his business hours.
Thanks for your patience. Kind regards. VIncent (for Jonathan)

Action taken: acknowledged client comments
Action plan: to Jona's attention
Notified: @Jonathan Avila

10:52 AM CEST

v_gay (IBM) changed Status from ***Waiting for IBM*** to ***IBM is working***.

10:28 AM CEST

Rina Forni (Customer)
Hello Jonathan,
sure I want an apar opened but I don't agree withe the FIN closing code.
What does "future release" mean exactly? If you reproduced the problem, you now know that it is not so hard having HCD misconfigured in large configuration and we are exposed to the problem NOW.
We are in the process of migrating to z/os 2.5, we could accept a fix for that release but not for an undetermined future release.
Also, even though the dev services will protect the system during the vary online process, this kind of misconfiguring problem could be more easily avoided if the HCD panels would have a SORT command available to ask for ordering the LCU list by CUADD: I think that this kind of function could be easily implemented on z/os 2.3 also.
Least but not last, either the final fix will involve also the HCD logic or we will ask for an RFE against HCD and in this second case having an apar number as root cause of the request may at the end cause someone to consider it.
thanks, Rina

10:28 AM CEST

Rina Forni (Customer) changed Status from ***Awaiting your feedback*** to ***Waiting for IBM***.

---

19 Jul 2022

07:18 PM CEST

Jonathan.Avila (IBM)

Hello Rina and Pier,
We were able to recreate the scenario with a forced mis-configuration similar for the issue in this case. There are plans to implement a fix in the next release however if you are willing to take a FIN apar please let us know.
Thank You,
Jonathan Avila,
DFSMS Support - Device Services
Action Taken: Inform customer of FIN apar
Action Planned: Await customer response
FUP: 7/26/2022

07:18 PM CEST

Jonathan.Avila (IBM) changed Status from ***IBM is working*** to ***Awaiting your feedback***.

---

13 Jul 2022

05:09 PM CEST

Jonathan.Avila (IBM)

Hello Rina,
We are still reviewing the process to collect documentation we will be in touch once we have instructions.
Thank you for your patience,
Jonathan Avila,
DFSMS Support - Device Services
Action Taken: Review Information for Documentation collection
Action Planned: Provide customer feedback
FUP: 7/20/2022

---

28 Jun 2022

10:07 PM CEST

Jonathan.Avila (IBM)

Hello Rina,
We are currently reviewing the details for the documentation to collect and we will be in touch soon.
Thank You,
Jonathan Avila,
DFSMS Support - Device Services
Action Taken: Review Information for Documentation collection
Action Planned: Provide customer feedback
FUP: 7/5/2022
10:07 PM CEST
Jonathan.Avila (IBM) changed Status from ***Waiting for IBM*** to ***IBM is working***.
09:26 PM CEST
ECUREP (IBM)
Mail from: Rina Forni <rina.forni@smeup.com>

Hello Vincent,

you are right, the correct word is “can”.

So, if your colleagues can help, we are quite sure we may try to reproduce the problem because we got trouble each time we misconfigured hcd.

Maybe it will not surface the same manner as this case, but two times we performed a mistake, two times we got troubles.

Ciao, Rina

Il giorno mar 28 giu 2022 alle 19:21 IBM Support <support@ecurep.ibm.com> ha scritto:

IBM Support
[servlet.ImageServer?id=0150z000002GllcAAC&oid=00D50000000c9MWEAY&lastMod=1592310594000]
Case notifications
Case Updated
28 June 2022, 19:20 PM CEST
TS008834715 Severity 2
SYA: CSAA861E FAILURE DETECTED IN MVS COMMON STORAGE TRACKING
IBM Support - 28 June 2022, 19:20 PM CEST

Hello Rina & Pier,

Sorry for my delay in replying here. I was out of the office yesterday. Thank you for the response, and a similar issue with the configuration issue, that case did not indicate overlays, but perhaps nothing critical was damaged at that time ether. I
assume your second word in the first sentence cant , is a typo. Based on your following statements, asking for the trace/meeting to discuss further what is needed, I think it should read can. I am going to involve my colleagues in Device Services as
the primary on this one which I was holding off on until you confirmed you'd be will to recreate the issue. They will respond next regarding what diagnostics can be provided.

Regards,

Vincent
09:26 PM CEST
ECUREP (IBM) changed Status from ***IBM is working*** to ***Waiting for IBM***.
07:20 PM CEST
vbonanno (IBM)
Hello Rina & Pier,
Sorry for my delay in replying here. I was out of the office yesterday. Thank you for the response, and a similar issue with the configuration issue, that case did not indicate overlays, but perhaps nothing critical was damaged at that time ether. I assume your second word in the first sentence ***cant*** , is a typo. Based on your following statements, asking for the trace/meeting to discuss further what is needed, I think it should read ***can***. I am going to involve my colleagues in Device Services as the primary on this one which I was holding off on until you confirmed you'd be will to recreate the issue. They will respond next regarding what diagnostics can be provided.
Regards,
Vincent

---

27 Jun 2022

05:27 PM CEST
Jack.Rothberg (IBM)
Action Taken: Update acknowledged, notified Vincent
Action Plan: hold for Vincents review
Notified: @VINCENT BONANNO
05:27 PM CEST
Jack.Rothberg (IBM) changed Status from ***Waiting for IBM*** to ***IBM is working***.
04:23 PM CEST
pierpaolo.mattavelli (Customer)
Hello Vincent,
we cant try to reproduce the issue in a DR environment, which kind of specific information (captured through MVS commands, traces, dumps) are you looking for?
In the past we experienced a similar issue, please refer to case TS004080538.
We are also available to attend a chat with you in order to plan a test and define the data we need to collect for you, please let us know.

Kind regards
Rina and Pier Paolo

---

23 Jun 2022

11:39 PM CEST
vbonanno (IBM)
Hello Pier,
We certainly do not intentionally want to delay diagnosis of this issue to another customer, but we are not able to reproduce this in house, even with the detailed descriptions provided. Being I am part of the VSM team, I can not be sure if this happens again it will damage a VSMP in ESQA, it could affect something else, so I don't have a good means of providing diagnostics here at this point either.
We were hesitant to ask a recreate on your part, since it involved a mis-configuration but, outside of that that we would not have a way to continue the pursuit here. I asked them and if you're willing to try to replicate what you did to cause this initially, they would work on providing additional diagnostics to continue pursuing this problem.
Please let us know if you'd like to continue in that regard or we can close this out.
Regards,
Vincent
11:22 AM CEST
liliya.radulova (IBM)
Hello Pier,
Thank you for the update, I will inform Vincent .
Regards,
Liliya
10:42 AM CEST
pierpaolo.mattavelli (Customer)
Hello,
thanks for your feedback.
I'm happy to close this case if you're happy to drop to your Customers a similar potential issue.
Please let me know, I'll act accordingly.
Kind regards
Pier Paolo Mattavelli

---

10 Jun 2022

03:42 PM CEST

vbonanno (IBM)
Hello Pier/Rina,
I've followed up with the device services team to check on their review and they attempted to drive the same configuration that you described, but it did not result in damage to common storage. Code review also has not revealed an issue here. At this point, there's not much more we can review to determine the root cause.
If you happen to experience another random overlay, we can continue further review.
Regards,
Vincent & Dirk
action taken:
followed up with device services for status and reported results to the customer, that we could not identify an issue resulting in an overlay from the configuration they specified.
action plan:
hold for additional requests/feedback from the customer.

---

2 Jun 2022

03:56 PM CEST
vbonanno (IBM)
action taken:
noted that the Devserv team is still discussing.
Action plan:
followup with Dirk on Status

---

26 Apr 2022

04:16 PM CEST
vbonanno (IBM)
Hello Rina,
Sorry you had trouble with the posts, there is 10,000 character limit per post, at least on our side, I suspect that you also ran into that. Breaking up the posts as you saw is the way around that.

Thank you for the details, I'll discuss with my colleague in Device Services and we'll let you know if we have any further questions. I hope you have a nice day.
Regards,
Vincent
03:51 PM CEST
Alexis.Ricci (IBM)

Action Taken:
Update noted
Action Plan:
notify Vincent of customer's response. Hold for his review.
02:33 PM CEST

Rina Forni (Customer)

Please read updates in the suggested order...
Part 1, part 2, part 3 and part 4.
Files already uploaded.....
regards, Rina and Pier Paolo
02:31 PM CEST

Rina Forni (Customer)

update part 4:
Different ports, different channels but the same LCU 47 of the same 75LBB11 box behind the definitions.
HCD accepted this definition without any warning because we have no way to define which ports belong to the same physical box.
The F280-F2FF addresses were defined as offline=no on all lpar and OS config.
The 8A80-8AFF addresses were defined only to SYA, SYT and SY4 systems and we sufferred problems only on these ones.
We performed the software activate on SYA and SYT systems belonging to parallel sysplex COPLEX and residing on the CSS1 of ZUPR01.
The hw activate was performed on SY2 system belonging to XCFPRD parallel sysplex and residing on CSS0 of ZUPR01.
We didn't receive any kind of error message during the hw Activate process on SY2.
Problem arised after F280 vary online on SYA and SYT and are in the dump you already got.
We forgot to delete 8A80 LCU before defining the F280 but at this time z/os doesn't have any mechanism in place that could help us in avoiding this error.
We are placing the s/n of the boxes in the comment line but there is no way to sort the LCU list by CUADD instead of CU number.
I'm going to upload an iodf compare report between the old iodf18 and the new iodf19.
Iodf19 is the iodf were we introduced the F280 LCU definition.
I will also send you an extract of the print iodf19 report regarding these two lcu.
Please let me know which other information you will need,
About IEECINIT/IECDINIT latest maintenance for R230:
Entry Type: MOD Zone Name: SYA0R1
Entry Name: IECDINIT Zone Type: TARGET
FMID: HDZ2230 LASTUPD: HDZ2230 TYPE=ADD

RMID: UJ90026 DISTLIB: AOSC5
Please let me know which other information you will need,
regards, Rina and Pier Paolo
02:31 PM CEST

Rina Forni (Customer)

update part 3:
Actually we decided to reuse the hw LCU to define new dasd that would be shared among systems.
But because of this sharing characteristics, we decided to use a particular device address range which was free in all the 12 CSS of the two z15: F280.
Also the channels and port we used to define the "new" F280 were different from the ones used for the "old" 8A80:

```
--------------------------------------------------------------------------------
- SY4                                  Row 1 of 12 More:     > -
- Command ===> _____________________________________________ Scroll ===> CSR   -
-                                                                              -
- Control unit number . : F280         DS8950F 75LBB10 SecB Condivisi         -
- Control unit type . . : 2107         Serial number . . . :                  -
-                                                                              -
- Connected switch.ports: 30.47 30.A9 30.BA (https://30.BA) 30.E6 31.47 31.A9 31.BA (https://31.BA) 31.E8     -
-                                                                              -
-                                                                              -
-                                                                              -
-                                                                              -
- ENTER to continue.                                                          -
-                                                                              -
-          -----------------Channel Path ID . Link Address------------------  -
- Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------  -
- ZDPR02.0  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZDPR02.1  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZDPR02.2  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZDPR02.3  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZDPR02.4  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZDPR02.5  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZUPR01.0  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZUPR01.1  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZUPR01.2  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZUPR01.3  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZUPR01.4  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
- ZUPR01.5  A0.47  A8.A9  7E.BA (https://7E.BA)  90.E6  AF.47  B9.A9  68.BA (https://68.BA)  83.E8    -
```

02:30 PM CEST

Rina Forni (Customer)

Update part 2:
The LCU we misconfigured is LSS 47 with ssid=BB47.
It was initially defined in hcd in august 2020 when we migrated to DS8950F with
cu number 8A80:

```
- SY4                                  Row 1 of 4 More:      > -
- Command ===> _____________________________________________ Scroll ===> CSR   -
-                                                            -
- Control unit number . : 8A80         ds8950F 75LBB10 SecB GM_SYA         -
- Control unit type . . : 2107         Serial number . . . :               -
-                                                            -
- Connected switch.ports: 30.0E 30.33 30.7E 31.0E 31.33 31.7E               -
-                                                            -
-                                                            -
-                                                            -
-                                                            -
- ENTER to continue.                                                 -
-                                                            -
-          ----------------Channel Path ID . Link Address----------------  -
- Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------  -
- ZDPR02.1  55.0E  73.0E  69.33  96.7E  77.33  A9.7E                  -
- ZDPR02.4  55.0E  73.0E  69.33  96.7E  77.33  A9.7E                  -
- ZUPR01.1  55.0E  73.0E  69.33  96.7E  77.33  A9.7E                  -
- ZUPR01.4  55.0E  73.0E  69.33  96.7E  77.33  A9.7E                  -
```

UA from 00 to 7F are base devices while UA 80 to FF are aliases.
Later on, the data residing on that old model 3 dasd were migrated to larger
capacity model and the physical dasd were removed from the box while the LCU and
device definition were retained in hcd to be eventually reused in the future.

02:29 PM CEST

Rina Forni (Customer)

update part 1:
I will try to give you a picture of Cedacri envinronment.
Cedacri has two cpu z15 hosting 19 z/os lpar belonging to different sysplexes.
Dasd storage data resides onto six box DS8950F.
When we activate a new i/o configuration, we issue the ACTIVATE IODF=xx,SOFT

command on every z/os lpar except for the last two lpar (one on each Z15)
where we issue the hw ACTIVATE IODF=xx command (eventually with the FORCE parameter if we perform any delete).
Due to the high number of system and dasd devices, we have many ficon channels for each box.
The box where we defined twice the same LCU has s/n 75LBB11 and it has the following ficon port attach:
box interface I0101 attached to FiconDir 30, port 0E
box interface I0000 attached to FiconDir 30, port 22
box interface I0111 attached to FiconDir 30, port 33
box interface I0010 attached to FiconDir 30, port 46
box interface I0200 attached to FiconDir 30, port 47
box interface I0301 attached to FiconDir 30, port 56
box interface I0311 attached to FiconDir 30, port 7E
box interface I0030 attached to FiconDir 30, port 82
box interface I0210 attached to FiconDir 30, port A9
box interface I0230 attached to FiconDir 30, port BA
box interface I0040 attached to FiconDir 30, port BB
box interface I0331 attached to FiconDir 30, port D9
box interface I0240 attached to FiconDir 30, port E6
box interface I0011 attached to FiconDir 31, port 0E
box interface I0100 attached to FiconDir 31, port 22
box interface I0031 attached to FiconDir 31, port 33
box interface I0110 attached to FiconDir 31, port 46
box interface I0300 attached to FiconDir 31, port 47
box interface I0231 attached to FiconDir 31, port 56
box interface I0241 attached to FiconDir 31, port 7E
box interface I0130 attached to FiconDir 31, port 82
box interface I0310 attached to FiconDir 31, port A9
box interface I0330 attached to FiconDir 31, port BA
box interface I0140 attached to FiconDir 31, port BB
box interface I0341 attached to FiconDir 31, port D9
box interface I0340 attached to FiconDir 31, port E8
box interface I0001 attached to FiconDir 50, port 07
box interface I0211 attached to FiconDir 51, port 07
02:28 PM CEST

Rina Forni (Customer)

Hello Vincent & Dirk,
we are trying to update the case but we are getting errors.
We will try to split the update in different part

01:58 PM CEST

liliya.radulova (IBM)

Hello Pier,
Thank you for the files.
I also noticed you did an empty update.
We are awaiting your answers to the questions from Vincent and Dirk.
Regards,
Liliya (for Vincent)

12:40 PM CEST

pierpaolo.mattavelli (Customer)

---

19 Apr 2022

03:36 PM CEST

vbonanno (IBM)

Hello Pier,
Sorry for the delay here, I was out of the office last week when my colleagues in Device Services reviewed further. To try to understand what configuration changes led to this, can you please help us understand in specific detail of what steps you performed that resulted in this misconfiguration that would help with possible recreate? Right now, there's no footprints that we can see thattell the sequence of events. For instance did you a few questions that came up
Did you delete the old SSSCB Device Table and add in the new table with incorrect data?
Can you explain how you defined these duplicate devices in the IODF, and the steps or sequence of events after?
Can you tell us if IEECINIT/IECDINIT are at the latest maintenance for R230? What ptf level are they are currently or f you've since applied PTFs since the problem occurrence, what ptf level were they at when you had this issue.?
Thanks in advance for your response, and back to you for now.
Regards,
Vincent & Dirk

---

1 Apr 2022

11:56 AM CEST

Torsten.Mueller (IBM)

I will inform Vincent about recent update.
Notified: @VINCENT BONANNO

11:46 AM CEST

pierpaolo.mattavelli (Customer)

Hello Vincent,
ascertained that the problem results from the overlapping of the devices, we are internally debating if, from a z/OS I/O management perspective, there are no self protecting mechanisms with the ability to prevent defining and bringing online devices belonging to two different LCU pointing to the same SSID.
I've now attached to the case SYA SMF 7x records and a RMF devices REPORT related to the overlapping addresses REPORTS(DEVICE(NMBR(8A80:8AFF,F280,F2FF))).
Report refers to the interval on which DEVICE F280 was bring ONLINE, this caused the abend of the ALLOCAS address space.
Please take a look and provide your considerations.
Many thanks, kind regards.
Pier Paolo

11:46 AM CEST

Torsten.Mueller (IBM)

Hi Pier,
we have received 2 new files for this case, but I can't find any comment7request about it.
Do they belong to this case?
Kind regards,
Torsten

---

28 Mar 2022

05:53 PM CEST

vbonanno (IBM)

Hello Pier,
Please see my colleague, Dirk, in DevServ support's previous update with their findings. I took a look at the SVC Dump provided, and at least at this point, Date: 03/22/2022 Time: 12:01:02.410679 , when the devices were brought online, the 'bad' vsmp' is still good at that point. From the overlay content it may have gone wrong shortly after then, as the device services like entry that is on top of the VSMP, at +8, 03/22/2022 12:01:09.600258 LOCAL X'DB3DDA71 5B402525' This was about 8 seconds after the dump started to be taken.
Regards,
Vincent
Action Taken:
Review of dump shows that it was taken on 3/22/2022 at 12:01:02.410679. Customer indicated this was when they brought the devices online..
From when devices brought online Dump

```
0A390000   E5E2D4D7   07CFF000   00001000   000000E2   | VSMP..0.......S |
0A390010   07A96310   DB363FC3   07ABC078   034C04E0   | .z.....C..{..<.\ |
0A390020   00000018   0CC71DDC   0821F070   DB3BC8E8   | .....G....0...HY |
0A390030   031152E8   08F76AD8   00000030   015041F2   | ...Y.7.Q.....&.2 |
```

From the overlay dump, taken: Date: 03/24/2022 Time: 10:44:35.378918 Local

```
0A390000   0735E8F8   07CFF000   DB3DCD08   21002525   | ..Y8..0..... |
0A390010   01030041   81CB4700   00000000   00000000   | ....a..... |
0A390020   00000018   0CC71DDC   0821F070   DB3BC8E8   | .....G....0...HY |
```

Taking a step back, and looking at the prior storage to the overlay: The storage at 0A38F040 was just obtained about a second earlier, at 12:01:01 versus dump time at 12:01:02.410679 by IOSVSLFD+0100 IN READ ONLY NUCLEUS

```
                                    Date       Time
ASID Job Name Id       St T Address  Length   Ret Addr MM/DD/YYYY HH:MM:SS
---- -------- -------- -- - -------- -------- -------- ---------- --------
   0000 *SYSTEM* ..... Ac S 0A38F040 00000FC0 0117D520 03/22/2022 12:01:01
==> 0000 *SYSTEM* ..... Ac S 0A390000 00001000 8152CAC0 02/03/2022 18:18:29
```

It's interesting to note that at the tail end of that xFC0 bytes, which would take us CLOSE to the damaged VSMP, we see, this IOS control block had pointers to UCBs, close to what we found in the overlay as part of the device table. I see many xFC0 areas being obtained around here, and do not see damage occurring.

```
0A38FDA0   00000000   0735E7A8   00000000   01000000   | ......Xy..... |       <------- 0735E7A8   (ucbs start w/ device F2A3)
0A38FDB0   098E9888   0A38FA08   00000000   80000000   | ..qh..... |
0A38FDC0.:0A38FF9F. LENGTH(X'01E0')--All bytes contain X'00'
0A38FFA0   00000000   0735E7A8   00000000   00000000   | ......Xy..... |
0A38FFB0.:0A38FFFF. LENGTH(X'50')--All bytes contain X'00'
0A390000   E5E2D4D7   07CFF000   00001000   000000E2   | VSMP..0.....S |
```

If I treat the overlay content +8 into the VSMP, as a timestamp: DB3DCD08 21002525, this converts to 03/22/2022 12:01:09.600258 LOCAL X'DB3DDA71 5B402525', versus when global data capture (common storage was dumped), 12:01:01.635330-03/22/2022 12:01:02.317347 . So I suspect sometime after the dump started was when the damage may have occurred from the information in the overlay content.
Action Plan:
Confirm Timestamp thoughts in the SSCBDEVE entry with Dirk. Report additional findings to customer / hold for feedback.
02:56 PM CEST

Dirk Gutschke (IBM)

Hello Pier,
here is Dirk from the DFSMS Device Services team.
regarding your comments:
We found a misconfiguration in HCD, the same DASD LCU with SSID BB47 and CUADD 47 has been defined twice in the same configuration with different FICONDIR ports and channels. Involved device ranges are 8A80-8AFF, F280-F2FF.
8A80 was already defined al IPL time while F280 was dinamically added (via ACTIVATE command) on 18th march. Is this the cause of the overlay?
Yes, I'm pretty sure that observed overlay is a direct result of the misconfiguration in HCD.
The data of the area that overlaid the VMSP ("overlaid vsmp") seems to be part of the STORAGE SUBSYSTEM STATUS CONTROL BLOCK structure.
More precise, the data seems to be a part of a 'DEVICE TABLE ENTRY (DESCRIBED BY STRUCTURE 'SSCBDEVE').

The UCB address

```
UCB 0735E8F8
```

within the "overlaid vsmp" data area: refers to

```
DEVICE 0.F2A5
```

which refers to

```
SERIAL=2107.996.IBM.750000000LBB11
SSID=BB47, LSS=47
CCA=25
```

but when I locate the entry for CCA=25 for SSID=BB47 in the DEVICE TABLE (at DVTP§=081821E0):

```
081821E0. E2E2C3C2 C4C80200 00000E20 00710020  20000020 F5F1F7F5 D3C2C2F1 F1000000 !SSCBDH.............5175LBB11...!
...
08182760. 0279C7A8 00000000 DB39188D 21002525  01030041 81CB4700 00000000 00000000 !.`Gy.................a.....!
```

then it refers me to a **different DEVN 8AA5 ?!?!**

```
UCB 0279C7A8 IS DEVICE 0.8AA5
```

but the serial#, SSID, LSS and CCA are the same as for DEVN F2A5

```
SERIAL=2107.996.IBM.750000000LBB11
SSID=BB47, LSS=47
CCA=25
```

That the reason why I stated that this overlay is a direct result of the mis-configuration in HCD.
But (at this point of time) I cannot explain why nor when the overlay has been occurred.
Would these information sufficient for you since is quite obviously that the overlay was been caused by the misconfiguration?
Your feedback is appreciated.
Kind regards, Dirk

ACTION TAKEN: reviewed DOCs; talked to Vincent
ACTION PLAN: await client feedback
02:41 PM CEST
vbonanno (IBM)
Hello Pier,
Thanks for the update and additional data. We'll review further. Thank you and you have a nice day as well.

Regards,
Vincent
Action Taken:
Update noted
Action Plan:
Review data.
01:30 PM CEST

pierpaolo.mattavelli (Customer)

Hello,
again on configuration activity, I've attached a DUMP taken by allocas at the time the new devices was brought online and a DEVSERV QD ouput command, here you can see the same devices with addresses starting with "8" and "F".
We hope this helps.
Have a nice day,
regards
Pier Paolo

---

25 Mar 2022

07:50 PM CET

vbonanno (IBM)

Hello Pier,
Thank you for sending in the EREP, in the interim, I have reached out to my colleagues in LOGGER and DEVSERV and they are reviewing the documentation further from their perspective. Your configuration activity is definitely interesting here. The UCB pointer which is +0 into the VSMP control block, represents a device number F2A5. It also looks like the overlay content resembles what an entry in the SSCBHD should be, but we do not see any other entries or the eyecatcher SSCBHD around the overlay.
We can't say exactly what is the cause of the overlay, but considering the content It does suggest something in the process with these devices is related. and the current focus.

```
UCBOB: 0735E8F8
CHAN..... F2A5
```

Regards,
Vincent
10:05 AM CET

Torsten.Mueller (IBM)

Hello Pier Paolo,
many thanks for the EREP and the additional information. Vincent will have a look when he is online in the afternoon.
Kind regards,
Torsten (for Vincent)

Notified: @VINCENT BONANNO

09:28 AM CET

pierpaolo.mattavelli (Customer)

Hello Vincent,
I uploaded new EREP reports from 08:00 as you requested, files are named EREPD08 (detail) and EREPS08 (summary).
We made an IPL yesterday night to reestablish the full system functionality.
Thanks a lot for your support,
kind regards
Pier Paolo
P.S.: We found a misconfiguration in HCD, the same DASD LCU with SSID BB47 and CUADD 47 has been defined twice in the same configuration with different FICONDIR ports and channels. Involved device ranges are 8A80-8AFF, F280-F2FF.
8A80 was already defined al IPL time while F280 was dinamically added (via ACTIVATE command) on 18th march. Is this the cause of the overlay?

12:05 AM CET

vbonanno (IBM)

Hello Pier,
I've been reviewing the documentation you provided, the CSAA861E , you reported is indicating that VSM's CSATRACKER function has been disabled. This has occurred, and is a result of an overlay. The control blocks (VSMP's) which represent the CSATRACKER data have been damaged. I can see, there has been corruption 1 VSMP control block, where the first 20 bytes have been damaged. Now that CSATRACKER is disabled, I would not expect you to see another abend, however, tracker cannot be enabled without an IPL. The EREP which covered about 5 minutes of time starting at 10;41, showed many abend1C5 (LOGGER abends) which occur before and after the damage detected in the VSMP, aside from this a few abend0C4s, none of these seem to show pointers to our corrupted data. From browsing the dump, I can see similar content of the overlay damage in device service control blocks.
At this time I do not see a connection to the overlay, but I am going to reach out to both logger and device services if they can review from their perspective and see if there is a connection to to the overlay.
Can you please send in more logrec, a few hours leading up to the time of the problem(10:45am) for further review.
Regards,
Vincent

12:01 AM CET

vbonanno (IBM)

Action Taken:
Review of the dump shows that it was taken for an abend0C4 occurring in a Virtual Storage Manager (VSM) freemain routine, IGVFSQA +DE at UA98724. This module is getting control for a freemain from Subpool xEF (SP239), an ESQA subpool.

```
0002 0001 009D71F8  SVC     78 00000000_0AFBA7AA  0000EF01 00000058 034C0510  Freemain                10:44:34.832229658 0042
                    07041000 80000000
0002 0001 009D71F8  PGM    010 00000000_015250AE  00040010 00000000          00040001 00000000 0001 0001 10:44:34.832233796 0042
                    04042000 80000000          69419800 <-        00000000
0000 044E 00913028  SSRV   133          00000000  00000103 00000830 2D8B97D0  Storage  Release           10:44:34.832244891 004B
                                       044E0000
0000 044E 00913028  PR     ...   0      2E6477AB 01269518                              044E
0000 044E 00913028  PC     ...   8      2E64764D             0030B          Storage  Obtain
0002 0001 009D71F8 *RCVY   PROG                  940C4000 00000010 00000000  00040001 00000000 0001 0001 10:44:34.832250824 0042
                                                 00000000
```

IGVFSQA was accessing GQE control blocks to remove the entry for the area being freed. I can see it tried to load R3 from R12+C. Register 69419070, being an invalid control block addresses, is ungetmained storage and causes the abend0C4-10.

```
 Symptom            Description
 -------            -----------
 PIDS/5752SC1CH      Program id: 5752SC1CH
 RIDS/NUCLEUS#L      Load module name: NUCLEUS
 RIDS/IGVFSQA        Csect name: IGVFSQA
 AB/S00C4            System abend code: 00C4
 PRCS/00000010       Abend reason code: 00000010
 REGS/90B9A          Register/PSW difference for R09:-0B9A
 RIDS/IGVRVSM#R      Recovery routine csect name: IGVRVSM
Time of Error Information

 PSW: 04042000 80000000 00000000 015250AE
 Instruction length: 04   Interrupt code: 0010
 Failing instruction text: 500C1E08 18AF5830 C00C1538
 Translation exception address: 00000000_69419800

 Breaking event address: 00000000_0152512C
 Registers 0-7
 GR: 034C0568 01030041 00000058 681E7C18  7F6C59E0 7F6C5A60 0000004C 025F9628
 AR: 00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000
 Registers 8-15
 GR: 034C0510 01525C48 01030041 00000006  69419070 7F6C5CC0 811363E0 02607758
 AR: 00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000
```

VSM control blocks appear to have been overlaid, and the overlay content, resulted in 69419070 being used as a GQEPTR. x0A390000 should be the xE2th VSMP control block. However, the header is overlaid, in fact the first 20 bytes appear to be damaged. +4 points to the next VSMP. When x0A390010 was used as a GQE, the storage address at x01030041 contains the value, x69419070, this was used then used as the GQEPTR.

```
overlaid vsmp
0A390000   0735E8F8   07CFF000   DB3DCD08   21002525   | ..Y8..0..... |
0A390010   01030041   81CB4700   00000000   00000000   | ....a..... |
0A390020   00000018   0CC71DDC   0821F070   DB3BC8E8   | .....G....0...HY |
0A390030   031152E8   08F76AD8   00000030   015041F2   | ...Y.7.Q.....&.2 |


good vsmp
07CFF000   E5E2D4D7   077EE000   00001000   000000E1   | VSMP.=\..... |
07CFF010   07ACBA60   DB400364   025F9A28   026F1300   | ...-. ...¬...?.. |
07CFF020   00000500   0192DB34   082DF3A0   DB3FB411   | .....k....3..... |
07CFF030   025F9A28   092FFB00   00000500   0192DB34   | .¬.....k.. |


where the bad GQEptr came from
01030041    694190   70405090   707CA7C8   00681E7C   |  .... &..@xH...@ |
         ^^^^^^^^^^
```

Reviewed system trace/logrec for abends that had addressability to xA390000, but did not see anything in registers or trace entries that showed a close connection to our overlaid vsmp. However, the logrec does show, recurring about a 1100 abend1C5, RSN804, that has been going on throughout the provided LOGREC and before / after the abend0C4 in VSM/overlay indication. The recurring abend1C5, appear
to be a recovery retry in IXGINLPA, re-driving the abend in H=01A3 P= 001A .
The LOGREC dataset covered errors from 10:41:59.63 - 10:46:59.84. The abend0C4 in VSM occurring at 10:44:35.378918. There were a few other abend0C4 in RMFGAT, but were not interesting, as not in key0 or had pointers to the overlaid storage in the registers.
SYMPTOM DESCRIPTION
------- -----------
PIDS/5752SCLOG PROGRAM ID: 5752SCLOG
RIDS/IXGINLPA#L LOAD MODULE NAME: IXGINLPA
RIDS/IXGF2BRW CSECT NAME: IXGF2BRW
AB/S01C5 SYSTEM ABEND CODE: 01C5
PRCS/00000804 ABEND REASON CODE: 00000804
TIME OF ERROR INFORMATION
PSW: 07040000 80000000 00000000 0D21F6E8

RECOVERY ENVIRONMENT
RECOVERY ROUTINE TYPE: FUNCTIONAL RECOVERY ROUTINE (FRR)
PSW AT ENTRY TO FRR: 070C0000 8D218130
Review of raw storage for the overlay content, I found very similar areas to the overlay content in the VSMP, which are for device services device tables, control block SSCBDH, However, I do not see any SSCBDH with entries with the same exact pattern

```
081821E0   E2E2C3C2   C4C80200   00000E20   00710020   | SSCBDH..... |
081821F0   20000020   F5F1F7F5   D3C2C2F1   F1000000   | ....5175LBB11... |
08182200   0279B5B0   00000000   DB39188D   21000A0A   | .`.............. |
       ptr toUCBs             timestamp?      xxxx varies
08182210   01030041   81CB4700   00000000   00000000   | ....a..... |
       same         same

08182220   0279B850   00000000   DB39188D   21000E0E   | .`.&..... |
08182230   01030041   81CB4700   00000000   00000000   | ....a..... |
08182240   0279C268   00000000   DB39188D   21001D1D   | .`B..... |
08182250   01030041   81CB4700   00000000   00000000   | ....a..... |
08182260   0279C508   00000000   DB39188D   21002121   | .`E..... |
08182270   01030041   81CB4700   00000000   00000000   | ....a..... |
08182280   0279BAF0   00000000   DB39188D   21001212   | .`.0..... |



overlaid vsmp:
0A390000   0735E8F8   07CFF000   DB3DCD08   21002525   | ..Y8..0..... |
       ptr toUCBs             timestamp?      xxxx varies
0A390010   01030041   81CB4700   00000000   00000000   | ....a..... |
       same         same
```

Action Plan:
Discuss abend1C5 with LOGGER support to see if they can determine any connection to the overlay
Discuss the SSCBDH related storage that looks to be the content that damaged the vsmp.

---

24 Mar 2022

11:54 PM CET
vbonanno (IBM)
.

05:27 PM CET

vbonanno (IBM)

Hello Pier,

Thank you for sending in the dump, logrec and syslog. I have begun review of the data you have sent in and will respond as soon as I can with details.

Regards,

Vincent

Action Taken:

Updates noted,

action Plan:

reviewing data.

03:52 PM CET

pierpaolo.mattavelli (Customer) created this case