

This is Google's cache of <http://forums11.itrc.hp.com/service/forums/questionanswer.do?threadId=1452274>. It is a snapshot of the page as it appeared on 14 Jun 2011 08:36:08 GMT. The [current page](#) could have changed in the meantime. [Learn more](#)

These search terms are highlighted: **crashconf dump space**

[Text-only version](#)

English

- » HP Home
- » Products & Services
- » Support & Drivers
- » Solutions
- » How to Buy

Search:

[» Forums advanced search](#)

» [Contact HP](#)

[IT Resource Center \(Forums\)](#) [All of HP \(US\)](#)

[IT Resource Center Forums](#) > [HP-UX](#) > [system administration](#)



hp-ux 11.31 **dump** device using San external LUN - This thread has been closed

» IT Resource Center

- » [Login](#)
- » [Register](#)
- » [My profile](#)

[Create a new message](#) [Receive e-mail notification if a new reply is posted](#)

[Reply to this message](#)

- » [Search knowledge base](#)
- » [Forums](#)
- » [Patch database](#)
- » [Download drivers, software and firmware](#)
- » [Warranty check](#)
- » [Support Case Manager](#)
- » [Software Update Manager](#)
- » [Training and Education](#)
- » [More maintenance and support options](#)

- » [Online help](#)
- » [Site map](#)

Author

[jose da silva](#)

Subject: hp-ux 11.31 dump device using San external LUN [Add to my favorites](#) **This thread has been closed**

Oct 17, 2010 20:11:21 GMT

Hello

I have successfully created secondary swap **space** on an external SAN LUN using the persistant devices. Now having several issues to configure an external lun for **DUMP**. The BL890c has 400GB of ram so /var/adm/crash is not large enough to hold a **dump** for analysis. Is it possible to redirect the **dump** " crash " to a different location ? What are the best practices?

Note: If you are the author of this question and wish to assign points to any of the answers, please login first. For more information on assigning points ,click [here](#)

Sort Answers By: [Date](#) or [Points](#)

[Chandrahasa s](#)

Expert in this area
This member has accumulated 500 or more points

Oct 18, 2010 02:44:41 GMT 8 pts

Hi,

I am facing server problem....

What are those several problems??

You can specify crash path on file

/etc/rc.config.d/savecrash

You need to specify path like

SAVECRASH_DIR=/var/adm/crash

Your memory size is 400GB so you also need to consider **dump** speed. HPux v3 came up with some new feature which will help to reduce **dump** time.

chandra

Assigning points will motivate each one in this forum

[jose da silva](#)

Oct 18, 2010 04:07:56 GMT N/A: Question Author

Hello Chandra

Thanks for your quick response. I think you have given me the correct hint. It seems I will have to assign /var/adm/crash path to a different location for example **/dump**. What I would like to know if I can assign a link between /var/adm/crash to **/dump** or should I specify the new **dump** path **/dump** in /etc/rc.config.d **/crashconf** ?

Member icons

- HP moderator
- Expert in this area

Member status

- ITRC Pro** *ITRC Pro*
250 points
- ITRC Graduate** *ITRC Graduate*
500 points
- ITRC Wizard** *ITRC Wizard*
1000 points
- ITRC Royalty** *ITRC Royalty*
2500 points
- ITRC Pharaoh** *ITRC Pharaoh*
7500 points
- Olympian** *Olympian*
20000 points
- 1-Star Olympian** *1-Star Olympian*
40000 points
- 2-Star Olympian** *2-Star Olympian*
80000 points

» [How to earn points](#)

» [Support forums](#)
[FAQs](#)

Question status

Magical answer



Message with a response that solved the author's question

Favorites status



Add to my favorites



Delete from my favorites

This thread has been closed
 Thread closed
 been closed

[Duncan Edmonstone](#) This member has accumulated 20000 or more points

Warm Regards,
 Jose

Oct 18, 2010 06:40:50 GMT 8 pts

Seems to be some confusion around terminology here... /var/adm/crash is not the **dump** device... it is a filesystems where dumps are written out to on reboot *following a **dump*** from the **dump** device.

The **dump** device is a raw partition that HPUX can write a **dump** to without having to worry about complex kernel functions like writing to filesystems etc... it is never a filesystem.

By default on most HPUX systems, the **dump** device is primary swap... you can check this using:

```
lvlnboot -v
```

HTH

Duncan

[Daniel Parkes](#) This member has accumulated 1000 or more points

Oct 18, 2010 09:27:42 GMT 8 pts

You can also check your crash devices using the **crashconf** command:

crashconf

Crash **dump** configuration has been changed since boot.

CLASS PAGES INCLUDED IN **DUMP** DESCRIPTION

```
-----
UNUSED 1919494 no, by default unused pages
USERPG 1051893 no, by default user process pages
BCACHE 409946 no, by default buffer cache pages
KCODE 7140 no, by default kernel code pages
USTACK 8640 yes, by default user process stacks
FSDATA 0 yes, by default file system metadata
KDDATA 740852 yes, by default kernel dynamic data
KSDATA 3098 yes, by default kernel static data
SUPERPG 35735 no, by default unused kernel super pages
```

Total pages on system: 4176798
 Total pages included in **dump**: 752590

Dump compressed: ON

DEVICE OFFSET(kB) SIZE (kB) LOGICAL VOL. NAME

```
-----
31:0x006000 1547124 4194300 64:0x000002 /dev/vg00/lvol2
31:0x006000 5741428 4194300 64:0x000010 /dev/vg00/lvswap2
-----
8388600
```

It's better not to use a symbolic link, you should change the dir in the /etc/rc.config.d/**crashconf** file

[jose da silva](#)

Oct 18, 2010 10:10:43 GMT N/A: Question Author

Hello Colleagues

Duncan you are correct I confused the issue between /var/adm/crash and **dump space**. I am going to give full details on the problem I am having which I have found a workaround.

The HP-UX 11.31 blade has a total of 400 GB of memory so I created in vg00 swap **space** of 50GB "vol2" and decided to do the following:-
 1. Create secondary swap **space** of 800 GB on an external SAN LUN " HITachi"
 2. Create a **dump space** of 400GB on an external SAN LUN.

I think the problem is that the Storage Administrator configured disk sizes of 70GB each,that is for swap "70GBx12disks=840GB" and for **Dump** "70GBx6disk=420GB". We created VG's with -p 32 -s 65535
 So when we create the volume groups and try to create the logical volumes with Strict and Contiguous allocation policies this process fails " Not enough free physical extents available. Failure possibly caused by Contiguous allocation

policy"

When I tested using only a 70GB lun with a size of 70GB disk " that is one disk only" the process passes no problems. So due to this fact it seems that for swap and **dump** it must be one LUN size consisting of one disk " NO EXTENDING of VG or LVs"

Please can you verify and give me advise on this problem. Perhaps I am missing something in HP-UX 11.31 version. Also since primary swap lvol2 was created with LV version 1 should we try to use LV version 2 and above?

Warm Regards,
Jose da Silva

[Daniel Parkes](#) This member has accumulated 1000 or more points

Oct 18, 2010 11:23:48 GMT 7 pts

Hi,

You don't need 800 gigs of swap **space**, at most if you didn't want to use pseudo swap, you could create a 400 gig swap device(it's always better to have several devices on different controllers/arrays with the same priority, than one big chunk).

If you use pseudo swap(default behaviour), you would be ok with much less **space**

[Daniel Parkes](#) This member has accumulated 1000 or more points

Oct 18, 2010 11:33:55 GMT 7 pts

Pseudo-Swap

Pseudo-swap **space** allows for the use of system memory as a third type of swap **space**. That is, HP-UX swap **space** can also consist of up to seven-eighths (87.5%) of system memory capacity.

For example, a computer with one GB of system memory and one GB of device and file system swap, can run up to 1.87 GB of processes. If any process attempts to grow or be created beyond this extended threshold, the process will fail.

When using pseudo-swap, since more processes can be created, the system load increases, causing more paging and deactivation activity.

By default, pseudo-swap **space** is configured to be available. If you do not wish to make use of it, you will need to re-set the tunable system parameter, `swappmem_on`, to 0 ("off"). (To modify a configurable parameter, see "Making Adjustments to Your System" in Chapter 1.)

[Chris Huys](#) This member has accumulated 1000 or more points

Oct 18, 2010 12:19:15 GMT 9 pts

Hi Jose,

Its not necessary to have a 400GB /var to "hold" a 400GB "full" crashdump.

The crashdump gets, when `COMPRESS=1` is set in `/etc/rc.config.d/savecrash`, gzipped when written from "physical **dump** device(s)" to `/var/adm/crash`.

Therefor the /var filesystem, would need to have 80GB, 20% of full crashdump, of free disk space available to save a "full" crashdump in `/var/adm/crash`.

Increase /var to f.e. 88GB, 80GB for **dump** and 8 GB for the rest of /var, and do a test with a full crashdump, i.e. change first with `crashconf` the parameters so that everything gets dumped, `#crashconf -i all;crashconf -v`, and then toc the system from the mp. `MP>CM>tc`.

NOTE: Default, when a system "crashes", only "selective" crashdump are dumped, which would only **dump** a fraction, check `crashconf -v` output, of the 400GB of RAM to physical dumpdevice(s), however in some cases, HP support will ask for a full crashdump and for this occasion, the system should indeed be "prepared" upfront.

For the "contiguous problem" of the question. Yes, I think, you cant have the **dump** logical volume, "cross" physical disk/lun boundaries, but on the other hand, you can for each of your 6 "70GB" "**dump**" "luns", create a separate "**dump**" "vol", i.e. `dumpvol1` of 70Gb/`dumpvol2` of 70G.. till `dumpvol6` of 70Gb, and add all this "dumpvolx" to `/etc/fstab`, f.e. `/dev/vg00/dump1 / dump defaults 0 0`, and also to `crashconf`, with `crashconf -s <dumpvol>`, also delete primary `dumpvol`, `/dev/vg00/lvol2` first from the list of `dumpvols` defined in `crashconf`, `crashconf -ds /dev/vg00/lvol2`. If then a crashdump occurs, the **dump** should be written consecutively(sp?) to each specific dumpdevice that it encounters in

crashconf.

NOTE: In HP-UX 11.31, it might even be that the dumpvolumes dont need to part anymore of /dev/vg00, well with vxvm this is the case, didnt test it with LVM. The dumpluns, where the dumpvolumes are created upon, however should have been made "lif'able" "upfront", through f.e. pvcreate -B /dev/rdisk/disk114;mkboot -e -l /dev/rdisk/disk114, to have them be part of the "persistent **dump** device list" (**crashconf** -v) , which is a absolute necessity, to have the dumpconfiguration "survive" reboots.

Greetz,
Chris

[Dennis Handly](#)

This member has accumulated 40000 or more points

Oct 18, 2010 15:04:40 GMT 8 pts

>Daniel: If you use pseudo swap (default behaviour), you would be ok with much less **space**

The default and only behavior for 11.31.

[Daniel Parkes](#) This member has accumulated 1000 or more points

Oct 18, 2010 15:35:21 GMT 7 pts

Thnx Dennis, I thought that was the case with 11.31 but wasn't sure.

[Chris Huys](#) This member has accumulated 1000 or more points

Oct 18, 2010 15:52:44 GMT 7 pts

> Thnx Dennis, I thought that was the case
> with 11.31 but wasn't sure.
It isnt its just dennis opinion. (and im offcourse of the opposite opinion) ;)

[jose da silva](#)

Oct 18, 2010 17:36:08 GMT N/A: Question Author

Thanks All for your inputs and guidelines.

Chris you stated the following which is very important:-

NOTE: In HP-UX 11.31, it might even be that the dumpvolumes dont need to part anymore of /dev/vg00, well with vxvm this is the case, didnt test it with LVM. The dumpluns, where the dumpvolumes are created upon, however should have been made "lif'able" "upfront", through f.e. pvcreate -B /dev/rdisk/disk114;mkboot -e -l /dev/rdisk/disk114, to have them be part of the "persistent **dump** device list" (**crashconf** -v) , which is a absolute necessity, to have the dumpconfiguration "survive" reboots.

Some of the docs stated the following HP-UX 11.31:

1. Select an unused disk example /dev/disk/disk52
2. pvcreate /dev/rdisk/disk52
3. vgcreate -p 32 -s 65535 /dev/vgdump
4. lvcreate -C y -r n -n dump_lvol /dev/vgdump
5. Add the new **dump** lvols in /etc/fstab as shown below:
(ie: /dev/vgdump/dump_lvol... **dump** defaults 0 0)
6. **crashconf** -a ; (add to **dump** configuration)
7. **crashconf** -s ; **crashconf** -v (For persistent **dump** across reboots)

Is the above correct, I had a look at some of the docs and it is not mentioned to pvcreate -B and mkboot cmds.Please advise

[Chris Huys](#) This member has accumulated 1000 or more points

Oct 18, 2010 21:25:49 GMT 8 pts

Hi Jose,

It doesnt look like its possible for the moment, with LVM, to have "persistent **dump** devices" "outside" of /dev/vg00.

NOTE: The "persistency" of the "**dump** devices", to have the dumpconfiguration survive reboots, is, imo, a "necessity" for dumpdevices outside of /dev/vg00..

From the latest HP-UX 11.31 **crashconf** patch, which does a update of the **crashconf** manpage.

PHCO_38958 11.31 **crashconf**(1M) cumulative patch

[..]

Defect

(QX:QXCR1000847727)

crashconf (1M) man page does not contain the information

"The logical volumes which are not part of the root volume

group cannot be configured as persistent **dump** devices".

Resolution:

crashconf (1M) man page is updated with the below information

"The logical volumes which are not part of the root volume group cannot be configured as persistent **dump** devices".

You could open a call with HP support and do a crosscheck and if the above is all "true", ask for the "LVM" feature to be added in a future patch.

For the rest, well you could always join the dark side, i.e. switch from lvm to vxvm as bootdisks and thus also for the dumpdevices, were it does work outside "rootdg" . ;)

on vxvm.

ronin # **crashconf** -v

Crash **dump** configuration has been changed since boot.

CLASS PAGES INCLUDED IN **DUMP** DESCRIPTION

- UNUSED 1638412 no, by default unused pages
- USERPG 145921 no, by default user process pages
- BCACHE 19943 no, by default buffer cache pages
- KCODE 11204 no, by default kernel code pages
- USTACK 1004 yes, by default user process stacks
- FSDATA 14 yes, by default file system metadata
- KDDATA 262472 yes, by default kernel dynamic data
- KSDATA 7955 yes, by default kernel static data
- SUPERPG 4744 no, by default unused kernel super pages

Total pages on system: 2091669

Total pages included in **dump**: 271445

Dump compressed: ON

Dump Parallel: ON

DEVICE OFFSET(kB) SIZE (kB) LOGICAL VOL. NAME

```

1:0x00000c 2350176 4194304 2:0x000001 /dev/vx/dsk/rootdg/swapvol
1:0x000048 515168 8388608 2:0x2892b8 /dev/vx/dsk/secdumpswapdg
/secdumpswapvol

```

12582912

Persistent **dump** device list:

/dev/vx/dsk/secdumpswapdg/secdumpswapvol

Greetz,
Chris

[VK2COT](#) Expert in

this area This member has accumulated 2500 or more points

Oct 18, 2010 21:27:30 GMT 10 pts

[Attachement is 363433.pptx](#)

Hello,

It is actually an excellent idea to separate **dump space** from boot volume group.

I wrote two presentations on this topics recently.

Note that compressed **dump** has ONE VERY CRITICAL REQUIREMENT:

Five CPUs per each **dump** unit

My presentation on swap/**dump** design is enclosed herewith.

I hope you find it useful.

I also wrote about my view and method for best practices for HP-UX server design and build. I can send it to this forum if there is any interest.

VK2COT

[jose da silva](#)

Oct 19, 2010 05:51:03 GMT N/A: Question Author

Hello V2COT

Thankyou for the valuable information, the slide presentations are excellent.

Can you also share with us your doc with your views and methods for best practices for the HP-UX server design and build.

In one of my previous responce I stated that the HP-UX 11v31 Admin Guide does not contain full examples on how to create the swap and **dump** spaces from scratch, they only refer to the lvcreate, **crashconf** and savecrash cmds it would be helpfull if it contained examples.

Warm Regards,
Jose

[VK2COT](#) Expert in
this area This
member has
accumulated 2500
or more points

Oct 19, 2010 10:35:52 GMT 10 pts

[Attachement is 363440.pptx](#)

Hello,

I am glad that you find the swap design document useful.

Here is the presentation on HP-UX build and good practices. It contains some of the things I deal with when I design and build HP-UX servers.

To add additional **dump** areas, the steps are quite simple. An example (I am not at my desk so the commands might be missing some flags):

```
# lvcreate -C y -r n -L 8192 -n lvdump2 /dev/vgextP01
```

```
# vi /etc/fstab
/dev/vgextP01/lvdump2 ... dump defaults 0 0
```

```
# crashconf -a
```

Cheers,

VK2COT

[jose da silva](#)

Oct 19, 2010 18:16:53 GMT N/A: Question Author

Hello All Concerned

Thanks for the information that I received which was helpfull. A special thanks to VK2COT for the slide presentations which gave us all a better understanding on the swap and **dump** settings for HP-UX 11v31.

Warm Regards to ALL
Jose da Silva

[jose da silva](#)Oct 19, 2010 18:19:30 GMT **Thread closed by author**

Thanks Once again to all involved. As per the information I have received I have configured the swap and **dump** spaces.

Warm Regards,
Jose

[Create a new message](#) [Receive e-mail notification if a new reply is posted](#)
[Reply to this message](#)

 [Printable version](#)

[Privacy statement](#)[Using this site means you accept its terms](#)

© 2011 Hewlett-Packard Development Company, L.P.