

Installing and Configuring the Compressed Dump Utility on HP-UX 11i v1

By: Euan McMaster, Senior Technology Consultant, BCS Customer Care Group



Introduction.....	2
Compressed Dump Overview.....	3
Supported System Configurations.....	5
Swap and Dump Device Considerations.....	6
Dump Time Comparisons With and Without Compression.....	7
Test System Configuration.....	7
Test Scenarios.....	7
For more information.....	8
Acknowledgments.....	8

Introduction

This paper discusses how to use the Compressed Dump utility (also known as CDUMP11i) on HP-UX 11i, and provides guidelines on when it might be appropriate to do so.

The Compressed Dump utility is primarily targeted for “large memory machines”; that is, those machines with more than four processors and physical memory greater than 2 GB.

The goal of the Compressed Dump feature is to increase the speed of the memory dump for HP-UX in the event of a system crash, so that dumps are taken faster and system availability is improved.

For an overview of the benefits see the Compressed Dump Overview section.

Dumps are compressed as long as certain conditions are met. These conditions are specified in the *HP-UX Compressed Dump* white paper, which you can access from the following URL:

http://docs.hp.com/en/5434/Cdump_WP.book.pdf

That white paper refers to HP-UX 11i v2 but the same conditions will also apply to HP-UX 11i v1.

The utility is not installed by default on HP-UX 11i v1; it became available in the June 2003 Software Pack:

<http://docs.hp.com/en/5187-2783/5187-2783.pdf>

The “Software Pack – Optional HP-UX 11i Core Enhancements” CD is included in the latest version of the HP-UX 11i v1 media kit.

You can download the entire Software Pack, or components, from the HP Software Depot at the following URL:

<http://h20293.www2.hp.com/portal/swdepot/displayProductInfo.do?productNumber=SWPACK>

Compressed Dump Overview

The Compressed Dump functionality affects dump configuration, dump-time console user-interface, dump file format, access library, and dump manipulation commands.

With the Compressed Dump feature, kernel dump and save times are reduced. The actual time reduction and dump space reduction depend on various factors like the throughput of the device configured for dump and the compressibility of memory.

With devices that give less than 15 MB/sec of sustained throughput for uncompressed dumps, enabling compression will result in a speed increase at least three times faster for dumping the default selection of page classes. On faster devices, the speed increase will be even greater. The disk space requirement should also be reduced by a factor of at least three for the default selection of page classes. The downtime and dump space requirements of the system are consequently reduced.

For example, on an HP-UX 11i v1 system that typically requires three hours to dump with the default dump utility, using the Compressed Dump option typically reduces the dump time to one hour, including the time for `savecrash` to copy the raw dump image to the crash dump directory.

The reduction in dump time observed during the tests that are detailed in the Dump Time Comparisons With and Without Compression section was a factor of five. In other lab tests, reduction times of greater than seven have been observed.

If a system crashes soon after booting, when most of the memory pages have not been touched by the operating system, and if the UNUSED class is also included in the dump, then Compressed Dump performance may be as slow as an uncompressed dump. This is because the UNUSED pages contain random data that could not be compressed.

At the time of a system crash, the dump subsystem examines the state of the system and its resources to determine whether it is possible to use compression. To reduce the time it takes for compression, additional processors are used to do compression in parallel with I/O. Also, since the dumped pages are compressed, additional data is maintained in the dump to check for the validity of the dump. These include checksums and the information required to uncompress the pages. If the dump device has adequate space to hold the compressed dump and enough processors are available to do compression consistent with I/O capability, then the system will proceed with the compressed dump.

After a system crash, you can change the compression choice by using the console. The compression choices are either a full dump or a selective dump.

When the system is dumping in compressed mode, dump progress will be updated on the console at least once every 15 seconds.

When `savecrash` runs during boot, the dump is copied into the crash dump directory.

`savecrash` will run in a shorter time when copying a compressed dump since the dump size will be smaller compared to an uncompressed dump.

The `savecrash` option `-z` is ignored since the dump is already compressed. The input of a chunksize smaller than the size of memory that belongs to one compression unit for compressed dump will also be ignored.

You can also analyze a compressed dump by using the latest versions of the kernel dump debug tools, `q4`, `adb`, and `kwdb`, linked with the latest `libcrash`.

Versions of these tools that support the new dump format are shown in the following table:

Tool	Version
q4	B.11.221 (printed on startup)
adb	11.23.01 (output of the \$B command in adb)
kwdb	2.1.5 (printed on startup)

You can download these or later versions of the tools from the Developer and Solution Partner Program Web site:

<http://h21007.www2.hp.com/dev/technologies/topic/1,2608,10301,00.html>

To use older versions of these tools, you can convert the dump into any of the older formats by using the `crashutil` command.

For example, the `crashutil -v CRASHDIR /var/adm/crash/crash.0 /var/adm/crash/crash.1` command will create `crash.1` in `CRASHDIR` format from `crash.0`, which is in `PARDIR` format. You can then use `crash.1` for debugging with versions of `q4`, `adb`, or `kwdb` that may be older than those shown in the previous table.

Supported System Configurations

The compression feature of the Compressed Dump utility is supported on PA-RISC based systems running HP-UX 11i v1 as long as they have more than four processors and at least 2 GB physical memory. On a system running vPars, each vPar must have more than four processors and must have at least 2 GB memory assigned to it.

The following additional requirements apply to PA-RISC systems running HP-UX 11i v1:

The patch baseline must include the following patches:

PHKL_24056

PHKL_28237

PHCO_28240

PHCO_28242

PHCO_28243

PHKL_28244

PHKL_28263

PHKL_28265

PHCO_29995

PHCO_30361

PHKL_30607

PHKL_33339

PHKL_34106

The vPar version must be at least V3.02 and the supporting patches must be installed.

Dump compression is not enabled by default; it must be enabled with the `crashconf -c` command.

Swap and Dump Device Considerations

The performance of the reboot after a dump is affected by the configuration of the swap and dump devices.

Dedicated dump devices will not shorten the time required to write from memory to the dump lvol, during the crash, but will shorten the reboot time. This is because the crash images are not at risk of being overwritten by page or swap activity, and `savecrash` can run in the background to save the crash files to the crash dump directory. Meanwhile, page and swap can be enabled at their normal point in the boot sequence and the boot can proceed in parallel.

However, if the dump device is also configured as one of the swap devices, the device cannot be enabled for paging until `savecrash` has finished saving the image from the device to the crash dump directory; therefore, the boot time will be longer. This extra time will be even greater if vPars are configured because multiple dump images may have to be saved.

As systems continue to be configured with more memory and, in many cases, more vPars, the difference in boot time will become more noticeable. Fortunately, storage space continues to reduce in price, and the optimum configuration of separate swap and dump devices becomes more affordable.

The reduction in reboot time achieved by configuring a separate dump device is likely to provide a worthwhile return on investment when system availability is a priority.

Dump Time Comparisons With and Without Compression

The following sections provide information on dump time comparisons with and without compression.

Test System Configuration

The test system configuration contained the following components:

- 64-way 380-GB SuperDome system PA-RISC 8800 CPUs
- HP-UX 11i v1 with the patches shown in the Supported System Configurations section
- 1 nPar, 2 vPars
- Running SX system exerciser to simulate a heavy load during testing
- Crash image size for each vPar: about 8 GB

Test Scenarios

The following table describes various test scenarios and the time it took to accomplish each test.

Scenario No.	Description	Time
1	1 nPar without compression	89 min
2	1 nPar with compression	17 min
3	1 nPar, 2 vPars, without compression	274 min
4	1 nPar, 2 vPars, with compression	82 min
5	1 nPar, 2 vPars, with compression and separate dump device	80 min

The times shown are the total time elapsed between the start of the crash and the point at which the system (with all partitions, where appropriate) is again available for use.

Note the additional time elapsed during test scenario 4, when swap and dump are configured to use the same device(s), compared with test scenario 5, when dedicated dump device is configured. This is the time taken for `savecrash` to write the dumps to the crash dump directory, before paging can be enabled and the boot can proceed. The additional time is proportional to the number of vPars. In the test scenarios, the crash images were quite small and there were only two vPars. In real-life situations the difference may be considerably greater.

There is no guarantee that the same relative savings will be made on customers' production systems as the compression algorithms depend on the pattern of data in memory. Nevertheless, a saving of three times or more can be expected.

For more information

HP-UX Compressed Dump A.01.01 Release Notes
First Edition: June 2003, 5187-2760, HP-UX 11i version 1 (B.11.11)
CD-ROM (Software Pack June 2003)

<http://docs.hp.com/en/5187-2760/index.html>

HP-UX Compressed Dump A.01.01 Release Note, Edition 2: E0905, 5991-2881

<http://docs.hp.com/en/5991-2881/ch01.html>

HP-UX 11i June 2003 Release Notes
Tenth Edition: June 2003, 5187-2783, (HP-UX 11i), CD-ROM (Instant Information)

<http://docs.hp.com/en/5187-2783/index.html>

Compressed Dump White paper Version 1.3

http://docs.hp.com/en/5434/Cdump_WP.book.pdf

HP-UX System Administration Tasks, Chapter 6. Managing Swap Space and Dump Areas

<http://docs.hp.com/en/B2355-90672/ch06.html>

Managing Systems and Workgroups, Chapter 6, Managing Swap and Dump

<http://docs.hp.com/en/B2355-90950/ch06s03.html>

HP-UX Virtual Partitions Release Notes
vPars A.03.02, December 2004

<http://docs.hp.com/en/6019/a.03.02.rls.notes.pdf>

Acknowledgments

The author would like to thank the following individuals who contributed in some way to the content of this white paper:

Bob Bowden
Brenda L. Dean
Geoff Gardner
John Harrigan
Bill Kirk
Wendy Rannenber

© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

October 2006

