

# Installing and Configuring the Compressed Dump Utility on HP-UX 11i V1

Euan McMaster  
Senior Technology Consultant  
BCS Customer Care Group



HP-UX, ServiceGuard and Integrity are all registered Trademarks of Hewlett-Packard Company

# Table of contents

|     |  |    |
|-----|--|----|
| 1   | Revision Control .....                                   | 3  |
| 2   | Introduction .....                                       | 4  |
| 3   | Compressed Dump Overview .....                           | 5  |
| 4   | Supported System Configurations .....                    | 7  |
| 5   | Swap and Dump Device Considerations.....                 | 8  |
| 6   | Dump Time Comparisons With and Without Compression ..... | 9  |
| 6.1 | Test System Configuration .....                          | 9  |
| 6.2 | Test Scenarios.....                                      | 9  |
| 7   | References .....   | 10 |
| 8   | Acknowledgements .....                                   | 11 |

# 1 Revision Control

| Date                       | Revision | Status          | Comment |
|----------------------------|----------|-----------------|---------|
| 23 <sup>rd</sup> June 2006 | 0.9      | Initial draft   |         |
| 28th July 2006             | 0.95     | Final draft     |         |
| 12th August 2006           | 1.0      | Initial release |         |

## 2 Introduction

This paper discusses the use of the Compressed Dump Utility (also known as CDUMP11i) on HP-UX 11i V1, and provides guidelines on when it might be appropriate to do so.

The Compressed Dump Utility is primarily targeted for “large memory machines”, i.e. those with more than 4 processors and physical memory greater than 2GB.

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.

An overview of the benefits is provided herein, in [Section3: Compressed Dump Overview.](#)

Dumps are compressed as long as certain conditions are met. These conditions are specified in the whitepaper at the following URL:

[http://docs.hp.com/en/5434/Cdump\\_WP.book.pdf](http://docs.hp.com/en/5434/Cdump_WP.book.pdf)

The above whitepaper refers throughout to HP-UX 11i V2 but the same conditions 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 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>

### 3 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 1.5MB/sec of sustained throughput for uncompressed dumps, enabling compression will result in a speed increase of at least 3X 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 3X for the default selection of page classes. The down-time and dump space requirements of the system are consequently reduced.

For example, on an HP-UX 11i V1 system that typically requires 3 hours to dump with the default dump utility, use of the Compressed Dump option typically reduces the dump time to 1 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 [Section 6 – Dump Time Comparisons With and Without Compression](#) was a factor of 5X. In other lab tests, a time reduction of 7X has been observed.

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

At the time of a system crash, the dump sub-system examines the state of the system and its resources to determine whether it is possible to use compression. To reduce the time taken for compression, additional processors are used to do compression in parallel with I/O. Also, since pages are dumped compressed, additional data is maintained in the dump to check for the validity of the dump. These include checksum and 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 compressed dump.

After a system crash, the choice for compression can be changed using the console, as can the choice of full or selective dump.

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

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

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

`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.

A compressed dump can be analyzed 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  |
|-------------------|--|
| <code>q4</code>   | B.11.221 (printed on startup)                                      |
| <code>adb</code>  | 11.23.01 (output of <code>\$B</code> command in <code>adb</code> ) |
| <code>kwdb</code> | 2.1.5 (printed on startup)   |

These, or later, versions of the tools can be downloaded from the developer & solution partner program website:

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

To use older versions of the same tools, the dump can be converted into any of the older formats using `crashutil`.

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

## 4 Supported System Configurations

The compression feature of the utility is supported on PA-RISC based systems running HP-UX 11i V1 as long as they have more than 4 processors and at least 2GB physical memory. On a system running VPARs, each VPAR must have more than 4 processors and at least 2GB 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 supporting patches must be installed.

Dump compression is not enabled by default; it must be enabled by ``crashconf -c``

## 5 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 and the boot time will thus 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 and 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.

## 6 Dump Time Comparisons With and Without Compression

### 6.1 Test System Configuration

- 64-way 380GB SuperDome system PA-RISC 8800 CPUs
- HP-UX 11i V1 with patches shown in [Section 4 – Supported System Configurations](#)
- 1 nPar, 2 vPars
- Running SX system exerciser to simulate a heavy load during testing
- Crash image size for each vPar ~8GB

### 6.2 Test Scenarios

| No. | Description   | Time    |
|-----|---|---------|
| 1   | single nPar without compression                                 | 89 min  |
| 2   | single nPar with compression                                    | 17 min  |
| 3   | single nPar, 2 vPars, without compression                       | 274 min |
| 4   | single nPar, 2 vPars, with compression                          | 82 min  |
| 5   | single 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 2 vPars. In real life situations the difference is likely to be considerably greater.

There is no guarantee that the same relative saving will be made on customers' production systems as the compression algorithms are dependent on the pattern of data in memory. Nevertheless, a saving of 3X, or more, can be expected.

## 7 References

### 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 Notes

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](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>

## 8 Acknowledgements

The author would like to thank the following, who contributed in some way to the content of this whitepaper:

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