

Linux on System z



Using the Dump Tools on Red Hat Enterprise Linux 7

Linux on System z



Using the Dump Tools on Red Hat Enterprise Linux 7

Note

Before using this information and the product it supports, read the information in “Notices” on page 71.

This edition applies to Red Hat Enterprise Linux 7 on IBM System z, and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2014.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this book	v	Chapter 9. Handling large dumps	35
Other relevant Linux on IBM System z publications	v	Compressing a dump using makedumpfile	36
Chapter 1. Tools overview	1	Compressing a dump using gzip and split	37
The kdump tool	2	Chapter 10. Sharing dump devices.	39
Stand-alone tools	3	Serialization and device locking	39
The VMDUMP tool	3	Sharing devices when dumping manually	40
Live-system dump	4	Sharing DASD devices on LPARs	40
Dump methods compared	4	Sharing DASD devices under z/VM	40
Chapter 2. Using kdump.	7	Sharing SCSI devices	40
How kdump works on System z.	7	Using attach and detach as locking mechanism	
Setting up kdump	9	under z/VM	41
How kdump is triggered	10	Sharing devices when dumping automatically.	41
Chapter 3. Using a DASD dump device	13	DASD (dump or dump_reipl panic action)	41
Installing the DASD dump tool.	13	DASD (vmcmd panic action)	41
Initiating a DASD dump	14	FCP-attached SCSI devices	42
Copying the dump from DASD with zgetdump	14	Sharing dump devices between different versions of	
Chapter 4. Using DASD devices for		Linux	42
multi-volume dump	17	Sharing dump resources with VMDUMP	43
Installing the multi-volume DASD dump tool	18	Appendix A. Examples for initiating	
Initiating a multi-volume DASD dump	19	dumps	45
Copying a multi-volume dump to a file	20	z/VM	45
Chapter 5. Using a tape dump device	21	Using kdump.	45
Installing the tape dump tool	21	Using DASD	45
Initiating a tape dump.	21	Using tape.	46
Tape display messages.	22	Using SCSI	46
Copying the dump from tape	22	Using VMDUMP	47
Preparing the dump tape.	22	Dumping NSSs	47
Using the zgetdump tool to copy the dump	23	HMC or SE	48
Checking whether a dump is valid, and printing		Testing automatic dump-on-panic	51
the dump header	23	Appendix B. Obtaining a dump with	
Chapter 6. Using a SCSI dump device	25	limited size	53
Installing the SCSI disk dump tool	25	Appendix C. Command summary	55
SCSI dump tool parameters	25	The zgetdump tool	55
Combined dump and target partition.	26	The dumpconf service.	59
Initiating a SCSI dump	26	The crash tool	62
Printing the SCSI dump header.	27	The vmconvert tool.	62
Chapter 7. Creating dumps on z/VM		The vmur tool	63
with VMDUMP.	29	Appendix D. Preparing for analyzing a	
Initiating a dump with VMDUMP.	29	dump	65
Copying the dump to Linux.	29	Appendix E. How to detect guest	
Chapter 8. Creating live-system dumps		relocation	67
with zgetdump	31	Accessibility	69
Creating a kernel dump on a live system	31	Notices	71
Opening a live-system dump with the crash tool	32	Trademarks	72

Index 73

About this book

This book describes tools for obtaining dumps of Linux for IBM® System z® instances running Red Hat Enterprise Linux 7. This book describes how to use DASD, tape, and SCSI dump devices, as well as how to use VMDUMP.

Unless stated otherwise, all z/VM® related information in this document assumes a current z/VM version, see www.ibm.com/vm/techinfo.

In this document, System z is taken to include all IBM mainframe systems supported by Red Hat Enterprise Linux 7 for System z. In particular, this includes IBM zEnterprise® BC12 (zBC12), IBM zEnterprise EC12 (zEC12), IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114) mainframes.

For Red Hat Enterprise Linux product documentation, including what is new, known issues, and frequently asked questions, see the Red Hat Enterprise Linux documentation Web site at

http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/

You can find the latest version of this document on developerWorks® at

www.ibm.com/developerworks/linux/linux390/documentation_red_hat.html

Authority

Most of the tasks described in this document require a user with root authority. In particular, writing to `procfs`, and writing to most of the described `sysfs` attributes requires root authority.

Throughout this document, it is assumed that you have root authority.

Other relevant Linux on IBM System z publications

Another Linux on IBM System z publication for Red Hat Enterprise Linux 7 is available on developerWorks.

You can find the latest versions of this publication at www.ibm.com/developerworks/linux/linux390/documentation_red_hat.html.

- *Device Drivers, Features, and Commands on Red Hat Enterprise Linux 7*, SC34-2710

For each of the following publications, the same web page points to the version that most closely reflects Red Hat Enterprise Linux 7:

- *How to Improve Performance with PAV*, SC33-8414
- *How to use FC-attached SCSI devices with Linux on System z*, SC33-8413
- *How to Set up a Terminal Server Environment on z/VM*, SC34-2596
- *libica Programmer's Reference*, SC34-2602
- *Secure Key Solution with the Common Cryptographic Architecture Application Programmer's Guide*, SC33-8294
- *Linux on System z Troubleshooting*, SC34-2612

Chapter 1. Tools overview

Different tools can be used for obtaining dumps for instances of Red Hat Enterprise Linux 7 running on IBM System z mainframes.

You can use the dump analysis tool **crash** to analyze a dump. Depending on your service contract, you might also want to send a dump to IBM support to be analyzed.

Table 1 summarizes the available dump tools:

Table 1. Dump tools summary

Dump aspect	kdump	DASD	Multi-volume DASD	SCSI	Tape	VMDUMP	Live-system dump with zgetdump
Environment	z/VM and LPAR	z/VM and LPAR	z/VM and LPAR	z/VM and LPAR	z/VM and LPAR	z/VM only	z/VM and LPAR
z/VM NSS	No	Yes (see note 3)	Yes (see note 3)	Yes (see note 3)	Yes (see note 3)	Yes	Yes
System size (See also "Dump size" on page 2)	Large	Small	Large	Large	Large	Small	Large
Speed	Fast	Fast	Fast	Fast	Slow	Slow	Fast
Medium	Any available medium	ECKD™ or FBA (see note 1) DASD	ECKD DASD	Linux file system	Tape cartridges	z/VM reader	Any available medium
Compression possible	While writing	No	No	While writing	Yes (See "Dump size" on page 2)	No	No
Dump filtering possible	While writing	When copying	When copying	When copying	When copying	When copying	No
Disruptive (see note 2)	Yes	Yes	Yes	Yes	Yes	No	No
Stand-alone	No	Yes	Yes	Yes	Yes	No	No

Note:

1. SCSI disks can be emulated as FBA disks. This dump method can, therefore, be used for SCSI-only z/VM installations.
2. In this context, disruptive means that the dump process kills a running operating system.
3. As of z/VM 6.3, you can use the nssdata IPL option to create a dump for a Linux instance that runs from a z/VM NSS.

Dump size

The dump size depends on the size of the system for which the dump is to be created. Except for `kdump`, all dump methods require persistent storage space to hold the kernel and user space of this system.

kdump

Initially uses the memory of the Linux instance for which a dump is to be created, and so supports any size. A persistent copy can be written to any medium of sufficient size. While writing, the dump size can be reduced through page filtering and compression.

DASD

Depends on the disk size. For example, ECKD model 27 provides 27 GB.

Multivolume DASD

Can be up to the combined size of 32 DASD partitions.

SCSI Depends on the capacity of the SCSI disk and which other data it contains.

Tape Depends on the tape drive. For example, IBM TotalStorage Enterprise Tape System 3592 supports large dumps and also offers hardware compression.

VMDUMP

Depends on the available spool space. The slow dump speed can lead to very long dump times for large dumps. Although technically possible, the slow dump speed makes VMDUMP unsuitable for large dumps.

zgetdump live-system dump

The dump can be written to any medium of sufficient size.

See Chapter 9, “Handling large dumps,” on page 35 for information specific to large dumps.

Note on device nodes

In all examples, the traditional device nodes for DASD, tape, and SCSI devices are used. You can also use the device nodes that `udev` creates for you.

The `kdump` tool

The `kdump` tool is made available through a Linux kernel and initial RAM disk that are preloaded in memory, along with a production system.

You do not have to install `kdump` on a dedicated dump device. The `kdump` system can access the memory that contains the dump of the production system through a `procf`s file.

Filtering out extraneous memory pages and compression can take place while the dump is written to persistent storage or transferred over a network. The smaller dump size can significantly reduce the write or transfer time, especially for large production systems.

Because `kdump` can write dumps through a network, existing file system facilities can be used to prevent multiple dumps from being written to the same storage space. Sharing space for dumps across an enterprise is possible without the more complex setups described in Chapter 10, “Sharing dump devices,” on page 39.

Stand-alone tools

Stand-alone tools are installed on a device on which you perform an IPL. Different tools are available depending on the device type.

Four stand-alone dump tools are shipped in the `s390utils` package as part of **zipl**:

- DASD dump tool for dumps on a single DASD device
- Multivolume DASD dump tool for dumps on a set of ECKD DASD devices
- Tape dump tool for dumps on (channel-attached) tape devices
- SCSI disk dump tool for dumps on SCSI disks

You need to install these tools on the *dump device*. A dump device is used to initiate a stand-alone dump by IPL-ing the device. It must have a stand-alone dump tool installed and should provide enough space for the dump.

Typically, the system operator initiates a dump after a system crash, but you can initiate a dump at any time. To initiate a dump, you must IPL the dump device. This is destructive, that is, the running Linux operating system is killed. The IPL process writes the system memory to the IPL device (DASD and tape) or directly to a file on a SCSI disk.

You can configure a dump device that is automatically used when a kernel panic occurs. For more information, see “The `dumpconf` service” on page 59.

For more information on **zipl**, refer to the **zipl** man page and to the **zipl** description in *Device Drivers, Features, and Commands on Red Hat Enterprise Linux 7*, SC34-2710. You can find the latest version of this document on developerWorks at: www.ibm.com/developerworks/linux/linux390/documentation_red_hat.html

Related tasks:

Chapter 3, “Using a DASD dump device,” on page 13

To use a DASD dump device, you need to install the stand-alone DASD dump tool and perform the dump process. Then, copy the dump to a file in a Linux file system.

Chapter 4, “Using DASD devices for multi-volume dump,” on page 17

You can handle large dumps, up to the combined size of 32 DASD partitions, by creating dumps across multiple volumes.

Chapter 5, “Using a tape dump device,” on page 21

You can use a channel-attached tape as a dump device. To use a tape, you need to install the stand-alone tape dump tool and perform the dump process. Then, copy the dump to a file in a Linux file system.

Chapter 6, “Using a SCSI dump device,” on page 25

You can use SCSI disks that are accessed through the `zfcp` device driver as dump devices. SCSI disk dumps are written as files in an existing file system on the dump partition. No copying is necessary.

The VMDUMP tool

The **VMDUMP** tool is a part of z/VM and does not need to be installed separately.

Dumping with **VMDUMP** is not destructive. If you dump an operating Linux instance, the instance continues running after the dump is completed.

VMDUMP can also create dumps for z/VM guests that use z/VM named saved systems (NSS).

Do not use **VMDUMP** to dump large z/VM guests; the dump process is very slow. Dumping 1 GB of storage can take up to 15 minutes depending on the used storage server and z/VM version.

For more information on **VMDUMP** see *z/VM CP Commands and Utilities Reference*, SC24-6175.

Related tasks:

Chapter 7, “Creating dumps on z/VM with VMDUMP,” on page 29

Use **VMDUMP** to create dumps on z/VM systems, using the z/VM reader as the dump medium.

Live-system dump

You can create a kernel dump from a live system without disruption.

Use the **zgetdump** tool that is shipped with the s390utils-base RPM to create a kernel dump while the Linux system continues running. No dump device must be prepared, because the `/dev/crash` device node is used to create the dump.

Dump methods compared

The process for preparing a dump device and obtaining a dump differs for the available dump methods.

Table 2. Comparing the dump methods.

Dump aspect	kdump	Stand-alone tools	VMDUMP	Live-system dump with zgetdump
Preparation	Reserve memory with the <code>crashkernel=</code> kernel parameter Load the <code>kdump</code> kernel and the initial RAM disk into the memory of the production system (<code>kexec</code>)	Write the stand-alone dump tool to the dump device (<code>zipl</code>) Define the panic shutdown action (<code>dumpconf</code>)	Define the panic shutdown action (<code>dumpconf</code>)	None
Dump trigger	Automatic: Kernel panic Initiated by operator: PSW restart	Automatic: Kernel panic Initiated by operator: IPL of the dump device	Automatic: Kernel panic Initiated by operator: z/VM CP VMDUMP command	Initiated by operator: zgetdump invocation
Initial dump space	Memory	Dump device	Spool device	Memory
Accessing the initial dump	Through <code>/proc/vmcore</code> from the <code>kdump</code> instance	Using zgetdump from a new Linux instance	Using vmur -c from a new Linux instance	Through <code>/dev/crash</code>

Table 2. Comparing the dump methods (continued).

Dump aspect	kdump	Stand-alone tools	VMDUMP	Live-system dump with zgetdump
Copying the initial dump to the final dump store (and releasing the initial dump space)	Copied from the kdump instance to any available storage	Copied from the new Linux instance to any available storage	Copied from the new Linux instance to any available storage	Copied from the current Linux instance to any available storage
Optional: Filtering the initial dump	Using <code>/proc/vmcore</code> and <code>makedumpfile</code> on the kdump instance	Using <code>zgetdump</code> and <code>makedumpfile</code> on the new Linux instance	Using <code>zgetdump</code> and <code>makedumpfile</code> on the new Linux instance	Not recommended

Chapter 2. Using kdump

You can use kdump to create system dumps for instances of Red Hat Enterprise Linux.

Advantages of kdump

kdump offers these advantages over other dump methods:

- While writing the dump, you can filter out extraneous pages and compress the dump, and so handle large dumps in a short time.
- When writing dumps over a network, you can use existing file system facilities to share dump space without special preparations.

Shortcomings of kdump

kdump has these drawbacks:

- kdump is not as reliable as the stand-alone dump tools. For critical systems, you can set up stand-alone dump tools as a backup, in addition to the kdump configuration (see “Failure recovery and backup tools” on page 9).
- kdump cannot dump a z/VM named saved system (NSS).
- For production systems that run in LPAR mode, kdump consumes memory (see “Memory consumption” on page 8).

How kdump works on System z

You can set up kdump according to your needs.

With kdump, you do not need to install a dump tool on the storage device that is to hold a future dump. Instead you use a kdump kernel, a Linux instance that controls the dump process.

The kdump kernel occupies a reserved memory area within the memory of the production system for which it is set up. The reserved memory area is defined with the `crashkernel=` kernel parameter. After the production system is started, the kdump kernel and its initial RAM disk (`initrd`) are loaded into the reserved memory area with the **kexec** tool.




Figure 1. Running production system with preloaded kdump kernel and initial RAM disk

At the beginning of the dump process, the reserved memory area is exchanged with the lower memory regions of the crashed production system. The kdump system is then started and runs entirely in the memory that has been exchanged with the reserved area. From the running kdump kernel, the memory of the crashed production system can be accessed as a virtual file, `/proc/vmcore`.




Figure 2. Running kdump kernel

This process is fast, because the kdump kernel is started from memory, and no dump data needs to be copied up to this stage. For Red Hat Enterprise Linux, the **makedumpfile** tool in the kdump `initrd` writes a filtered and compressed version of the dump to a file on persistent storage, locally or over a network. Again, this saves time, because the dump is reduced in size while it is written or transferred.

By default, kdump `initrd` automatically IPLs the production system after the dump is written.

Memory consumption

Although each Linux instance must be defined with additional memory for kdump, the total memory consumption for your z/VM installation does not increase considerably.

On most architectures, the inactive kdump system consumes the entire memory that is reserved with the `crashkernel=` kernel parameter.

For Linux on z/VM, only the kdump image and its initial RAM disk consume actual memory. The remaining reserved memory is withheld by the z/VM hypervisor until it is required in exchange for the lower memory region of the crashed production system.

Because the kdump image and initial RAM disk are not used during regular operations, z/VM swaps them out of memory some time after IPL. Thereafter, no real memory is occupied for kdump until it is booted to handle a dump.

For Linux in LPAR mode, the reserved memory area consumes real memory.

Failure recovery and backup tools

If kdump fails, stand-alone dump tools or VMDUMP can be used as backup tools. Backup tools are, typically, set up only for vital production systems.

Because of being preloaded into memory, there is a small chance that parts of kdump are overwritten by malfunctioning kernel functions. The kdump kernel is, therefore, booted only if a checksum assures the integrity of the kdump kernel and initial RAM disk. This failure can be recovered automatically by setting up a backup dump tool with the **dumpconf** service or through a backup dump that is initiated by a user. See “The dumpconf service” on page 59.

A second possible failure is the kdump system itself crashing during the dump process. This failure occurs, for example, if the reserved memory area is too small for the kdump kernel and user space. For this failure, initiate a backup dump, which captures data for both the crashed production system and the crashed kdump kernel. You can separate this data with the **zgetdump --select** option. See “The zgetdump tool” on page 55.

Related reference:

“The dumpconf service” on page 59

The **dumpconf** service configures the action to be taken if a kernel panic or PSW restart occurs.

“The zgetdump tool” on page 55

The **zgetdump** tool reads or converts a dump.

Setting up kdump

Red Hat Enterprise Linux provides several ways of setting up kdump.

About this task

You can choose between the following methods of setting up kdump:

- The firstboot utility: Basic tool to be used after first boot. For a configuration example, see the chapter on Firstboot in the *Red Hat Enterprise Linux Installation Guide*.
- The Kernel Dump Configuration utility **system-config-kdump**: Graphical tool with more configuration options. For a configuration example, see the chapter on kdump in the *Red Hat Enterprise Linux Deployment Guide*.

- Manually using the configuration file `/etc/kdump.conf`. For a configuration example, see the chapter on `kdump` in the *Red Hat Enterprise Linux Deployment Guide*.

What to do next

As a backup, you can set up a stand-alone dump tool in addition to `kdump`. See “The `dumpconf` service” on page 59 about how to run a backup tool automatically, if `kdump` fails.

Related reference:

“The `dumpconf` service” on page 59

The **dumpconf** service configures the action to be taken if a kernel panic or PSW restart occurs.

How kdump is triggered

A kernel panic automatically triggers the dump process with `kdump`. When your Linux system does not respond and `kdump` is not triggered automatically, depending on your system environment, there are additional methods for triggering the dump process.

About this task

With `kdump` installed, a kernel panic or PSW restart trigger `kdump` rather than the shutdown actions defined in `/sys/firmware`. The definitions in `/sys/firmware` are used only if an integrity check for `kdump` fails (see also “Failure recovery and backup tools” on page 9 and “The `dumpconf` service” on page 59).

Procedure

Use one of the methods according to your environment:

- For Linux in LPAR mode: Run the **PSW restart** task on the HMC. See “HMC or SE” on page 48 for details.
- For Linux on z/VM: Run the z/VM CP **system restart** command. For example, issue this command from a 3270 terminal:

```
#cp system restart
```

- For Linux on z/VM: Configure the z/VM watchdog to trigger `kdump`. Set **system restart** as the z/VM CP command to be issued if the watchdog detects that the Linux instance failed. See *Device Drivers, Features, and Commands on Red Hat Enterprise Linux 7*, SC34-2710 about how to configure the z/VM watchdog.

Results

The dump process loads the `kdump` kernel. Depending on the `kdump` configuration `/proc/vmcore` is copied and filtered by the `kdump` initrd, and then the production system is rebooted.

What to do next

Verify that your production system is up and running again. Send the created dump to a service organization.

Related reference:

“The dumpconf service” on page 59

The **dumpconf** service configures the action to be taken if a kernel panic or PSW restart occurs.

Chapter 3. Using a DASD dump device

To use a DASD dump device, you need to install the stand-alone DASD dump tool and perform the dump process. Then, copy the dump to a file in a Linux file system.

DASD dumps are written directly to a DASD partition that was not formatted with a file system. The following DASD types are supported:

- ECKD DASDs
 - 3380
 - 3390
- FBA DASDs

Installing the DASD dump tool

Install the DASD dump tool on an unused DASD partition. Memory dumps are written to this partition.

Before you begin

You need an unused DASD partition with enough space (memory size + 10 MB) to hold the system memory. If the system memory exceeds the capacity of a single DASD partition, use the multivolume dump tool, see Chapter 4, “Using DASD devices for multi-volume dump,” on page 17.

About this task

The examples assume that `/dev/dasdc` is the dump device and that you want to dump to the first partition `/dev/dasdc1`.

The steps that you need to perform for installing the DASD dump tool depend on your type of DASD, ECKD or FBA:

- If you are using a DASD of type ECKD, perform all three of the following steps.
- If you are using a DASD of type FBA, skip steps 1 and 2 and perform step 3 only.

Procedure

1. (ECKD only) Format your DASD with **dasdfmt**. A block size of 4 KB is recommended. For example:

```
# dasdfmt -f /dev/dasdc -b 4096
```

2. (ECKD only) Create a partition with **fdasd**. The partition must be sufficiently large (the memory size + 10 MB). For example:

```
# fdasd /dev/dasdc
```

3. Install the dump tool with the **zipl** command. Specify the dump device on the command line. For example:

```
# zipl -d /dev/dasdc1
```

Note: When you use a DASD of type ECKD that is formatted with the traditional Linux disk layout `ldl`, the dump tool must be reinstalled with `zip1` after each dump.

Initiating a DASD dump

You can initiate a dump from a DASD device.

Procedure

To obtain a dump with the DASD dump tool, perform the following main steps:

1. Stop all CPUs.
2. Store status on the IPL CPU.
3. IPL the dump tool on the IPL CPU.

Note: Do not clear storage!

The dump process can take several minutes depending on the device type you are using and the amount of system memory. After the dump completes, the IPL CPU should go into disabled wait.

The following PSW indicates that the dump process completed successfully:
(64-bit) PSW: 00020000 80000000 00000000 00000000

Any other disabled wait PSW indicates an error.

After the dump tool is IPLed, messages that indicate the progress of the dump are written to the console:

```
Dumping 64 bit OS
00000032 / 00000256 MB
00000064 / 00000256 MB
00000096 / 00000256 MB
00000128 / 00000256 MB
00000160 / 00000256 MB
00000192 / 00000256 MB
00000224 / 00000256 MB
00000256 / 00000256 MB
Dump successful
```

Results

You can IPL Linux again.

See Appendix A, “Examples for initiating dumps,” on page 45 for more details.

Copying the dump from DASD with `zgetdump`

You can copy a DASD dump to a file system by using the `zgetdump` tool.

About this task

By default, the `zgetdump` tool takes the dump device as input and writes its contents to standard output. To write the dump to a file system, you must redirect the output to a file.

Procedure

Assuming that the dump is on DASD device `/dev/dasdc1` and you want to copy it to a file named `dump_file`:

```
# zgetdump /dev/dasdc1 > dump_file
```

What to do next

You can use **zgetdump** to display information about the dump. See “Checking whether a DASD dump is valid and printing the dump header” on page 58 for an example.

For general information about **zgetdump**, see “The zgetdump tool” on page 55 or the man page.

Chapter 4. Using DASD devices for multi-volume dump

You can handle large dumps, up to the combined size of 32 DASD partitions, by creating dumps across multiple volumes.

Before you begin

Prepare a set of ECKD DASD devices for a multivolume dump, install the stand-alone dump tool on each DASD device that is involved. Then, perform the dump process, and copy the dump to a file in a Linux file system.

About this task

You can specify up to 32 partitions on ECKD DASD volumes for a multivolume dump. The dump tool is installed on each volume involved. The volumes must be:

- In subchannel set 0.
- Formatted with the compatible disk layout (cdl, the default option when you use the **dasdfmt** command.)

You can use any block size, even mixed block sizes. However, to speed up the dump process and to reduce wasted disk space, use block size 4096.

For example, Figure 3 shows three DASD volumes, *dasdb*, *dasdc*, and *dasdd*, with four partitions that are selected to contain the dump. To earmark the partition for dump, a dump signature is written to each partition.

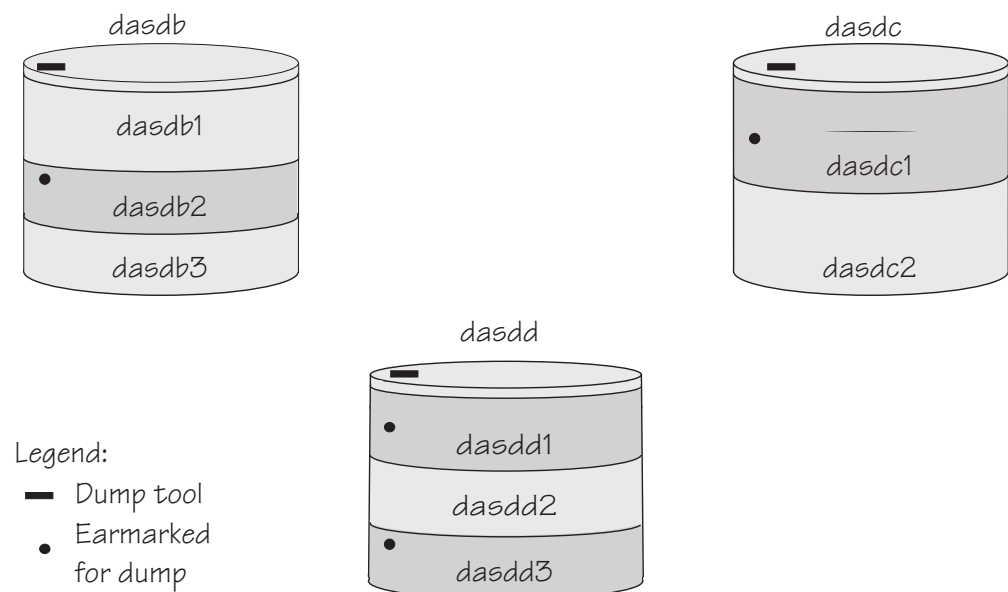


Figure 3. Three DASD volumes with four partitions for a multivolume dump

The partitions need to be listed in a configuration file, for example:

```
/dev/dasdb2
/dev/dasdc1
/dev/dasdd1
/dev/dasdd3
```

You can define a maximum of three partitions on one DASD. All three volumes are prepared for IPL; regardless of which you use the result is the same.

The following sections take you through the entire process of creating a multivolume dump.

Installing the multi-volume DASD dump tool

This example shows how to perform the dump process on two partitions, /dev/dasdc1 and /dev/dasdd1, which reside on ECKD volumes /dev/dasdc and /dev/dasdd.

About this task

Assume that the corresponding bus IDs (as displayed by **lsdasd**) are 0.0.4711 and 0.0.4712, so the respective device numbers are 4711 and 4712.

Procedure

1. Format both dump volumes with **dasdfmt**. Specify **cdl** (compatible disk layout), which is the default. Preferably, use a block size of 4 KB:

```
# dasdfmt -f /dev/dasdc -b 4096
# dasdfmt -f /dev/dasdd -b 4096
```

2. Create the partitions with **fdasd**. The sum of the partition sizes must be sufficiently large (the memory size + 10 MB):

```
# fdasd /dev/dasdc
# fdasd /dev/dasdd
```

3. Create a file named `sample_dump_conf` containing the device nodes of the two partitions, separated by one or more line feed characters (0x0a). The file's contents are as follows:

```
/dev/dasdc1
/dev/dasdd1
```

4. Prepare the volumes using the **zipl** command. Specify the dump list on the command line:

```
# zipl -M sample_dump_conf
Dump target: 2 partitions with a total size of 1234 MB.
Warning: All information on the following partitions will be lost!
/dev/dasdc1
/dev/dasdd1
Do you want to continue creating multi-volume dump partitions (y/n)?
```

Results

Now the two volumes /dev/dasdc and /dev/dasdd with device numbers 4711 and 4712 are prepared for a multivolume dump. Use the **-device** option of **zgetdump** to display information about these volumes:

```
# zgetdump -d /dev/dasdc
Dump device info:
Dump tool.....: Multi-volume DASD dump tool
Version.....: 2
Architecture.....: s390x (64 bit)
Dump size limit...: none
Force specified...: no

Volume 0: 0.0.4711 (online/valid)
Volume 1: 0.0.4712 (online/valid)
```

During **zipl** processing both partitions were earmarked for dump with a valid dump signature. The dump signature ceases to be valid when data other than dump data is written to the partition. For example, writing a file system to the partition overwrites the dump signature. Before writing memory to a partition, the dump tool checks the partition's signature and exits if the signature is invalid. Thus any data inadvertently written to the partition is protected.

You can circumvent this protection, for example, if you want to use a swap space partition for dumping, by using the **zipl** command with the **--force** option. This option inhibits the dump signature check, and any data on the device is overwritten. Exercise great caution when using the force option.

The **zipl** command also takes a size specification, see Appendix B, “Obtaining a dump with limited size,” on page 53. For more details on the **zipl** command, see *Device Drivers, Features, and Commands on Red Hat Enterprise Linux 7*, SC34-2710.

Initiating a multi-volume DASD dump

After preparing the DASD volumes, you can initiate a multi-volume dump by performing an IPL from one of the prepared volumes.

Procedure

To obtain a dump with the multivolume DASD dump tool, perform the following steps:

1. Stop all CPUs.
2. Store status on the IPL CPU.
3. IPL the dump tool using one of the prepared volumes, either 4711 or 4712.

Note: Do not clear storage!

The dump process can take several minutes depending on each volume's block size and the amount of system memory. After the dump has completed, the IPL CPU should go into disabled wait.

The following PSW indicates that the dump process has completed successfully:
(64-bit) PSW: 00020000 80000000 00000000 00000000

Any other disabled wait PSW indicates an error.

After the dump tool is IPLed, messages that indicate the progress of the dump are written to the console:

```
Dumping 64 bit OS
Dumping to: 4711
00000128 / 00001024 MB
00000256 / 00001024 MB
00000384 / 00001024 MB
```

```
00000512 / 00001024 MB
Dumping to: 4712
00000640 / 00001024 MB
00000768 / 00001024 MB
00000896 / 00001024 MB
00001024 / 00001024 MB
Dump successful
```

Results

You can IPL Linux again.

Copying a multi-volume dump to a file

Use the **zgetdump** command to copy the multi-volume dump.

About this task

This example assumes that the two volumes `/dev/dasdc` and `/dev/dasdd` (with device numbers 4711 and 4712) contain the dump. Dump data is spread along partitions `/dev/dasdc1` and `/dev/dasdd1`.

Procedure

Use **zgetdump** without any options to copy the dump parts to a file:

```
# zgetdump /dev/dasdc > multi_volume_dump_file
Format Info:
Source: s390mv
Target: s390
Copying dump:
00000000 / 00001024 MB
00000171 / 00001024 MB
00000341 / 00001024 MB
00000512 / 00001024 MB
00000683 / 00001024 MB
00000853 / 00001024 MB
00001024 / 00001024 MB
Success: Dump has been copied
```

If you want to check only the validity of the multivolume dump rather than copying it to a file, use the `-info` option with **zgetdump**. See “Checking whether a DASD dump is valid and printing the dump header” on page 58 for an example.

Chapter 5. Using a tape dump device

You can use a channel-attached tape as a dump device. To use a tape, you need to install the stand-alone tape dump tool and perform the dump process. Then, copy the dump to a file in a Linux file system.

The following tape devices are supported:

- 3480
- 3490
- 3590
- 3592

The following sections take you through the entire process of creating a dump on a tape device.

Installing the tape dump tool

Install the tape dump tool on the tape that is to hold the dump.

Before you begin

Have enough empty tapes ready to hold the system memory (memory size + 10 MB).

About this task

The examples assume that `/dev/ntibm0` is the tape device that you want to dump to.

Procedure

Perform these steps to install the tape dump tool:

1. Insert an empty dump cartridge into your tape device.
2. Ensure that the tape is rewound.
3. Install the dump tool with the `zip1` command. Specify the dump device on the command line. For example:

```
# zip1 -d /dev/ntibm0
```

Initiating a tape dump

Initiate a tape dump by performing an IPL on the IPL CPU.

Procedure

To obtain a dump with the tape dump tool, perform the following main steps:

1. Ensure that the tape is rewound.
2. Stop all CPUs.
3. Store status on the IPL CPU.
4. IPL the dump tool on the IPL CPU.

Note: Do not clear storage!

The dump tool writes the number of dumped MB to the tape drive message display.

The dump process can take several minutes, depending on the device type you are using and the amount of system memory available. When the dump is complete, the message `dump*end` is displayed and the IPL CPU goes into disabled wait.

The following PSW indicates that the dump was taken successfully:

(64-bit) PSW: 00020000 80000000 00000000 00000000

Any other disabled wait PSW indicates an error.

After the dump tool is IPLed, messages that indicate the progress of the dump are written to the console:

```
Dumping 64 bit OS
00000032 / 00000256 MB
00000064 / 00000256 MB
00000096 / 00000256 MB
00000128 / 00000256 MB
00000160 / 00000256 MB
00000192 / 00000256 MB
00000224 / 00000256 MB
00000256 / 00000256 MB
Dump successful
```

5. You can IPL Linux again.

What to do next

See Appendix A, “Examples for initiating dumps,” on page 45 for more details.

Tape display messages

Messages might be shown on the tape display.

Messages

number

The number of MB dumped.

`dump*end`

The dump process ended successfully.

Copying the dump from tape

You can copy a tape dump to a file system with the `zgetdump` tool.

Before you begin

The `mt` utility must be installed.

Preparing the dump tape

You need to rewind the tape, and find the correct position on the tape to start copying from.

About this task

Use the `mt` tool to manipulate the tape.

Procedure

1. Rewind the tape.

For example:

```
# mt -f /dev/ntibm0 rewind
```

2. Skip the first file on the tape (this file is the dump tool itself).

For example:

```
# mt -f /dev/ntibm0 fsf
```

Using the `zgetdump` tool to copy the dump

Use the `zgetdump` tool to copy the dump file from the tape to a file system.

Before you begin

The tape must be in the correct position (see “Preparing the dump tape” on page 22).

About this task

By default, the `zgetdump` tool takes the dump device as input and writes its contents to standard output. To write the memory dump to a file system, you must redirect the output to a file.

The example assumes that the dump is on tape device `/dev/ntibm0`.

Procedure

Copy the dump from tape to a file named `dump_file` in the file system:

```
# zgetdump /dev/ntibm0 > dump_file
```

For general information on `zgetdump`, see “The `zgetdump` tool” on page 55 or the man page.

Checking whether a dump is valid, and printing the dump header

To check whether a dump is valid, use the `zgetdump` command with the `-i` option.

Procedure

1. Ensure that the volume is loaded.
2. Skip the first file on the tape (this file is the dump tool itself):

```
# mt -f /dev/ntibm0 fsf
```

3. Issue the `zgetdump` command with the `-i` option:

```
# zgetdump -i /dev/ntibm0
```

The `zgetdump` command goes through the dump until it reaches the end. See also “Using `zgetdump` to copy a tape dump” on page 57

Chapter 6. Using a SCSI dump device

You can use SCSI disks that are accessed through the `zfcpx` device driver as dump devices. SCSI disk dumps are written as files in an existing file system on the dump partition. No copying is necessary.

Installing the SCSI disk dump tool

You install the SCSI dump tool with the `zipl` command.

Before you begin

- The kernel-kdump RPM (named `kernel-kdump-3.10-xx.e17.s390x.rpm`) must be installed on your system.
- The dump directory needs enough free space (memory size + 10 MB) to hold the system memory.

About this task

The SCSI dump tool (also referred to as the SCSI Linux System Dumper, or SD) is written to one partition, referred to here as the *target partition*. The memory dump can be written to a second partition, the *dump partition*, provided it is on the same physical disk. Only the target partition needs to be mounted when `zipl` is run. In a single-partition configuration, the target partition is also the dump partition.

SCSI dump tool parameters

When you install the SCSI disk dump tool, the following parameters can be specified in a 'parameters' line in the `zipl` configuration file or specified in the `zipl` command line with the `-P` option

Parameters

`dump_dir=<directory>`

Path to the directory (relative to the root of the dump partition) to which the dump file is to be written. This directory is specified with a leading slash. The directory must exist when the dump is initiated.

For example, if the dump partition is mounted as `/dumps`, and the parameter `dump_dir=/mydumps` is defined, the dump directory would be accessed as `/dumps/mydumps`.

The default is `/` (the root directory of the partition).

`dump_mode=interactive | auto`

Action that is taken if there is no room on the file system for the new dump file. **interactive** prompts the user to confirm that the memory dump with the lowest number is to be deleted. **auto** automatically deletes this file.

The default is **interactive**.

In rare cases, you might want to complement or overwrite the SCSI dump tool parameters that were configured with `zipl`. For example, you might want to change the dump mode setting when you initiate the dump. How you specify such parameters depends on whether your Linux instance runs in LPAR mode or as a z/VM guest. For more information, see the SCSI examples in Appendix A, "Examples for initiating dumps," on page 45.

Combined dump and target partition

A single partition on a SCSI device can be used as both the dump partition and target partition.

About this task

This example assumes that `/dev/sda` is a SCSI device that contains no data and is to be used exclusively as a dump device. Because no other data is to be stored on the device, a single partition is created that serves as both dump and target partition.

Procedure

1. Create a single partition with **fdisk**, using the PC-BIOS layout:

For example:

```
# fdisk /dev/sda
```

The created partition is `/dev/sda1`.

2. Format this partition with the xfs or ext4 file system.

For example:

```
# mkfs.ext4 /dev/sda1
```

3. Mount the partition at a mount point of your choice and create a subdirectory to hold the dump files.

For example:

```
# mount /dev/sda1 /dumps  
# mkdir /dumps/mydumps
```

4. Install the dump tool with the **zipl** command. Specify the dump device on the command line.

For example:

```
# zipl -D /dev/sda1 -t /dumps -P "dump_dir=/mydumps"
```

5. Unmount the file system:

```
# umount /dumps
```

Results

When you IPL `/dev/sda1` using boot program selector 1 or 0 (default), the dump is written to directory `mydumps` on partition 1 of `/dev/sda`. The boot program selector is located on the load panel, see Figure 6 on page 51 for an example.

Initiating a SCSI dump

To initiate the dump, IPL the SCSI dump tool with the SCSI dump load type.

About this task

The dump process can take several minutes depending on the device type you are using and the amount of system memory. The dump progress and any error messages are reported on the operating system messages console.

Procedure

IPL the SCSI dump tool.

See Appendix A, “Examples for initiating dumps,” on page 45 for more details.

Results

The dump process creates a new dump file in the dump directory. All dumps are named `dump.<n>`, where `<n>` is the dump number. A new dump receives the next highest dump number out of all dumps in the dump directory (see the `dump_dir` parameter under “SCSI dump tool parameters” on page 25).

For example, if there are already two dump files named `dump.0` and `dump.1` in the dump directory, the new dump is named `dump.2`.

When the dump completes successfully, you can IPL Linux again.

You do not need to convert the dump or copy it to a different medium. To access the dumps, mount the dump partition.

Printing the SCSI dump header

To print the dump file header, use `zgetdump` with the `-i` option.

Procedure

Specify the `zgetdump` command with the `-i` option:

```
# zgetdump -i dump.0
General dump info:
Dump format.....: lkcd
Version.....: 8
System arch.....: s390x (64 bit)
CPU count (online): 2
CPU count (real)...: 2
Dump memory range..: 1024 MB

Memory map:
0000000000000000 - 000000003fffffff (1024 MB)
```

Chapter 7. Creating dumps on z/VM with VMDUMP

Use VMDUMP to create dumps on z/VM systems, using the z/VM reader as the dump medium.

Before you begin

Do not use VMDUMP to dump large z/VM guests; the dump process is very slow. Dumping 1 GB of storage can take up to 15 minutes depending on the used storage server and z/VM version.

About this task

This section describes how to create a dump with VMDUMP, how to transfer the dump to Linux, and how to convert the z/VM dump to a convenient format. VMDUMP does not need to be installed separately.

The following sections take you through the process of creating a dump with VMDUMP.

Initiating a dump with VMDUMP

Start the dump VMDUMP process with the CP VMDUMP command.

Procedure

Issue the following command from the 3270 console of the z/VM guest virtual machine:

```
#CP VMDUMP
```

Results

z/VM CP temporarily stops the z/VM guest virtual machine and creates a dump file. The dump file is stored in the reader of the z/VM guest virtual machine. After the dump is complete, the Linux on z/VM instance continues operating.

You can use the **TO** option of the VMDUMP command to direct the dump to the reader of another guest virtual machine of the same z/VM system.

Example

To write the dump to the reader of z/VM guest virtual machine linux02 issue:

```
#CP VMDUMP TO LINUX02
```

For more information about VMDUMP refer to *z/VM CP Commands and Utilities Reference*, SC24-6175.

Copying the dump to Linux

Copy the dump from the z/VM reader using the **vmur** command.

Procedure

1. Find the spool ID of the **VMDUMP** spool file in the output of the **vmur li** command:

```
# vmur li
ORIGINID FILE CLASS RECORDS CPY HOLD DATE TIME NAME TYPE DIST
T6360025 0463 V DMP 00020222 001 NONE 06/11 15:07:42 VMDUMP FILE T6360025
```

In the example the required **VMDUMP** file spool ID is 463.

2. Copy the dump into your Linux file system using the **vmur receive** command. To convert the dump into a format that can be processed with the Linux dump analysis tool **crash**, convert the dump using the **--convert** option:

```
# vmur rec 463 -c myvmdump
vmdump information:
architecture: 64 bit (big)
storage.....: 256 MB
date.....: Thu Feb  5 08:39:48 2009
cpus.....: 1
256 of 256 |#####| 100%
```

Results

The created file, named myvmdump, can now be used as input to **crash**.

Chapter 8. Creating live-system dumps with zgetdump

If you require a kernel dump of a Linux instance, but no downtime is acceptable, you can create a kernel dump from a live system without disruption.

Because the Linux system continues running while the dump is written, and kernel data structures are changing during the dump process, the resulting dump contains inconsistencies. The faster the dump process completes, the fewer inconsistencies the resulting live-system dump will contain. Therefore, run the dump process with the highest acceptable priority.

You can change the scheduling priority with the `nice` command. For example, use `nice -n -20` to set the highest possible priority.

Creating a kernel dump on a live system

You can create non-disruptive kernel dumps on a running Linux system with the `zgetdump` tool.

Before you begin

- The dump directory needs enough free space (memory size + 10 MB) to hold the system memory.
- Ensure that during the dump process, no memory hotplug or CPU hotplug is performed.
- If applicable, stop the `cpuplugd` service by issuing the command:

```
# service cpuplugd stop
```

- Load the crash kernel module by issuing the command:

```
# modprobe crash
```

Procedure

1. Optional: Use the `-i` option to print information for the currently running Linux image:

```
# zgetdump -i /dev/crash
General dump info:
  Dump format.....: devmem
  Dump method.....: live
  UTS node name.....: mylnxsys
  UTS kernel release.: 3.3.0
  UTS kernel version.: #68 SMP PREEMPT Thu Mar 22 10:21:30 CET 2012
  System arch.....: s390x (64 bit)
  Dump memory range..: 896 MB

Memory map:
0000000000000000 - 0000000037ffffff (896 MB)
```

2. Create a dump from a live system by specifying `/dev/crash` as input dump and redirecting the output to a dump file. Run the dump process with a high priority.

```
# nice -n -20 zgetdump /dev/crash > dump.s390
```

Optionally, you can also specify a target dump format with the **-f** option:

```
# zgetdump /dev/crash -f elf > dump.elf
```

3. Optional: Print information for the live-system dump. Use the **-i** option to print information for live-system dumps that are generated by **zgetdump**:

```
# zgetdump -i dump.elf
General dump info:
Dump format.....: elf
Version.....: 1
Dump method.....: live
UTS node name.....: mylnxsys
UTS kernel release.: 3.3.0
UTS kernel version.: #68 SMP PREEMPT Thu Mar 22 10:21:30 CET 2012
System arch.....: s390x (64 bit)
Dump memory range..: 896 MB

Memory map:
0000000000000000 - 0000000037ffffff (896 MB)
```

The value "live" in the **Dump method** field indicates that this is a dump from a live system.

Example

```
# nice -n -20 zgetdump /dev/crash -f elf > dump.elf
Format Info:
Source: devmem
Target: elf
Copying dump:
00000000 / 00000896 MB
00000149 / 00000896 MB
...
00000747 / 00000896 MB
00000896 / 00000896 MB
Success: Dump has been copied
```

What to do next

After you create a dump from a live system, you can work with crash, see “Opening a live-system dump with the crash tool.”

After the live dump has been copied to a file system, you can compress it with **makedumpfile**. Note that the dump level must not be greater than 1 because of the dump inconsistencies.

For example:

```
# makedumpfile dump.elf -c -d 1 dump.kdump
```

Opening a live-system dump with the crash tool

Inconsistencies in a kernel dump from a live system can cause some crash commands to fail.

Procedure

- Use the **crash** command to find information about whether a dump is from a live system. This information is displayed in the startup messages, or when you use the **sys** command:


```
# crash dump.elf /usr/lib/debug/lib/modules/3.10.<xx>.e17.s390x/vmlinux
...
  KERNEL: /usr/lib/debug/lib/modules/3.10.<xx>.e17.s390x/vmlinux
  DUMPFILE: dump.elf [LIVE DUMP]
  CPUS: 6
...
crash> sys | grep DUMPFILE
...
  DUMPFILE: dump.elf [LIVE DUMP]
...
```

The tag [LIVE DUMP] informs you that the dump contains inconsistencies.

- Detect whether a dump is from a live system by using the **help -p** command:

```
# crash> help -p | grep flags2
flags2: 40 (LIVE_DUMP)
```

- Use the `--minimal` option if the crash tool fails to start because of inconsistent data structures in the kernel dump. With this option, crash tolerates a degree of inconsistency. However, only a subset of crash commands is then available:

```
# crash --minimal dump.elf /usr/lib/debug/lib/modules/3.10.<xx>.e17.s390x/vmlinux
...
NOTE: minimal mode commands: log, dis, rd, sym, eval, set and exit
```

Chapter 9. Handling large dumps

Methods exist for handling memory dumps that are especially large (greater than 10 GB in size).

Before you begin

The preferred method for handling memory dumps of large production systems is using `kdump`. With `kdump`, you do not need to set up a dedicated dump device with a dump tool for each individual system. Instead, you set aside storage space to receive any dumps from across your installation. If you want to set up a backup dump method for a critical system with a large memory, the information here applies when you use `kdump`.

About this task

Large memory dumps present a challenge as they:

- Take up a large amount of disk space
- Take a long time when dumping
- Use considerable network bandwidth when they are sent to the service organization.

Note: Sometimes you can re-create the problem on a test system with less memory, which makes the handling much easier. Take this option into account before you create a large memory dump.

Procedure

Complete these steps to prepare and process a large memory dump.

1. Choose a dump device. If you want to dump a system with a large memory footprint, you must prepare a dump device that is large enough. You can use the following dump devices for large dumps:

Single-volume DASD

- 3390 model 9 (up to 45 GB)
- 3390 model A (up to 180 GB)

Multivolume DASD

Up to 32 DASDs are possible.

- 32 x 3390 model 9 (up to 1.4 TB)
- 32 x 3390 model A (up to 5.7 TB)

z/VM FBA emulated SCSI dump disk

FBA disks can be defined with the CP command `SET EDEVICE`. These disks can be used as single-volume DASD dump disks. The SCSI disk size depends on your storage server setup.

SCSI dump

The SCSI disk size depends on your storage server setup. The ext4 file system dump size limit is 16 TB and for xfs it is 8 EB.

Dump on 3592 channel-attached tape drive

Cartridges with up to 4 TB capacity.

- Do not use **VMDUMP** for large systems because this dump method is very slow.
2. Estimate the dump time. The dump speed depends on your environment, for example your SAN setup and your storage server. Assume about 100 MB per second dump speed on DASDs or SCSI disks and a system with 50 GB memory. The dump will take about eight minutes. Do a test dump on your system to determine the dump speed for it. Then, you have an indication of how long a dump takes in case of emergency.
 3. Reduce the dump size. For transferring dumps in a short amount of time to a service organization, it is often useful to reduce the dump size or split the dump into several parts for easier and faster transmission. To reduce the dump, choose one of these methods:
 - “Compressing a dump using makedumpfile”
 - “Compressing a dump using gzip and split” on page 37
 4. Send the dump.

Compressing a dump using makedumpfile

Use the **makedumpfile** tool to compress s390 dumps and exclude memory pages that are not needed for analysis. Alternatively, you can use the **gzip** and **split** commands.

About this task

Compressing the dump substantially reduces the size of dump files and the amount of time needed to transmit them from one location to another. For Red Hat Enterprise Linux, the **makedumpfile** tool is included in the **kexec-tools** RPM that you can install, for example, with **yum install kexec-tools**. Because **makedumpfile** expects as input dump files in ELF format, you first have to transform your s390 format dump to ELF format. This is best done by mounting the dump using the **zgetdump** command.

Procedure

1. Mount the dump in ELF format by performing one of these steps:
 - To mount a DASD dump from the partition `/dev/dasdb1` to `/mnt`, issue:

```
# zgetdump -m -f elf /dev/dasdb1 /mnt
```
 - To mount a SCSI dump from file `dump.0` to `/mnt`, issue:

```
# zgetdump -m -f elf dump.0 /mnt
```
2. Locate the `vmlinux` file in the `debuginfo` RPM. After mounting the dump in ELF format with **zgetdump**, the dump is available in the file named `/mnt/dump.elf`. In order to use **makedumpfile** with dump level greater than one, you also need the `vmlinux` file that contains necessary debug information. You find this file in the kernel `debuginfo` RPM. Issue the following commands (the `xx` in the example must be replaced by the appropriate kernel version that caused the dump):

```
# rpm -qlp kernel-debuginfo-3.10-xx.e17.s390x.rpm | grep vmlinux
```
3. Extract the `vmlinux` file to `./usr/lib/debug/lib/modules/3.10-xx.e17.s390x/`. Issue the following command:

```
# rpm2cpio kernel-debuginfo-3.10-xx.e17.s390x.rpm | cpio -idv *vmlinux*
./usr/lib/debug/lib/modules/3.10-xx.e17.s390x/vmlinux
1079519 blocks
```

4. Use the **-d** (dump level) option of **makedumpfile** to specify which pages to exclude from the dump. See the man page for **makedumpfile** for a description of the dump level and other options of **makedumpfile**.

This example compresses the dump file named `/mnt/dump.elf` (**-c** option) and excludes pages that are typically not needed to analyze a kernel problem. Excluded pages are: pages containing only zeroes, pages used to cache file contents (cache, cache private), pages belonging to user spaces processes, and free pages (maximum dump level 31):

```
# makedumpfile -c -d 31 -x vmlinux /mnt/dump.elf dump.kdump
```

The newly created file, named `dump.kdump` should be much smaller than the original file, named `dump.elf`. Until your kernel problem is resolved, it is recommended to keep the original dump file. This will enable you to reduce the dump level, if it turns out that the pages that had been excluded are still needed for problem determination.

5. For initial problem analysis, you can also extract the kernel log with **makedumpfile**, and send it to your service organization:

```
# makedumpfile --dump-dmesg -x vmlinux /mnt/dump.elf kernel.log
```

What to do next

After you have used **makedumpfile**, you can unmount the dump:

```
# zgetdump -u /mnt
```

Compressing a dump using gzip and split

Use the **gzip** and **split** commands to compress the dump and split it into parts. Alternatively, you can use the **makedumpfile** command.

Procedure

1. Compress the dump and split it into parts of one GB by using the **gzip** and **split** commands.

- For a DASD dump:

```
# zgetdump /dev/dasdd1 | gzip | split -b 1G
```

- For a tape dump:

```
# mt -f /dev/ntibm0 rewind
# mt -f /dev/ntibm0 fsf
# zgetdump /dev/ntibm0 | gzip | split -b 1G
```

- For a SCSI dump:

```
# cat /mnt/dump.0 | gzip | split -b 1G
```

This creates several compressed files in your current directory:

```
# ls  
# xaa xab xac xad xae
```

2. Create md5 sums of parts:

```
# md5sum * > dump.md5
```

3. Upload the parts together with the MD5 information to the service organization.

4. The receiver (the service organization) must do the following:

a. Verify md5 sums:

```
# cd dumpdir  
# md5sum -c dump.md5  
xaa: OK  
xab: OK  
...
```

b. Merge parts and uncompress the dump:

```
# cat x* | gunzip -c > dump
```

Chapter 10. Sharing dump devices

For reasons of economy, you might want to share dump devices rather than setting up a dedicated dump device for each Linux instance.

This section applies to sharing dump devices that have been set up with stand-alone dump tools.

With `kdump`, you can transmit the dump through a network and use existing mechanisms to prevent conflicts when concurrently writing multiple dumps to a shared persistent storage space. `VMDUMP` uses `z/VM` resources to hold the initial dump and the integrity of each dump is handled by the `z/VM` system.

Serialization and device locking

To share devices, some kind of serialization is needed to prevent two systems from dumping at the same time and thus corrupting the dumps.

Either the involved operators must prevent concurrent dumps manually, or, in some cases, available system mechanisms can be used to prevent this. It is possible in many cases to use a pool of devices for sharing. For the sake of simplicity, most of the following examples use only one dump device.

Possible serialization mechanisms:

External

Operators must find an external way to ensure serialization manually.

Link Exclusive write for minidisk is used as a locking mechanism (see “Sharing DASD devices under `z/VM`” on page 40).

Attach Attach and detach is used as locking mechanism (see “Using attach and detach as locking mechanism under `z/VM`” on page 41).

`vmcmd`

Use the `vmcmd` panic action (see “DASD (`vmcmd` panic action)” on page 41).

Alternatively, use no serialization and take the risk that dumps are overwritten, see “DASD (`dump` or `dump_reipl` panic action)” on page 41).

Table 3 shows the serialization methods available for different system configurations.

Table 3. *Serialization of dump devices overview*

	DASD		SCSI	
	<code>z/VM</code>	LPAR	<code>z/VM</code>	LPAR
Manual dump	link, attach, external	external	attach, external	external
Automatic dump	overwrite, <code>vmcmd</code>	overwrite	N/A	N/A

Sharing devices when dumping manually

In the following sections, it is assumed that you start the dump process manually, without using automatic dump on panic.

Sharing DASD devices on LPARs

Configure your IOCDs so that all LPARs that want to share the dump device can access the DASD device. There is no system mechanism available for serialization. Exclusive access must be ensured manually by the involved system operators.

Sharing DASD devices under z/VM

Under z/VM, DASD devices can be shared if they are defined as sharable minidisks for a **NOLOG** user.

About this task

Exclusive access can be guaranteed by the **link** CP command using the exclusive write mode. Because with this mode only one DASD can be linked to one z/VM guest virtual machine at the same time, the dump device will be locked for other systems until it is detached.

Procedure

To create a dump after a system crash, perform these steps:

1. To link the dump device, issue a command of the form:

```
#cp link <disk owner> <vdev1> <vdev2> EW
```

where

- *<disk owner>* is the user ID in the system directory whose entry is to be searched for device *<vdev1>*.
 - *<vdev1>* is the specified user's virtual device number.
 - *<vdev2>* is the virtual device number that is to be assigned to the device for your virtual machine configuration.
2. Create the dump using device *<vdev2>*
 3. Reboot your Linux system.
 4. On your Linux system, set dump device *<vdev2>* online.
 5. On your Linux system, copy the dump using **zgetdump**.
 6. On your Linux system, set dump device *<vdev2>* offline.
 7. Detach the dump device:

```
#cp detach <vdev2>
```

Results

The dump DASD is free again and can be used by other systems.

Sharing SCSI devices

You can share SCSI devices for dumping from multiple Linux systems.

If you want to share FCP attached SCSI disks for dump, they must be accessible through your SAN on all Linux systems that want to use the dump device. The involved operators must ensure manually that two dumps are not taking place at the same time. Otherwise, if multiple Linux systems write to the shared dump device at the same time, both the dump file and the file system on the dump device might be damaged.

Using attach and detach as locking mechanism under z/VM

For your shared dump devices, you can use attach and detach as a locking mechanism.

At any time, only one z/VM guest virtual machine can attach a device. A Linux instance that runs as a guest of a class B z/VM guest virtual machine can lock a shared dump device by attaching it. If you use one single FCP adapter for memory dump on all systems, attach and detach can also be used as locking mechanism for SCSI dump.

Sharing devices when dumping automatically

You can configure a dump to be created automatically should a kernel panic occur.

About this task

The automatic dump on panic can be configured in `/etc/sysconfig/dumpconf` (see “The dumpconf service” on page 59).

DASD (dump or dump_reipl panic action)

It is possible to share DASD devices for automatic dump on panic, but there is no serialization mechanism available.

About this task

As there is no serialization mechanism available, two systems dumping at the same time might corrupt the dumps. Normally, system crashes are quite rare and therefore the chance of corrupt dumps is low, but you have to consider carefully if this is an acceptable risk. Such a dump setup is a trade-off between reliability and resource expenses. You have to consider the likelihood of two concurrent system crashes and the business impact of losing a dump.

Procedure

To share DASDs under z/VM, you must use minidisks that are linked in access mode multiple-write (MW) to all systems where you want to configure dump on panic.

DASD (vmcmd panic action)

You can specify up to five CP commands in a configuration file. These commands run if a kernel panic occurs.

Before you begin

Define minidisks 4e1 and 4e2 with disk owner user **SHARDISK** and prepare them as dump DASDs.

About this task

With z/VM, you can use the panic action **vmcmd** in `/etc/sysconfig/dumpconf` to specify up to five commands that are run in case of a kernel panic. You can use this mechanism to implement locking through the exclusive link or attach method.

In this example, assume that we want to link either 4e1 or 4e2 as device number 5000 and then create the dump using device 5000. The first free DASD will be linked. If both devices are already linked to other z/VM guest virtual machines, the system will stop without creating a dump.

Procedure

The corresponding configuration for `/etc/sysconfig/dumpconf` looks like this:

```
ON_PANIC=vmcmd
VMCMD_1="LINK SHARDISK 4E1 5000 EW"
VMCMD_2="LINK SHARDISK 4E2 5000 EW"
VMCMD_3="STORE STATUS"
VMCMD_4="IPL 5000"
```

Results

After the dump process has finished, you must perform an IPL on the Linux system manually, copy the dump, and detach the disk 5000.

Compared to “DASD (dump or dump_reipl panic action)” on page 41, this option has the advantage that you cannot get corrupted dumps and you can use more than one dump device. It has the disadvantage that automatic re-IPL is not possible.

FCP-attached SCSI devices

Device sharing for automatic dumps is risky when using FCP-attached devices.

For automatic dump on a FCP-attached SCSI device, device sharing should not be used. Otherwise, if multiple Linux systems write to the shared dump device at the same time, you may not only corrupt the dump file but also the file system on the dump device.

Sharing dump devices between different versions of Linux

Do not share dump devices between Linux installations with different major releases.

For example, you should not share dump devices between Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6.

You can share dump devices between Linux installations with different service levels. Prepare the dump device with the **zipl** tool from the lowest service level. For example, if you have systems with Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7, you should prepare your dump device using the **zipl** tool from Red Hat Enterprise Linux 6. Newer tools such as **zgetdump** or dump analysis tools such as **crash** always can process dumps that have been created with older **zipl** versions. The other way around might work, but it is not guaranteed to work.

Sharing dump resources with VMDUMP

You can use **VMDUMP** concurrently on multiple z/VM guest virtual machines.

Note that the dump speed is slow and therefore is best for very small systems only. The shared resource here is the z/VM spool space. You have to ensure that it has enough space to hold multiple dumps created by **VMDUMP**.

Appendix A. Examples for initiating dumps

You can initiate dumps from different control points, such as the z/VM 3270 console or the HMC.

z/VM

You can initiate dumps from z/VM using `kdump`, a DASD device, tape, a SCSI device, or `VMDUMP`.

About this task

The following examples assume the 64-bit mode. Corresponding 31-bit examples would have a different PSW but be the same otherwise.

Using `kdump`

With `kdump`, you do not need a dump device to initiate the dump.

Before you begin

Your Linux instance must be set up for `kdump` as described in “Setting up `kdump`” on page 9.

Procedure

Issue the **system restart** z/VM CP command. For example from a 3270 terminal emulation for the Linux instance to be dumped:

```
#cp system restart
```

Boot messages for the `kdump` kernel indicate that the dump process started.

Using DASD

You can initiate a dump from a DASD device.

Example

If 193 is the dump device:

```
#cp cpu all stop  
#cp store status  
#cp i 193
```

On z/VM, a three-processor machine in this example, you will see messages about the disabled wait:

```
01: The virtual machine is placed in CP mode due to a SIGP stop from CPU 00.  
02: The virtual machine is placed in CP mode due to a SIGP stop from CPU 00.  
"CP entered; disabled wait PSW 00020000 80000000 00000000 00000000"
```

You can now IPL your Linux instance and resume operations.

Using tape

This example shows how you can initiate a dump from z/VM using tape.

Procedure

If 193 is the tape device:

1. Rewind the tape:

```
#cp rewind 193
```

2. Stop all CPUs:

```
#cp cpu all stop
```

3. Store status:

```
#cp store status
```

4. IPL the tape device:

```
#cp i 193
```

Results

On z/VM, a three-processor machine in this example, you will see messages about the disabled wait:

```
01: The virtual machine is placed in CP mode due to a SIGP stop from CPU 00.  
02: The virtual machine is placed in CP mode due to a SIGP stop from CPU 00.  
"CP entered; disabled wait PSW 00020000 80000000 00000000 00000000"
```

You can now IPL your Linux instance and resume operations.

Using SCSI

Initiating a dump using a SCSI disk.

Before you begin

SCSI dump from z/VM is supported as of z/VM 5.4.

About this task

Assume your SCSI dump disk has the following parameters:

- WWPN: 4712076300ce93a7
- LUN: 4712000000000000
- FCP adapter device number: 4711
- Boot program selector: 3

Results

Messages on the operating system console will show when the dump process is finished.

Example

```
#cp set dumpdev portname 47120763 00ce93a7 lun 47120000 00000000 bootprog 3
#cp ipl 4711 dump
```

What to do next

You can now IPL your Linux instance and resume operations.

In rare cases, you might want to overwrite or complement the existing SCSI dump tools parameters that have been configured with **zipl**. For example, you might want to change the dump mode setting. You can use a command of this form to specify SCSI dump tools parameters to be concatenated to the existing parameters:

```
#cp set dumpdev scpdata '<parameters>'
```

Enter this command before entering the **IPL** command.

In contrast to SCSI IPL configurations, where you can use a leading equal sign to replace all kernel parameters, you cannot use a leading equal sign to replace all SCSI dump tool parameters. Specifying the parameters with a leading equal sign causes the dump to fail.

Using VMDUMP

You can initiate a dump under z/VM by using VMDUMP.

Procedure

To initialize a dump with **VMDUMP**, issue this command from the console of your z/VM guest virtual machine:

```
#cp vmdump
```

Results

Dumping does not force you to perform an IPL. If the Linux instance ran as required before dumping, it continues running after the dump is completed.

Dumping NSSs

You can create dumps from NSSs with VMDUMP. As of z/VM 6.3 you can also use stand-alone dump with the nssdata IPL option to dump NSSs.

Procedure

- To initiate a DASD or tape dump for NSS, issue this command from the console of your z/VM guest virtual machine:

```
#cp i <devno> nssdata
```

where *<devno>* is the device number for the dump device.

- To initiate a SCSI dump for NSS, issue this command from the console of your z/VM guest virtual machine:

```
#cp i <devno> dump nssdata
```

where *<devno>* is the FCP adapter device number.

HMC or SE

You can initiate a dump process on an LPAR from an HMC (Hardware Management Console) or SE (Support Element).

About this task

The following description refers to an HMC, but the steps also apply to an SE. The steps are similar for DASD, tape, and SCSI. Differences are noted where applicable. You cannot initiate a dump with **VMDUMP** from the HMC or SE.

Procedure

1. In the left navigation pane of the HMC, expand Systems Management and Servers and select the mainframe system you want to work with. A table of LPARs is displayed in the upper content area on the right.
2. Select the LPAR for which you want to initiate the dump.
3. In the **Tasks** area, expand **Recovery**. Proceed according to your dump device:
 - If you are using **kdump**, click **PSW restart**. This initiates the dump process. Skip the remaining procedure, no further steps are required.
 - If you are dumping to DASD or tape, click **Stop all** in the **Recovery** list to stop all CPUs. Confirm when you are prompted to do so.
 - If you are dumping to a SCSI disk, skip this step and proceed with step 4 on page 49

Figure 4 on page 49 shows an example of an HMC with a selected mainframe system and LPAR. The **Load**, **PSW restart**, and **Stop all** tasks can be seen in the expanded **Recovery** list.




Figure 4. HMC with the Load, PSW restart, and Stop all tasks

4. Click **Load** in the **Recovery** list to display the Load panel.

For a dump to DASD or tape:

- Select **Load type** "Normal".
- Select the **Store status** check box.
- Type the device number of the dump device into the **Load address** field.

Figure 5 on page 50 shows a Load panel with all entries and selections required to start the dump process for a DASD or tape dump device.

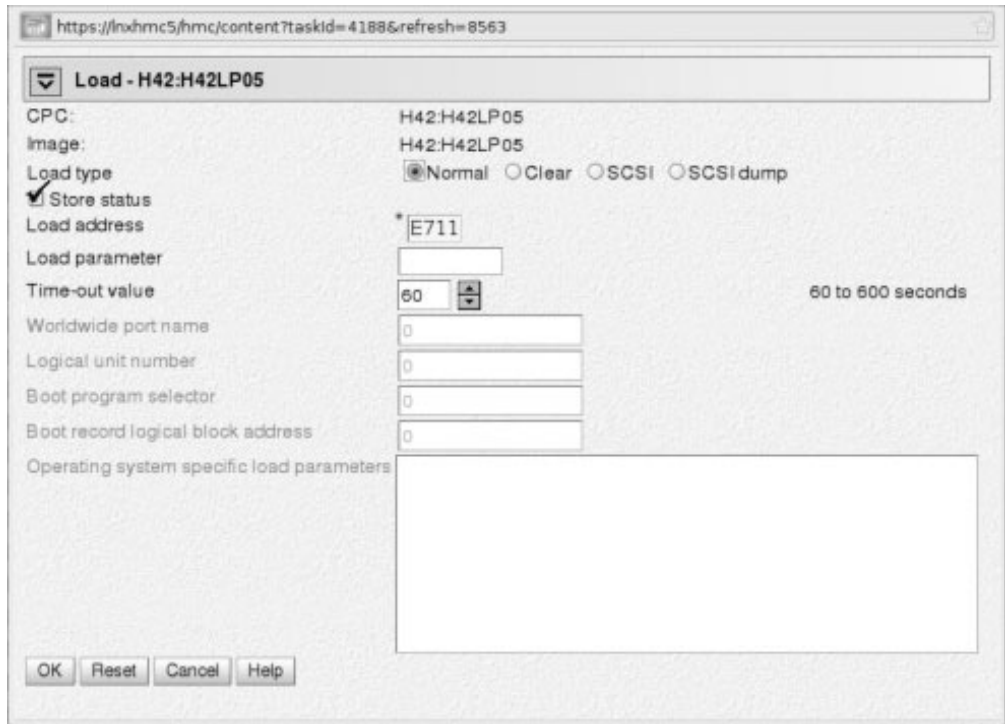


Figure 5. Load panel for dumping to DASD or tape

For a dump to SCSI disk:

- a. Select **Load type** "SCSI dump".
- b. Type the device number of the FCP adapter for the SCSI disk into the **Load address** field.
- c. Type the World Wide Port name of the SCSI disk into the **World wide port name** field.
- d. Type the Logical Unit Number of the SCSI disk into the **Logical unit number** field.
- e. Type the configuration number of the dump IPL configuration in the **Boot program selector** field.

The 'configuration number' defines the IPL or dump configuration which is to be IPLed. The numbering starts with 1 and is related to the menu of IPL/dump entries in the **zipl** configuration file for the SCSI disk. Configuration number 0 specifies the default configuration.

- f. Accept the defaults for the remaining fields.

In rare cases, you might want to overwrite or complement the existing SCSI dump tools parameters that have been configured with **zipl**. For example, you might want to change the dump mode setting. In the **Operating system specific load parameters** field, you can specify SCSI dump tools parameters to be concatenated to the existing parameters.

In contrast to SCSI IPL configurations, where you can use a leading equal sign to replace all kernel parameters, you cannot use a leading equal sign to replace all SCSI dump tool parameters. Specifying the parameters with a leading equal sign causes the dump to fail.

Figure 6 on page 51 shows a Load panel with all entries and selections required to start the SCSI dump process.

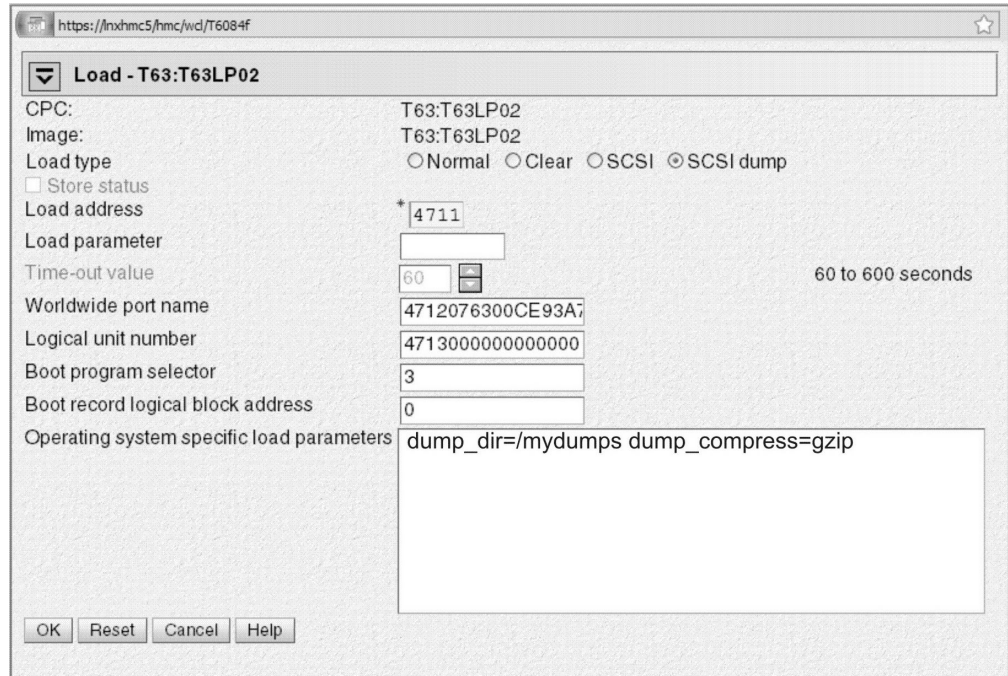


Figure 6. Load panel with enabled SCASI feature for dumping to SCASI disk

5. Click **OK** to start the dump process.
6. Wait until the dump process completes. Click the **Operating System Messages** icon for progress and error information.

Results

When the dump has completed successfully for a stand-alone dump tool, you can IPL Linux again. When using kdump, the re-IPL is typically done automatically by the kdump initrd after the dump has been copied.

Testing automatic dump-on-panic

Cause a kernel panic to confirm that your dump configuration is set up to automatically create a dump if a kernel panic occurs.

Before you begin

You need a Linux instance with active magic sysrequest functions.

Procedure

Crash the kernel with a forced kernel panic.

If your method for triggering the magic sysrequest function is:	Enter:
A command on the 3270 terminal or line-mode terminal on the HMC	^~c
A command on the hvc0 terminal device	<code>Ctrl+o c</code>
Writing to procfs	echo c > /proc/sysrq-trigger

Note: `Ctrl+0` means pressing 0 while holding down the control key. See *Device Drivers, Features, and Commands on Red Hat Enterprise Linux 7*, SC34-2710 for more details about the magic sysrequest functions.

Results

The production system crashes. If kdump is set up correctly, the kdump kernel is booted and the dump can be accessed through `/proc/vmcore`.

Appendix B. Obtaining a dump with limited size

The `mem` kernel parameter can make Linux use less memory than is available to it. A dump of such a Linux system does not need to include the unused memory. You can use the `zipl size` option to limit the amount of memory that is dumped.

About this task

The `size` option is available for all `zipl` based dumps: DASD, tape, and SCSI, in command-line mode or in configuration-file mode. The `size` option is appended to the dump device specification with a comma as separator.

The value is a decimal number that can optionally be suffixed with K for kilobytes, M for megabytes, or G for gigabytes. Values specified in byte or kilobyte are rounded to the next megabyte boundary.

Be sure not to make the dump size smaller than the amount of memory actually used by the system to be dumped. Limiting the dump size to less than the amount of used memory results in an incomplete dump.

Example

The following command prepares a DASD dump device for a dump that is limited to 100 megabyte:

```
# zipl -d /dev/dasdc1,100M
```

Appendix C. Command summary

The descriptions of the commands contain only the relevant options and parameters, for a full description refer to the man pages.

- The zgetdump tool
- The dumpconf service
- The crash tool
- The vmconvert tool
- “The vmur tool” on page 63

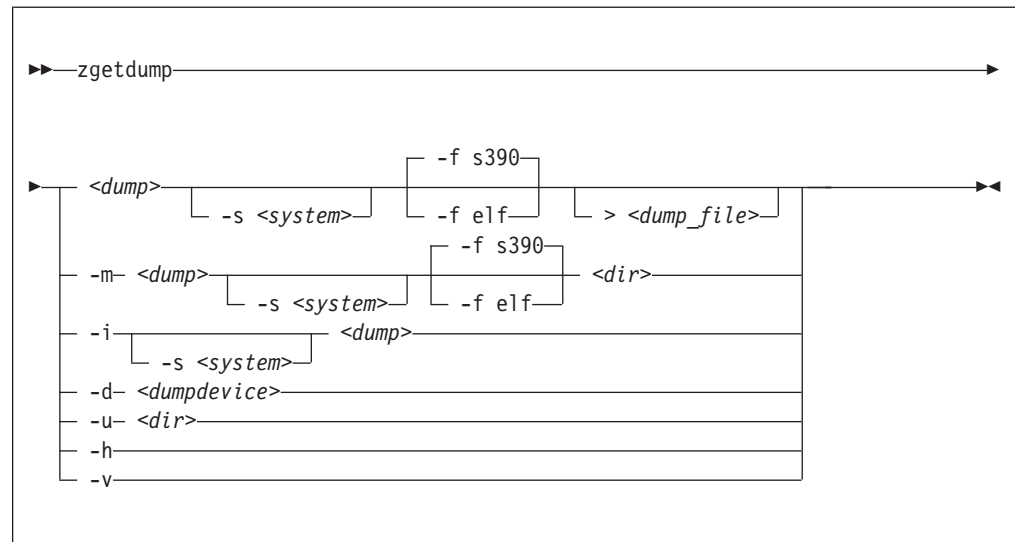
The zgetdump tool

The **zgetdump** tool reads or converts a dump.

The dump can be located either on a dump device or on a file system. The dump content is written to standard output, unless you redirect it to a specific file. You can also mount the dump content, print dump information, or check whether a DASD device contains a valid dump tool.

Before you begin: Mounting is implemented with "fuse" (file system in user space). Therefore the fuse kernel module must to be loaded before you can use the **--mount** option.

zgetdump syntax



Parameters

<dump>

is the file, DASD device or partition, or tape device node where the dump is located:

- Regular dump file (for example /testdir/dump.0)
- DASD partition device node (for example /dev/dasdc1)
- DASD device node for multivolume dump (for example /dev/dasdc)

- Tape device node (for example /dev/ntibm0)

Note: For a DASD multivolume dump it is sufficient to specify only one of the multivolume DASDs as *<dump>*.

<dump_file>

Is the file to which the output is redirected. The default is standard output.

<dumpdevice>

Specifies the dump device for the **-d** option. The device node of the DASD device, for example /dev/dasdb.

-s <system> or --select <system>

for dumps that capture two systems, selects the system of interest. This option is mandatory when accessing the dump of a crashed kdump instance, but returns an error if applied to a regular dump.

A dump can contain data for a crashed production system and for a crashed kdump system. A dump like this is created if a stand-alone dump tool is used to create a dump for a kdump instance that crashed while creating a dump for a previously crashed production system. *<system>* can be:

prod

to select the data for the crashed production system.

kdump

to select the data for the kdump instance that crashed while creating a dump for the previously crashed production system.

-m <dump> <dir> or --mount <dump> <dir>

Mounts the *<dump>* to mount point *<dir>* and generates a virtual target dump file instead of writing the content to standard output. The virtual dump file is named *dump.FMT*, where *FMT* is the name of the specified dump format (see the **--fmt** option).

-u <dir> or --umount <dir>

Unmounts the dump that is mounted at mount point *<dir>*. You can specify the dump itself instead of the directory, for example /dev/dasdd1. This option is a wrapper for **fusermount -u**.

-i <dump> or --info <dump>

Displays the dump header information from the dump and performs a validity check.

-d <dumpdevice> or --device <dumpdevice>

Checks whether the specified ECKD or FBA device contains a valid dump tool and prints information about it.

-f <format> or --fmt <format>

Uses the specified target dump format *<format>* when writing or mounting the dump. The following target dump formats are supported:

elf Executable and Linking Format core dump (64 bit only)

s390 S/390® dump (default)

-h or --help

Displays the help information for the command.

-v or --version

Displays the version information for the command.

Using zgetdump to copy a dump

Assuming that the dump is on DASD partition /dev/dasdb1 and that you want to copy it to a file named dump_file:

```
# zgetdump /dev/dasdb1 > dump_file
```

Using zgetdump to transfer a dump with ssh

Assuming that the dump is on DASD partition /dev/dasdd1 and that you want to transfer it to a file on another system with ssh:

```
# zgetdump /dev/dasdd1 | ssh user@host "cat > dump_file_on_target_host"
```

Using zgetdump to transfer a dump with FTP

Follow these steps to transfer a dump with FTP:

1. Establish an FTP session with the target host and log in.
2. To transfer a file in binary mode, enter the FTP **binary** command:

```
ftp> binary
```

3. To send the dump file to the host issue a command of the following form:

```
ftp> put |"zgetdump /dev/dasdb1" <dump_file_on_target_host>
```

Using zgetdump to copy a multi-volume dump

Assuming that the dump is on DASD devices /dev/dasdc and /dev/dasdd spread along partitions /dev/dasdc1 and /dev/dasdd1, and that you want to copy it to a file named multi_volume_dump_file:

```
# zgetdump /dev/dasdc > multi_volume_dump_file
```

For an example of the output from this command, see Chapter 4, "Using DASD devices for multi-volume dump," on page 17.

Using zgetdump to copy a tape dump

Assuming that the tape device is /dev/ntimb0:

```
# zgetdump /dev/ntimb0 > dump_file
Format Info:
Source: s390tape
Target: s390

Copying dump:
00000000 / 00001024 MB
00000171 / 00001024 MB
00000341 / 00001024 MB
00000512 / 00001024 MB
00000683 / 00001024 MB
00000853 / 00001024 MB
00001024 / 00001024 MB

Success: Dump has been copied
```

Using zgetdump to create a dump from a live system

To store an ELF-format dump from a live system in a file called `dump.elf`, issue:

```
# nice -n -20 zgetdump /dev/crash -f elf > dump.elf
```

For an example of the output from this command, see “Creating a kernel dump on a live system” on page 31.

Checking whether a tape dump is valid, and printing the dump header

Assuming that the tape device is `/dev/ntibm0`:

```
# zgetdump -i /dev/ntibm0
Checking tape, this can take a while...
General dump info:
Dump format.....: s390tape
Version.....: 5
Dump created.....: Mon, 14 Jan 2013 17:26:46 +0200
Dump ended.....: Mon, 14 Jan 2013 17:27:58 +0200
Dump CPU ID.....: ff00012320948000
UTS kernel release.: 3.10.0-123.el7.s390x
UTS kernel version.: #1 SMP Mon Nov 19 16:52:53 EST 2012
Build arch.....: s390x (64 bit)
System arch.....: s390x (64 bit)
CPU count (online): 2
CPU count (real)...: 2
Dump memory range..: 1024 MB
Real memory range..: 1024 MB

Memory map:
0000000000000000 - 000000003fffffff (1024 MB)
```

Checking whether a DASD dump is valid and printing the dump header

Assuming that the dump is on a partition, `part1`, of a DASD device `/dev/dasdb1`:

```
# zgetdump -i /dev/dasdb1
General dump info:
Dump format.....: s390
Version.....: 5
Dump created.....: Mon, 10 May 2010 17:32:36 +0200
Dump ended.....: Mon, 10 May 2010 17:32:48 +0200
Dump CPU ID.....: ff00012320948000
Build arch.....: s390x (64 bit)
System arch.....: s390x (64 bit)
CPU count (online): 2
CPU count (real)...: 2
Dump memory range..: 1024 MB
Real memory range..: 1024 MB

Memory map:
0000000000000000 - 000000003fffffff (1024 MB)
```

Checking whether a device contains a valid dump record

Checking DASD device `/dev/dasda`, which is a valid dump device:

```
# zgetdump -d /dev/dasdb
Dump device info:
Dump tool.....: Single-volume DASD dump tool
Version.....: 2
Architecture.....: s390x (64 bit)
DASD type.....: ECKD
Dump size limit...: none
```

Checking DASD device `/dev/dasdc`, which is not a valid dump device:

```
# zgetdump -d /dev/dasdc
zgetdump: No dump tool found on "/dev/dasdc"
```

Using the mount option

Mounting is useful for multivolume DASD dumps. After a multivolume dump has been mounted, it is shown as a single dump file that can be accessed directly with dump processing tools such as **crash**.

The following example mounts a multivolume DASD dump as an ELF dump, processes it with **crash**, and unmounts it with **zgetdump**:

```
# zgetdump -m -f elf /dev/dasdx /dumps
# crash vmlinux /dumps/dump.elf
# zgetdump -u /dumps
```

Mounting can also be useful when you want to process the dump with a tool that cannot read the original dump format. To do this, mount the dump and specify the required target dump format with the `--fmt` option.

Selecting data from a dump that includes a crashed kdump

The following example mounts dump data for a crashed production system from a DASD backup dump for a failed kdump (see “Failure recovery and backup tools” on page 9 for details).

```
# zgetdump -s prod -m /dev/dasdb1 /mnt
```

The dumpconf service

The **dumpconf** service configures the action to be taken if a kernel panic or PSW restart occurs.


The service is installed as a script under `/etc/init.d/dumpconf` and reads the configuration file `/etc/sysconfig/dumpconf`.

Note: kdump does not depend on **dumpconf** and can neither be enabled nor disabled with **dumpconf**. If kdump has been set up for your production system, dump tools as configured with **dumpconf** are used only if the integrity check for kdump fails. With kdump set up, you can use **dumpconf** to enable or disable backup dump tools. See also “Failure recovery and backup tools” on page 9.

To enable the **dumpconf** service, issue:

```
# chkconfig --add dumpconf
```

dumpconf service syntax



Parameters

start

Enable configuration defined in `/etc/sysconfig/dumpconf`.

stop

Disable the **dumpconf** service.

status

Show current configuration status of the **dumpconf** service.

-h or --help

Display short usage text on console. To view the man page, enter **man dumpconf**.

-v or --version

Display version number on console, and exit.

Keywords for the configuration file

ON_PANIC

Shutdown action to be taken if a kernel panic or PSW restart occurs. Possible values are:

dump Dump Linux and stop system.

reipl Reboot Linux.

dump_reipl

Dump Linux and reboot system.

vmcmd

Execute specified CP commands and stop system.

stop Stop Linux (default).

DELAY_MINUTES

The number of minutes that the activation of the **dumpconf** service is to be delayed. The default is zero.

Using **reipl** or **dump_reipl** actions with **ON_PANIC** can lead to the system looping with alternating IPLs and crashes. Use **DELAY_MINUTES** to prevent such a loop. **DELAY_MINUTES** delays activating the specified panic action for a newly started system. When the specified time has elapsed, the **dumpconf** service activates the specified panic action. This action is taken should the system subsequently crash. If the system crashes before the time has elapsed, the previously defined action is taken. If no previous action has been defined, the default action (STOP) is performed.

VMCMD_<X>

Specifies a CP command, <X> is a number from one to five. You can specify up to five CP commands that are executed in case of a kernel panic or PSW restart. Note that z/VM commands, device addresses, and names of z/VM guest virtual machines must be uppercase.

DUMP_TYPE

Type of dump device. Possible values are **ccw** and **fcp**.

DEVICE

Device number of dump device.

WWPN

WWPN for SCSI dump device.

LUN LUN for SCSI dump device.

BOOTPROG

Boot program selector

BR_LBA

Boot record logical block address.

Example configuration files for the dumpconf service

- Example configuration for a CCW dump device (DASD) using **reipl** after dump and **DELAY_MINUTES**:

```
ON_PANIC=dump_reipl
DUMP_TYPE=ccw
DEVICE=0.0.4714
DELAY_MINUTES=5
```

- Example configuration for FCP dump device (SCSI disk):

```
ON_PANIC=dump
DUMP_TYPE=fcp
DEVICE=0.0.4711
WWPN=0x5005076303004712
LUN=0x4713000000000000
BOOTPROG=0
BR_LBA=0
```

- Example configuration for re-IPL if a kernel panic or PSW restart occurs:

```
ON_PANIC=reipl
```

- Example of sending a message to the z/VM guest virtual machine "MASTER", executing a **CP VMDUMP** command, and rebooting from device 4711 if a kernel panic or PSW restart occurs:

```
ON_PANIC=vmcmd
VMCMD_1="MSG MASTER Starting VMDUMP"
VMCMD_2="VMDUMP"
VMCMD_3="IPL 4711"
```

Note that z/VM commands, device addresses, and names of z/VM guest virtual machines must be uppercase.

Examples for using the dumpconf service

Use the **dumpconf** service to enable and disable the configuration.

- To enable the configuration:

```
# service dumpconf start
ccw dump device configured. "dump" on panic configured.
```

- To display the status:

```
# service dumpconf status
type....: ccw
device..: 0.0.4714
on_panic: dump
```

- To disable dump on panic:

```
# service dumpconf stop
Dump on panic is disabled now
```

- To display the status again and check that the status is now stopped.

```
# service dumpconf status
on_panic: stop
```

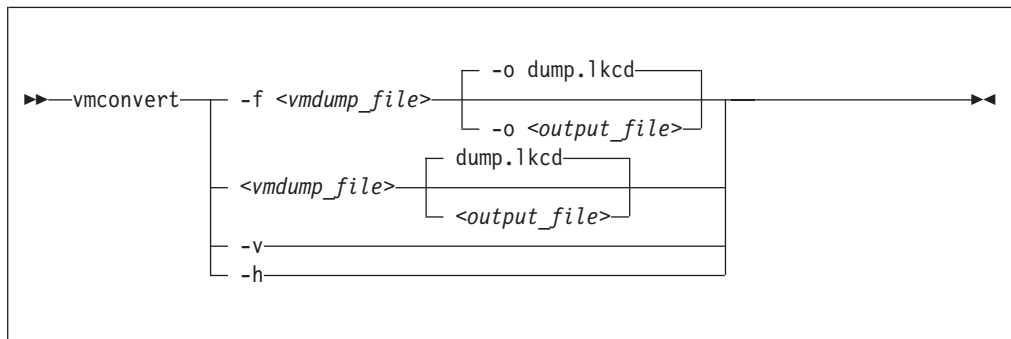
The crash tool

The **crash** tool is a GPL-licensed tool maintained by Red Hat. For more details see the tool online help.

The vmconvert tool

The **vmconvert** tool converts a dump that was created with **VMDUMP** into a file that can be analyzed with **crash**.

vmconvert syntax



Parameters

<vmdump_file> or **-f <vmdump_file>** or **--file <vmdump_file>**
Specifies the **VMDUMP** created dump file to be converted.

<output_file> or **-o <output_file>** or **--output <output_file>**
Specifies the name of the dump file to be created. The default is **dump.1kcd**.

-v or --version

Displays the tool version.

-h or --help

Displays the help information for the command.

Example

To convert a **VMDUMP**-created dump file `vmdump1` into a dump file `dump1.lkcd` that can be processed with **crash** issue:

```
# vmconvert -f vmdump1 -o dump1.lkcd
```

You can also use positional parameters:

```
# vmconvert vm.dump lkcd.dump
vmdump information:
architecture: 32 bit
date.....: Fri Feb 18 11:06:45 2005
storage.....: 16 MB
cpus.....: 6
 16 of 16 |#####| 100%
'lkcd.dump' has been written successfully.
```

The vmur tool

The **vmur** command can receive a **VMDUMP** file from the z/VM reader and convert it into a file that can be analyzed with **crash**.

Issue a command of the following form:

```
# vmur receive -c <spool ID> <dump file name>
```

Parameters

<spool ID>

Specifies the **VMDUMP** file spool ID.

<dump file name>

Specifies the name of the output file to receive the reader spool file's data.

For more details, see the **vmur** man page and *Device Drivers, Features, and Commands on Red Hat Enterprise Linux 7, SC34-2710*

Example

To receive and convert a **VMDUMP** spool file with spool ID 463 to a file named `dump.lkcd` on the Linux file system in the current working directory:

```
# vmur rec -c 463 dump.lkcd
```

Appendix D. Preparing for analyzing a dump

To analyze your dump with **crash**, additional files are required.

If you need to send your dump for analysis, it might be good to include these additional files with the dump file. Your distribution typically provides the additional files in RPMs.

If a Red Hat Enterprise Linux 7 dump is to be analyzed with **crash**, include:

- **vmlinux** (full): Contains addresses of kernel symbols and datatype debug information

Red Hat Enterprise Linux 7 debug files

The Red Hat Enterprise Linux 7 debug file is:

Table 4. Red Hat Enterprise Linux 7 debug file name

Debug file	Path
vmlinux (full)	/usr/lib/debug/lib/modules/3.10-xx.e17.s390x/vmlinux

The RPM that contains this file is:

Table 5. Red Hat Enterprise Linux 7 debuginfo RPM name

Red Hat Enterprise Linux version	RPM
Red Hat Enterprise Linux 7	kernel-debuginfo-3.10-xx.e17.s390x.rpm

Appendix E. How to detect guest relocation

Information about guest relocations is stored in the s390 debug feature (s390dbf). You can access this information in a kernel dump or from a running Linux instance.

About this task

You can detect if a Linux instance was moved to another z/VM guest virtual machine or LPAR. One available mechanism for guest relocation is z/VM Single System Image (SSI).

You can access the s390 debug feature **lgr** from a live system or with the **crash** tool from a kernel dump. When the debug feature contains only one entry, no relocation was detected and the entry identifies the boot virtual guest machine. For each detected relocation one additional entry is written.

Procedure

Choose the method that suits your purpose:

- Use the **crash** tool to read from a kernel dump. Issue a command of this form:

```
# crash <vmlinux files> <dump file>
crash> s390dbf lgr hex_ascii
```

- Use the **cat** command on a live system to read the debugfs entry for **lgr**. Issue a command of this form:

```
# cat /sys/kernel/debug/s390dbf/lgr/hex_ascii
```

Example

Assume that one relocation of the z/VM guest virtual machine ZVMGUEST was detected from a z/VM in LPAR VM000A to a z/VM in LPAR VM000B. You can see this relocation in the kernel dump:

```
# crash vmlinux dump
crash> s390dbf lgr hex_ascii
00 01317816806:277332 3 - 00 .. | ... IBM2817000000000000EAA1402 ... VM000A ... ZVMGUEST
00 01317866806:277332 3 - 00 .. | ... IBM2817000000000000EAA1402 ... VM000B ... ZVMGUEST
```

Alternatively, you can see such a relocation from a running system:

```
# cat /sys/kernel/debug/s390dbf/lgr/hex_ascii
00 01317816806:277332 3 - 00 .. | ... IBM2817000000000000EAA1402 ... VM000A ... ZVMGUEST
00 01317866806:277332 3 - 00 .. | ... IBM2817000000000000EAA1402 ... VM000B ... ZVMGUEST
```

For more information about the complete s390dbf record, see the struct `os_info` definition in the Linux kernel source code.

Accessibility

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Documentation accessibility

The Linux on System z publications are in Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience difficulties when you use the PDF file and want to request a Web-based format for this publication, use the Reader' Comments form in the back of this publication, send an email to eservdoc@de.ibm.com, or write to:

IBM Deutschland Research & Development GmbH
Information Development
Department 3282
Schoenaicher Strasse 220
71032 Boeblingen
Germany

In the request, be sure to include the publication number and title.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility at

www.ibm.com/able

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Adobe is either a registered trademark or trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Index

Special characters

- /etc/init.d/dumpconf
 - dumpconf configuration file 59
- /etc/kdump.conf
 - configuration file for kdump 9

A

- accessibility 69
- analyzing dump
 - preparing for 65
- attach and detach
 - use as locking mechanism 41
- authority
 - root v
- automatic dump
 - dump-on-panic 51
 - sharing devices 41
- automatic dump-on-panic 51

C

- commands
 - crash 62
 - dumpconf 59
 - summary 55
 - vmconvert 62
 - vmur 63
 - zgetdump 55
- compressing dump
 - using gzip and split 37
 - using makedumpfile 36
- configuration file
 - dumpconf 59
 - kdump 9
- crash 62
 - live dump, opening 32
 - preparing for analyzing a dump with 65
- crashkernel=
 - kernel parameter 7
- creating dumps
 - tools overview 1

D

- DASD
 - vmcmd panic action 41
- DASD devices
 - on LPARs, sharing 40
 - under z/VM, sharing 40
 - using as dump device 13
 - using for multi-volume dump 17
- DASD devices, sharing 40
- DASD dump
 - initiating 14
- DASD dump tool 18
 - installing 13

- DASD multi-volume dump
 - starting 19
- detach and attach
 - use as locking mechanism 41
- developerWorks 3
- device locking 39
- dump
 - compressing using gzip and split 37
 - compressing using makedumpfile 36
 - copy from DASD with zgetdump
 - command 14
 - copy multi-volume dump from DASD
 - with zgetdump command 20
 - copy tape from 22
 - copy with zgetdump command 23
 - limited size 53
 - live-system 4
 - multi-volume 17
 - preparing for analyzing 65
 - sharing space for 2
 - starting a multi-volume DASD 19
 - tape, checking if valid 23
 - tape, initializing 21
- dump devices
 - DASD 13, 41
 - DASD, shared, on LPAR 40
 - DASD, shared, under z/VM 40
 - definition 3
 - SCSI 25
 - SCSI, shared 41
 - sharing 39
 - sharing between different Linux versions 42
 - sharing with VMDUMP 43
 - sharing, when dumping
 - automatically 41
 - sharing, when dumping manually 40
 - tape 21
- dump methods
 - comparison 4
- dump panic action 41
- dump tool
 - crash 62
 - DASD, installing 13
 - tape, installing 21
- dump tools
 - dumpconf service 59
 - multi-volume, DASD 18
 - stand-alone 3
 - stand-alone tape 21
 - summary 55
 - vmconvert 62
 - VMDUMP 3, 29
 - vmur 63
 - zgetdump 55
- dump tools overview 1
- dump_reipl panic action 41
- dumpconf service 59

F

- FCP-attached SCSI devices 42
- firstboot utility
 - kdump setup 9

G

- guest relocation 67
- gzip command 37

H

- handling large dumps 35

I

- initiate dumps 10, 45
 - from z/VM 45
 - HMC or SE 48
 - using a DASD device 45
 - using VMDUMP 47
 - using VMDUMP for NSS 47
- initiating a dump 29
- initiating a dump to tape 46
- initiating dump
 - using SCSI 46
- initiating dumps
 - examples 45

K

- kdump
 - advantages and disadvantages 7
 - comparison with other dump methods 4
 - initiate 10
 - initiating 45
 - introduction 2
 - setup 9
 - testing automatic dump-on-panic 51
- kdump kernel 7
- kernel dump
 - creating from live system 31
- Kernel Dump Configuration utility
 - kdump setup 9

L

- large dump
 - handling 35
 - multi-volume 17
- limit amount of memory dumped 53
- Linux versions
 - sharing dump devices between 42
- live dump
 - open with crash 32
- live system
 - creating dump 31

- live-system dump 4
- locking mechanism
 - under z/VM 41
- LPARs
 - sharing DASD devices on 40

M

- makedumpfile 36
- memory
 - reserved for kdump kernel 7
- memory dump
 - compressing 36
 - compressing, using gzip and split 37
- messages
 - tape display 22
- multi-volume dump
 - starting 19
- multi-volume dumps 17
 - DASD tool 18

N

- non-disruptive dump
 - using zgetdump 31

O

- opening live-system dump
 - using crash 32

P

- parameters
 - SCSI dump tool 25
- profs
 - required authority v

R

- relocation
 - guest 67
- root authority v

S

- s390utils 3
- SCSI devices
 - for dumping, shared 41
 - used for automatic dumping 42
- SCSI dump
 - initiating 27
 - printing the dump header 27
 - single partition 26
- SCSI dump device 25
- SCSI dump tool
 - installing 25
 - parameters 25
- SCSI Linux System Dumper 25
- SD
 - See SCSI Linux System Dumper
- serialization 39
- single partition
 - used for SCSI dump 26

- split command 37
- summary
 - commands for dumps 55
- system restart
 - z/VM CP command 45

T

- tape 46
 - copy dump from 22
 - display messages 22
 - use for dumping, preparing 22
 - using as dump device 21
- tape dump
 - checking if valid 23
 - initializing 21
- tape dump tool
 - installing 21
- testing 51
- tools for creating dumps 1
- transfer time
 - reducing with kdump 2

U

- using kdump 10, 45

V

- vmcmd panic action 41
- vmconvert 62
- VMDUMP 29
 - comparison with other dump
 - methods 4
 - copying dump 30
 - initiate dump process 29
 - introduction 3
 - sharing dump resources with 43
- vmur 63
 - using to copy VMDUMP dump 30

Z

- z/VM 40
 - z/VM CP command
 - system restart 45
- zgetdump 55
 - comparison with other dump
 - methods 4
 - create a dump from a live system 31
- zgetdump tool 14, 20, 23
- zipl
 - size option 53

Readers' Comments — We'd Like to Hear from You

Linux on System z
Using the Dump Tools
on Red Hat Enterprise Linux 7

Publication No. SC34-2711-00

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:

- Send your comments to the address on the reverse side of this form.
- Send your comments via email to: eservdoc@de.ibm.com

If you would like a response from IBM, please fill in the following information:

Name

Address

Company or Organization

Phone No.

Email address



Fold and Tape

Please do not staple

Fold and Tape

PLACE
POSTAGE
STAMP
HERE

IBM Deutschland Research & Development GmbH
Information Development
Department 3248
Schoenaicher Strasse 220
71032 Boeblingen
Germany

Fold and Tape

Please do not staple

Fold and Tape



SC34-2711-00

