

Linux on Z and LinuxONE  
Development stream (Kernel 5.2)

*Using the Dump Tools*



**Note**

Before using this information and the product it supports, read the information in [“Notices” on page 61.](#)

This edition applies to the Linux on Z Development stream for kernel 5.2, s390-tools version 2.7.1, and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2004, 2019.**

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

---

# Contents

<b>Summary of changes.....</b>	<b>v</b>
Updates for kernel 4.17 to 5.2.....	v
Updates for kernel 4.16.....	v
Updates for kernel 4.15.....	v
<b>About this document.....</b>	<b>vii</b>
Linux on Z and LinuxONE publications.....	vii
Device nodes used in this publication.....	viii
<b>Chapter 1. Planning for dumps.....</b>	<b>1</b>
Tool overview.....	1
Maximum dump size by tool.....	3
Dump methods compared.....	4
Automatic dumping.....	5
Dump tool version dependencies.....	5
<b>Chapter 2. Using kdump.....</b>	<b>7</b>
How kdump works on Z.....	7
Setting up kdump.....	9
Initiating a dump.....	10
Accessing the dump.....	12
<b>Chapter 3. Using a DASD dump device.....</b>	<b>15</b>
Installing the DASD dump tool.....	15
Initiating a DASD dump.....	16
Copying the dump from DASD with zgetdump.....	19
<b>Chapter 4. Using DASD devices for multivolume dump.....</b>	<b>21</b>
Installing the multivolume DASD dump tool.....	22
Initiating a multivolume DASD dump.....	23
Copying a multivolume dump to a file.....	24
<b>Chapter 5. Using a tape dump device.....</b>	<b>25</b>
Installing the tape dump tool.....	25
Initiating a tape dump.....	25
Copying the dump from tape.....	28
<b>Chapter 6. Using a SCSI disk dump device.....</b>	<b>31</b>
Installing the SCSI disk dump tool .....	31
Initiating a SCSI dump.....	32
Copying the dump from SCSI disks with zgetdump.....	35
Printing the SCSI dump header.....	35
<b>Chapter 7. Using VMDUMP.....</b>	<b>37</b>
Initiating a dump with VMDUMP .....	37
Copying the dump to Linux.....	38
<b>Chapter 8. Creating live-system dumps with zgetdump.....</b>	<b>39</b>
Creating a kernel dump on a live system.....	39

Opening a live-system dump with the crash tool.....	40
<b>Chapter 9. Processing dumps.....</b>	<b>43</b>
Reducing dump size.....	43
Sending a dump to IBM Support.....	46
<b>Appendix A. Obtaining a dump with limited size.....</b>	<b>47</b>
<b>Appendix B. Command summary.....</b>	<b>49</b>
The zgetdump tool.....	49
The dumpconf service.....	54
The crash tool.....	56
The vmconvert tool.....	57
The vmur tool.....	57
<b>Accessibility.....</b>	<b>59</b>
<b>Notices.....</b>	<b>61</b>
Trademarks.....	61
<b>Index.....</b>	<b>63</b>

# Summary of changes

---

This revision reflects changes to the Development stream for kernel 5.2.

## Updates for kernel 4.17 to 5.2

---

This edition (SC33-8412-19) contains changes related to kernels 4.17 to 5.2.

### New information

- The kexec tool now provides an option for checking the validity of a kdump kernel, see [“Setting up kdump”](#) on page 9.
- You can now analyze dumps of KASLR kernels with crash, see [“The crash tool”](#) on page 56

### Changed Information

- None.

This revision also includes maintenance and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

### Deleted Information

- None.

## Updates for kernel 4.16

---

This edition (SC33-8412-18) contains changes related to the kernel 4.16 release.

### New information

- Stand-alone DASD dump tools now drop empty memory areas on dumping to disk, see [“The zgetdump tool”](#) on page 49.

### Changed Information

- None.

This revision also includes maintenance and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

### Deleted Information

- None.

## Updates for kernel 4.15

---

This edition (SC33-8412-17) contains changes related to kernels 4.11 - 4.15.

### New information

- Systemd is now used in place of System V, see [“The dumpconf service”](#) on page 54.

**Changed Information**

- None.

This revision also includes maintenance and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

**Deleted Information**

- None.

## About this document

---

This document describes tools for obtaining dumps of Linux on IBM Z® (Z) or IBM LinuxONE (LinuxONE) instances. It describes how to use DASD, channel-attached tape, and SCSI disks, as well as how to use kdump and VMDUMP.

For details about IBM® tested Linux environments, see [www.ibm.com/systems/z/os/linux/resources/testedplatforms.html](http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html).

Unless stated otherwise, all IBM z/VM® related information in this document assumes a current z/VM version, see [www.vm.ibm.com/techinfo/lpmigr/vmleos.html](http://www.vm.ibm.com/techinfo/lpmigr/vmleos.html).

You can find the latest version of this publication on the IBM Knowledge Center at:

[www.ibm.com/developerworks/linux/linux390/documentation\\_dev.html](http://www.ibm.com/developerworks/linux/linux390/documentation_dev.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.

## Linux on Z and LinuxONE publications

---

You can find publications for Linux on Z and LinuxONE on IBM Knowledge Center.

These publications are available on IBM Knowledge Center at [www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz\\_r\\_lib.html](http://www.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_lib.html)

- *Device Drivers, Features, and Commands*
- *Using the Dump Tools*
- *Running Docker Containers on IBM Z*, SC34-2781
- *KVM Virtual Server Quick Start*, SC34-2753
- *KVM Virtual Server Management*, SC34-2752
- *How to use FC-attached SCSI devices with Linux on z Systems®*, SC33-8413
- *libica Programmer's Reference*, SC34-2602
- *Exploiting Enterprise PKCS #11 using openCryptoki*, SC34-2713
- *Secure Key Solution with the Common Cryptographic Architecture Application Programmer's Guide*, SC33-8294
- *Pervasive Encryption for Data Volumes*, SC34-2782
- *How to set an AES master key*, SC34-7712
- *Troubleshooting*, SC34-2612
- *Kernel Messages*, SC34-2599
- *How to Improve Performance with PAV*, SC33-8414
- *How to Set up a Terminal Server Environment on z/VM*, SC34-2596

## Device nodes used in this publication

---

There can be multiple device nodes for the same device.

The DASD and tape examples in this publication use the standard device nodes. The SCSI examples use multipath device nodes. If you are using a Linux distribution that provides udev, you can also use the device nodes that udev creates for you. See your distribution documentation to find out which nodes are available.



# Chapter 1. Planning for dumps

Be prepared before disaster strikes! Consider what dump method you want to use, what size dumps you need to handle, and what possibilities exist to limit the size or spread the dump over several devices.

## Tool overview

Different tools can be used for obtaining dumps of Linux on Z instances.

**Note:** As of IBM z13®, simultaneous multithreading is available for Linux in LPAR mode. You must use dump tools that support multithreading on such LPARs. See [“Dump tool version dependencies”](#) on page 5.

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 on page 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
System size (see note “1” on page 1)	Large	Small	Large	Large	Large	Small	Large
Speed	Fast	Fast	Fast	Fast	Slow	Slow	Fast
Medium	Any available medium	ECKD or FBA DASD (see note “2” on page 2)	ECKD DASD	SCSI partition	Tape cartridges	z/VM reader	Any available medium
Compression possible	While writing	No	No	No	Yes (see note “1” on page 1)	No	No
Dump filtering possible	While writing	When copying	When copying	When copying	When copying	When copying	No
Disruptive (see note “3” on page 2)	Yes	Yes	Yes	Yes	Yes	No	No
Stand-alone	No	Yes	Yes	Yes	Yes	No	No

**Note:**

1. For dump system sizes, see also [“Maximum dump size by tool”](#) on page 3.

2. SCSI disks can be emulated as FBA disks. This dump method can, therefore, be used for SCSI-only z/VM installations.
3. In this context, disruptive means that the dump process kills a running operating system.

## kdump

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 procfs 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.

IPLing the production system again clears the dump in memory, so the dump must be saved to persistent storage before a subsequent re-IPL.

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.

## 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 s390-tools package as part of the **zipl** package:

- 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 must provide enough space for the dump. For Linux on z/VM, the dump device must be on subchannel set 0. For Linux in LPAR mode, the device can be on any subchannel set.

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 device, channel-attached tape, or 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 54](#).

All examples for installing stand-alone tools by using a **zipl** configuration file assume that `/etc/zipl.conf` is used as the configuration file and that `/etc/zipl.conf` is the default configuration file.

For more information on **zipl**, refer to the **zipl** man page and to the **zipl** description in *Device Drivers, Features, and Commands*, SC33-8411. You can find the latest version of this document on developerWorks® at:

[www.ibm.com/developerworks/linux/linux390/documentation\\_dev.html](http://www.ibm.com/developerworks/linux/linux390/documentation_dev.html)

## 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*, SC24-6268.

## Live-system dump

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

Use the **zgetdump** tool that is shipped with the s390-tools package to create a kernel dump while the Linux system continues running. No dump device must be prepared, because the `/dev/mem` device node is used to create the dump.

## Maximum dump size by tool

---

The dump size depends on the size of the system for which the dump is to be created.

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 A provides several hundreds of GB, depending on the storage server model. To write a large memory dump to single volume DASD, you can use:

- 3390 model 9 (up to 45 GB)
- 3390 model A (up to 1 TB)

### **Multivolume DASD**

Can be up to the combined size of 32 DASD partitions. To write a large memory dump to multiple DASD volumes, you can use up to 32 DASDs:

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

### **SCSI**

Depends on the capacity of the SCSI disk and which other data it contains. To write a large memory dump to SCSI disk, you can use:

#### **z/VM emulated FBA device that represents a real SCSI 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 disk**

The SCSI disk size depends on your storage server setup. For SCSI dump partitions greater than 2 TB, you must use the GPT disk layout.

### **Channel-attached tape**

Depends on the tape drive. For example, IBM TotalStorage™ Enterprise Tape System 3592 supports large dumps and also offers hardware compression. For large memory dumps, cartridges with up to 10 TB capacity are available.

### **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 [“Reducing dump size” on page 43](#) for information specific to large dumps.

## 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 kdump kernel and the initial RAM disk into the memory of the production system Use <b>kexec</b> or  <pre>systemctl start kdump</pre>	Write the stand-alone dump tool to the dump device ( <b>zipl</b> )  Define the panic shutdown action ( <b>dumpconf</b> )	Define the panic shutdown action ( <b>dumpconf</b> )	None
Dump trigger	<b>Automatic:</b> Kernel panic  <b>Initiated by operator:</b> PSW restart	<b>Automatic:</b> Kernel panic  <b>Initiated by operator:</b> IPL of the dump device	<b>Automatic:</b> Kernel panic  <b>Initiated by operator:</b> z/VM CP <b>VMDUMP</b> command	<b>Initiated by operator:</b> <b>zgetdump</b> invocation
Initial dump space	Memory	Dump device	Spool device	Memory
Accessing the initial dump	Through <code>/proc/vmcore</code> from the kdump instance (automatically done by kdump initrd)	Using <b>zgetdump</b> from a new Linux instance	Using <b>vmux -c</b> from a new Linux instance	Through <code>/dev/mem</code>
Copying the initial dump to the final dump store (and releasing the initial dump space)	Copied from the kdump instance to any available storage (automatically done by kdump initrd)	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 <b>makedumpfile</b> on the kdump instance (automatically done by kdump initrd)	Using <b>zgetdump</b> and <b>makedumpfile</b> on the new Linux instance	Using <b>zgetdump</b> and <b>makedumpfile</b> on the new Linux instance	Not recommended

# Automatic dumping

You can configure a dump device that is automatically used when a kernel panic occurs.

If you set up kdump, a kernel panic or PSW restart automatically triggers a dump. You can also use dumpconf to set up automatic dumping for standalone tools, or for your backup dump solution.


See “The dumpconf service” on page 54 for how to set up dumpconf. Once you have set up the automation, you can cause a kernel panic to test the configuration.

# Dump tool version dependencies

Use dump tools with multithreading support for Linux instances that run on LPARs with enabled multithreading.

Table 3 on page 5 shows how you can verify the version of the different tools.

Table 3. Tool versions for multithreading	
Dump tool	Your tools support multithreading if:
DASD dump tool Tape dump tool	The command <code>zipl -v</code> displays a version of 1.30.0 or later, and the command <code>zgetdump -d &lt;dump device&gt;</code> displays either the extended dump tool or a dump tool version of 4 or later.
SCSI dump tool	The Linux kernel that is used for the SCSI dump tool is version 4.0 or later.
kdump	The kdump kernel version is 4.0 or later.

 **Attention:** If you must IPL a Linux kernel with a version earlier than 4.0 on an LPAR that has previously run with multithreading enabled, use `IPL clear`. Otherwise any memory dumps that you create for this Linux kernel might be flawed.



---

## Chapter 2. Using kdump

You can use kdump to create system dumps for instances of Linux on Z.

### Before you begin

For using kdump, you require:

- A Linux kernel image that has been compiled with the common code kernel configuration options CONFIG\_KEXEC, CONFIG\_CRASH\_DUMP, and CONFIG\_PROC\_VMCORE. If you want to enable validity checks for the kdump kernel, you also need the CONFIG\_KEXEC\_FILE option. This requirement applies to both, the production system for which kdump is set up and to the kdump image.
- An initial kdump RAM disk if required by the kdump kernel image.
- A kexec package with kdump support.
- A System z9® or later mainframe system.

### 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 cannot be used for issues that occur before kdump is initialized, for example, for early boot problems. For such cases, use a standalone dump tool.
- 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).
- For production systems that run in LPAR mode, kdump consumes memory (see [“Memory consumption”](#) on page 8)

---

## How kdump works on 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 init service loads the kdump kernel and its initial RAM disk (initrd) 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 was 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. Typically, the **makedumpfile** tool is now used to write a filtered and compressed version of the dump to a file on persistent storage, locally or over a network. Again, this method saves time because the dump is reduced in size while it is written or transferred.

After the dump is written, the production system can be IPLed again.

### 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 amount of 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 54](#).

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 49](#).

## Setting up kdump

---

Before you can use kdump, you must load the kdump kernel into a reserved memory area.

### Before you begin

- The steps that are described here are typically done for you by your distribution or through distribution-specific configuration tools.
- You need a kdump kernel image and initial RAM disk.
- The kexec-tools package must be installed.

### About this task

The reserved memory area must be sufficiently large to accommodate the kdump system, including user space, when it is booted. If too little memory is reserved, kdump itself crashes if booted (see [“Failure recovery and backup tools” on page 9](#)).

### Procedure

1. For your production system, code the `crashkernel=` kernel parameter according to this syntax:

```
crashkernel=<size>@<offset>
```

where `<size>` specifies the amount of memory to be reserved and `<offset>` the beginning of the memory range. The values are integers, optionally followed by K for kilobyte, M for megabyte, or G for gigabyte. The values are adjusted to multiples of 1 MB.

The specified memory area, `<offset>` through `<offset> + <size>`, and a corresponding memory area 0 through `<size>`, must both be available and must both not contain any memory holes.

If you omit `@<offset>`, a suitable offset is chosen for you. If you specify an offset, it must be greater than the size of the reserved memory.

For example, the following specification reserves 128 MB for the kdump kernel at an automatically selected suitable offset.

```
crashkernel=128M
```

The following specification reserves 128 MB for the kdump kernel at an offset of 256 MB.

```
crashkernel=128M@256M
```

Optionally, you can make the specification of reserved memory size and offset dependent on the size of the available memory. See `Documentation/kernelparameters.txt` in the Linux source tree for details about how to do this and about further details of the `crashkernel=` kernel parameter.

2. Boot your production system.
3. Optional: Issue the following command to confirm that a memory area was reserved for the kdump kernel:

```
# cat /proc/iomem
```

The command output must include a memory range for `Crash kernel`.

4. Load the kdump kernel with the **kexec** command according to this syntax:

```
# kexec -a -p <image> --initrd <initrd> --command-line "<kparms>"
```

where:

**<image>**

specifies the kdump kernel image.

**<initrd>**

specifies the initial RAM disk of the kdump kernel. This specification can be omitted if the kdump kernel does not require an initial RAM disk.

**<kparms>**

specifies kernel parameters for the kdump kernel.

**-a**

specifies that kexec should use an advanced kexec load function, if available. If the advanced function is not available, kexec falls back to the default load function. This is the preferred method of calling kexec.

**-s**

is an optional parameter that uses an advanced kexec load function. This function enables the current kernel to check the validity of the kdump kernel before loading it.

#### Example:

```
# kexec -a /boot/kdump.image --initrd /boot/kdump.initrd \
--command-line="dasd=eb90 root=/dev/ram0 maxcpus=1"
```

### Results

A kernel panic or PSW restart now triggers an automatic dump process with kdump.

### 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 54](#) about how to run a backup tool automatically, if kdump fails.

## Initiating a dump

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 54](#)).

To trigger kdump, use one of the methods according to your environment:

- For Linux on z/VM: Run the z/VM CP **system restart** command. See [“z/VM guest example” on page 11](#) for an example.
- For Linux in LPAR mode: Run the **PSW restart** task on the HMC. See [“LPAR example” on page 11](#) for an example.

**Tip:** On both z/VM and LPAR you can use the diag288 watchdog to trigger kdump. The default setup triggers the correct actions. For more details about the diag288 watchdog, see *Device Drivers, Features, and Commands*, SC33-8411

## Results

After kdump is triggered, first kernel messages from the booting kdump kernel and later dump progress messages are issued. The messages are written to the Operating System Messages applet of the HMC for LPAR and to the 3270 terminal for z/VM.

## What to do next

Because the dump is initially held in memory, you must process the dump before IPLing your production system. A new IPL clears the dump. On distributions, kdump scripts copy the dump and reboot automatically.

## z/VM guest example

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

### Before you begin

Your Linux instance must have been 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, issue:

```
#cp system restart
```

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

## LPAR example

You can initiate a kdump 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.

### Procedure

1. In the 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 content area.
2. Select the LPAR for which you want to initiate the dump.
3. In the **Tasks** area, click **PSW restart**, which is located in the **Recovery** section. This initiates the dump process.  
Figure 3 on page 12 shows an example of an HMC with a selected mainframe system and LPAR. The **PSW restart** task can be seen in the **Tasks** area.

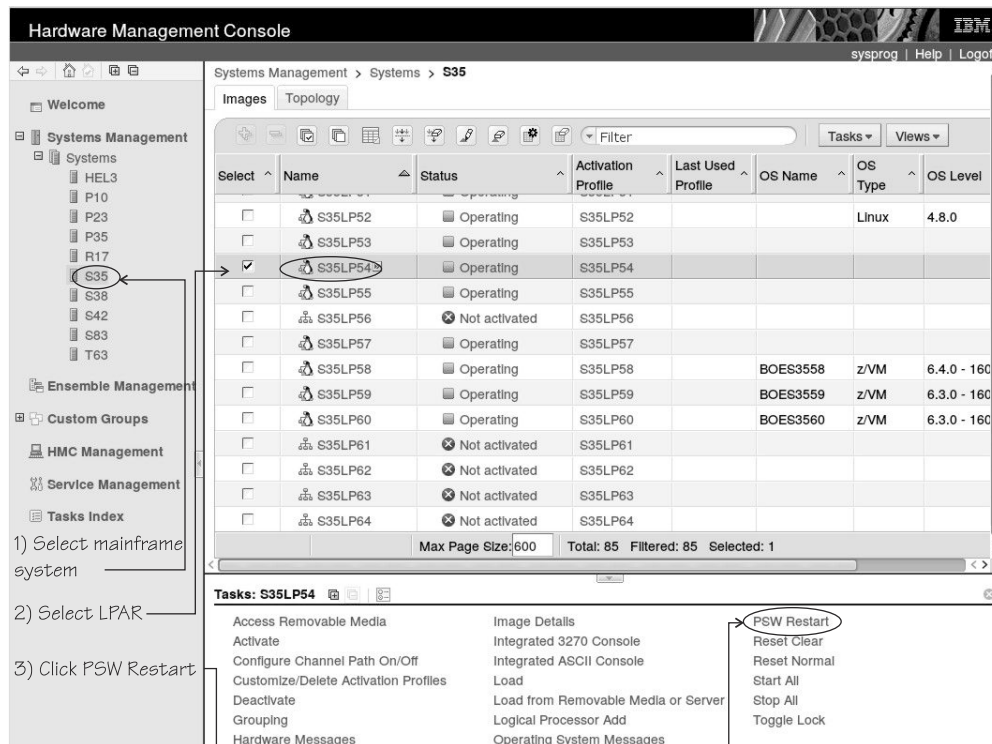


Figure 3. HMC with the **PSW restart** task

- Wait until the dump process completes. Click the **Operating System Messages** icon for progress and error information.

## Results

kdump automatically collects the dump and reboots Linux.

## Accessing the dump

After the kdump kernel has started, the dump can be accessed and copied from memory to a file on persistent storage. Typically, this process is automated through tools in an initial RAM disk that is provided with your distribution.

### Before you begin

The user space in the kdump distribution initial RAM disk is typically set up to perform the following actions for you:

- Reading the dump of the production system
- Filtering and compressing the dump
- Saving the dump to a file on persistent storage
- Rebooting the production system

The steps that follow describe how to perform these tasks from the command line.

### About this task

On the running kdump kernel, the dump can be accessed through two virtual files:

#### **/proc/vmcore**

represents the dump in Executable and Linkable Format (ELF) core format and includes memory, CPU register, and vmcoreinfo information.

## **/dev/oldmem**

represents the dump in form of an unstructured linear memory.

The following description uses the more convenient ELF format at `/proc/vmcore`.

### **Procedure**

1. Optional: Use **zgetdump** to display information about the dump.

```
# zgetdump -i /proc/vmcore
```

2. Optional: Analyze the dump at `/proc/vmcore` with the **crash** tool, as usual. This means that you cannot IPL your production system until you have finished the analysis.

3. Copy the dump to persistent storage.

The following examples illustrate some of the options:

This example copies the entire dump to a device that is available at `/dumps/dump.elf` in the Linux file system:

```
# cp /proc/vmcore /dumps/dump.elf
```

This example uses `scp` to copy the entire dump to a file `/dumps/dump.elf` in the file system of another system, `dumpstore`:

```
# scp /proc/vmcore user@dumpstore:/dumps/dump.elf
```

This example uses **makedumpfile** to copy a compressed and filtered version of the dump to a file `/dumps/dump.kdump` in the Linux file system:

```
# makedumpfile -c -d 31 /proc/vmcore /dumps/dump.kdump
```

This example uses **makedumpfile** to copy a compressed and filtered version of the dump to a file `/dumps/dump.kdump` in the file system of another system, `dumpstore`:

```
# makedumpfile -F -c -d 31 /proc/vmcore | \
ssh user@dumpstore "cat > /dumps/dump.kdump"
```

The **makedumpfile -c** option compresses the dump, the **-F** option converts it to the flat format required for transferring the file, and **-d** filters unwanted data from the dump. The number that follows **-d** must be in the range 0 through 31. The number represents a bit mask that specifies which page types to filter out.

For more details, see the man page for **makedumpfile**.

4. Issue `reboot`, to start the production system again.

### **Results**

You can now analyze the dump with **crash**.



---

## 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 is 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. 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 multivolume dump,” on page 21](#).

#### About this task

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

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

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

#### Procedure

1. (ECKD only) Format your DASD with **dasdfmt**.

Use a block size of 4 KB.

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 using the **zipl** command.

You can specify the dump device on the command line or use a configuration file.

Command line example:

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

Configuration file example:

a. Edit `/etc/zipl.conf` to add the following lines:

```
[dump_dasd]
dumpto=/dev/dasdc1
```

b. Issue:

```
# zipl dump_dasd
```

**Note:** When using an ECKD-type DASD formatted with the traditional Linux disk layout `ldl`, the dump tool must be reinstalled using **zipl** 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 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
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.

## z/VM guest example

In this example, a dump to DASD device 193 is initiated from z/VM.

### 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.

## LPAR examples

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

### HMC or SE web interface

In this example, a Linux instance running on an LPAR is dumped to a DASD device. The dump is initiated from an HMC or SE.

### About this task

The following description refers to an HMC, but the steps also apply to an SE.

### Procedure

1. In the 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 content area.
2. Select the LPAR for which you want to initiate the dump.
3. In the **Tasks** area, click **Stop all** (which is located in the **Recovery** section) to stop all CPUs. Confirm when you are prompted to do so.
4. In the **Tasks** area, click **Load** to display the Load panel.  
Figure 4 on page 17 shows an example of an HMC with a selected mainframe system and LPAR. The **Load** and **Stop all** tasks can be seen in the Tasks area.

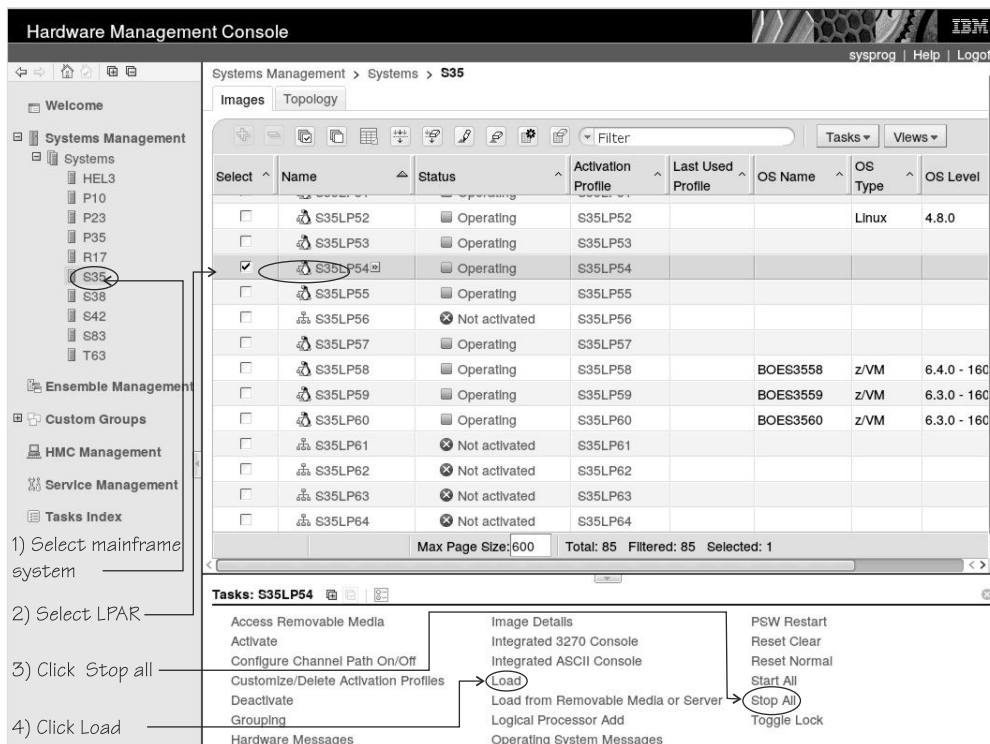


Figure 4. HMC with the **Load** and **Stop all** tasks

5. Select **Load type** "Normal".

6. Select the **Store status** check box.

7. Type the device number of the dump device into the **Load address** field.

To IPL from a subchannel set other than 0, specify five digits: The subchannel set ID followed by the device number, for example 1E711. Figure 5 on page 18 shows a Load panel with all entries and selections required to start the dump process for a DASD or tape dump device.

The figure shows a graphical user interface window titled "Load - S35:S35LP54". It contains several fields and controls for configuring a dump process. The "CPC:" field is set to "S35:S35LP54". The "Image:" field is also set to "S35:S35LP54". Under "Load type", there are four radio buttons: "Normal" (selected), "Clear", "SCSI", and "SCSI dump". The "Store status" checkbox is checked. The "Load address" field contains "+1E711". The "Load parameter" field is empty. The "Time-out value" is set to "60" with a dropdown arrow, and a note "60 to 600 seconds" is displayed. Below these are four input fields for "Worldwide port name", "Logical unit number", "Boot program selector", and "Boot record logical block address", all containing the value "0". A large text area for "Operating system specific load parameters" is at the bottom. At the very bottom are four buttons: "OK", "Reset", "Cancel", and "Help".

Figure 5. Load panel for dumping to DASD

8. Click **OK** to start the dump process.

9. 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.

## HMC or SE remote API using snipl

You can trigger HMC or SE activities remotely by using the **snipl** command.

### Before you begin

As of **snipl** version 2.1.9 the **snipl** command can be used for dump handling. The required setup for **snipl** usage and further details are described in *Device Drivers, Features, and Commands*, SC33-8411. You can dump to a DASD disk.

### About this task

For example, assume that you have a **snipl** configuration file `/etc/snipl.conf` containing the following specifications:

```
server=myse9.example.com
image=LPARLNX1
image=LPARLNX2

server=myse5.example.com
image=LPRLN05
```

Further assume that you have prepared a dump DASD (in this example with device number 5199) with the **zipl** tool.

### Procedure

Use the following **snipl** commands to write a memory dump of LPARLNX1 to the prepared DASD:

### 1. Stop the CPUs:

```
# snipl LPARLNX1 --stop
Server myse9.example.com from config file /etc/snipl.conf is used
processing.....
LPARLNX1: acknowledged.
```

### 2. IPL the dump tool on DASD 5199, prepared with the dump tool:

```
# snipl LPARLNX1 --load -A 5199 --storestatus
Server myse9.example.com from config file /etc/snipl.conf is used
processing.....
LPARLNX1: acknowledged.
```

To IPL from a subchannel set other than 0, specify five digits: The subchannel set ID followed by the device number, for example 15199.

### 3. Monitor the dump progress:

```
# snipl LPARLNX1 --dialog
LPARLNX1: acknowledged.
Starting operating system messages interaction for
partition LPARLNX1 (Ctrl-D to abort):
00000128 / 00001024 MB
...
00000896 / 00001024 MB
00001024 / 00001024 MB
Dump successful
```

## Automatic dump example

On both z/VM and LPAR, you can use the dumpconf service to set up automatic dumping. In this example, a dump is automatically triggered when a kernel panic occurs.

### About this task

Use the dumpconf service to set up automatic dumping. To set up dumping, edit the configuration file `/etc/sysconfig/dumpconf`.

### Example

Example configuration for a CCW dump device (DASD):

```
ON_PANIC=dump
DUMP_TYPE=ccw
DEVICE=0.0.4714
```

For details on how to set up dumpconf, see [“The dumpconf service” on page 54](#).

## Copying the dump from DASD with zgetdump

You can copy a DASD dump to a file system 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.elf`:

```
# zgetdump /dev/dasdc1 > dump.elf
```

### 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 52](#) for an example.

For general information about **zgetdump**, see [“The zgetdump tool” on page 49](#) or the man page.

## Chapter 4. Using DASD devices for multivolume dump

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

### Before you begin


You need to prepare a set of ECKD DASD devices for a multivolume dump, install the stand-alone dump tool on each DASD device involved, 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 formatted with the compatible disk layout (cdl, the default option when using the **dasdfmt** command.)

You must specify block size 4096 for **dasdfmt**.

For example, [Figure 6 on page 21](#) shows three DASD volumes, dasdb, dasdc, and dasdd, with four partitions selected to contain the dump. To earmark the partition for dump, a dump signature is written to each partition.



*Figure 6. 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 will take you through the entire process of creating a multivolume dump.

## Installing the multivolume 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 device bus-IDs (as displayed by **lsdasd**) are 0.0.4711 and 0.1.4712.

### Procedure

1. Format both dump volumes with **dasdfmt**.

The command shown uses the default cdl (compatible disk layout) and specifies a block size of 4KB.

```
# 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 `mvdump.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.

You can specify the dump list on the command line or use the zipl configuration file.

Command line example:

```
# zipl -M mvdump.conf
Dump target: 2 partitions with a total size of 100000 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)?
```

zipl configuration file example:

- a. Copy `mvdump.conf` to `/boot/mvdump.conf` and edit `/etc/zipl.conf` to add the following lines:

```
[multi_volume_dump]
mvdump=/boot/mvdump.conf
```

- b. Issue:

```
# zipl multi_volume_dump
```

### Results

Now the two volumes `/dev/dasdc` and `/dev/dasdd` with device bus-IDs 0.0.4711 and 0.1.4712 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 (extended)
  Version.....: 1
  Architecture.....: s390x (64 bit)
  Dump size limit...: none
  Force specified...: no

  Volume 0: 0.0.4711 (online/valid)
  Volume 1: 0.1.4712 (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 A, “Obtaining a dump with limited size,” on page 47. For more details about the **zipl** command, see *Device Drivers, Features, and Commands*, SC33-8411.

## Initiating a multivolume DASD dump

After preparing the DASD volumes, you can initiate a multivolume 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 main steps:

1. Stop all CPUs.
2. Store status on the IPL CPU.
3. IPL the dump tool using one of the prepared volumes, either 0.0.4711 or 0.1.4712.

**Note:** Do not clear storage!

For an example performed for a single DASD on the HMC, see “HMC or SE web interface” on page 17. For an example performed for a single DASD on z/VM, see “z/VM guest example” on page 16.

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: 0.0.4711
00000128 / 00001024 MB
00000256 / 00001024 MB
00000384 / 00001024 MB
00000512 / 00001024 MB
Dumping to: 0.1.4712
00000640 / 00001024 MB
00000768 / 00001024 MB
00000896 / 00001024 MB
```

```
00001024 / 00001024 MB
Dump successful
```

4. You can IPL Linux again.

## Copying a multivolume dump to a file

---

Use the **zgetdump** command to copy the multivolume dump.

### About this task

This example assumes that the two volumes `/dev/dasdc` and `/dev/dasdd` (with bus IDs 0.0.4711 and 0.1.4712) 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 > dump.elf
Format Info:
Source: s390mv_ext
Target: elf

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 only check 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 52](#) 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 you want to dump to.

#### Procedure

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

```
# zipl -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 should go 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
```

Messages might be shown on the tape display.

***number***

The number of MB dumped.

***dump\*end***

The dump process ended successfully.

5. You can IPL Linux again.

## z/VM guest example

You can initiate a dump to tape from z/VM.

### Example

If 193 is the tape device:

```
#cp rewind 193
#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.

## LPAR example

You can initiate a dump to tape 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.

### Procedure

1. In the 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 content area.

2. Select the LPAR for which you want to initiate the dump.
3. In the **Tasks** area, click **Stop all** to stop all CPUs. Confirm when you are prompted to do so.
4. In the **Tasks** area, click **Load** to display the Load panel.

Figure 7 on page 27 shows an example of an HMC with a selected mainframe system and LPAR. The **Load** and **Stop all** tasks can be seen in the **Tasks** area.

The screenshot shows the IBM Hardware Management Console (HMC) interface. On the left, the 'Systems Management' tree is expanded, showing a list of systems (HEL3, P10, P23, P35, R17, S35, S38, S42, S83, T63). The system 'S35' is selected. The main area displays a table of LPARs for system 'S35'. The table has columns: Select, Name, Status, Activation Profile, Last Used Profile, OS Name, OS Type, and OS Level. The LPAR 'S35LP54' is selected. Below the table, the 'Tasks' panel for 'S35LP54' is shown, containing various task options. Annotations with numbered arrows point to specific elements: 1) 'Select mainframe system' points to 'S35' in the left tree; 2) 'Select LPAR' points to 'S35LP54' in the table; 3) 'Click Stop all' points to the 'Stop All' button in the tasks panel; 4) 'Click Load' points to the 'Load' button in the tasks panel.

Select	Name	Status	Activation Profile	Last Used Profile	OS Name	OS Type	OS Level
<input type="checkbox"/>	S35LP52	Operating	S35LP52			Linux	4.8.0
<input type="checkbox"/>	S35LP53	Operating	S35LP53				
<input checked="" type="checkbox"/>	S35LP54	Operating	S35LP54				
<input type="checkbox"/>	S35LP55	Operating	S35LP55				
<input type="checkbox"/>	S35LP56	Not activated	S35LP56				
<input type="checkbox"/>	S35LP57	Operating	S35LP57				
<input type="checkbox"/>	S35LP58	Operating	S35LP58		BOES3558	z/VM	6.4.0 - 160
<input type="checkbox"/>	S35LP59	Operating	S35LP59		BOES3559	z/VM	6.3.0 - 160
<input type="checkbox"/>	S35LP60	Operating	S35LP60		BOES3560	z/VM	6.3.0 - 160
<input type="checkbox"/>	S35LP61	Not activated	S35LP61				
<input type="checkbox"/>	S35LP62	Not activated	S35LP62				
<input type="checkbox"/>	S35LP63	Not activated	S35LP63				
<input type="checkbox"/>	S35LP64	Not activated	S35LP64				

Tasks: S35LP54

- Access Removable Media
- Activate
- Configure Channel Path On/Off
- Customize/Delete Activation Profiles
- Deactivate
- Grouping
- Hardware Messages
- Image Details
- Integrated 3270 Console
- Integrated ASCII Console
- Load
- Load from Removable Media or Server
- Logical Processor Add
- Operating System Messages
- PSW Restart
- Reset Clear
- Reset Normal
- Start All
- Stop All
- Toggle Lock

Figure 7. HMC with the **Load** and **Stop all** tasks

5. Select **Load type** "Normal".
6. Select the **Store status** check box.
7. Type the device number of the dump device into the Load address field. To IPL from a subchannel set other than 0, specify five digits: The subchannel set ID followed by the device number, for example 1E711.

Figure 8 on page 28 shows a Load panel with all entries and selections required to start the dump process for a channel-attached tape dump device.

Figure 8. Load panel for dumping to tape

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

## Results

When the dump has completed successfully, you can IPL Linux again.

## Copying the dump from tape

You can copy a tape dump to a file system using the **zgetdump** tool.

### Before you begin

You must have installed the **mt** utility.

## 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 is the dump tool itself).

For example:

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

## Using the **zgetdump** tool to copy the dump

Use the **zgetdump** tools 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 28).

### 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.

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

### Procedure

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

```
# zgetdump /dev/ntibm0 > dump.elf
```

For general information on **zgetdump**, see [“The zgetdump tool”](#) on page 49 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 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 51.



---

## Chapter 6. Using a SCSI disk dump device

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

---

### Installing the SCSI disk dump tool

You install the SCSI disk standalone dump tool with the **zipl** command.

#### Before you begin

The dump partition needs enough free space (memory size + 10 MB) to hold the system memory.

The following examples assume that a SCSI dump disk is accessed through the device mapper multipath device node `/dev/mapper/36005076303ffd40100000000000020c0-part1` and is IPLed with the following parameters:

- devno: 0.0.4711
- wwpn: 0x4712076300ce93a7
- lun: 0x4712000000000000

#### Example

A partition on a SCSI disk is used as dump partition.

This example assumes that `/dev/mapper/36005076303ffd40100000000000020c0` is the dump device, and that you want to dump to the first partition, `/dev/mapper/36005076303ffd40100000000000020c0-part1`. Always use multipath devices instead of single path SCSI disk device nodes, if possible.

1. Create a partition with **fdisk** or **parted**, using the PC-BIOS or GPT layout.

For example:

```
# fdisk /dev/mapper/36005076303ffd40100000000000020c0
```

2. Install the dump tool by using the **zipl** command. You can specify the dump device on the command line or use a configuration file.

#### Command line example:

```
# zipl -d /dev/mapper/36005076303ffd40100000000000020c0-part1
```

#### Configuration file example:

- a. Edit `/etc/zipl.conf` to add the following lines:

```
[scsidump]
dumppto=/dev/mapper/36005076303ffd40100000000000020c0-part1
```

- b. Issue **zipl**:

```
# zipl scsidump
```

When you perform an IPL from `/dev/mapper/36005076303ffd40100000000000020c0` by using boot program selector 1 or 0 (default), the memory dump is written directly to partition 1 of `/dev/mapper/36005076303ffd40100000000000020c0`. The boot program selector is located on the load panel, see [Figure 10 on page 34](#) for an example.

## Initiating a SCSI dump

To initiate the dump, IPL the SCSI dump tool using 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.

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

```
Linux System Dumper starting

Version 3.0 (64 bit)
Linux version ...

Dump start at:..
Fri, 9 Dec 2016 13:22:55 +0000

Dump parameters:..
devno.....: 0.0.4711
wwpn.....: 0x4712076300ce93a7
lun.....: 0x4712000000000000
conf.....: 0
partition: /dev/sda1

Writing dump:
  0 of 6144 MB   0.0%   0 MB/s unknown ETA
2489 of 6144 MB 40.5% 277 MB/s 0:00:13 ETA
5114 of 6144 MB 83.2% 269 MB/s 0:00:03 ETA
6144 of 6144 MB 100.0% 267 MB/s 0:00:00 ETA

Dump successful.
```

### Results

The dump process copies the dump to the dump partition. When the dump completes successfully, you can IPL Linux again. You can then extract the dump from the dump partition into a file.

## z/VM guest example

You can initiate a dump to a SCSI disk using z/VM.

### About this task

Assume your SCSI dump disk has the following parameters:

- WWPN: 4712076300ce93a7
- LUN: 4712000000000000
- FCP adapter device number: 4711

### 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
#cp ipl 4711 dump
```



## What to do next

You can now IPL your Linux instance and resume operations.

## LPAR examples

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

### HMC or SE web interface

You can initiate a dump process on an LPAR from an HMC or SE web interface.

### About this task

The following description refers to an HMC, but the steps also apply to an SE.

### Procedure

1. In the 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 content area.

2. Select the LPAR for which you want to initiate the dump.

3. In the **Tasks** area, click **Load** to display the Load panel.

Figure 9 on page 33 shows an example of an HMC with a selected mainframe system and LPAR. The **Load** task can be seen in the **Tasks** area.

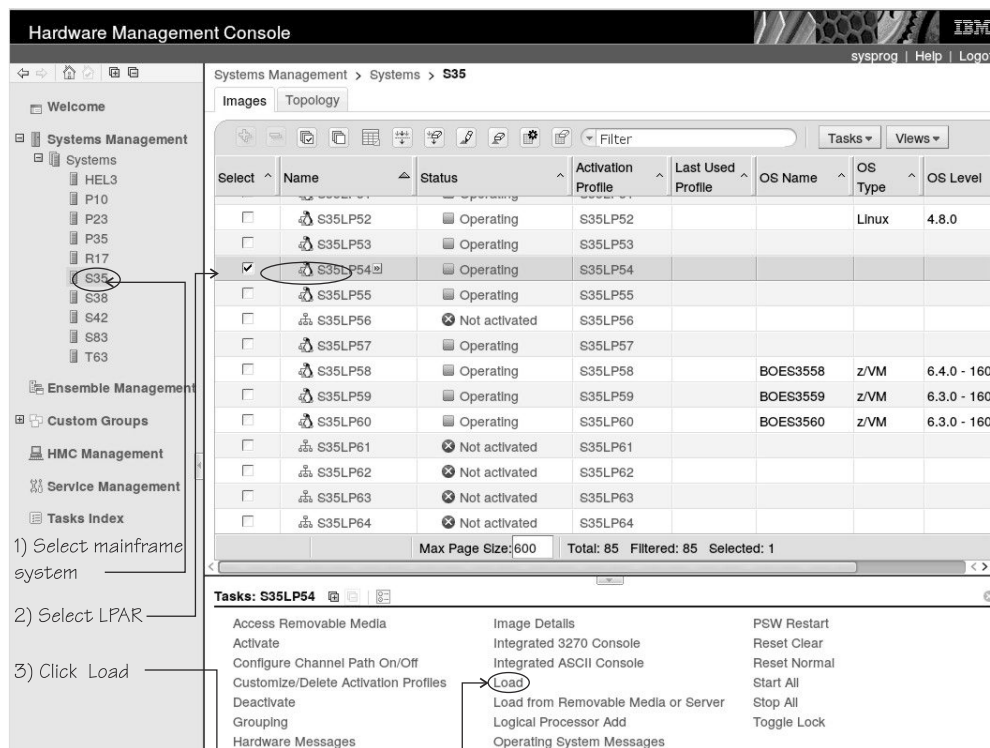


Figure 9. HMC with the **Load** task

4. Select **Load type** "SCSI dump".
5. Type the device number of the FCP adapter for the SCSI disk into the **Load address** field.
6. Type the World Wide Port name of the SCSI disk into the **World wide port name** field.
7. Type the Logical Unit Number of the SCSI disk into the **Logical unit number** field.
8. Type 0 in the Boot program selector field.
9. Accept the defaults for the remaining fields.

Figure 10 on page 34 shows a Load panel with all entries and selections required to start the SCSI dump process.

Figure 10. Load panel with enabled SCSI feature for dumping to SCSI disk

10. Click **OK** to start the dump process.
11. 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.

## HMC or SE remote API using snipl

You can trigger HMC or SE activities remotely by using the **snipl** command.

### About this task

As of snipl version 2.1.9 the **snipl** command can be used for dump handling. The required setup for **snipl** usage and further details are described in *Device Drivers, Features, and Commands*, SC33-8411. You can dump to a SCSI disk.

Assume you want to write a memory dump to a SCSI disk with:

- A WWPN of 4712076300ce93a7
- A LUN of 4712000000000000
- An FCP adapter with device number 4711

### Procedure

To write the dump, issue a **snipl** command like the following:

```
# snipl LPRLNX05 --scsidump -A 4711 --wwpn_scsiload 4712076300ce93a7 --lun_scsiload
4712000000000000
Server myse5.example.com from config file /etc/snipl.conf is used
processing...
LPRLNX05: acknowledged.
```

## Automatic dump example

On both z/VM and LPAR, you can use the dumpconf service to set up automatic dumping. In this example, a dump is automatically triggered when a kernel panic occurs.

### About this task

To set up dumping using dumpconf, edit the configuration file `/etc/sysconfig/dumpconf`.

## Example

Example configuration for an FCP dump device (SCSI disk), where the disk has device bus ID 0.0.4711, the WWPN is 0x4712076300ce93a7, and the LUN is 0x4712000000000000:

```
ON_PANIC=dump
DUMP_TYPE=fc
DEVICE=0.0.4711
WWPN=0x4712076300ce93a7
LUN=0x4712000000000000
BOOTPROG=0
BR_LBA=0
```

For details on how to set up dumpconf, see [“The dumpconf service” on page 54](#).

## Copying the dump from SCSI disks with zgetdump

You can copy a SCSI dump to a file system 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 SCSI disk `/dev/mapper/36005076303ffd40100000000000020c0-part1` and you want to copy it to a file named `dump.elf`:

```
# zgetdump /dev/mapper/36005076303ffd40100000000000020c0-part1 > dump.elf
```

### What to do next

You can use **zgetdump** to display information about the dump. See [“Checking whether a SCSI dump is valid and printing the dump header” on page 52](#) for an example. For general information about **zgetdump**, see [“The zgetdump tool” on page 49](#) or the man page.

## 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 /dev/mapper/36005076303ffd40100000000000020c0-part1
General dump info:
  Dump format.....: elf
  Version.....: 1
  UTS node name.....: mylnxsys
  UTS kernel release.: 3.3.0
  UTS kernel version.: #68 SMP PREEMPT Thu Mar 20 10:21:30 CET 2014
  System arch.....: s390x (64 bit)
  CPU count (online): 3
  Dump memory range.: 768 MB

Memory map:
  0000000000000000 - 000000002ffffff (768 MB)
```



---

## Chapter 7. Using VMDUMP

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

### About this task

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.

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.

---

### Initiating a dump with VMDUMP

Start the 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 T0 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*, SC24-6268.

### Example

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.

## 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, 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

- A Linux kernel image that was compiled with the common code kernel configuration option `CONFIG_STRICT_DEVMEM=n`.
- 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`.

#### Procedure

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

```
# zgetdump -i /dev/mem
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.: 512 MB

Memory map:
  0000000000000000 - 000000001fffffffff (512 MB)
```

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

```
# nice -n -20 zgetdump /dev/mem > 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.: 512 MB

Memory map:
  0000000000000000 - 000000001fffffffff (512 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/mem > dump.elf
Format Info:
Source: devmem
Target: elf
Copying dump:
00000000 / 00000512 MB
00000110 / 00000512 MB
...
00000512 / 00000512 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”](#) on page 40.

## 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 vmlinux vmlinux.debug
...
  KERNEL: /boot/vmlinux
  DUMPFILE: /mnt/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 vmlinux vmlinux.debug
...
NOTE: minimal mode commands: log, dis, rd, sym, eval, set and exit
```



---

## Chapter 9. Processing dumps

You can copy and transfer the dump file to another system, reduce the dump size, and send the reduced dump to IBM Support.

### Procedure

- To copy and transfer an LPAR dump on DASD, tape, or SCSI dump to another system, use the **zgetdump** command, see [“The zgetdump tool” on page 49](#).

With kdump, you can transmit the dump through a network. Use existing mechanisms to prevent conflicts when concurrently writing multiple dumps to a shared persistent storage space.

- To analyze the dump, you can use crash, see [“The crash tool” on page 56](#).
- To receive a VMDUMP file, use the **vmux** command, see [“The vmur tool” on page 57](#)

---

## Reducing dump size

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

### Before you begin

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

### About this task

Large 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 being sent to the service organization.

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

### Procedure

1. Choose a dump device.

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

- Single-volume DASD
- Multivolume DASD, up to 32 DASDs are possible.
- z/VM emulated FBA device that represents a real SCSI 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 disk. The SCSI disk size depends on your storage server setup. For SCSI dump partitions greater than 2 TB, you must use the GPT disk layout.
- 3592 channel-attached tape drive, cartridges with up to 10 TB capacity are available.

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. Assuming about 100 MB per second dump speed on DASDs or SCSI disks and a system with 50 GB memory, the dump takes 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” on page 44](#)
- [“Compressing a dump using gzip and split” on page 45](#)

4. Send the dump.

## Compressing a dump using makedumpfile

Use the **makedumpfile** tool (version 1.3.7 or higher) can be used 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. 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 the partition `/dev/mapper/36005076303ffd40100000000000020c0-part1` to `/mnt`, issue:

```
# zgetdump -m -f elf /dev/mapper/36005076303ffd40100000000000020c0-part1 /mnt
```

2. Create a file with a filtered and compressed version of the dump.

Use the **makedumpfile -d** (dump level) option to excludes pages that are typically not needed to analyze a kernel problem. For dump level 31, pages containing only zeroes, pages used to cache file contents (cache, cache private), pages belonging to user-space processes, and free pages are all excluded.

See the man page for **makedumpfile** for a description of the dump level and other options of **makedumpfile**.

The following command accesses a dump at `/mnt/dump.elf` filters it with dump level 31, compresses it, and writes it to a file `/dumps/dump.kdump`:

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

You might want to retain a copy of the original dump file until the problem is resolved. This reserves the option to create further copies at different dump levels should any of the excluded pages be required for problem determination.

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

```
# makedumpfile --dump-dmesg /mnt/dump.elf /dumps/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 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 /dumps/mydumps/dump.0 | gzip | split -b 1G
```

This will create 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
```

## Sending a dump to IBM Support

---

After compressing the dump, you can transfer it using FTP or an HTTPS upload server.

### **Before you begin**

You need to have opened a problem management record (PMR) before you send the data. You need the PMR number to correctly name the uploaded files.

### **Procedure**

See Enhanced Customer Data Repository at [www.ibm.com/support/pages/enhanced-customer-data-repository-ecurep-send-data.html](http://www.ibm.com/support/pages/enhanced-customer-data-repository-ecurep-send-data.html) for detailed instructions.

---

## Appendix A. Obtaining a dump with limited size

The `mem kernel` parameter can make Linux use less memory than is available to it. A dump of a Linux system like this 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

This section does not apply to `kdump`.

The **size** option is available for all **zipl** based dumps except SCSI: DASD and tape, 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
```

An equivalent section in a configuration file could look like this:

```
[dump1]
dumpto=/dev/dasdc1,100M
```





## Appendix B. Command summary

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

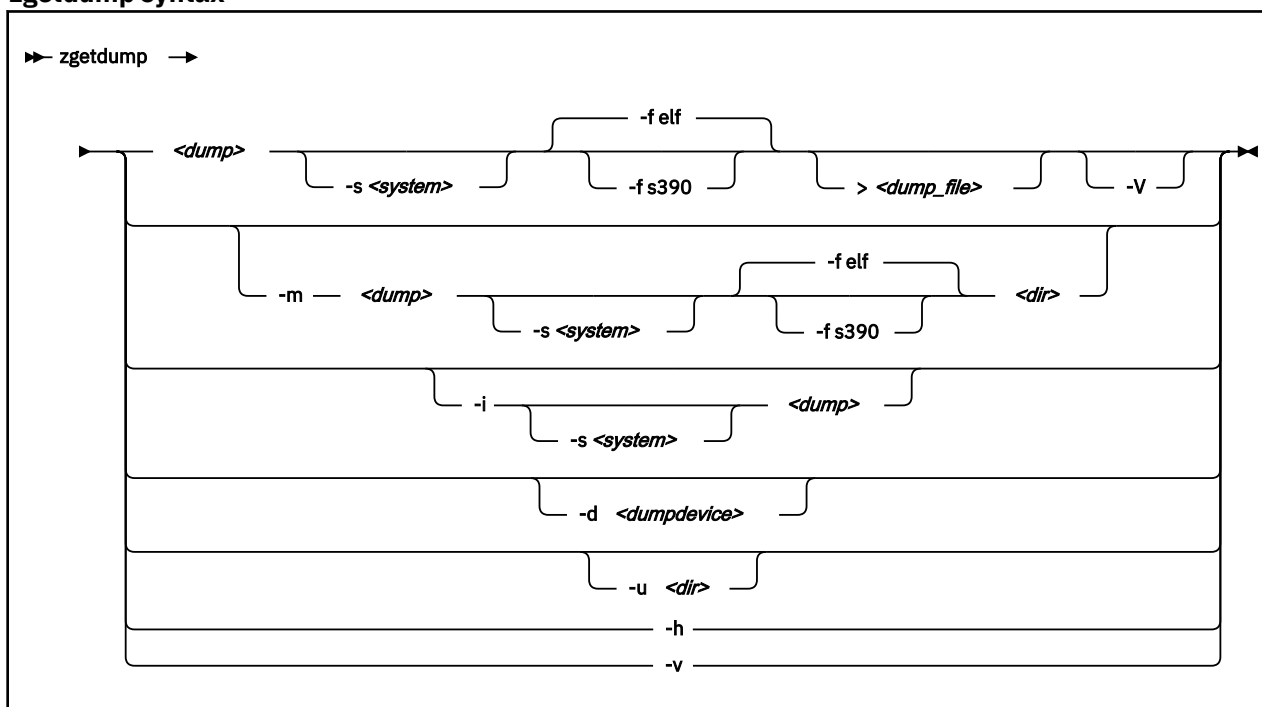
- [The zgetdump tool](#)
- [The dumpconf service](#)
- [The crash tool](#)
- [The vmconvert tool](#)
- [“The vmur tool” on page 57](#)

### The zgetdump tool

The **zgetdump** tool copies a source dump into a target dump with a configurable dump format. The source dump can be located either on a dump device or on a file system. The source 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, check whether a DASD device contains a valid dump tool, or create a non-disruptive kernel dump on a live system.

**Before you begin:** Mounting is implemented with fuse (file system in user space). Therefore the fuse kernel module must be loaded before you can use the **-m** option. Load the module, for example, with **modprobe fuse**.

#### zgetdump syntax



#### Parameters

##### <dump>

is the file, DASD device or partition, multipath partition of a SCSI disk, channel-attached tape device, or live system device node where the source 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)
- Device mapper multipath partition device node of a SCSI disk (for example /dev/mapper/36005076303ffd40100000000000020c0-part1)
- Tape device node (for example /dev/ntibm0)
- Device node for live system (/dev/mem)

**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, or a multipath device node of a SCSI disk, for example /dev/mapper/36005076303ffd40100000000000020c0.

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

for dumps that capture two systems, selects the system of interest. To check whether a dump contains two systems, use **zgetdump -i**. The **-s** 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.

**all**

to select the data for both prod and kdump.

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

Mounts the source dump *<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, FBA, or SCSI disk 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 (default)

**s390**

S/390® dump

**-V or --verbose**

Shows the detailed memory-map layout when printing the dump header information for s390\_ext and ELF dump formats.

**-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.elf:

```
# zgetdump /dev/dasdb1 > dump.elf
```

**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.elf"
```

**Using zgetdump to transfer a dump with FTP**

Assuming that you want to use FTP to transfer a dump to a file, dump.elf, on another system:

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 FTP host issue:

```
ftp> put |"zgetdump /dev/dasdb1" dump.elf
```

**Using zgetdump to copy a multivolume 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 dump.elf:

```
# zgetdump /dev/dasdc > dump.elf
```

For an example of the output from this command, see [Chapter 4, “Using DASD devices for multivolume dump,”](#) on page 21.

**Using zgetdump to copy a tape dump**

Assuming that the tape device is /dev/ntibm0:

```
# zgetdump /dev/ntibm0 > dump.elf
```

Format Info:

Source: s390tape

Target: elf

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/mem > dump.elf
```

For an example of the output from this command, see [“Creating a kernel dump on a live system”](#) on page 39.

## 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, 10 May 2010 17:26:46 +0200
  Dump ended.....: Mon, 10 May 2010 17:27:58 +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 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_ext
  Version.....: 1
  Dump created.....: Mon, 09 Jul 2018 20:00:52 +0200
  Dump ended.....: Mon, 09 Jul 2018 20:00:55 +0200
  Dump CPU ID.....: ff3fec8729648000
  UTS node name.....: s8360020
  UTS kernel release.: 4.16.0
  UTS kernel version.: #1 SMP Fri Jun 22 10:14:22 CEST 2018
  Build arch.....: s390x (64 bit)
  System arch.....: s390x (64 bit)
  CPU count (online): 2
  CPU count (real)...: 2
  Dump memory range..: 2048 MB
  Real memory range..: 2048 MB
  Dump file size.....: 493 MB
```

```
Memory map:
0000000000000000 - 000000007fffffff (2048 MB)
```

## Checking whether a SCSI dump is valid and printing the dump header

Assuming that the dump is on the first partition of a SCSI disk, for example `/dev/mapper/36005076303ffd40100000000000020c0-part1`:

```
# zgetdump -i /dev/mapper/36005076303ffd40100000000000020c0-part1
General dump info:
  Dump format.....: elf
  Version.....: 1
  UTS node name.....: r3545010
  UTS kernel release.: 3.12.25-2-default
  UTS kernel version.: #1 SMP Mon Jul 28 12:18:48 UTC 2014
  System arch.....: s390x (64 bit)
  CPU count (online).: 3
  Dump memory range..: 1024 MB

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

## Checking whether a DASD device contains a valid dump record

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

```
# zgetdump -d /dev/dasdb
Dump device info:
  Dump tool.....: Single-volume DASD dump tool (extended)
  Version.....: 1
  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"
```

## Checking whether a SCSI disk contains a valid dump record

Checking SCSI multipath device /dev/mapper/36005076303ffd40100000000000020c0, which is a valid dump device:

```
# zgetdump -d /dev/mapper/36005076303ffd40100000000000020c0
Dump device info:
  Dump tool.....: Single-volume SCSI dump tool
  Version.....: 1
  Architecture.....: s390x (64 bit)

Partition info:
  Partition number...: 1
  Maximum dump size..: 20473 MB
```

Checking SCSI multipath device /dev/mapper/36005076307ffc5e300000000000084cf, which is not a valid dump device:

```
# zgetdump -d /dev/mapper/36005076307ffc5e300000000000084cf
zgetdump: No dump tool found on "/dev/mapper/36005076307ffc5e300000000000084cf"
```

## 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 source DASD dump as an ELF dump, processes it with **crash**, and unmounts it with **zgetdump**:

```
# zgetdump -m /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. Use the **--fmt** option to mount the dump with a format other than the default format.

## 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
```

## Checking whether a dump has captured two systems

A dump can contain data from two systems. To check for this use **zgetdump -i**, for example, assuming that the previous dump example contains both a dump from the production system and a kdump kernel dump:

```
# zgetdump -i /dev/mapper/36005076303ffd40100000000000020c0-part1
zgetdump: The dump contains "kdump" and "production system"
Access "production system" with "-s prod"
Access "kdump" with "-s kdump"
Access the complete dump with "-s all"
Send both dumps to your service organization
```

## The dumpconf service

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

The service is installed as systemd unit 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.

### dumpconf systemd service syntax



### Parameters

#### start

enables the configuration that is defined in `/etc/sysconfig/dumpconf`.

#### stop

stops the **dumpconf** service.

#### status

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

### Keywords for the configuration file

#### ON\_PANIC

Shutdown action to be taken if a kernel panic or PSW restart occur. 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 **dumpconf** 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, **dumpconf** 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 eight. You can specify up to eight CP commands that are executed in case of a kernel panic or PSW restart. 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 `fcpx`.

**DEVICE**

Bus ID of the dump device.

**WWPN**

WWPN for SCSI disk.

**LUN**

LUN for SCSI disk.

**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=fcpx
DEVICE=0.0.4711
WWPN=0x4712076300ce93a7
LUN=0x4712000000000000
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"
```

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:

```
# systemctl start dumpconf
```

- To display the status:

```
# systemctl status dumpconf
| dumpconf.service - Configure dump on panic
  Loaded: loaded (/usr/lib/systemd/system/dumpconf.service...)
  Active: active (exited) since Mon 2017-11-06 20:47:15 CET; 1s ago
  Process: 12200 ExecStart=/lib/s390-tools/dumpconf start (code=exited, status=0/SUCCESS)
 Main PID: 12200 (code=exited, status=0/SUCCESS)

Nov 06 20:47:15 s3545002 systemd[1]: Starting Configure dump on panic ...
Nov 06 20:47:15 s3545002 dumpconf[12200]: dump_reipl on panic configured: Using ccw...
Nov 06 20:47:15 s3545002 systemd[1]: Started Configure dump on panic ...
```

- To disable dump on panic:

```
# systemctl stop dumpconf
```

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

```
# systemctl status dumpconf
| dumpconf.service - Configure dump on panic
  Loaded: loaded (/usr/lib/systemd/system/dumpconf.service...)
  Active: inactive (dead)
  ...
Nov 06 20:49:33 s3545002 systemd[1]: Stopping Configure dump on panic ...
Nov 06 20:49:33 s3545002 dumpconf[12217]: Dump on panic is disabled now
Nov 06 20:49:33 s3545002 systemd[1]: Stopped Configure dump on panic ...
```

## The crash tool

The **crash** tool is a GPL-licensed tool maintained by Red Hat.

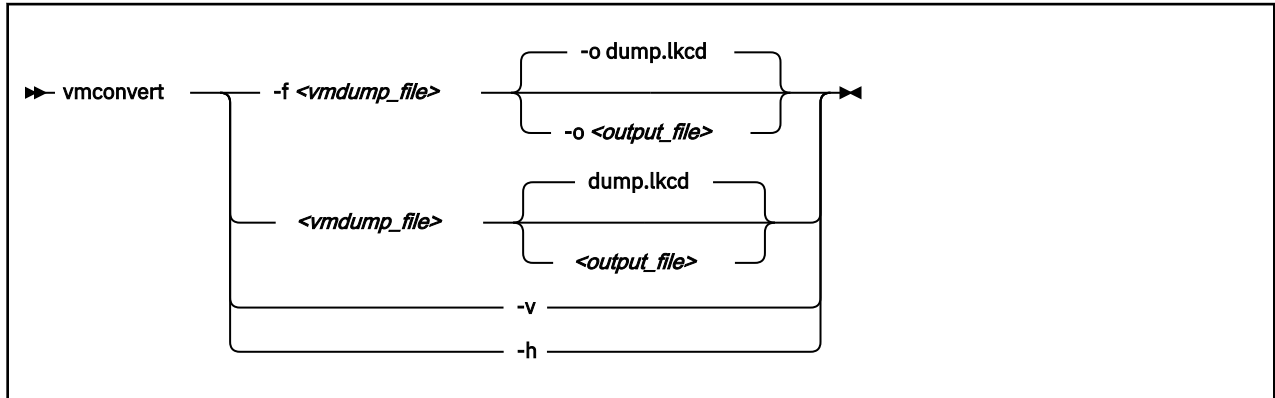
As of crash 7.2.6, you can analyze dumps of kernels with enabled KASLR support. Use crash with the **--kaslr auto** option. KASLR requires that the dump contains vmcoreinfo, which is always included with kdump. For all other dump types, such as VMDUMP, stand-alone dumps, and qemu dumps, convert the dump to an ELF dump before using crash. To convert the dump, use the command **zgetdump -f elf**.



## 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.lkcd.

**-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 data of the reader spool file.

For more details, see the **vmur** man page and *Device Drivers, Features, and Commands*, SC33-8411.

## Example

The vmur commands access the reader device, which has to be online. To set it online, it needs to be freed from cio\_ignore. Example:

```
# chzdev -ae c
Configuring devices in the active configuration only
Generic CCW device 0.0.000c configured
```

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
```

# 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 Z and LinuxONE 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 send an email to [eservdoc@de.ibm.com](mailto: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](http://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](http://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](http://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.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

---

# Index

## A

accessibility [59](#)

## C

commands

summary [49](#)

zgetdump [49](#)

compress memory dump [45](#)

compressing a memory dump [44](#)

crash

live dump, opening [40](#)

crash tool [56](#)

## D

DASD

using as dump device [15](#)

DASD devices

using for multivolume dump [21](#)

DASD dump

initiating [16](#)

DASD dump tool

installing [15](#)

DASD multivolume dump

starting [23](#)

developerWorks [2](#)

device nodes [viii](#)

dump

copy from DASD with zgetdump command [19](#)

copy from SCSI with zgetdump command [35](#)

copy multivolume dump from DASD with zgetdump command [24](#)

copy tape from [28](#)

copy with zgetdump command [29](#)

limited size [47](#)

live-system [3](#)

multivolume [21](#)

starting a multivolume DASD [23](#)

tape, checking if valid [29](#)

tape, initializing [25](#)

dump device

DASD [15](#)

definition [2](#)

tape [25](#)

dump devices

SCSI [31](#)

dump methods

comparison [4](#)

dump tool

crash [56](#)

DASD, installing [15](#)

tape, installing [25](#)

dump tools

dumpconf service [54](#)

multivolume, DASD [22](#)

dump tools (*continued*)

stand-alone [2](#)

stand-alone tape [25](#)

summary [49](#)

version dependencies [5](#)

vmconvert [57](#)

VMDUMP [2](#), [37](#)

vmur command [57](#)

zgetdump [49](#)

dump tools overview [1](#)

dumpconf service [54](#)

## E

ECuRep [46](#)

## G

gzip command [45](#)

## H

handling large dumps [43](#)

## I

initiate dumps

HMC or SE [11](#), [17](#), [19](#), [26](#), [33](#), [34](#)

using a DASD device [16](#)

using tape [26](#)

using VMDUMP [37](#)

initiating a dump [37](#)

initiating dump

using SCSI [32](#)

## K

kdump

accessing the dump [12](#)

advantages and disadvantages [7](#)

comparison with other dump methods [4](#)

initiate [10](#)

initiating [11](#)

introduction [2](#)

on IBM Z [7](#)

kernel dump

creating from live system [39](#)

## L

large dump

handling [43](#)

multivolume [21](#)

limit amount of memory dumped [47](#)

live dump

open with crash [40](#)

live system  
    creating dump [39](#)  
live-system dump [3](#)

## M

makedumpfile [44](#)  
memory dump  
    compressing [44](#)  
messages  
    tape display [25](#)  
multivolume dump  
    starting [23](#)  
multivolume dumps  
    DASD tool [22](#)

## N

non-disruptive dump  
    using zgetdump [39](#)

## P

PMR [46](#)

## R

remotely triggering a dump to DASD [18](#)  
remotely triggering a dump to SCSI [34](#)

## S

SCSI dump  
    initiating [32](#)  
    printing the dump header [35](#)  
    single partition [31](#)  
SCSI dump device [31](#)  
SCSI dump tool  
    installing [31](#)  
sending a dump to IBM [46](#)  
single partition  
    used for SCSI dump [31](#)  
split command [45](#)  
standard device nodes [viii](#)  
summary  
    commands for dumps [49](#)  
system restart  
    z/VM CP command [11](#)

## T

tape  
    copy dump from [28](#)  
    display messages [25](#)  
    use for dumping, preparing [28](#)  
    using as dump device [25](#)  
tape dump  
    checking if valid [29](#)  
    initializing [25](#)  
tape dump tool  
    installing [25](#)  
tools for creating dumps [1](#)

trigger a dump to DASD remotely [18](#)  
trigger a dump to SCSI remotely [34](#)

## U

using kdump [10](#), [11](#)

## V

vmconvert [57](#)  
VMDUMP  
    comparison with other dump methods [4](#)  
    copying dump [38](#)  
    initiate dump process [37](#)  
    introduction [2](#)  
vmur command  
    use to copy VMDUMP dump [38](#)

## Z

z/VM CP command  
    system restart [11](#)  
zgetdump  
    comparison with other dump methods [4](#)  
    create a dump from a live system [39](#)  
zgetdump tool [19](#), [24](#), [29](#), [35](#)  
zipl  
    size option [47](#)







SC33-8412-19

