RHEL: Crash kernel dumps configuration and analysis on RHEL 6

Article Number: 122 | Rating: Unrated | Last Updated: Sat, Jun 2, 2018 8:37 AM

RHEL: Crash kernel dumps configuration and analysis on RHEL 6

```
Kernel dumps may provide invaluable insights when debugging serious
issues.

1.- Install following RPMs:

         kexec-tools
         crash

    To be able to analyze crash dumps, following packages should be
also installed:

         kernel-debuginfo-common
         kernel-debuginfo


2.- Append "crashkernel=128M@16M" to the kernel parameters in
/boot/grub/grub.conf
    (see next points to adjust the size of crashkernel):

    kernel /vmlinuz-2.6.18-238.el5 ro root=/dev/rvg/rootlv nodmraid
rhgb quiet crashkernel=128M@16M

3.- Reboot the system
```

```
4.- Check that the crash kernel has been loaded:

         # cat /proc/iomem | grep Crash kernel
         01000000-08ffffff : Crash kernel


5.- Configure kdump to dump to:

   - Either locally; add following lines to /etc/kdump.conf:

         path /var/crash
         core_collector makedumpfile -d 31 -c


*** Note: This config can be done also by running:

# system-config-kdump

Please check the option box "Enable kdump" at the top of the Dialog.

Next, you have to define the memory to reserve for Kdump In the
dialog you see the memory information
for your system and the usable memory for Kdump. On most systems a
value of "128MB" Kdump memory
should be enough.

Finally, you need to define a location where to store the dump file.
You have the choice between
'file', 'nfs', 'ssh', 'raw', 'ext2', and 'ext3'. This setup is
straight forward, please configure
the kdump as it fit’s best into your environment. The simplest
configuration for the location is
"file:///var/crash".


   - Or to a remote server; add following lines to /etc/kdump.conf:
         net root@<kdump_remote_server>
         core_collector makedumpfile -d 31 -c


6.- Propagate SSH keys so that the vmcore could be sent via scp
without the need to enter any password
    (Take this point into account only if dumping to a remote
server):

         # service kdump propagate


7.- Configure kdump to start automatically

         # chkconfig kdump on
         # service kdump start


8.- Sync all filesystems:

         # sync


9.- Provoke a kernel panic with:

         # echo "1" > /proc/sys/kernel/sysrq
         # echo "c" > /proc/sysrq-trigger


10.- Now the crash kernel should get booted and on the remote system
a vmcore should get created
     under /var/crash:

         # tree /var/crash
         /var/crash
         |-- 192.168.12.227-2010-01-21-20:16:16
         `-- vmcore.flat
```

```
11.- The vmcore.flat needs to be processed in order to analyze the
core dump via the crash utility:

         # cat "vmcore.flat" | makedumpfile -R "/tmp/vmcore"
         The dumpfile is saved to /tmp/vmcore.
         makedumpfile Completed.



12.- Now you may analyze the vmcore with the crash utility:

         # crash
/usr/lib/debug/lib/modules/2.6.18-128.1.10.el5/vmlinux /tmp/vmcore

         crash 4.0-8.9.1.el5
         Copyright (C) 2002, 2003, 2004, 2005, 2006, 2007, 2008,
2009  Red Hat, Inc.
         Copyright (C) 2004, 2005, 2006  IBM Corporation
         Copyright (C) 1999-2006  Hewlett-Packard Co
         Copyright (C) 2005, 2006  Fujitsu Limited
         Copyright (C) 2006, 2007  VA Linux Systems Japan K.K.
         Copyright (C) 2005  NEC Corporation
         Copyright (C) 1999, 2002, 2007  Silicon Graphics, Inc.
         Copyright (C) 1999, 2000, 2001, 2002  Mission Critical
Linux, Inc.
         This program is free software, covered by the GNU General
Public License,
         and you are welcome to change it and/or distribute copies of
it under
         certain conditions.  Enter "help copying" to see the
conditions.
         This program has absolutely no warranty.  Enter "help
warranty" for details.

         GNU gdb 6.1
         Copyright 2004 Free Software Foundation, Inc.
         GDB is free software, covered by the GNU General Public
License, and you are
```

```
         welcome to change it and/or distribute copies of it under
certain conditions.
         Type "show copying" to see the conditions.
         There is absolutely no warranty for GDB.  Type "show
warranty" for details.
         This GDB was configured as "x86_64-unknown-linux-gnu"...

          KERNEL:
/usr/lib/debug/lib/modules/2.6.18-128.1.10.el5/vmlinux
          DUMPFILE: /tmp/vmcore  [PARTIAL DUMP]
          CPUS: 1
          DATE: Thu Jan 21 20:21:20 2010
          UPTIME: 00:03:10
          LOAD AVERAGE: 1.09, 0.46, 0.17
          TASKS: 445
          NODENAME: vrhel03
          RELEASE: 2.6.18-128.1.10.el5
          VERSION: #1 SMP Wed Apr 29 13:53:08 EDT 2009
          MACHINE: x86_64  (2666 Mhz)
          MEMORY: 1 GB
          PANIC: "SysRq : Trigger a crashdump"
          PID: 7835
          COMMAND: "bash"
          TASK: ffff81040699d0c0  [THREAD_INFO: ffff8103fed24000]
          CPU: 1
          STATE: TASK_RUNNING (SYSRQ)

         crash>



------------------------------------------------------------------
------
    The kdump procedure
------------------------------------------------------------------
------

1.- The normal kernel is booted with crashkernel=... as a kernel
option, reserving some memory for
```

```
    the kdump kernel. The memory reserved by the crashkernel
parameter is not available to the
    normal kernel during regular operation.  It is reserved for later
use by the kdump kernel.
2.- The system panics.
3.- The kdump kernel is booted using kexec, it uses the memory area
that was reserved via the
    crashkernel parameter.
4.- The normal kernel's memory is captured into a vmcore.

*** Note: Not reserving enough memory for the kdump kernel can lead
to the kdump operation failing.

*** Warning: Unless the system has enough memory, the Kernel Dump
Configuration utility will not
             start and you will be presented with an error message.



------------------------------------------------------------------------
------
    Configuring crashkernel on RHEL6.0 and RHEL6.1 kernels
------------------------------------------------------------------------
------

Some mappings of ram and appropriate crashkernel values:

ram size        crashkernel parameter        ram / crashkernel factor
>0GB            128MB                        15
>2GB            256MB                        23
>6GB            512MB                        15
>8GB            768MB                        31

The last column contains a ram/crashkernel factor.

The table is covered by the following crashkernel configuration:

    crashkernel=0M-2G:128M,2G-6G:256M,6G-8G:512M,8G-:768M
```

```
For servers with more RAM it is recommended to compute the
crashkernel parameter using the factors
that have been observed so far: 15 to stay on a safe side (maybe
wasting memory), using a factor
of 20 should also work. Please also note that the maximum size of RAM
that should be reserved here is 896M.



---------------------------------------------------------------------
------
    Configuring crashkernel on RHEL6.2 (and later) kernels
---------------------------------------------------------------------
------

Starting with RHEL6.2 kernels crashkernel=auto should be used. The
kernel will automatically reserve
an appropriate amount of memory for the kdump kernel.

Additionally some improvements have been made in the RHEL6.2 kernel
which have reduced the overall
memory requirements of kdump.

The amount of memory reserved for the kdump kernel can be estimated
with the following scheme:

    base memory to be reserved = 128MB
    an additional 64MB added for each TB of physical RAM present in
the system. So
    for example if a system has 1TB of memory 192MB (128MB + 64MB)
will be reserved.

*** Note: It is recommended to verify that kdump is working on all
systems after installation of
          all applications. The memory reserved by crashkernel=auto
takes only typical RHEL
          configurations into account. If 3rd party modules are used
more memory might have to be
          reserved. Thus, if a testdump fails it is a good strategy
```

```
to verify if it works with
          crashkernel=768M@0M and if it does do further debugging of
the memory requirements
          using the debug_mem_level option in /etc/kdump.conf.

*** Note: crashkernel=auto will only reserve memory on systems with
4GB or more physical memory.
          If the system has less than 4GB of memory the memory must
be reserved in explicitly
          configuring crashkernel=128M. Since RHEL6.3GA
(kernel-2.6.32-279.el6) this limit has
          been lowered to 2GB.

*** Warning: You need to take care that you have enough disk space on
the configured location.
```

Posted - Sat, Jun 2, 2018 8:37 AM. This article has been viewed 4927 times.

Online URL: http://kb.ictbanking.net/article.php?id=122