

Relax and Recover:

“The OpenSource [mksysb] for Linux on Power.”

Sébastien Chabrolles

s.chabrolles@fr.ibm.com

Power System Linux Center

IBM Montpellier Client Center

May 17th, 2017



When adversity strikes ...




“If anything can go wrong... it will.” (Murphy’s law)




**KEEP
CALM
RELAX
AND
RECOVER**

My Backup Server was there !



Backup Strategy and Disaster Recovery

“What should I do if a disaster strikes”

- **Backing up your server are necessary ... but not enough**

- Backups must be externalized.
- Backup strategy must be defined (what to backup, how, when ...).
- Disaster Recovery plan must be defined.

- **What is Disaster Recovery?**

The process by which a business function is restored to the normal, steady state after a disaster.

- **What is Business Continuity?**

The way that a business function will operate after a disaster, until such time as the normal, steady state is restored.



Disaster Recovery Plan (DRP)

- DRP addresses need to recover from an emergency with minimum impact to the enterprise.
 - Protects enterprise from major services failure
 - Minimizes risk to enterprise from delays in providing services
 - Guarantees reliability of standby systems by testing and simulation
 - Minimizes personnel decision-making required during disaster recovery
- Backups of data are necessary! but not enough in case of losing the complete Operating System.
 - Reinstalling the OS from scratch takes hours
 - Restoring the backups a few more hours
 - Fine-tuning of configurations takes days
 - Even months later issues pop up!
- It is absolute necessary to foresee an inventory of hardware and software and use **tool to quickly rebuild completely a system** :
 - **From scratch**
 - **OS + configuration + data** (*as it was before*)
 - **Mostly Automated.** (*speedup + avoid human error*)



UNIX System Recovery Tools



Flash Archive (up to Solaris 10)



HP Ignite-UX



IBM mksysb

*OS recovery tools are
part of most of the
UNIX since 90's !*

What is the mksysb command?



```
#man mksysb
mksysb Command
Purpose:
Creates an installable image of the root volume group either
in a file or onto a bootable tape.
```

Source: AIX mksysb man page

Translation:

Provides System Administrators the assurance of a BMR (*Bare-Metal Recovery*) solution, a provisioning tool, or way to roll back a system after a bad update.

=> Keeps SysAdmins from losing their jobs when all hell breaks loose.

1. Bootable : *(hd5 / spot)*
2. LVM + FS Layout : *image.data*
3. System Backup : *result of AIX **backup** command*

- Stored in bootable tape or NIM server images.
- Could be restored on dissimilar POWER HW.

Linux System Recovery Tools ?



*That's my
Dad!*



*“Only wimps use tape backup:
REAL men just upload their
important stuff on ftp and let the
rest of the world mirror it.”*

– Linus Torvalds –



Linux System Recovery Tools ?

Go to the community ...

Image Your Hard Drive using dd

Submitted by [sandip](#) on Fri, 02/11/2005 – 21:15

- Boot from the live cdrom distribution.
- Switch to root.
- Make sure **NO partitions** are mounted from the source hard drive.

1- Mount the external HD.

```
# mount -t vfat /dev/sda1 /mnt/sda1
```

2- Back up the drive.

```
# dd if=/dev/hda conv=sync,noerror bs=64K | gzip -c >  
/mnt/sda1/hda.img.gz
```

3- To restore your system:



```
# gunzip -c /mnt/sda1/hda.img.gz | dd of=/dev/hda  
conv=sync,noerror bs=64K
```

Limitations:

- Full Backup Only
- Full restore is the only option
- Backup cannot be made “Live”.
Need to reboot the server on DVD rescue to backup the server.
- Restore on the same Hardware.

Why Linux doesn't provide a tool like "mksysb" by default ?

... simply because "Linux is not UNIX" ;-)



System Architecture

IBM (Power)

RHEL 6: Red Hat has certified **939** different servers from **42** vendors

(x86), (x86_64), (Itanium II), (Power), (SPARC), (ARM), (System z)



Disk Partitioning

Logical Volumes

Logical Volumes, Partitions, Software RAID, Raw Disk, **Any block device!**



Filesystem Type

JFS, JFS2

Ext2,3,4, btrfs, Reiserfs, VFAT, XFS, JFS, OCFS2



Bootloader

bosboot

LILO, GRUB, GRUB2, ELILO, yaboot, lilo (SUSEppc), EXTLINUX ...

As with many Linux products...options abound

- Some “Linux compatible” Backup tools :

- SEP Sesam
- Relax and Recover*
- Storix*
- Cristie*
- IBM TSM
- HP DataProtector
- Symantec Netbackup
- Duplicity
- CommVault Galaxy
- EMC Networker
- Bacula
- Bareos
- Rsync
- GNU tar

* Close to mksysb functionalities



Relax and Recover – (*aka ReaR*)

- **OpenSource Disaster Recovery Solution for Linux**

<http://relax-and-recover.org/>

<https://github.com/rear/rear>

GPL 3 Software – Developers in Germany and Belgium

- 100% Bash script – no GUI and no dependencies
- Utilize kernel, modules, binaries of host (kernel ≥ 2.6)
- Support any combination of SW/HW RAID, LVM
- Internal backup on CIFS, NFS ...
- Boot media on CD/DVD, USB key and LAN (PXE)

- Successor of mkCDrec



ReaR – Main Principle

- **Which backup mechanism to use?**
 - Internal : GNU tar, rsync
 - External : Bacula, Commercial backup program
- **Where will the backups reside?**
 - external USB disk, tape, local spare disk
 - Remote network location:
 - NFS share
 - CIFS share
 - sshfs (*via fuse*)
- **How shall we start the rescue image?**
 - CDROM (ISO image)
 - tape (OBDR)
 - USB disk
 - Network (PXE)



Installing Relax and Recover

- With your preferred package manager.

- Rear is included in several Linux official repo's (Fedora, EPEL and SLES)

(It is now shipped in RedHat and Sles HA extension.)

- `yum install rear`
- `zypper install rear`
- `apt-get install rear`

- But not always the latest stable version available.

- Directly from the source repository.

- Github : <https://github.com/rear/rear>



```
# git clone https://github.com/rear/rear
# cd rear
(# git checkout rear-2.00)
# make
# make install
```

- Possibility to generate packages :

- rpm package : `make rpm`
- debian package : `make deb`

More information <https://github.com/rear/rear/blob/master/README.adoc>



This repository Search Pull requests Issues Gist

rear / rear Watch 47 Unstar 170 Fork 97

Code Issues 50 Pull requests 1 Projects 0 Wiki Pulse Graphs

Relax-and-Recover - Linux bare metal disaster recovery and system migration solution (cfr. mksysb, ignite) <http://relax-and-recover.org/>

2,186 commits 2 branches 38 releases 60 contributors GPL-3.0

Branch: master New pull request Create new file Upload files Find file Clone or download

gdha committed on GitHub Merge pull request #1034 from rear/dev Latest commit 28199db 5 days ago

.github	Update CONTRIBUTING.md	2 months ago
doc	adding the final release notes (we hope at least)	7 days ago
etc	Removed site.conf per request by gdha	a year ago
packaging	Tagging release 1.19 into packaging files	7 days ago
usr	Tagging release 1.19 into packaging files	7 days ago
.gitignore	Support having a var dir somewhere deep down in our tree. Fixes #792	7 months ago
.travis.yml	Ensure new commits are tested at GitHub/TravisCI	10 months ago
AUTHORS	Rebrand Rear as Relax-and-Recover where possible	4 years ago
COPYING	Changing GPLv2 to GPLv3.	10 months ago
Makefile	fixed a few more suse spelling typos	28 days ago
README.adoc	Corrected misspelled forms to Relax-and-Recover (issue1008).	25 days ago

ReaR – [command] usage

- Rear is very easy to use : `rear [options] <command>`

```
Available options:
-h --help          usage information
-c DIR            alternative config directory; instead of /etc/rear
-d              debug mode; log debug messages
-D             debugscript mode; log every function call (via 'set -x')
--debugscripts SET same as -d -v -D but debugscript mode with 'set -SET'
-r KERNEL       kernel version to use; current: '3.10.0-327.22.2.el7.ppc64le'
-s            simulation mode; show what scripts rear would include
-S          step-by-step mode; acknowledge each script individually
-v          verbose mode; show more output
-V --version  version information
```



- After initial configuration, a simple “`rear mkbackup`” will be sufficient to backing up your system.

```
List of commands:
checklayout  check if the disk layout has changed
dump        dump configuration and system information
format      format and label media for use with rear
mkbackup    create rescue media and backup system
mkbackuponly backup system without creating rescue media
mkrescue    create rescue media only
recover     recover the system
validate    submit validation information
Use 'rear -v help' for more advanced commands.
```



ReaR – backup process



Linux server to Backup

All those steps are done "online".

rear mkbackup

1 rear mkrescue



Create a **Bootable** Media including:


- Disk / Raid / LVM / FS layout.
- Network configuration.

Boot media can be sent to a remote location.

2 rear mkbackuponly




Backup Data with possibility to send it to a remote location.



Remote Backup Server

ReaR – mkrescue

- “**rear mkrescue**” command creates a bootable media with :
 - all the information needed to recreate the server FS structure
 - Set of useful tool for recovery (fdisk, partprobe ...).
 - IP info => allow network recovery.
- Format can be: ISO, USB, ODBR (tape), PXE (network)
 - Use **OUTPUT** variable to set rescue device type.
 - Use **OUTPUT_URL** to set its location (can be remote with nfs, sshfs etc..)
if unset, **OUTPUT_URL = BACKUP_URL** (see **mkbackuponly**)



Remote server



OUTPUT=ISO



OUTPUT=PXE



OUTPUT=ODBR



OUTPUT=USB



```
OUTPUT_URL=nfs://<IP>/PATH
OUTPUT_URL=ftp://<IP>/PATH
OUTPUT_URL=cifs://<IP>/PATH
OUTPUT_URL=sshfs://<IP>/PATH
```

ReaR – mkbackuponly

- “**rear mkbackuponly**” command starts the **“Real Data” backup**.
- Backup could be :
 - Internal : Managed by REAR with OS tool like TAR.
 - External : Managed by Third-Party backup product like : IBM Spectrum Protect (TSM), NetBackup, HP Protect etc ...
- Internal Backup are defined via BACKUP variable
 - BACKUP = NETFS uses TAR with compression (gzip by default) to store backup in a remote server.
 - BACKUP = RSYNC uses rsync to copy files to backup to a remote server.
- BACKUP_URL defines the remote location of backup. The following remote storage protocol are supported:
 - BACKUP_URL=file:///directory/path/
 - BACKUP_URL=tape:///dev/nst0
 - BACKUP_URL=nfs://nfs-server-name/share/path
 - BACKUP_URL=cifs://cifs-server-name/share/path
 - BACKUP_URL=sshfs://root@server/export/archives
- More information here : <https://github.com/rear/rear/blob/master/doc/user-guide/03-configuration.adoc>



ReaR – configuration file (*simple example*)

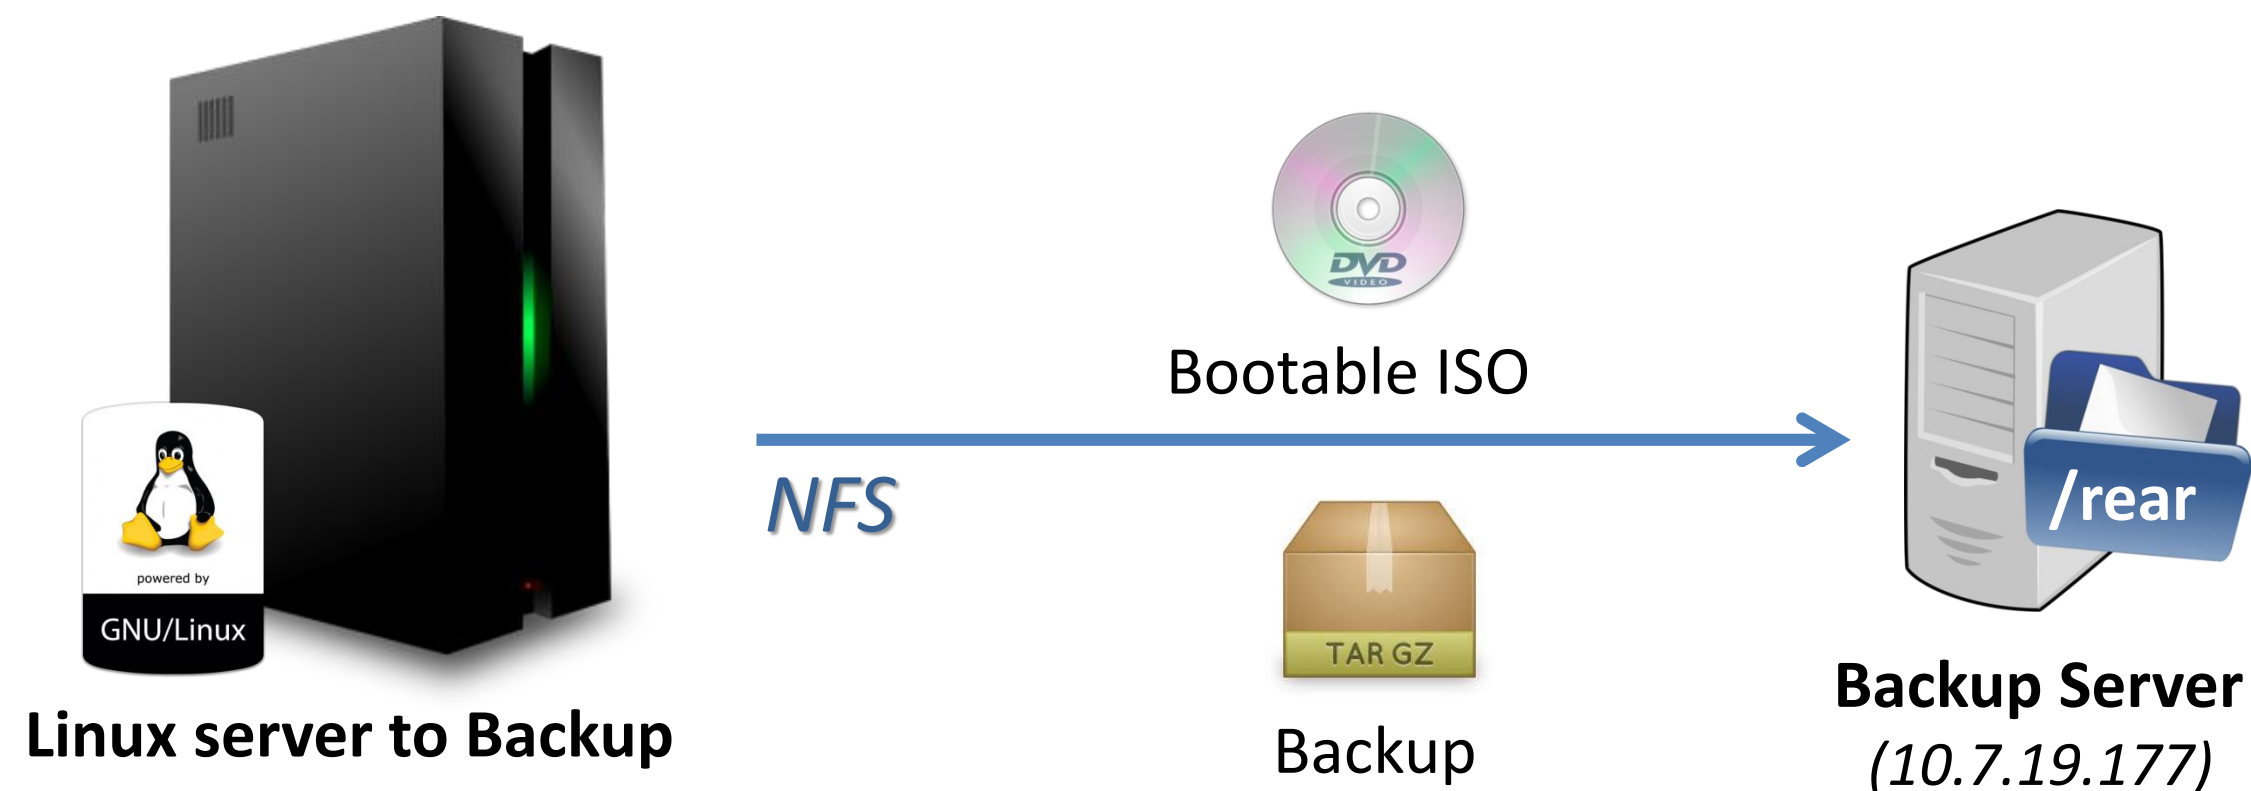
- Define your Rear setting in `/etc/rear/local.conf` (or `/etc/rear/site.conf`)
 - `site.conf` parameter variable will be overwritten by `local.conf`

```
### write the rescue initramfs to USB and update the USB bootloader
OUTPUT=ISO

### create a backup using the internal NETFS method, using 'tar'
BACKUP=NETFS

### write both rescue image and backup to the device labeled BACKUP_URL
BACKUP_URL=nfs://10.7.19.177/rear

### Activate SSH with the following root password for rescue
SSH_ROOT_PASSWORD="reardemo"
```



ReaR – Recover workflow

1. Boot on the REAR-rescue media

```
Relax-and-Recover 1.19-git201610141054 / 2016-10-14
Relax-and-Recover comes with ABSOLUTELY NO WARRANTY; for details see
the GNU General Public License at: http://www.gnu.org/licenses/gpl.html
Host rhel72LE-176 using Backup NETFS and Output ISO
Build date: Sat, 15 Oct 2016 14:01:05 +0200

Red Hat Enterprise Linux Server 7.2 (Maipo)
Kernel 3.10.0-327.22.2.el7.ppc64le on an ppc64le

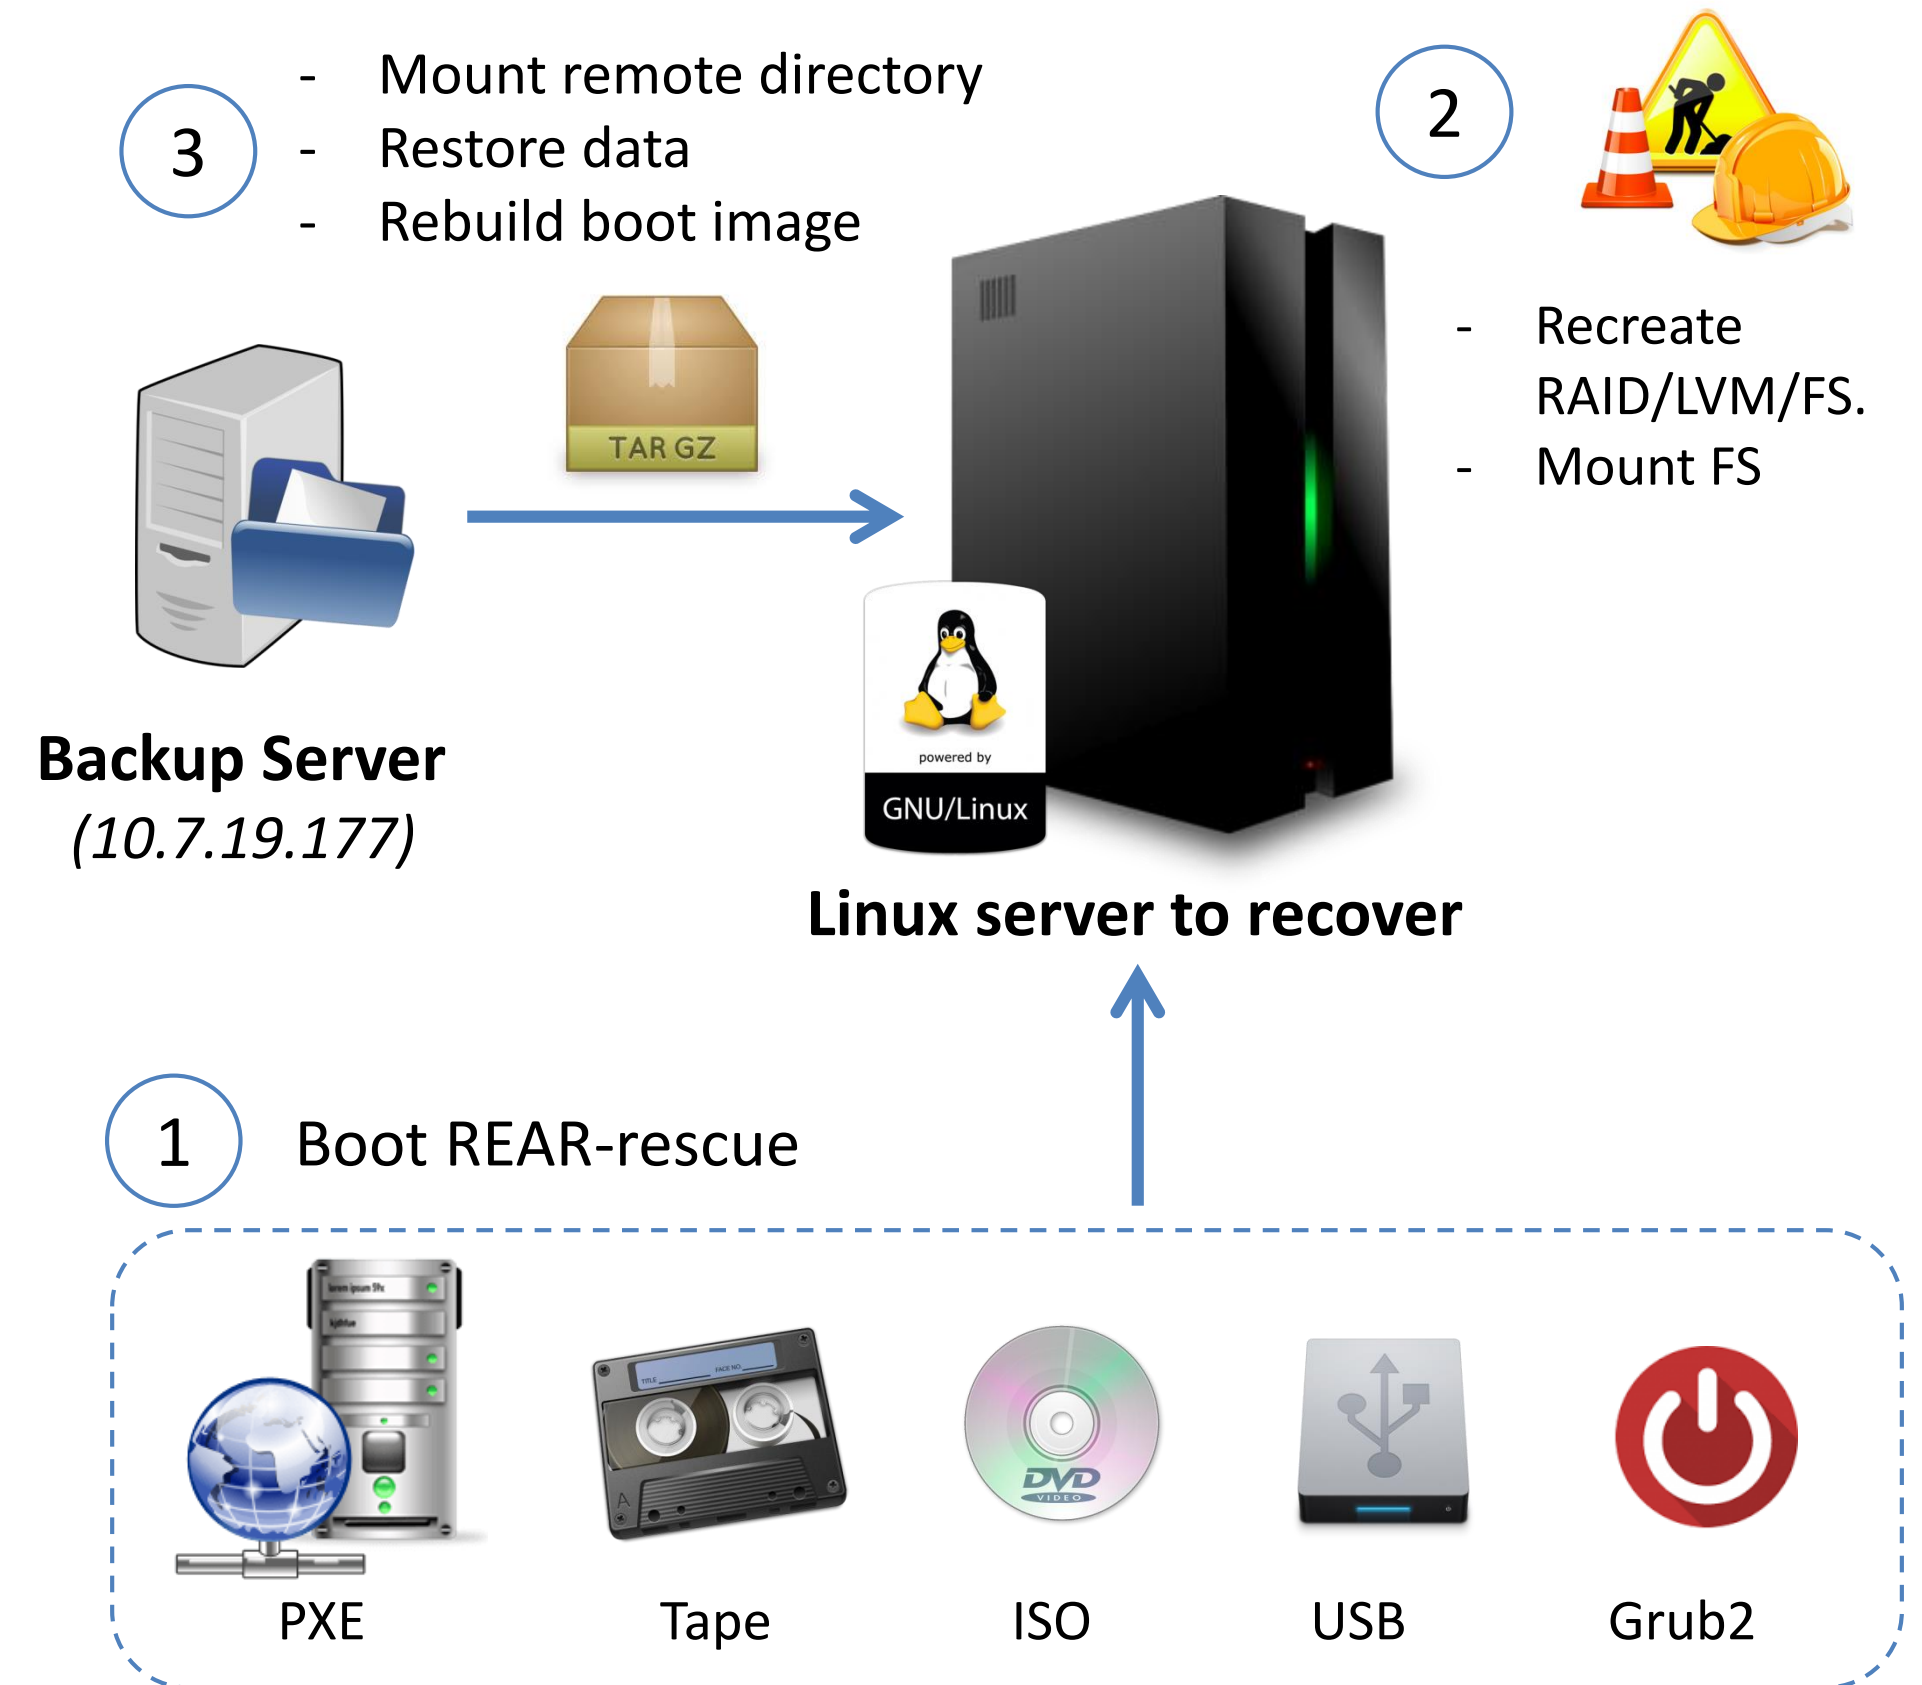
rhel72LE-176 login: root

Welcome to Relax-and-Recover. Run "rear recover" to restore your system !
RESCUE rhel72LE-176:~ #
```

2. Run “rear recover”


- ✓ Check HW, disk/FS layout versus backup info
- ✓ Recreate FS layout (partition/RAID/LVM/FS)
- ✓ Mount remote FS for backup restoration
- ✓ Restore the backup data
- ✓ Rebuild initrd / bootloader if needed.

3. Inspect result (mounted in /mnt/local) & Reboot.



Relax & Recover Demo 1


Simple Backup / Restore using ReaR with NetFS

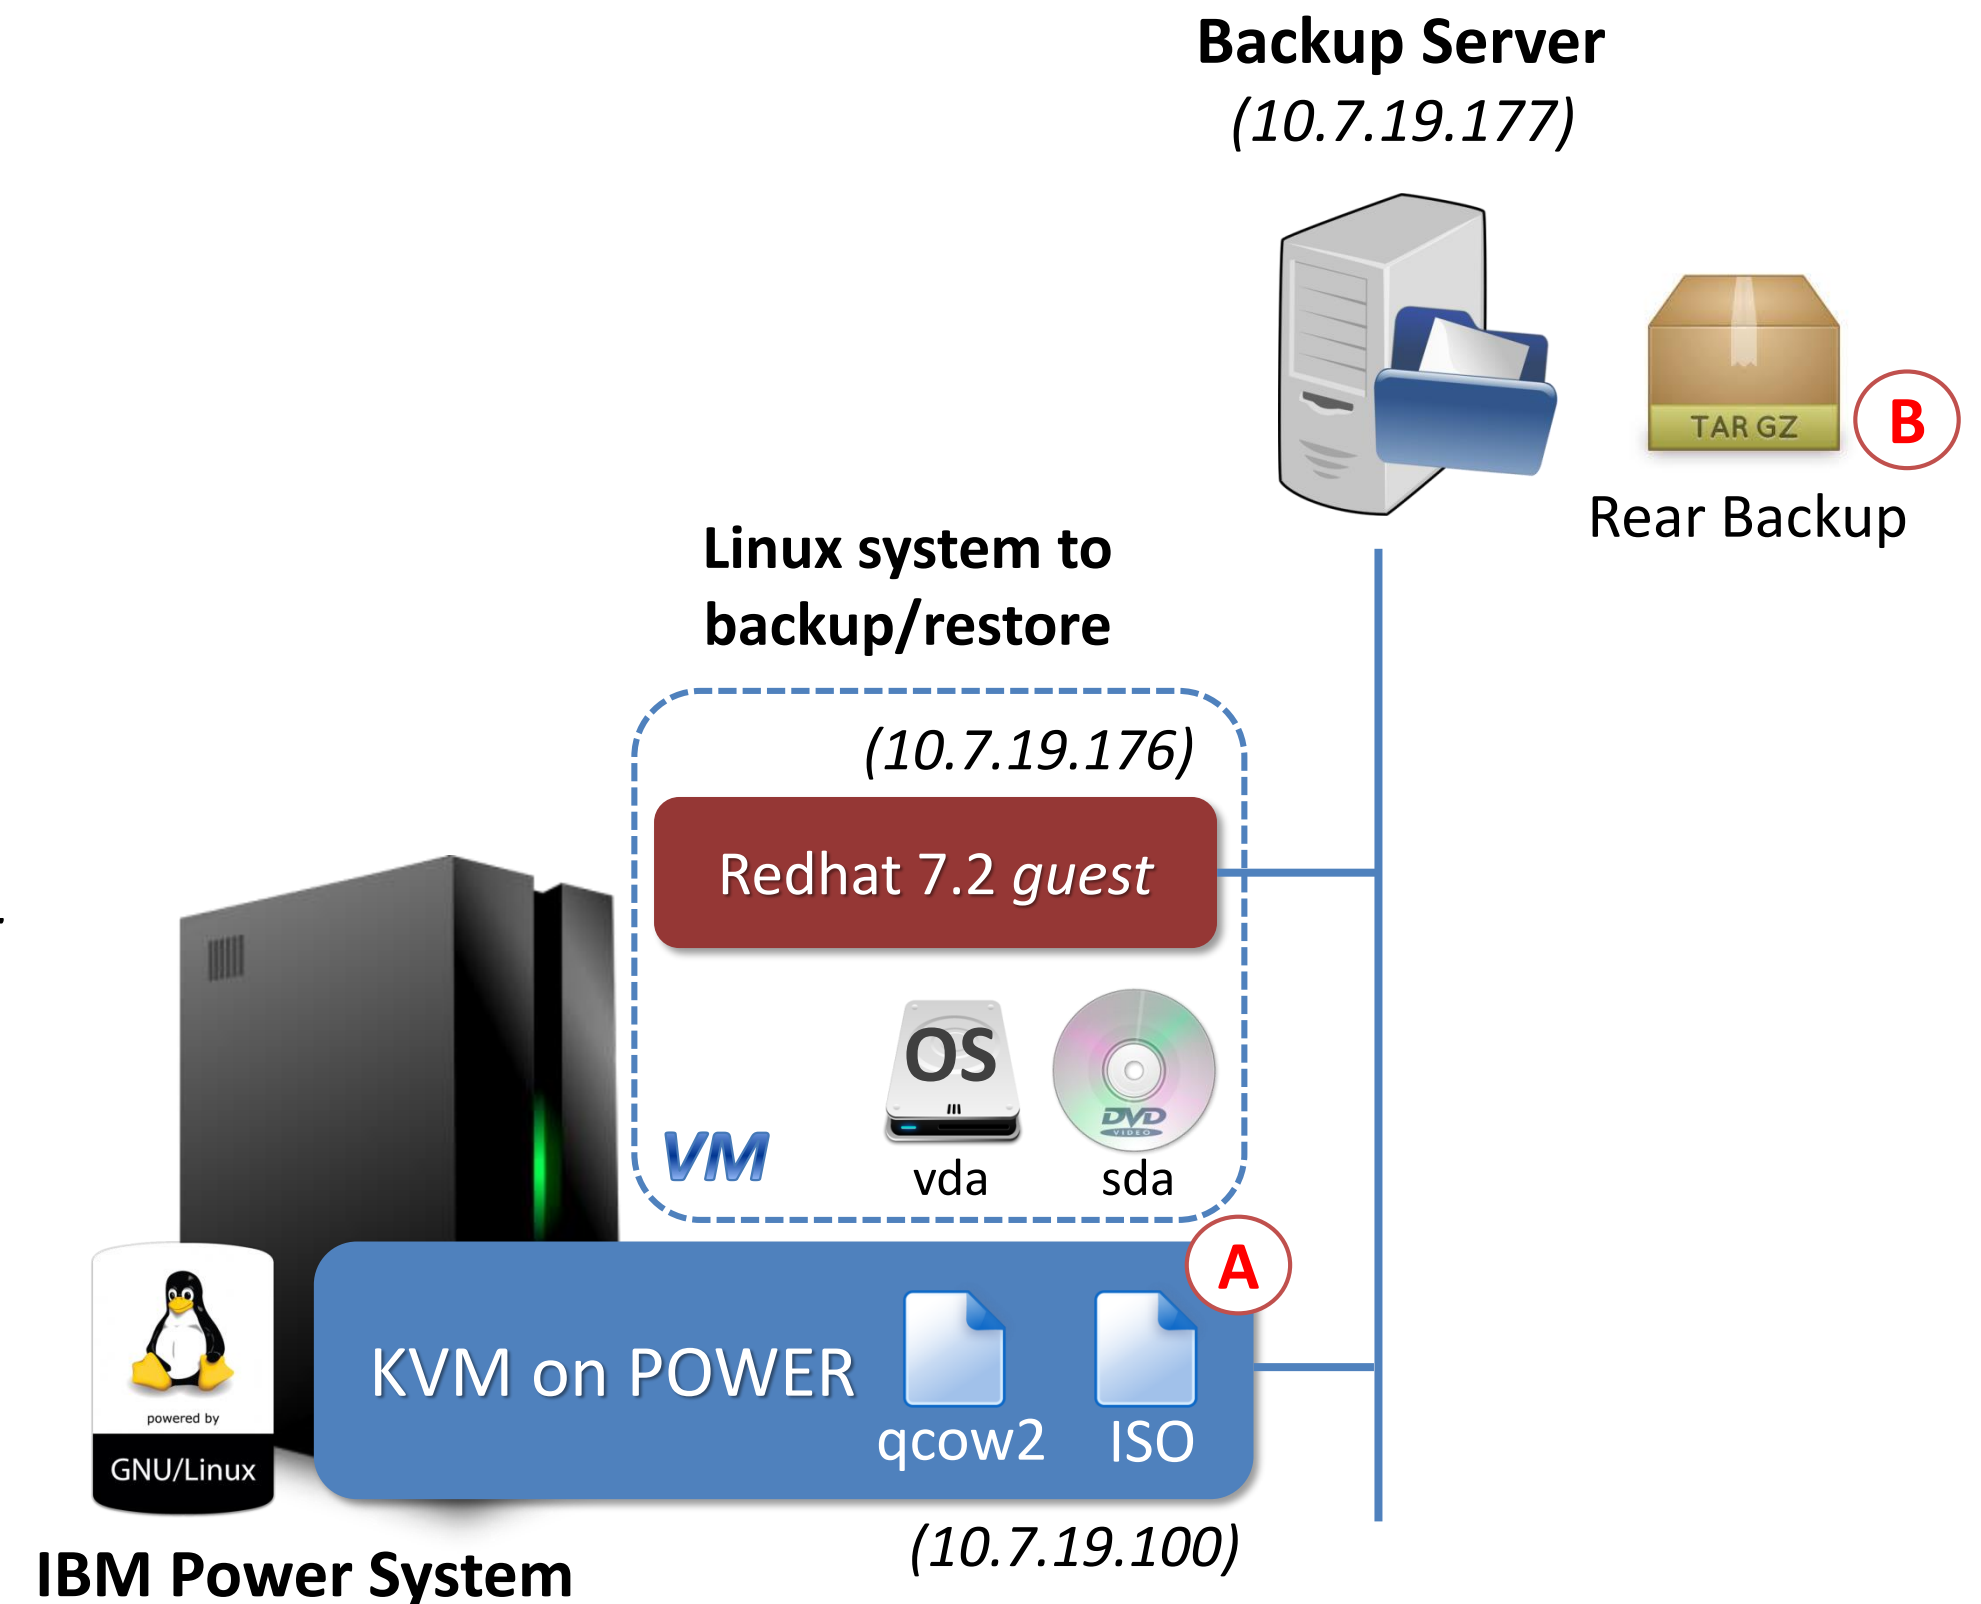


ReaR demo 1 – Simple System Backup/Restore (10 min)

Objective: Simple backup of a KVM guest using :

- ISO image as bootable Rescue-DVD.
- NFS server to store the Backup.

1. Use “`rear -v mkbackup`” to create:
 - A. A bootable DVD on the KVM host
 - B. A Full tar backup on a remote server via NFS.
2. Do something **BAD** to BREAK your system (*prevent it to reboot properly.*)
3. Reboot the system ... **FAILED !!!** 
4. Boot on a ReaR rescue bootable device.
5. Recover the system: “`rear -v recover`”
6. Reboot again on the real system.



ReaR demo 1: "local.conf" used *(for reference)*

```
# Default is to create Relax-and-Recover rescue media as ISO image
# set OUTPUT to change that
# set BACKUP to activate an automated (backup and) restore of your data
# Possible configuration values can be found in /usr/share/rear/conf/default.conf
#
# This file (local.conf) is intended for manual configuration. For configuration
# through packages and other automated means we recommend creating a new
# file named site.conf next to this file and to leave the local.conf as it is.
# Our packages will never ship with a site.conf.

### write the rescue initramfs to USB and update the USB bootloader
OUTPUT=ISO

### create a backup using the internal NETFS method, using 'tar'
BACKUP=NETFS

### write both rescue image and backup to the device labeled BACKUP_URL
BACKUP_URL=nfs://10.7.19.177/rear

### Activate SSH with the following root password for rescue
SSH_ROOT_PASSWORD="reardemo"
```

I need more “enterprise ready” features

This is interesting ... but ...

- I got hundreds of server to manage
 - ⇒ I don't want to spend my “precious” time in copying or managing DVD “rescue media”
 - ⇒ I would need help to migrate to migrate on **NEW hardware**.



- I already have a “Enterprise backup solution”.
 - ⇒ I don't want to backup twice
 - ⇒ And my Backup solution has better feature like “deduplication” or point in time restore..



- My company is “unique” and doing very “special stuff”...
 - ⇒ I need a “special tool” to backup my “special systems”



TAILORMADE SOLUTION

Setting up a “Net Boot” server for Linux on Power

PXE boot is not “natively” available on Linux on Power, but “grub2” provides a network-boot with similar functionality

1. Install and enable tftp service on your server (example for RedHat here under)

- `yum install tftp-server`
- enable tftp server by setting “`disable=no`” in “`/etc/xinetd.d/tftp`”
- `systemctl start xinetd && systemctl enable xinetd`
- Update your Firewall configuration if needed (`firewall-cmd -add-service=tftp --permanent`)

2. Create a **grub2 network boot** directory in your **tftpboot** directory.

- `grub2-mknetdir --net-directory=/var/lib/tftpboot`
- ⇒ This will create a “**boot**” directory with all the files needed for a Power network boot.
- ⇒ THIS OPERATION **MUST BE RUN ON A LINUX ON POWER !!!!**


*This is where
the magic happens*

3. You will have to update your **dhcpd.conf** to use **/boot/grub2/powerpc-ieee1275/core.elf** as filename

```
#
# DHCP Server Configuration file.
# see /usr/share/doc/dhcp*/dhcpd.conf.example
# see dhcpd.conf(5) man page
#

subnet 10.7.19.0 netmask 255.255.255.0 {
  allow bootp;
  next-server 10.7.19.177; # tftpserver ip address
  filename "/boot/grub2/powerpc-ieee1275/core.elf";
  option routers 10.7.19.254;

  host sles11sap-144 {
    hardware ethernet 1a:f4:ea:94:64:0c;
    fixed-address 10.7.19.144;
  }
}
```



Network boot server
(DHCP + BOOTP + TFTP)

ReaR configuration for Netboot (*POWER*)

- ReaR capability to generate “grub styled” PXE configuration will be available in ReaR 2.1.
(You can still test it by getting latest code from [github](#) or use a package version > `rear-2.00-1.git201705041626.ppc64le.rpm`)
- Use **OUTPUT=PX**E and **PXE_CONFIG_GRUB_STYLE=y** in your **local.conf** file to enable it.
- You also have to specify URLs where ReaR can put files generated for recovery:
 - **PXE_TFTP_URL**: location used to store “**rear-kernel**” and “**rear-initrd**”
⇒ it must point to the root of your tftp server (`tftpboot dir`)
 - **PXE_CONFIG_URL**: location used to store the dedicated “**grub.cfg**” file which contain information about location of the “**rear-kernel**” and “**rear-initrd**” file.
⇒ It must point to `(tftpboot dir)/boot/grub2/powerpc-ieee1275`
⇒ During boot sequence, tftp client looks for a file named `grub.cfg-<MAC_ADDRESS>` or `grub.cfg-<IP_HEX>`

Extract of local.conf with PXE configuration

```
# Using PXE/NetBoot Server Rescue image
OUTPUT=PX
PXE_CONFIG_GRUB_STYLE=y
PXE_TFTP_URL="nfs://10.7.19.177/var/lib/tftpboot"
PXE_CONFIG_URL="nfs://10.7.19.177/var/lib/tftpboot/boot/grub2/powerpc-ieee1275"
```

ReaR – Backup Software Integration

ReaR provides simple integrated full backup.

ReaR is really focus on disaster recovery and not backup.

- It could be integrated with common backup software to delegate “*file backup*” to a real backup infrastructure:
 - Backup software: Data storage and retrieval
 - ReaR: Recover system layout and make it work again
 - ReaR orchestrates the overall process and use the backup software to restore the backup data.

⇒ Use the best tool for the job.

- Enabled by setting “**BACKUP**” variable.
- Currently compatible with :

Tivoli Storage Manager (BACKUP=TSM)

HP Data Protector (BACKUP=DP)

Symantec NetBacakup (BACKUP=NBU)

Galaxy 5, 6, and 7 (BACKUP=GALAXY)

Galaxy 10 [Commvault Simpana] (BACKUP=GALAXY10)

Bacula (BACKUP=BACULA)

Bareos (BACKUP=BAREOS) (*A fork of Bacula*)

Rsync Backup Made Easy (BACKUP=RBME)

Duplicity/Duply (BACKUP=DUPLICITY)

EMC Networker, also known as Legato (BACKUP=NSR)

SEP Sesam (BACKUP=SESAM)

Borg Backup (BACKUP=BORG)

FDR/Upstream (BACKUP=FDRUPSTREAM)

Novastor NovaBACKUP DC (BACKUP=NBKDC)

Centralized Backup



ReaR – Backup Software Integration – *(tips)*

ReaR is really focus on disaster recovery and not backup. *(repeat)*

- Don't run "`rear mkbbackup`" => use your external backup tool with scheduling policy.
- Use "`rear mkrescue`" to create the bootable rescue device.

⇒ **Don't forget to update it regularly (run `rear mkrescue`) if you**

- **change your disk / FS layout (creating or resizing LV and FS)**
- **modify backup client configuration files.**


TIPS:

"`rear checklayout`" command determines if there is any change in the fs layout since the last "`rear mkrescue`".

- Return 0 => no change
- Return 1 => change: need to regenerate rescue device.

⇒ You can simply use the following cron schedule task to automate this.

```
#cat /etc/cron.d/rear
30 1 * * * root /usr/sbin/rear checklayout || /usr/sbin/rear mkrescue
```



ReaR – Migration, Restoring to a different HW.

- Enable recovery on dissimilar hardware.
(that is not the same as the original system **but still the same arch**).

For example:

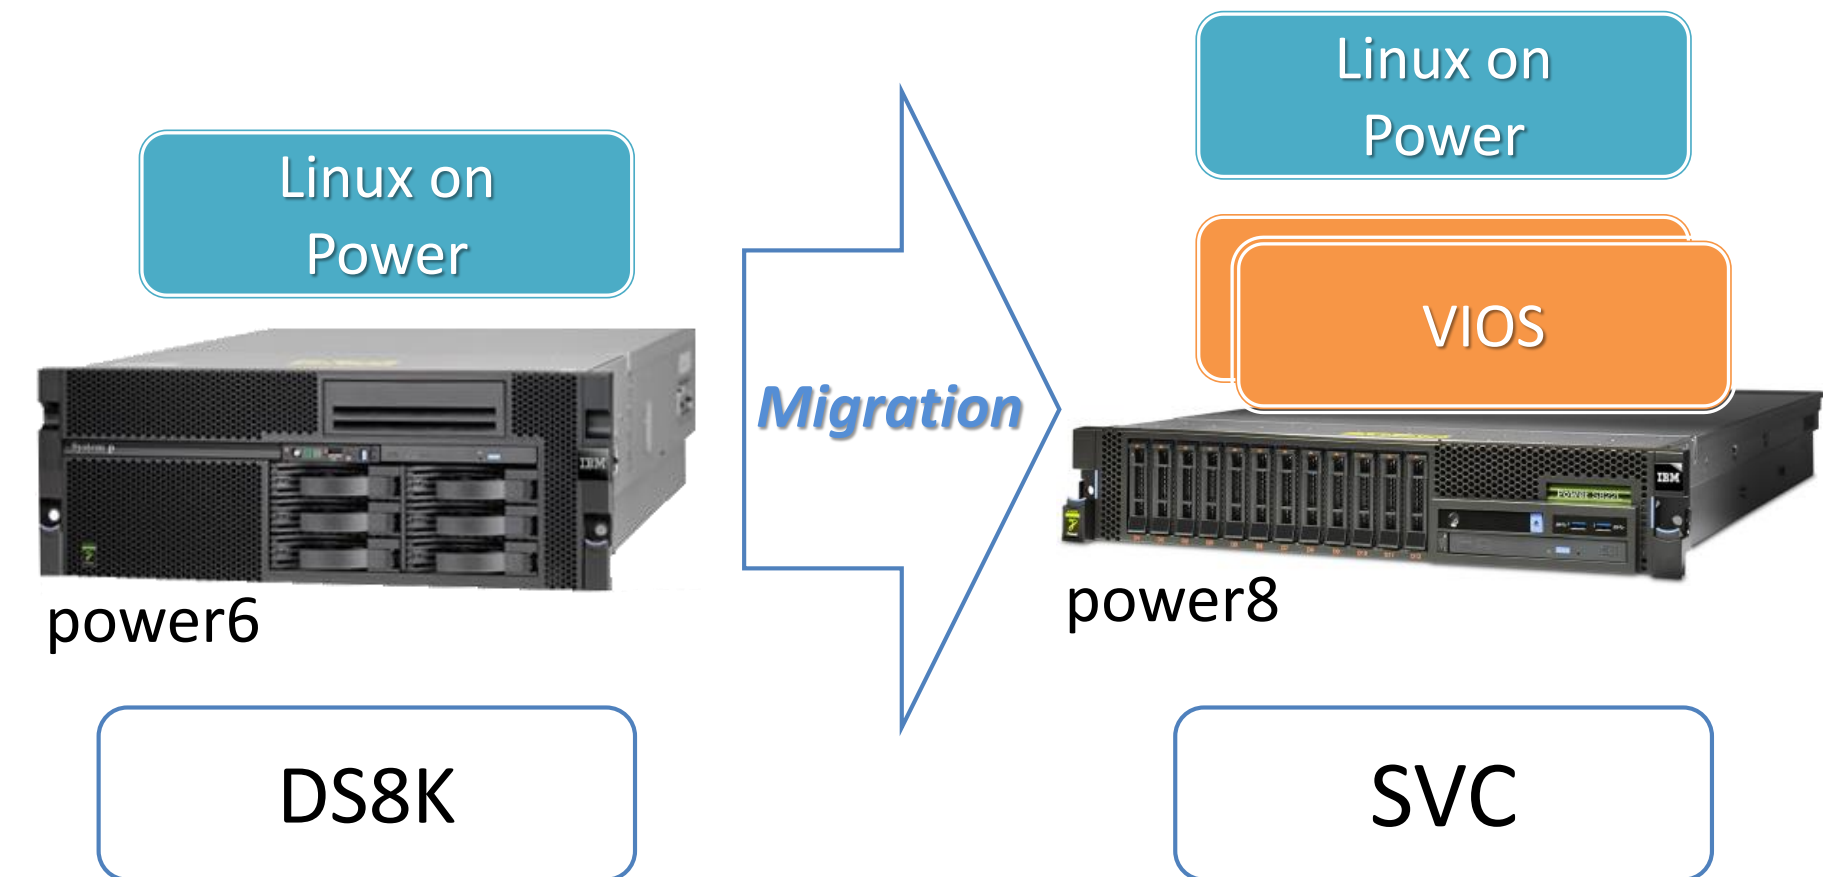
- P2V, V2P, V2V, P2P
- From POWER7 to POWER8
- From VIOS vscsi to Direct SAN
- From PowerKVM to PowerVM
- ...

- **Network :**

- network and storage drivers are adjusted
- remap network MAC addresses
- use another IP address, or using dhcp via templates or from kernel command line

- **Disks :**

- map hard disks if they do not match (e.g. hda -> sda)
- rebuild the initial ramdisk if needed (for new storage drivers)
- migration to SAN storage (Experimental)




Example, moving from :

- Power6 to Power8
- DS8K disk to SVC
- Dedicated IO to VIO:
 - * virtual Network
 - * virtual fiber (NPIV)

Relax & Recover Demo 2

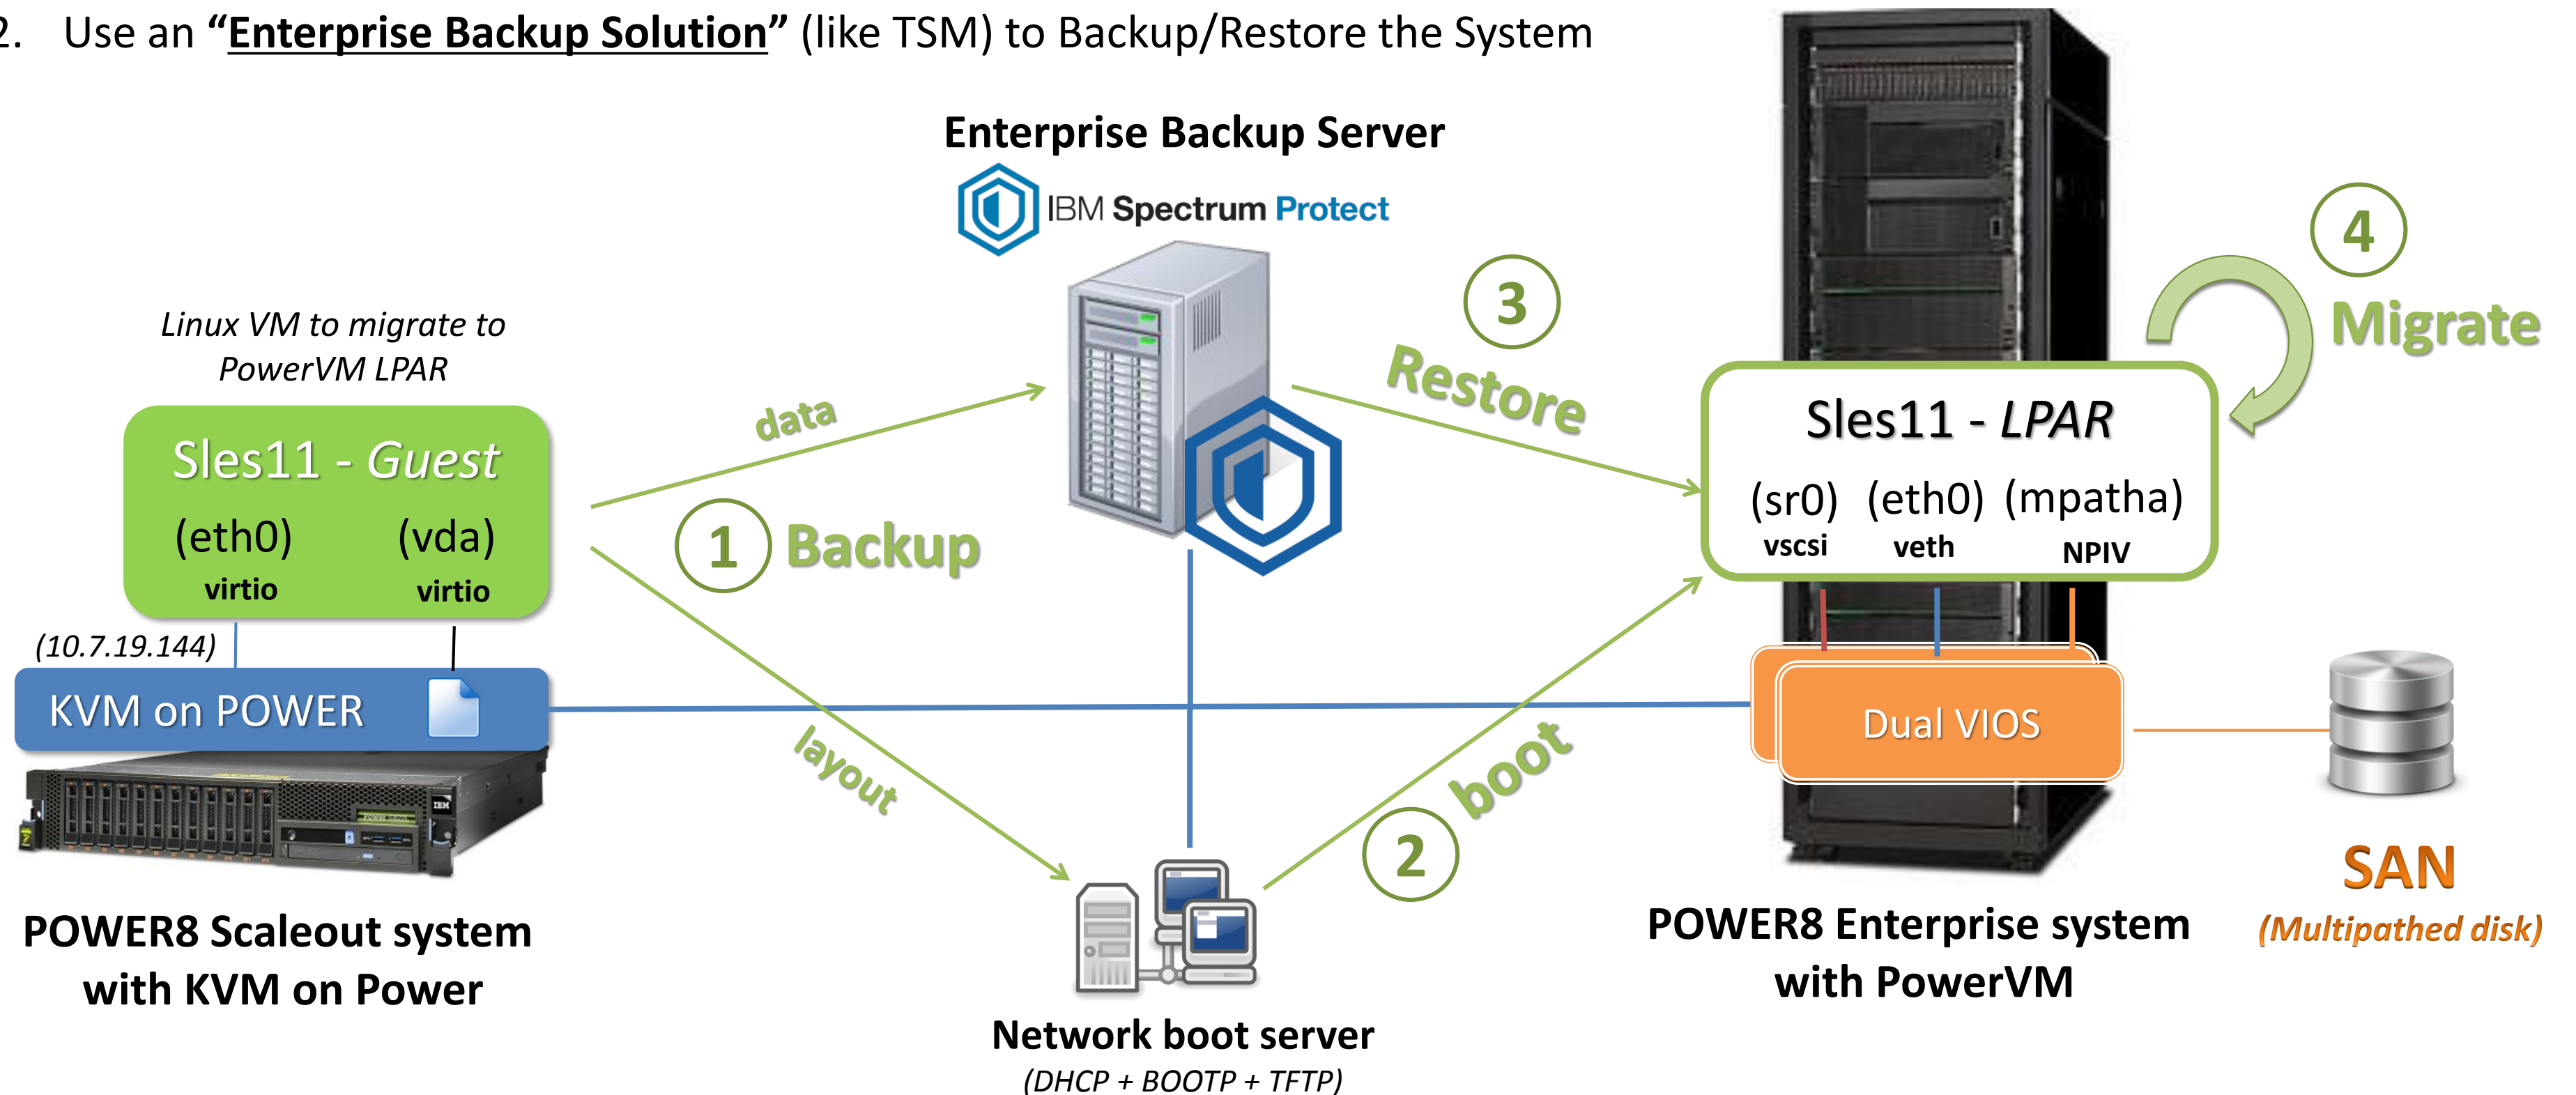
1. ReaR rescue over the Network
2. ReaR integration with TSM
3. Restoration on dissimilar HW



ReaR demo 2 – KVM guest to PowerVM LPAR Migration. (15 min)

Objectives:

1. Migrate a Sles11 **KVM** guest to **PowerVM LPAR**
 - From virtio network to virtual ethernet (*New MAC !!*)
 - From virtio filebacked disk to NPIV SAN based (*vda -> mpatha*)
 - New device driver (*ibmvscsi, ibmveth, ibmvfc*) => need to rebuild “**initrd**” ramdisk.
2. Use an “**Enterprise Backup Solution**” (like TSM) to Backup/Restore the System



ReaR demo 2: "local.conf" used (for reference)

- By default ReaR exclude multipath device (**disable this**)
 - Our target system boot over SAN disk
- This will force multipath configuration during block device probing

- Use PXE to store rescue media
- PXE with grub2 mode (POWER compatible)
- **PXE_TFTP_URL** must be your tftpboot root dir.
- **OUTPUT_PREFIX_PXE** store rescue data in this subdir.
- **PXE_CONFIG_URL**: where client will find boot cmd to use
- **PXE_CREATE_LINKS=IP** : Use IP suffix instead of MAC (because we want to migrate to a different HW)

- PowerVM can have trouble during boot if "initrd" size is too BIG
Using a better compression like lzma (.xz) can help
- Reduce to the minimum the TSM client binary to copy to the ReaR rescue image => reduce initrd size.

```
# Default is to create Relax-and-Recover rescue media as ISO image
# set OUTPUT to change that
# set BACKUP to activate an automated (backup and) restore of your data
# Possible configuration values can be found in /usr/share/rear/conf/default.conf
#
# This file (local.conf) is intended for manual configuration. For configuration
# through packages and other automated means we recommend creating a new
# file named site.conf next to this file and to leave the local.conf as it is.
# Our packages will never ship with a site.conf.

# Prepare for Migration to PowerVM with BootOnSAN
AUTOEXCLUDE_MULTIPATH=n
BOOT_OVER_SAN=y

# Using PXE/NetBoot Server Rescue image
OUTPUT=PXE

PXE_CONFIG_GRUB_STYLE=y
PXE_TFTP_URL="nfs://10.7.19.177/var/lib/tftpboot"
OUTPUT_PREFIX_PXE=rear/$HOSTNAME

PXE_CONFIG_URL="nfs://10.7.19.177/var/lib/tftpboot/boot/grub2/powerpc-ieee1275"
PXE_CREATE_LINKS=IP
PXE_REMOVE_OLD_LINKS=y

# Backup TSM
BACKUP=TSM

# Reduce initrd for PowerVM
REAR_INITRD_COMPRESSION=lzma

COPY_AS_IS_TSM=( /etc/adsm/TSM.PWD /opt/tivoli/tsm/client/ba/bin/dsmc
/opt/tivoli/tsm/client/ba/bin/tsmbench_incl excl /opt/tivoli/tsm/client/ba/bin/dsm.sys
/opt/tivoli/tsm/client/ba/bin/dsm.opt /opt/tivoli/tsm/client/api/bin64/libgpfs.so
/opt/tivoli/tsm/client/api/bin64/libdmapi.so /opt/tivoli/tsm/client/ba/bin/EN_US/dsmclientV3.cat
/usr/local/ibm/gsk8* )
```


ReaR Keywords

- **Enterprise Solution**
 - Fully automated.
 - Seamless integration with most popular Enterprise Backup Solution.
- **Open Source (GPL)**
 - Source code included (100% bash).
- **Modular Design**
 - Easily extendible with new functions.
- **Scalability**
 - One solution for all Linux systems, unlimited scale-out.
- **Usability**
 - Documentation, community & commercial support.



Highly Customizable

Want to participate to the ReaR community ?



- ReaR code is hosted on github.



<https://github.com/rear/rear>

- Possibility to raise issues or propose modifications.
- Since 2015 (rear 1.18), some IBMers contribute to improve POWER architecture support:
 - Multipathing support
 - PowerVM Bootlist generation after recover
 - PPC64 with Yaboot support
 - PPC64LE with grub2 support
 - Ubuntu support
 - PowerKVM guest support
 - PowerNV support
 - PowerKVM <=> PowerVM Migration
 - IBM TSM integration
 - PowerVM/KVM netboot support (*grub2 netboot*)

Commit	Description	Time
gdha	Merge pull request #1034 from rear/dev	Latest commit 28199db 5 days ago
.github	Update CONTRIBUTING.md	2 months ago
doc	adding the final release notes (we hope at least)	7 days ago
etc	Removed site.conf per request by gdha	a year ago
packaging	Tagging release 1.19 into packaging files	7 days ago
usr	Tagging release 1.19 into packaging files	7 days ago
.gitignore	Support having a var dir somewhere deep down in our tree. Fixes #792	7 months ago
.travis.yml	Ensure new commits are tested at GitHub/TravisCI	10 months ago
AUTHORS	Rebrand Rear as Relax-and-Recover where possible	4 years ago
COPYING	Changing GPLv2 to GPLv3.	10 months ago
Makefile	fixed a few more suse spelling typos	28 days ago
README.adoc	Corrected misspelled forms to Relax-and-Recover (issue1008).	25 days ago



ReaR Core Team:

- Gratien D'haese (*ReaR main author*)
- 4 main Developers (*1 from Suse*)
- + **60 contributors** :
- At least "3 known" IBMers

Thanks to Jason Furmaneck (IBM USA), Masanori Mitsugi (IBM JP), Sebastien Chabrolles (IBM FR)

POWER test Matrix: <https://github.com/rear/rear/wiki/Test-Matrix-rear-1.19#hardware-vendors>

ReaR and POWER: support matrix & known issue

- **ReaR version and POWER support**

Even if ReaR supports POWER arch since more than 2 years now, a lot of issues were corrected recently (especially regarding “multipathing “support).

If the version provided to the distribution doesn't work for you, I really encourage you to test the latest version available on [github](#) (*master branch*)

Linux OS Version	PowerVM	PowerNV (baremetal)	KVM Guest	Migration*
RedHat 6 (ppc64)	OK	N/A	OK	OK
RedHat 7 (ppc64/ppc64le)	OK	OK	OK	OK
Suse SLES 11 SP4 (ppc64)	OK	N/A	OK	OK
Suse SLE 12 SP2 (ppc64le)	OK	N/A	OK	OK
Ubuntu 16.04 (ppc64le)	OK	OK	OK	<i>not tested yet</i>

Test performed based on rear-2.00-1.git201705111237 (rear-2.1-devel)

- **PowerVM firmware & linux kernel 256MB RMA limit.**

- kernel and drivers need more and more place.
initrd in rescue media (like RedHat rescue or ReaR) can be bigger than 100MB.
=> **PowerVM LPAR cannot load firmware + kernel + initrd ...**

New Linux kernel should switch to a 512MB RMA size.

This should be included in RHEL 7.4 and a backport for RHEL 7.3 is in discussion.

BUGZILLA

Workaround: Use “`REAR_INITRD_COMPRESSION=lzma`” in your “`local.conf`” file.

It should reduce your initrd size by half, but use more CPU time to be generated.

```
PowerVM RHEL7.3
OF stdout device is: /vdevice/vty@30000000
Preparing to boot Linux version 3.10.0-514.el7.ppc64 (mockbuild@ppc-
021.build.eng.bos.redhat.com) (gcc version 4.8.5 20150623 (Red Hat 4.8.5-11)
(GCC) ) #1 SMP Wed Oct 19 11:30:41 EDT 2016
Detected machine type: 0000000000000101
Max number of cores passed to firmware: 256 (NR_CPUS = 2048)
Calling ibm,client-architecture-support... done
command line: BOOT_IMAGE=/vmlinuz-0-rescue-3a0bde33ffc1407985be3a777b9748cf
root=/dev/mapper/rhel_d--zswap--rearrh-root ro crashkernel=auto
rd.lvm.lv=rhel_d-zswap-rearrh/root rd.lvm.lv=rhel_d-zswap-rearrh/swap
memory layout at init:
memory_limit : 0000000000000000 (16 MB aligned)
alloc_bottom : 00000000f210000
alloc_top    : 0000000010000000
alloc_top_hi : 0000000010000000
rmo_top     : 0000000010000000
ram_top     : 0000000010000000
Could not allocate memory for RTAS
EXIT called ok
>
```

Need Support ?



Gratiën D'haese, “*Relax-and-Recover*” author and founder of “*IT3 Consultants*”, provides IT services, consulting and support around ReaR and Disaster Recovery.



- **Relax-and-Recover (Rear) Consultancy Services:**

- Help with writing & implementing DR policy
- Designing & configuring rear as a central recovery solution
- Proof of Concept
- Training Session / Workshop

- **Relax-and-Recover (Rear) Support Services :**



1. Limited Support Contract
2. Business Hours Support Contract
3. 24x7 Support Contract

Website: <http://it3.be/rear-support/index.html>



Non OpenSource: Linux System Recovery Tools



- Founded in 1999 by the author of IBM's sysback
- Bare-metal recovery for Linux since 2002
- TSM Integration since 2007
- Bare-metal recovery for Solaris since 2008



Non OpenSource: Linux System Recovery Tools



System recovery and
recovery assurance.



<http://www.cristie.com/>

Enhance your server protection planning with Power Systems and Cristie Recovery Suite. We're proud to support system recovery, migration and replication for Linux and AIX on Power Systems.

[Read more](#)

Thank you for your attention ...

Any Question ?




IBM

Sébastien Chabrolles
s.chabrolles@fr.ibm.com

*Power Systems Linux Center
IBM Systems
IBM Client Center Montpellier
France*



Montpellier “Power Systems Linux Center” Capabilities & Focus



Power Systems Linux Center for Europes

Strategic Focus on:

- Big Data
- Cloud
- Mobile
- Opensource DBs

Providing Support for :

- Access to platforms (HW + SW + Support)
- Customer architecture design
- General Developer Resources Support
- ISV Resources Support
- Education / Training



Capabilities

Talk and Teach : Customer, BP, CSI & ISV Briefings, Demos, Videos

Design: Pre-sales Customer support, Customer Consultancy, Architecture Design Workshops

Prove : Linux on Power Benchmark & PoCs, Remote Power Linux Platform access

+ **Second level of support** for technical IIC, IMTs (Infrastructure, Virtualization, OS, Compilers, Certification programs, ...)



Philippe Chonavel
Open Source &
Accelerated Solutions
Center Manager




Alain Roy
Engagement Leader
Contact: a2roy@fr.ibm.com



Sébastien Chabrolles
Technical Leader
Linux on Power Specialist
(Performance, Virtualization)



Julien Limodin
Linux on Power Specialist
(Middleware knowledge,
Mobile & Java skills, Cloud)



Fabrice Moyen
Linux on Power Specialist
(Performance, Virtualization,
SAP Hana)



Christophe Menichetti
Power Architect
(Big Data knowledge /
Competitive knowledge)

Notices and Disclaimers

Copyright © 2016 by International Business Machines Corporation (IBM). No part of this document may be reproduced or transmitted in any form without written permission from IBM.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IN NO EVENT SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY. IBM products and services are warranted according to the terms and conditions of the agreements under which they are provided.

IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”

Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law

Notices and Disclaimers Con't.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. IBM EXPRESSLY DISCLAIMS ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, ibm.com, Aspera®, Bluemix, Blueworks Live, CICS, Clearcase, Cognos®, DOORS®, Emptoris®, Enterprise Document Management System™, FASP®, FileNet®, Global Business Services®, Global Technology Services®, IBM ExperienceOne™, IBM SmartCloud®, IBM Social Business®, Information on Demand, ILOG, Maximo®, MQIntegrator®, MQSeries®, Netcool®, OMEGAMON, OpenPower, PureAnalytics™, PureApplication®, pureCluster™, PureCoverage®, PureData®, PureExperience®, PureFlex®, pureQuery®, pureScale®, PureSystems®, QRadar®, Rational®, Rhapsody®, Smarter Commerce®, SoDA, SPSS, Sterling Commerce®, StoredIQ, Tealeaf®, Tivoli®, Trusteer®, Unica®, urban{code}®, Watson, WebSphere®, Worklight®, X-Force® and System z® Z/OS, are trademarks of International Business Machines Corporation, 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.