

Packages - Bug #2942

Spaces in filenames giving problems while installing

2020-12-08 05:55 PM - infinite_recursion

Status:	info needed	% Done:	0%
Priority:	bug		
Assignee:			
Category:			
Description			
iptables lz4 zlib			
armv7h filenames of these packages have spaces in them. There should be : instead of a space there.			

History

#1 - 2020-12-08 06:12 PM - oaken-source

are you sure about that? on my end, there's no spaces in the pkgnames:

<https://repo.parabola.nu/pool/alarm/iptables-1:1.6.2-3-armv7h.pkg.tar.xz>

although, there is something weird about these packages. I notice that the pseudo-simlink on the repo browser does not work.

#2 - 2020-12-08 06:22 PM - infinite_recursion

Absolutely, press 'download from mirror' at the link below

<https://www.parabola.nu/packages/core/armv7h/iptables/>

There's a space whereas in other packages, I didn't see any. I'm facing an error while pacstrap base libelogind in above 3 packages.

#3 - 2020-12-08 07:20 PM - oaken-source

all of these packages come from archlinuxarm though. we should probably check if upstream has the same problem.

#4 - 2020-12-08 07:45 PM - grizzlyuser

It seems to be not related to architecture and can be reproduced on Arch too for a package that has an epoch in PKGBUILD and thus in version too:

1. Go to: <https://www.archlinux.org/packages/community/any/mythes-hu/>
2. Click 'Download From Mirror' and download the file.
3. Check if the file has a space in the name.

I don't know what other tools do when download the package, but Firefox-based browsers save it with space in the name. Looks like the web browser does some name sanitizing, or simply has a bug when a file has a colon in the name.
Couldn't reproduce when installing with pacstrap: it's saved to pacman cache with correct name.

In case of iptables, it looks like armv7h package name extension is .xz, while for other architectures it's .zst. Also the response headers differ when downloading. For example, from Parabola website 'content-type' for .xz packages is 'application/x-xz', while for .zst it's 'text/plain', which looks wrong. From Arch website 'Content-Type' (notice different case of header name) looks to be always 'application/octet-stream' regardless of the extension. Which points out some differences in how servers respond.

[infinite_recursion](#), what tool did you use to download those files before running pacstrap? And where in the filesystem were they saved with incorrect names? It's a general rule to provide as much information in the bug report as possible. Like detailed steps to reproduce, error logs, etc.

#5 - 2020-12-09 05:23 AM - infinite_recursion

```
$ sudo pacstrap -C /usr/share/pacman/defaults/pacman.conf.armv7h harddisk/ base libelogind
==> Creating install root at harddisk/
==> Installing packages to harddisk/
:: Synchronizing package databases...
nonsystemd is up to date
libre                               331.0 KiB   86.0 KiB/s  00:04 [#####] 100%
kernels is up to date
core                                163.8 KiB   75.8 KiB/s  00:02 [#####] 100%
extra                               2.0 MiB    648 KiB/s  00:03 [#####] 100%
```

```
community is up to date
pcr                               359.4 KiB   401 KiB/s 00:01 [#####] 100%
resolving dependencies...
:: There are 2 providers available for libudev:
:: Repository nonsystemd
    1) eudev-libudev
:: Repository libre
    2) systemd-libudev

Enter a number (default=1):
:: There are 3 providers available for init:
:: Repository nonsystemd
    1) openrc-init  2) sysvinit
:: Repository libre
    3) systemd-sysvcompat

Enter a number (default=1):
:: There are 2 providers available for libdbus-1.so=3-32:
:: Repository nonsystemd
    1) dbus
:: Repository core
    2) dbus

Enter a number (default=1):
looking for conflicting packages...
warning: dependency cycle detected:
warning: ca-certificates-utils will be installed before its coreutils dependency
warning: dependency cycle detected:
warning: pinentry will be installed before its libcap dependency
warning: dependency cycle detected:
warning: openrc will be installed before its openrc-init dependency

Packages (117) acl-2.2.53-3 archlinux-keyring-20201028-1 archlinux32-keyring-20201117-1.0
archlinuxarm-keyring-20140119-1 attr-2.4.48-3 audit-2.8.5-6 bash-5.0.018-2 bzip2-1.0.8-4
ca-certificates-20181109-4 ca-certificates-mozilla-3.59-1 ca-certificates-utils-20181109-4
coreutils-8.32-1 curl-7.73.0-1 db-5.3.28-5 dbus-1.12.16-5.nonsystemd1 e2fsprogs-1.45.6-3
eudev-libudev-3.2.9-1 expat-2.2.10-2 file-5.39-1 filesystem-2019.10-1.par1.nonsystemd1
findutils-4.7.0-2 gawk-5.1.0-1 gcc-libs-10.2.0-1 gdbm-1.18.1-3 gettext-0.21-1
glib2-2.66.2-1 glibc-2.32-2 gmp-6.2.0-1 gnupg-2.2.24-1 gnutls-3.6.15-1 gpgme-1.15.0-1
grep-3.6-1 gzip-1.10-3 hwids-20200813.1-1 iana-etc-20201012-1 icu-67.1-1
iproute2-5.9.0-1 iptables-1:1.8.6-1 iputils-20200821-1 keyutils-1.6.3-1 kmod-27-2
krb5-1.18.2-1 less-563-1 libarchive-3.4.3-1 libassuan-2.5.4-1 libcap-2.45-2
libcap-ng-0.8-1 libcroco-0.6.13-2 libelf-0.181-1 libffi-3.3-4 libgcrypt-1.8.7-1
libgpg-error-1.39-1 libidn2-2.3.0-1 libksba-1.4.0-2 libldap-2.4.54-1 libmnl-1.0.4-3
libnetfilter_conntrack-1.0.8-1 libnfnetlink-1.0.1-4 libnftnl-1.1.8-1 libnghttp2-1.41.0-1
libnl-3.5.0-2 libp11-kit-0.23.20-4.nonsystemd1 libpcap-1.9.1-4 libpsl-0.21.1-1
libsasl-2.1.27-3 libsecret-0.20.4-1 libssh2-1.9.0-2 libtasn1-4.16.0-1 libtirpc-1.2.6-1
libunistring-0.9.10-3 libx11-1.7.0-1 libxau-1.0.9-3 libxcb-1.14-1 libxcrypt-4.4.17-1
libxdmcp-1.1.3-3 libxml2-2.9.10-6 licenses-20200427-1.parabola2
linux-libre-api-headers-5.3.1_gnu-1 lz4-1:1.9.2-2 mpfr-4.1.0-1 ncurses-6.2-1 nettle-3.6-1
npth-1.6-2 openrc-0.42.1-6 openrc-init-0.42.1-6 openssl-1.1.1.h-1 opensysusers-0.5.1-2
opentmpfiles-0.2-2 p11-kit-0.23.20-4.nonsystemd1 pacman-5.2.1-4.parabola3
pacman-mirrorlist-20201202-1.parabola2 pam-1.5.0-1 pambase-20200721.1-2.nonsystemd1
parabola-keyring-20201125-1 pciutils-3.7.0-1 pcre-8.44-1 perl-5.32.0-3 pinentry-1.1.0-5
procps-ng-3.3.16-2 psmisc-23.3-2 readline-8.0.004-1 sed-4.8-1 shadow-4.8.1-4
sqlite-3.33.0-2 tar-1.32-3 tzdata-2020d-1 util-linux-2.35.1-1.nonsystemd1
util-linux-libs-2.36.1-3 xcb-proto-1.14.1-1 xorgproto-2020.1-1 xz-5.2.5-1
your-freedom-20201207-1 your-initfreedom-20200617-1 zlib-1:1.2.11-4 zstd-1.4.5-1
base-2-2.parabolal.nonsystemd2 libelogind-243.7-2

Total Download Size:      0.53 MiB
Total Installed Size:    447.81 MiB


:: Proceed with installation? [Y/n]
:: Retrieving packages...
error: could not open file harddisk//var/cache/pacman/pkg/lz4-1:1.9.2-2-armv7h.pkg.tar.xz.part: Invalid argument
nt
error: could not open file harddisk//var/cache/pacman/pkg/lz4-1:1.9.2-2-armv7h.pkg.tar.xz.part: Invalid argument
nt
error: could not open file harddisk//var/cache/pacman/pkg/lz4-1:1.9.2-2-armv7h.pkg.tar.xz.part: Invalid argument
nt
error: could not open file harddisk//var/cache/pacman/pkg/lz4-1:1.9.2-2-armv7h.pkg.tar.xz.part: Invalid argument
nt
error: could not open file harddisk//var/cache/pacman/pkg/lz4-1:1.9.2-2-armv7h.pkg.tar.xz.part: Invalid argument
nt
```



```
rgument
warning: failed to retrieve some files
error: failed to commit transaction (failed to retrieve some files)
Errors occurred, no packages were upgraded.
==> ERROR: Failed to install packages to new root
```

harddisk/ is the mounted sd card.

Yes, did the download. It has space in the name.

pacstrap automatically downloads, saves at places and installs. They're saved at the above shown directories in prompts.

#6 - 2020-12-09 05:42 AM - bill-auger

the errors in that last post are probably unrelated - those look like orphans of incomplete downloads - i would delete those

```
# rm /mnt/var/cache/pacman/pkg/*.part
```

#7 - 2020-12-09 05:50 AM - infinite_recursion

There are no part files in that directory. My guess is that it downloads part files(doesn't save it exactly, maybe it is in RAM. Don't know the exact procedure for download.), can't process it further because pacman's internal package name has : and part file doesn't, so it throws an error.

#8 - 2020-12-09 05:58 AM - infinite_recursion

I downloaded them, renamed and tried to move but : is an invalid character for filenames. Maybe there should be something else instead of : . Could you guys reproduce the error?

#9 - 2020-12-09 08:00 AM - bill-auger

i beleive that iceweasel is misbehaving in that way; but that is completely unrelated to pacman/pacstrap - you can see that those *.part filenames do not contain a space - i am quite certain that the completed download would also not contain any spaces - so that log is probably not related to the original topic

i am still confused what exactly is the problem - what leads you to beleive that there are packages on your system with spaces in the file name? - can you show an error message which demonstrates the original problem?

also, you did not answer grizzlyuser's question

did you download some packages using a web browser?

which files did you download manually?

why did you download *any* files manually?

there is a wiki article describing how to bootstrap ARM from x86 - did you follow that?

https://wiki.parabola.nu/Building_armv7h_packages_on_a_x86_system#Static_QEMU_.28option_1.29

#10 - 2020-12-09 09:36 AM - infinite_recursion

I answered grizzlyuser. pacstrap downloads it automatically. The ones I downloaded that were with space are from iceweasel.

I was trying to move the downloaded files to /cache/pacman/pkg/ dir and then execute pacstrap to skip the downloading part but I can't move something with ':' in it's name.

Yes, I followed those steps. I've done this before. Haven't faced issues.

Freedora isn't available for armv7. guix and parabola are the only 2 dists for libre armv7 OS installation with parabola being the easier one to install.

#11 - 2020-12-09 09:44 AM - bill-auger

i can not reproduce any such error with pacstrap - its more of a subtle question that probably no one has asked before - "is it possible to pacstrap ARM/nonsystemd from a x86/nonsystemd host" - we discussed that on IRC, and it probably is not possible in the direct way - the wiki guide i linked above requires a systemd service - probably a service file would need to be made for openrc - there are other ways to accomplish it though, such as booting the parabola ARM QEMU image with QEMU, with the real target disk passed-through to the machine

if you can offer "steps to reproduce" the error, it would be important to address; but otherwise i think we should close this ticket, or categorize it as a feature request - if there is any "strictly a bug" in this BR, it would be "iceweasel mangles downloaded filenames"; but i dont think that is a bug either,

just undefined behavior - i did some experiments nonetheless

in the rsync repo itself, the filename is:
mythes-hu-1:1.0-1-any.pkg.tar.xz

':', '/', and '@' are technically reserved characters for URLs -
i believe that it is undefined behavior for these (or spaces) to
exist in URLs - ie: serving such files defers it to the client
to handle them in any arbitrary way - the conventional solution
has always been to simply avoid having those chars (or spaces)
in any filenames which are accessible via URL; or the least, for
the server to process the filename upon delivery to clients -
iceweasel correctly presents the URL with an escaped "URL
entity" (%3A) - eg:
<http://quantum-mirror.hu/mirrors/pub/archlinux/pool/community/mythes-hu-1%3A1.0-1-any.pkg.tar.xz>

... but downloading with iceweasel produces a file with a space
in the filename:

```
$ ls mythes-hu-*  
'mythes-hu-1 1.0-1-any.pkg.tar.xz'
```

this could be seen as a bug in both the client and/or the server
IMHO; but again i believe that there is no standard protocol
defined for this - as the standard assumes that admins will
take steps to avoid it, if they want to ensure that clients will
not mangle the filename arbitrarily; it is really not a bug in
any client, even if the client rejected the download entirely -
i tried with a few more clients - the behavior is inconsistent,
as i expected

epiphany also mutates the filename, producing:
mythes-hu-1_1.0-1-any.pkg.tar.xz

wget, elinks, and qutebrowser handle this as expected, producing:
mythes-hu-1:1.0-1-any.pkg.tar.xz

of those, IMHO the choice of the space char made by mozilla is
the worst of all possible ASCII chars to choose; but its really
not a bug; just a poor design decision - i suppose that we could
patch that to behave like elinks and qutebrowser and/or ask
mozilla to do so

#12 - 2020-12-09 09:58 AM - infinite_recursion

I tried with "wget <https://mirror.grapentin.org/parabola/core/os/armv7h/lz4-1%3A1.9.2-2-armv7h.pkg.tar.xz>"
Works fine.

I checked all the packages here <https://mirror.grapentin.org/parabola/core/os/armv7h/>
I'm facing problems in every package having ':' that has to be installed. Others work fine. Was this something introduced recently?

#13 - 2020-12-09 10:10 AM - bill-auger

- Status changed from unconfirmed to info needed

we would need you to offer some explicit "steps to reproduce" the error - no one is able to reproduce the behavior you see, given only the vague
description in this BR - it can not be considered as a bug until then

i tried this from an installed systemd system, a systemd LiveISO, and an openrc LiveISO - in all cases, all packages were downloaded successfully - it
fails while executing the install hooks; but that is not related to any package filenames - pacstrap would fail before the hooks stage, if there were a
problem with any of the packages

```
# pacstrap -C /usr/share/pacman/defaults/pacman.conf.armv7h /mnt base libelogind
```

there is clearly more information about your system, and your procedure that is unknown - for example, using
/usr/share/pacman/defaults/pacman.conf.armv7h would not have either nonsystemd or kernels enabled - at the very least, you must have modified
that - more information would be needed to investigate this any further

#14 - 2020-12-09 10:26 AM - infinite_recursion

Forget about base libelogind. Just try this single package lz4. I'm having problems. So what info do I send.

Information is far from vague. It is absolutely obvious that ':' is creating this error and the 10 prompts show different attempts to get the file from 10
mirrors. I don't get where the part files come from? I don't have any in cache/pacman/pkg directory.

Where is pacstrap script in parabola?

Like the one below?

<https://git.archlinux.org/arch-install-scripts.git/tree/pacstrap.in>

#15 - 2020-12-09 10:35 AM - bill-auger

it is not obvious - you have not shown any error message which would indicate that pacstrap has downloaded a file with a space in its filename; and you have not given instructions for how someone else may try to see a similar error for one's-self - you gave only one example command - that command does not produce the same result on any parabola system on which i tried it; and it does not indicate a file with a space in the name

maybe something is borked with your system, maybe your target disk has a peculiar format; but even that is not clear, because you have not shown evidence of the problem, as you described it - those broken downloads could indicate some problem; but probably not the problem which you are describing - we dont even know which variant of parabola that your host system is running

one simple thing i can suggest is to download the parabola systemd LiveISO, and try the wiki instructions from the clean live system - if that fails, at least then we would know that it is not a problem with the host system; because that works for me

for example, an excellent bug report may look something like this:

system info:

parabola x86_64/nonsystemd

steps to reproduce:

1. download parabola ISO systemd/CLI 2020-08
2. boot ISO
3. enable [nonsystemd] and [kernels] in /usr/share/pacman/defaults/pacman.conf.armv7h
4. format the target disk with EXT4
5. mount the target disk on /mnt
6. run this command: # pacstrap -C /usr/share/pacman/defaults/pacman.conf.armv7h /mnt lz4

expected result: pacstrap should install the armv7h 'lz4' package into the chroot successfully

actual result:

```
# pacstrap -C /usr/share/pacman/defaults/pacman.conf.armv7h /mnt lz4
==> Creating install root at /mnt
==> Installing packages to /mnt
:: Synchronizing package databases...
nonsystemd is up to date
libre is up to date
kernels is up to date
core is up to date
extra is up to date
community is up to date
pcr is up to date
resolving dependencies...
looking for conflicting packages...
```

```
Packages (6) filesystem-2020.08.21-1.parabola1  glibc-2.32-2  iana-etc-20201012-1  linux-libre-api-headers-5.3
.l1_gnu-1
            tzdata-2020d-1  lz4-1:1.9.2-2
```

Total Installed Size: 47.79 MiB

:: Proceed with installation? [Y/n]

```
(6/6) checking keys in keyring [#####]
#####] 100%
(6/6) checking package integrity [#####]
#####] 100%
(6/6) loading package files [#####]
#####] 100%
(6/6) checking for file conflicts [#####]
#####] 100%
(6/6) checking available disk space [#####]
#####] 100%
:: Processing package changes...
(1/6) installing linux-libre-api-headers [#####]
#####] 100%
(2/6) installing tzdata [#####]
#####] 100%
(3/6) installing iana-etc [#####]
#####] 100%
(4/6) installing filesystem [#####]
#####] 100%
```

```
(5/6) installing glibc [#####] 100%
Optional dependencies for glibc
    gd: for memusagestat
(6/6) installing lz4 [#####] 100%
call to execv failed (Exec format error)
error: command failed to execute correctly
```

#16 - 2020-12-10 12:58 PM - infinite_recursion

Could you guys try "pacstrap -C /usr/share/pacman/defaults/pacman.conf.armv7h /mnt lz4" on a FAT32 partition? /mnt should be FAT32

#17 - 2020-12-10 01:23 PM - eschwartz

Don't use the crippled vfat filesystem format, it will eat your filenames.

#18 - 2020-12-10 01:29 PM - eschwartz

See also https://bugs.archlinux.org/task/24270?string=fat32&project=3&search_in_comments=1&search_in_details=1

#19 - 2020-12-10 01:37 PM - infinite_recursion

Ok. I used FAT32 because it was written parabola doesn't support ext4 fully at time of writing the ARMv7 installation guide. Please close this. Thanks.
If FAT32 is supported by parabola, this is a bug.