



Retrace Server

Michal Toman <mtoman@redhat.com>

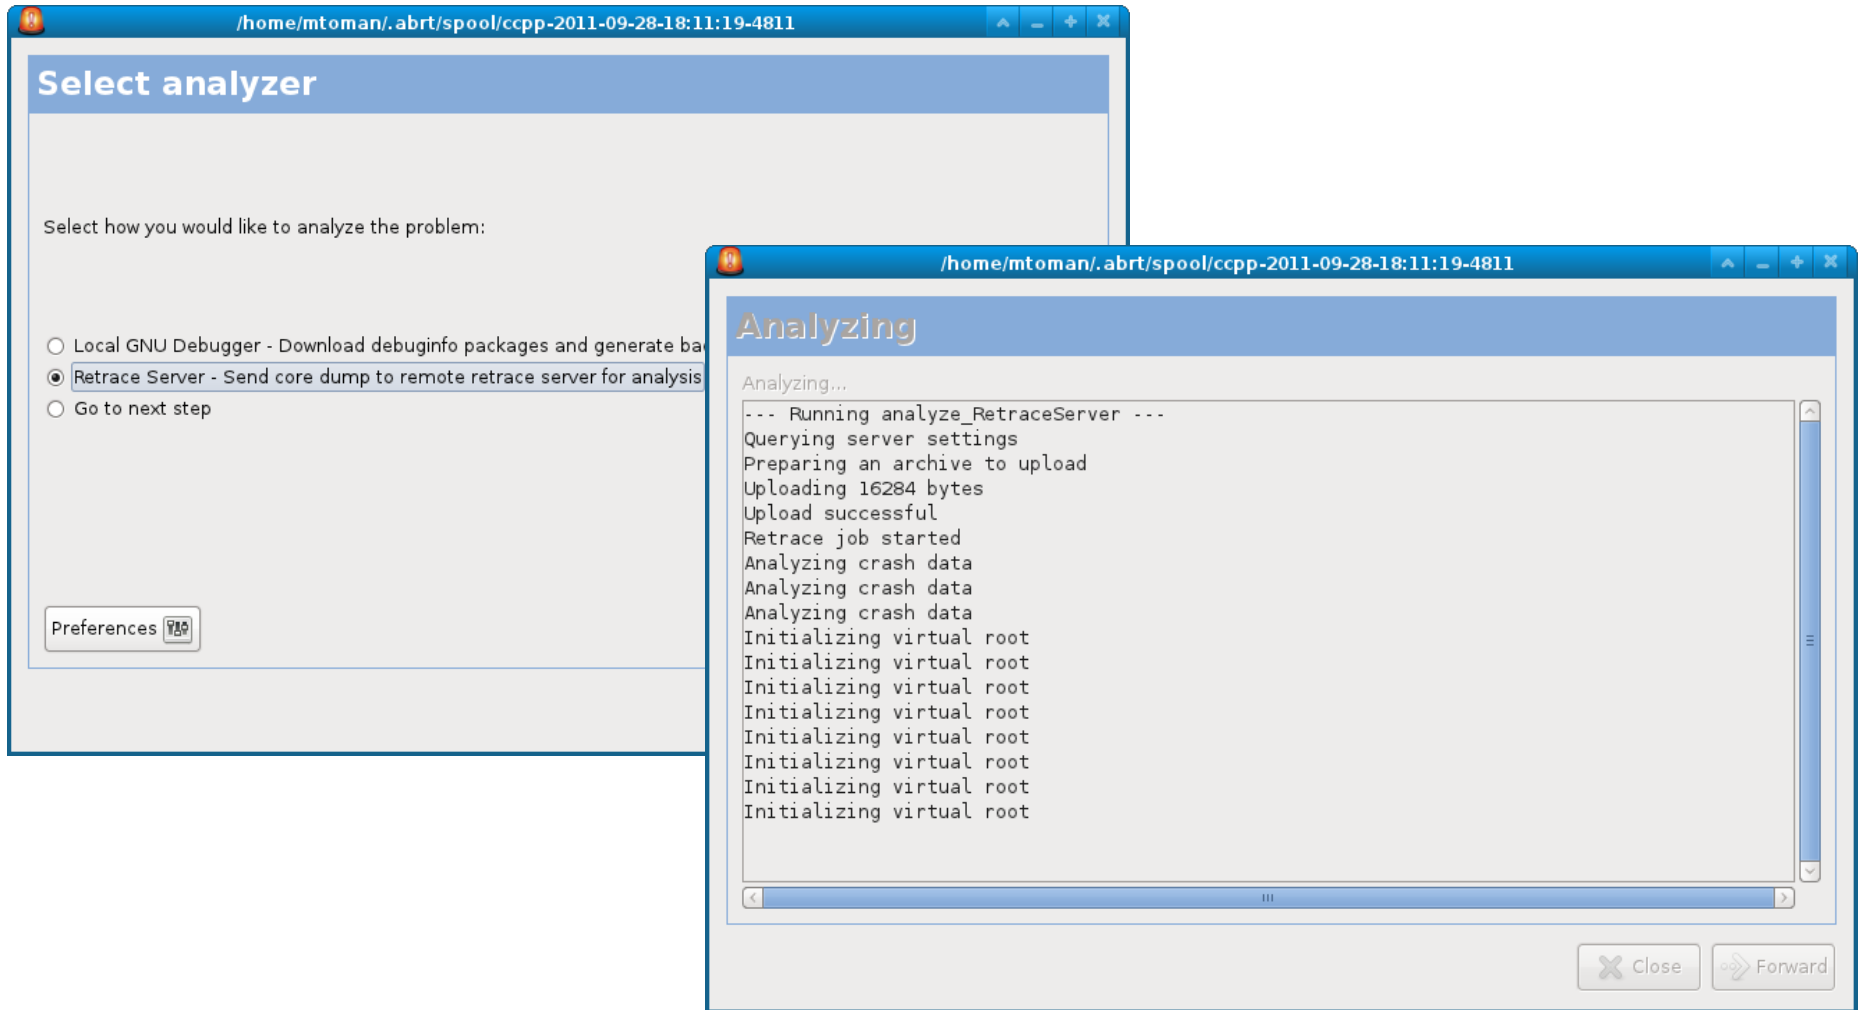
FUDCon Milan, 2011-10-01

Agenda

- What is Retrace Server
- Why to use
- How does it work
- Current status in Fedora
- Future plans



What is Retrace Server

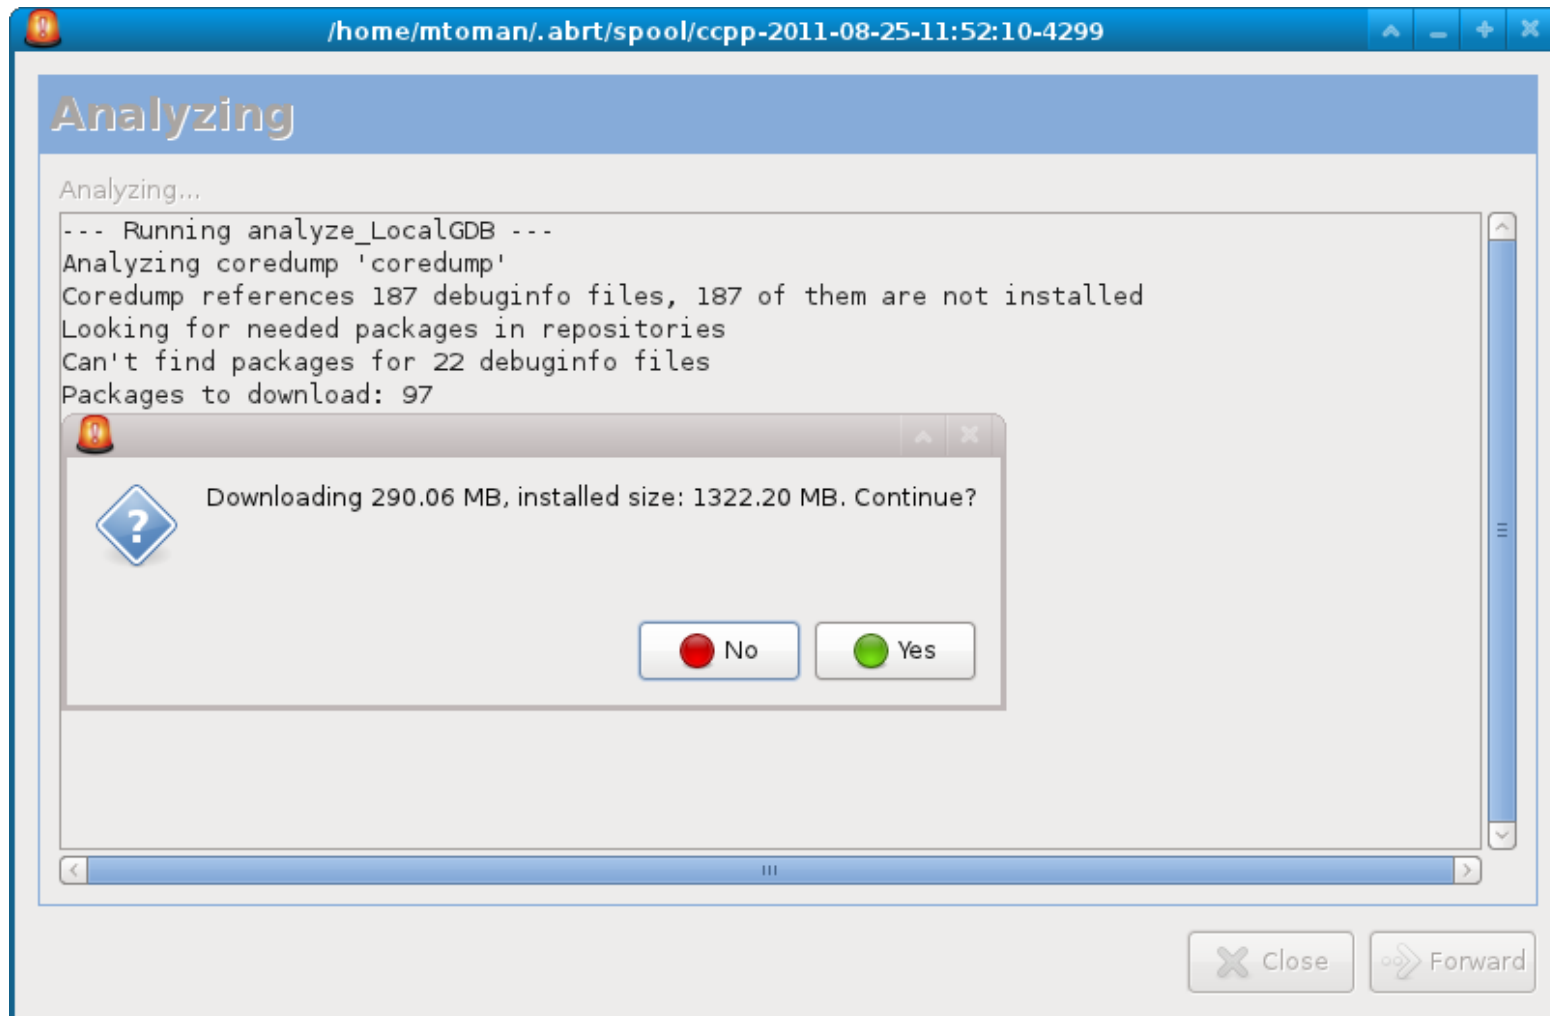


What is Retrace Server

- Analyzer for C/C++ coredumps
- Started in September 2010, part of ABRT project
- Stand-alone service, HTTP API
- Part of Fedora infrastructure
- Retrace Server is NOT a backtrace analyzer!



Why to use Retrace Server



Why to use Retrace Server

- No need to install debuginfo packages locally
 - HDD space (SSD users)
 - Network bandwidth
 - Time
 - Unneeded data
- Higher chance to get useful backtrace
 - Unavailable debuginfo
 - Uninstalled applications
 - Updated applications



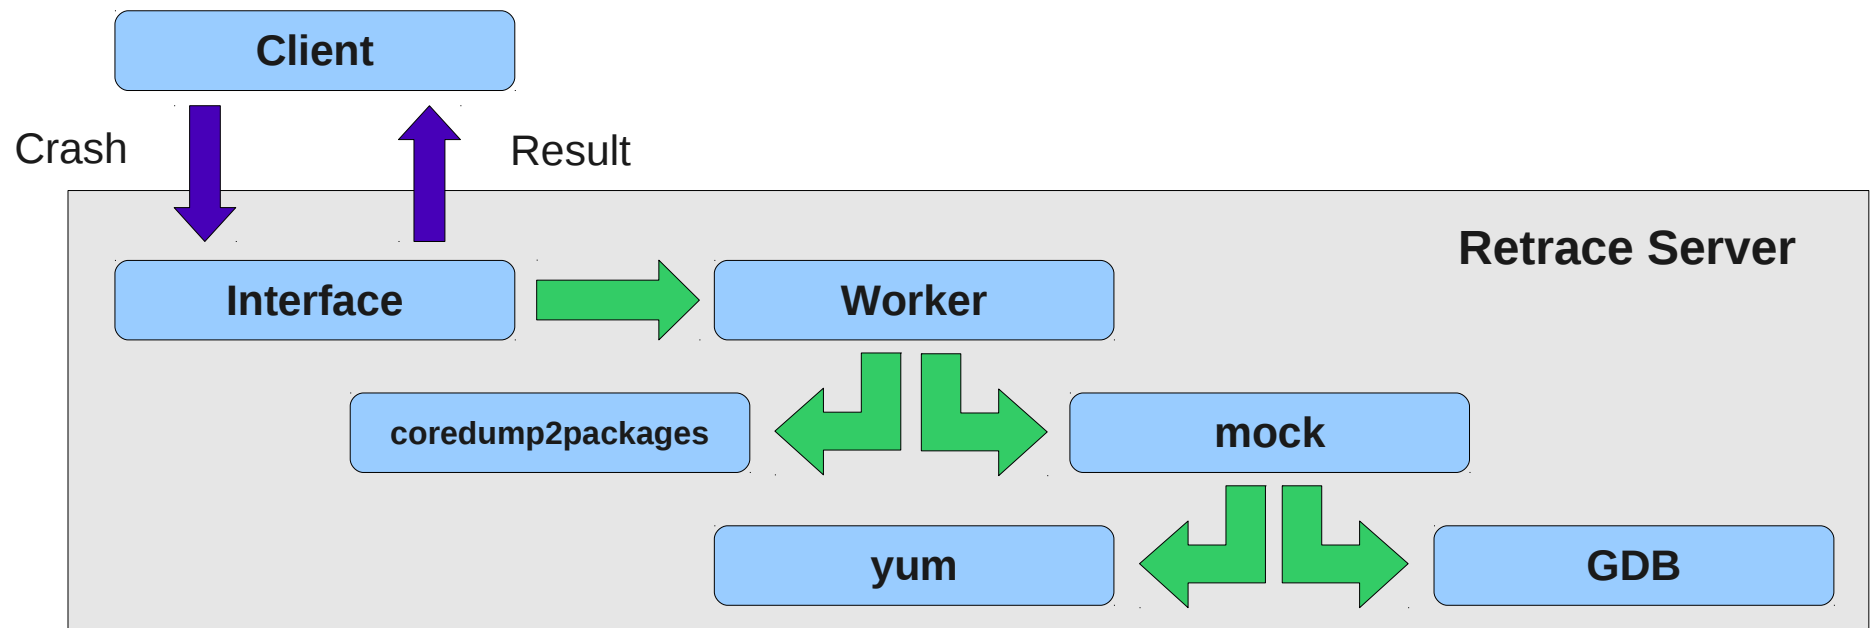
How does it work

- Set of Python/WSGI scripts
- HTTPS – coredumps may contain sensitive data
- 4 main parts
 - Interface (httpd, mod_wsgi)
 - Worker (mock, gdb)
 - Repository synchronizer (yum, rsync)
 - Garbage collector

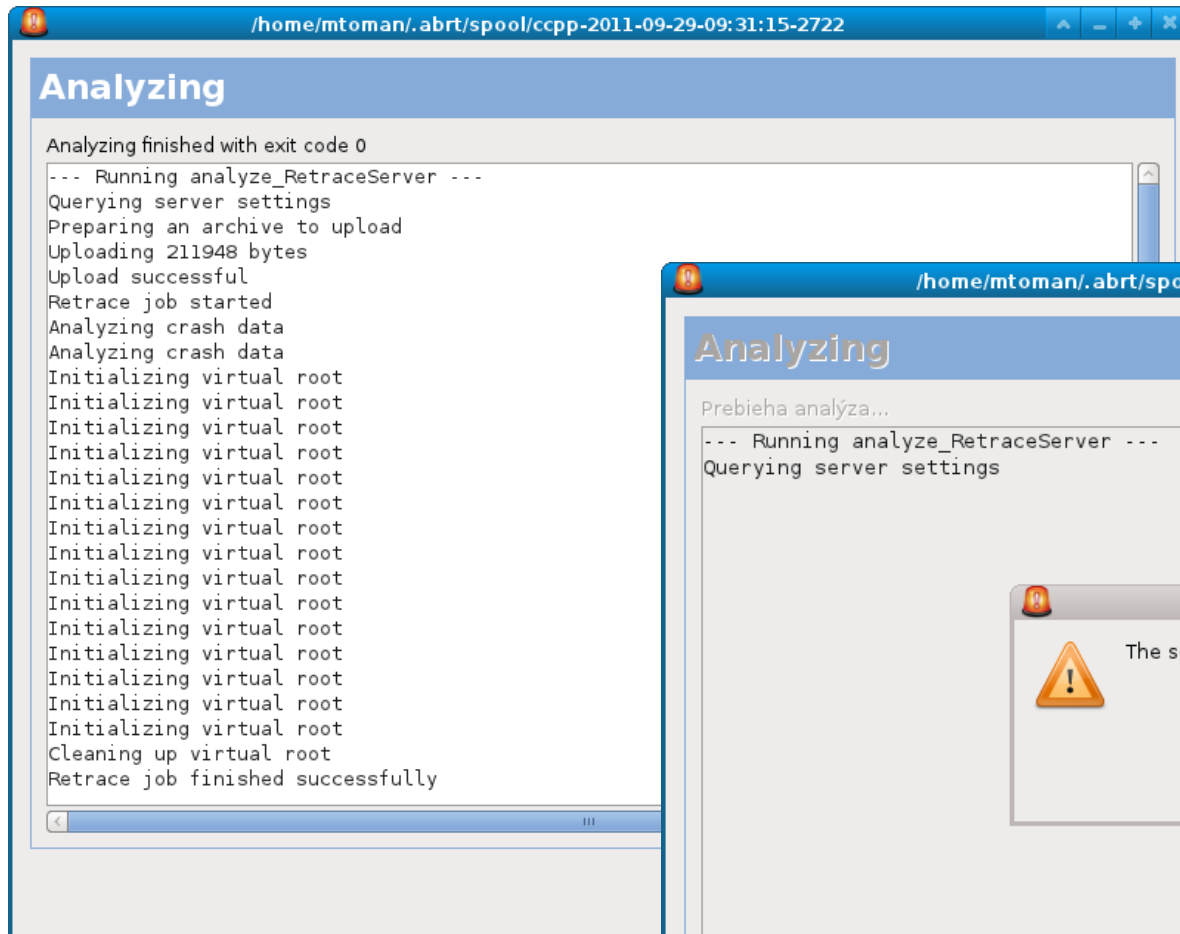


How does it work

- Client uploads the crash to the server
- Retrace Server prepares a chroot and runs GDB
- Client periodically asks for status
- Once finished client downloads the result

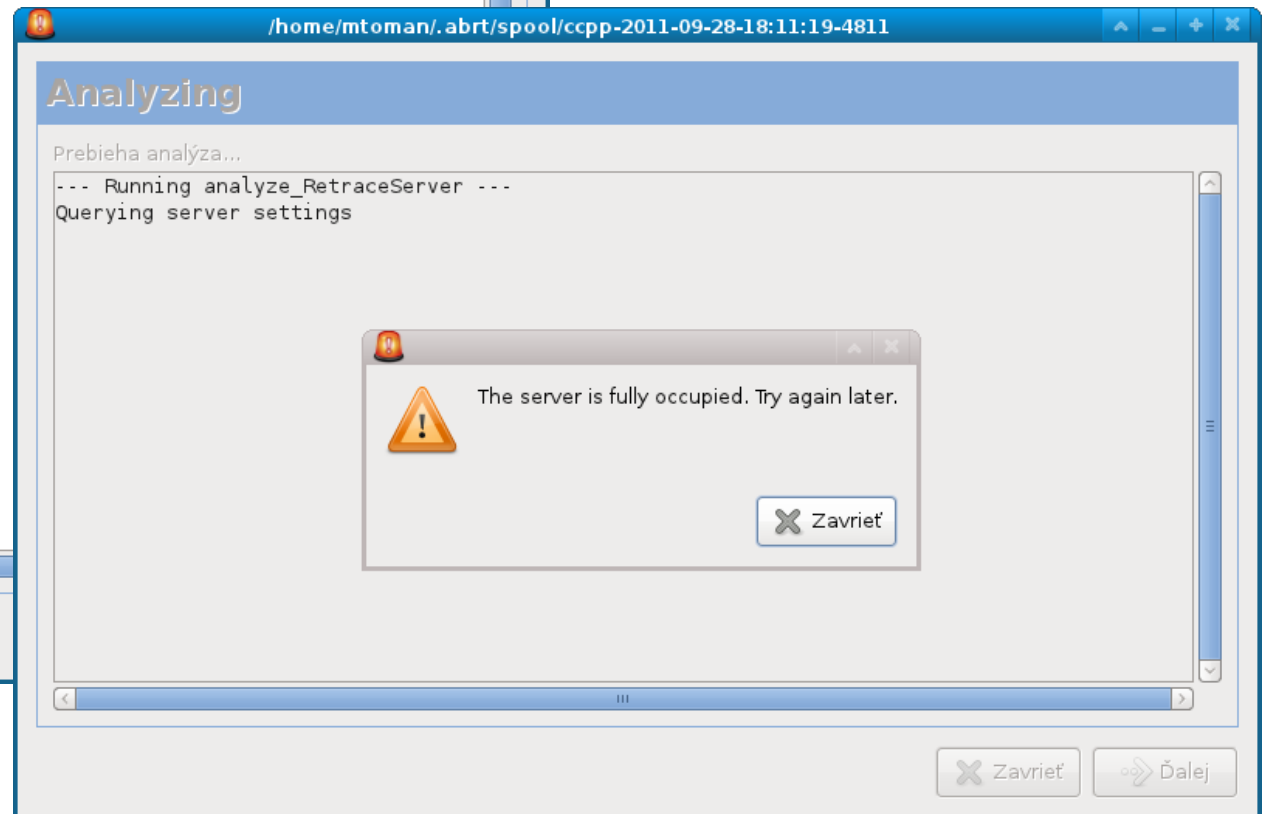


Current status in Fedora



The terminal window shows the output of the Retrace analyze_RetraceServer command. The process starts with querying server settings, preparing an archive, and uploading 211948 bytes. It then starts a Retrace job, analyzes crash data, and initializes a virtual root 15 times. After cleaning up the virtual root, the job finishes successfully with exit code 0.

```
Analizing finished with exit code 0
--- Running analyze_RetraceServer ---
Querying server settings
Preparing an archive to upload
Uploading 211948 bytes
Upload successful
Retrace job started
Analyzing crash data
Analyzing crash data
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Initializing virtual root
Cleaning up virtual root
Retrace job finished successfully
```



The terminal window shows the output of the Retrace analyze_RetraceServer command. It starts with querying server settings, but then an error dialog box appears: "The server is fully occupied. Try again later." The dialog box has a "Zavriet" (Close) button. The terminal window also has "Zavriet" and "Dalej" (Next) buttons at the bottom.

```
Prebieha analýza...
--- Running analyze_RetraceServer ---
Querying server settings
```

The server is fully occupied. Try again later.

Zavriet

Zavriet Dalej




Current status in Fedora

- retrace01.fedoraproject.org
 - Active (alias to retrace.fedoraproject.org)
 - Virtual machine, not enough storage
 - Often full (max 3 parallel jobs)
 - Only supports F15
- retrace02.fedoraproject.org
 - Physical machine
 - Supports F15, F16, rawhide
 - To replace retrace01 soon (before F16 release)




Current status in Fedora



Current status in Fedora

Most retraced packages

2011-08-25 - 2011-09-29



Current status in Fedora

- Known problems
 - Debuginfo packages
 - Completely missing
 - Partially generated
 - Wrong symlinks
 - yum ignores specific versions
 - Speed



Future plans

- Get as close to 100% reliability as possible
- Be able to retrace vmcores
- Involve statistics
- Optimize speed
- <https://fedorahosted.org/abrt/wiki/AbrtRetraceServer>





Q&A

mtoman@redhat.com

kklic@redhat.com

crash-catcher@lists.fedorahosted.org