

Update Accelerator revisited

Post Reply Search this topic... 91 posts 1 2 3 4 5 ... 7

Update Accelerator revisited

by BeBiMa March 27th, 2015, 11:39 pm As announced some days before, I've made some enhancements to Updatexlrator.

Main goal Updatexlrator has grown more and more. There were added new update sources (software vendors), the URLs of existing sources were updated too.

This produced the "If mountains of Axel", not easy to maintain and to verify. Each addition increases complexity. Another aspect of complexity is the abstraction of a set of update URLs to a regular expression.

Therefore some time ago the idea was born to simplify this process:
- Addition are not made in the source, but in a description file for these new sources
- The user must not compute a regular expression for the URLs. Only the URLs must be given, the regex is computed by IPfire
- Not all sources are used by all users. Thus a kind of selection should be possible.
- The decision whether the updatexlrator is responsible for handling the URL should be quick. Some sort of history of update requests could help to speed up the process.

Realization (till now) The update sources (vendors) are described by files in directory /var/ipfire/updatexlrator/sources. Filenames is the vendor, e.g. Microsoft. Each file contains positive matches, negative matches and the source type.

How to find them? Updatexlrator (file /usr/sbin/updatexlrator) contains if statements with the schematic (=~ means "matches", != means "doesn't match")

```
CODE: SELECT ALL
if ( $url =~ m@<exp 1>@i |
    $url =~ m@<exp n>@i
&&($url != m@<exp n+1>@i |
    ...
    $url != m@<exp m>@i ) {
    $url = check( ...,$url, ..., <vendors>, <type> )
}
```

From this code sequence <exp 1> ... <exp n> are the positive matches and (if existent) <exp n+1> ... <exp m> are the negative matches, <type> is either \$unique or \$mirror.

Thus the descriptive file contains a couple of lines adhering to the syntax <command>,<expression> with the meaning

- <command>:
o p ... regex that should match as whole (^<expression>\$)
o P ... regex that should match (without ^\$)
o n ... regex that should not match as whole (^<expression>\$)
o N ... regex that should not match (without ^\$)
o a ... URL that should match (accept)
o d ... URL that should not match (deny)
o t ... type of source (unique, mirror)

- <expression>:
o rexp for p,P,n,N commands (our positive and negative matches from above)
o URL for a,d commands
o type for t command (u=unique, m=mirror)

The p and n commands take the <exp i> from our code segment (without ^ and \$ brackets for ease of use!). These are the "normal" commands (till now). Commands a and d allow to add plain URLs (no knowledge of regex necessary, I hope).

This database of defined vendors is compiled to data structure holding one regex for positive match and negative match for each vendor. The condition can be reduced by rules of logic to

```
CODE: SELECT ALL
if ( ( $url =~ <p_match> ) & ( $url != <n_match> ) ) {
    check( ..... );
}
```

This is iterated over all know vendors. The compilation is done by call of

```
CODE: SELECT ALL
/var/ipfire/updatexlrator/bin/sources.pl
```

(This is inspired by the processing for language files in IPfire). This script generates a list of known vendors also (file /var/ipfire/updatexlrator/sources-used).

If you add new vendor descriptions, you should add an image file in /srv/web/ipfire/html/images/updbooste also. The web interface will honor it.

For names I tried to follow the rule: vendor names begin with upper case, the image file names contain the vendor in lowercase. This is the state found in the running system.

For acceleration of the search process the list of known vendors is resorted by the updatexlrator process: If an URL yields an update request, the corresponding vendor is put at top of the list. This seldomly or never used vendors slowly drop to the bottom.

Further one can disable vendors, which should not be cached or aren't used at all by moving the definition file from /var/ipfire/updatexlrator/sources to /var/ipfire/updatexlrator/sources_disabled and rerunning the compilation process.

The file /var/ipfire/updatexlrator/sources_used represents a somewhat recent state of the internal list, therefore it is sorted according to the strategy from above, but not alphabetical. It is located in the tmpfs /var/log/rdd and saved/restored at restart.

While there are several updatexlrator processes (defined in proxy configuration), each process has its own local copy. Further each process sends each vendor found to a server process, which sorts the system wide list sources_used. A synchronisation of local lists and system wide list is done at restart of updatexlrator only (yet).

Installation for test Download the attached file, transfer it to your IPfire system. Unpack and install with

```
CODE: SELECT ALL
tar xzf updx_new.<version>.tgz
cd updx_new
./updx_new install
```

restart updatexlrator

Uhhh, a bit lengthy this explanation, but the process for realization wasn't that short. 😊

Please test and ask.

-Bernhard

P.S. There are some other things to optimize for this module. I'm working on it. This should be a first version only.


P.P.S. Thx to Nöman, who initiated the "revisit" to the updatexlrator "construction site" and added a bunch of new vendors. I hope I've integrated all without errors. (viewtopic.php?p=508&t=12868)

EDIT: correct installer. v 0.9-2 (28-03-15 19:00) installer creates perl sub dir, deleted ESET

EDIT: new version v 1.0-0 (16-04-15 18:15) disable by moving to directory sources_disabled, some new vendors v 1.0-1 (19-04-15 20:11) typos corrected. Corrected WUI v 1.0-2 (20-04-15 00:25) Processing of HTTP error responses in download, TODO: check in updxlrator v 1.0-3 (21-04-15 13:45) Corrected addition to an empty repository, illegal end of line chars in source definition v 1.0-3a (22-04-15 02:15) fixed rights / ownership of files v 1.1-0 (13-05-15 19:03) New database for meta data, new structure of installer (see post http://forum.ipfire.org/viewtopic.php?f=935&p84935) v 1.1.-0a fixed typos 😊 v 1.1-1 (15-05-15 23:56) Directories in repository for vendors are generated if not existent yet.

ATTACHMENTS updx_new.1.1-1.tgz MD5: 7487adec4e3e4701af27cbe5361a7de (90.59 KiB) Downloaded 1115 times updx_new.1.1-0a.tgz MD5: 5b5976c15ec088505c51685bc5f78f9e (90.47 KiB) Downloaded 360 times updx_new.1.1-0.tgz MD5: 86532efa16dc9139e5772a500c9855d8 (90.46 KiB) Downloaded 359 times updx_new.1.0-3a.tgz (82.27 KiB) Downloaded 400 times updx_new.1.0-3.tgz (82.19 KiB) Downloaded 401 times

Last edited by BeBiMa on May 15th, 2015, 10:00 pm, edited 20 times in total.



Re: Update Accelerator revisited by Nöman March 28th, 2015, 3:36 am I'm having issues with the installer script. It keeps giving several errors after line 46.

Re: Update Accelerator revisited by BeBiMa March 28th, 2015, 9:09 am Sorry, didn't check this piece of software 😊 I've added the corrected version to the top post.



Re: Update Accelerator revisited by Nöman March 28th, 2015, 4:32 pm Thank you, most of the installer issues seem fixed. There is only one thing. You need to make a directory called /usr/lib/perl5/5.12.3/Regexp/ before you try to copy Assemble.pm to it.

Re: Update Accelerator revisited by Nöman March 28th, 2015, 5:00 pm I've tested update accelerator after I made the new directory, It works well. I also noticed something. In the maintenance page, there is still ESET as a supported vendor. I removed it in the last revision because ESET uses authentication which the update accelerator couldn't use.

Re: Update Accelerator revisited by BeBiMa March 28th, 2015, 6:06 pm Nöman wrote: Thank you, most of the installer issues seem fixed. There is only one thing. You need to make a directory called /usr/lib/perl5/5.12.3/Regexp/ before you try to copy Assemble.pm to it. I also noticed something. In the maintenance page, there is still ESET as a supported vendor. I removed it in the last revision because ESET uses authentication which the update accelerator couldn't use

Thx. Done. BTW: The file /var/ipfire/updatexlrator is not really updated by the compilation process. New vendors are added, but deleted vendors are not deleted (yet). Just remove the file before recompilation or delete the vendor with an editor.




Re: Update Accelerator revisited by BeBiMa April 11th, 2015, 2:22 pm Some state about the work. I'll try to eliminate the ugly issues with weird file names. This is a bit more complex than expected. All files belonging to update accelerator must be inspected. Having problems to update my Avira Free in the last days, I detected a problem with source definitions. Many sources/vendors are marked as "mirror" at the moment (in standard accelerator and extensions from Nöman), in most cases this doesn't really matter, but for Avira it did!

The classifications used are:
- unique ... file is identified by the URL as a whole (a.b/v1/prog.exe, a.b/v2/prog.exe)
- mirror ... file is identified by its name (a.b/upd/prog_v1.exe, a.b/upd/prog_v2.exe)

I can't really check, whether all classifications are right. Therefore it would be nice to get reports from users of the various update sources. (Avira is corrected 😊)

- Bernhard




Re: Update Accelerator revisited by BeBiMa April 11th, 2015, 2:22 pm Some state about the work. I'll try to eliminate the ugly issues with weird file names. This is a bit more complex than expected. All files belonging to update accelerator must be inspected. Having problems to update my Avira Free in the last days, I detected a problem with source definitions. Many sources/vendors are marked as "mirror" at the moment (in standard accelerator and extensions from Nöman), in most cases this doesn't really matter, but for Avira it did!

The classifications used are:
- unique ... file is identified by the URL as a whole (a.b/v1/prog.exe, a.b/v2/prog.exe)
- mirror ... file is identified by its name (a.b/upd/prog_v1.exe, a.b/upd/prog_v2.exe)

I can't really check, whether all classifications are right. Therefore it would be nice to get reports from users of the various update sources. (Avira is corrected 😊)

- Bernhard



Re: Update Accelerator revisited by Nöman April 11th, 2015, 3:51 pm Hi, I have been trying adding two sources to update accelerator. The software are Libreoffice and Openoffice. The issue is that Libreoffice seems to use metadata to redirect to a mirror. Openoffice redirects to a sourceforge page which uses mirrors. What would the URL source look like?

Re: Update Accelerator revisited by BeBiMa April 11th, 2015, 4:34 pm What kind of updates do you want to cache? Libreoffice usually does an update as a whole. The user can fetch it from the Libreoffice sites directly (possibly redirected) or from one of the mirrors. I personally prefer to get it as a torrent (an unpredictable bunch of "mirrors"). EDIT: An inspection of the download pages of Libreoffice shows URLs adhering to the RegExp ^http://(.*/libreoffice/|stable/|(.*/))LibreOffice_(.*)\.msi\$|tar.gz\$ The file name seems to be unambiguous (e.g. LibreOffice_4.4.2_Win_x86.msi), thus we can classify as "mirror". The endings must be completed for all supported OSs.




Re: Update Accelerator revisited by Nöman April 11th, 2015, 8:55 pm Thank you. I corrected the unique vs mirror entries. I made Avira support Professional Updates - I think it has different definitions between the free and paid versions. I added Libreoffice support.

Re: Update Accelerator revisited by bloater99 April 13th, 2015, 2:25 pm LibreOffice does not have an automated update system as far as I know. You either have to manually apply the update or create a custom .msi and do it through Group Policy. I suppose, though, you can still cache the download so others who are also manually downloading can get it from the cache.

Re: Update Accelerator revisited by bloater99 April 13th, 2015, 4:37 pm I have not tested this as my only IPfire install is in a production setting. But I do use the default update accelerator and wondered if in your version have you separated Firefox and Thunderbird? Currently it just says Mozilla with no indication of which Mozilla package. I don't know if it's possible, but it would be nice to some day be able to see separate entries for Firefox and Thunderbird 😊

Re: Update Accelerator revisited by BeBiMa April 13th, 2015, 5:10 pm If you can separate the download URL of Firefox and Thunderbird, no problem. My concept allows you define "packages" (now named "vendors" or "sources") with their dedicated set of URLs. Therefore if you want to divide the Mozilla repository into two, it is your choice. But I think we'll support the big bunch "Mozilla" only. About the LibreOffice you are right. But Update Accelerator is not a cache for automated download only, but a cache for multiple downloads in the local network. BTW: If it would cache "updates" only, there weren't the problem with blanks in file name. The "prototype example" is the setup file for Firefox. 😊



Re: Update Accelerator revisited by BeBiMa April 16th, 2015, 4:29 pm Finally, the first version which implements most ideas. Main change is that vendors are disabled/enabled by moving between directories sources and sources_disabled. Directory sources_disabled is only a storage for unused vendors. All vendors in sources are active! First trial to handle a repository distributed over several drives/partitions. Calculations for free/used disk space should work. Output in web interface isn't really nice (yet). EDIT: "Beautifulified" updatexlrator.cgi and description in wiki coming soon.



Re: Update Accelerator revisited by Nöman April 17th, 2015, 10:03 pm I tested the new version on a clean version of IPFire. Here is what I've found so far: The install script has a typo: (Don't forget to run /var/ipfire/updatexlrator/bin/sources.pl again!) The database seems to use the old, original sources rather than the new ones There seems to be an error with permissions. Many files seem to have their permissions changed. when I try to start the script I get:

```
CODE: SELECT ALL
/usr/sbin/updxlrator
Name "Updxsrc::VendorPipe" used only once: possible typo at /usr/sbin/updxlrator line 100.
Name "Updxsrc::match_all" used only once: possible typo at /usr/sbin/updxlrator line 125.
Name "Updxsrc::srcVendors" used only once: possible typo at /usr/sbin/updxlrator line 81.
Name "Updxsrc::srvVendors" used only once: possible typo at /usr/sbin/updxlrator line 97.
Can't locate /root/work/usrc in @INC (@INC contains: /usr/lib/perl5/site_perl/5.12.3/i586-linux-thread-multi /usr/lib/p
```

Last edited by Nöman on April 20th, 2015, 12:17 am, edited 1 time in total.

Post Reply 91 posts 1 2 3 4 5 ... 7

Return to "Development" Delete cookies All times are UTC


Update Accelerator revisited

Post Reply Search this topic... 91 posts

Re: Update Accelerator revisited

by BeBiMa » April 19th, 2015, 4:20 pm

Thank you very much. There were some typos and "optimizations" in the files. A fixed version is ready for download in the top post.




BeBiMa Profile: Posts: 2842, Joined: July 30th, 2011, 12:55 pm, Location: Mannheim

Re: Update Accelerator revisited

by BeBiMa » April 19th, 2015, 6:13 pm

Forgot the WUI file. Please download the updated version 1.0-1



BeBiMa Profile: Posts: 2842, Joined: July 30th, 2011, 12:55 pm, Location: Mannheim

Re: Update Accelerator revisited

by Nöman » April 19th, 2015, 7:51 pm

I tried the new version. It seems to pickup the new sources, but it keeps giving me these errors. /var/log/updatedxlrator/download.log

```
2015-04-19 15:19:20 [5876] Retrieving file for local cache: vlc-plugin-notify_2.1.4-0ubuntu14.04.1_amd64.deb
2015-04-19 15:19:20 [5909] Retrieving file for local cache: vlc-plugin-pulse_2.1.4-0ubuntu14.04.1_amd64.deb
2015-04-19 15:19:20 [5876] Download finished with result code: ERROR
2015-04-19 15:19:20 [5909] Download finished with result code: ERROR
```

I get .info files in /var/updatecache/download, and the webinterface is empty. Example File Contents: REMOTESIZE=0 CFMIRROR=1 STATUS=1 REMOTETIME=0 FILENAME=vlc_2.1.4-0ubuntu14.04.1_amd64.deb SRCURL=http://us.archive.ubuntu.com/ubuntu/pool/universe/v/vlc/vlc_2.1.4-0ubuntu14.04.1_amd64.deb VENDORID=linux

Nöman Profile: Posts: 299, Joined: July 26th, 2013, 1:56 am


Re: Update Accelerator revisited

by BeBiMa » April 19th, 2015, 9:33 pm

This source doesn't exit. Unfortunately the HTTP response wasn't analysed yet. My new Update Accelerator uses the remaining .info file comes from my change of the condition "download of file running" (was: updatefile exists, now: updatefile.info exists). Should be corrected with version 1.0-2.

The check in updxxlrator itself should be changed also. Is on my TODO list.

Last edited by BeBiMa on April 20th, 2015, 9:30 am, edited 1 time in total.



BeBiMa Profile: Posts: 2842, Joined: July 30th, 2011, 12:55 pm, Location: Mannheim

Re: Update Accelerator revisited

by Nöman » April 20th, 2015, 12:13 am

The logs say now that it detects the download. It just won't download the file. When you test it on your system does it download?


Nöman Profile: Posts: 299, Joined: July 26th, 2013, 1:56 am

Re: Update Accelerator revisited

by BeBiMa » April 20th, 2015, 9:29 am

If I try to download the file, I get an error "404 Not Found". This error is documented in download.log and the download is not started. At the moment the error is detected in program download, because the detection in module updxxlrator is missing. updxxlrator logs to cache.log, download logs to download.log.

Maybe there is a problem with a .info file in the directory /var/updatecache/download/linux. If it still exists, just delete it. My new Update Accelerator uses the existence of a foo.info file as condition "download for foo in progress". If this condition is true download does nothing more. Thus the invalid state (no download, but .info exists) isn't cleared.



BeBiMa Profile: Posts: 2842, Joined: July 30th, 2011, 12:55 pm, Location: Mannheim

Re: Update Accelerator revisited

by Nöman » April 20th, 2015, 11:13 am

I meant, is the update accelerator on your system able to cache any files. I keep trying to download different sources I know work. I keep just getting logs when the update accelerator recognizes a download. I just doesn't seem to do anything else. No .info files anymore nor download.log. When I install it on a fresh version of IPFire, I install the package and then configure the webproxy and restart the update accelerator. Is this the right way to install?


Nöman Profile: Posts: 299, Joined: July 26th, 2013, 1:56 am

Re: Update Accelerator revisited

by BeBiMa » April 20th, 2015, 11:34 am

Hi, surelyly Update Accelerator caches files on my system. Without this, I wouldn't have published the package. I'll check the sources in the package against my installation. And yes, the way you install it, should work (if you restart Update Accelerator by WUI). Perhaps there is a problem with rights and owners.

EDIT: Files are identical.




BeBiMa Profile: Posts: 2842, Joined: July 30th, 2011, 12:55 pm, Location: Mannheim

Re: Update Accelerator revisited

by BeBiMa » April 21st, 2015, 2:17 pm

There was a problem with the download of files, if there was no directory in the repository for the vendorid yet. /var/updatecache) but the partition of /var/updatecache/<vendorid>. With a non-existent directory the function return "garbage". This is fixed now (version 1.0-3). For non-existent directories the repository is chosen. This is the basic directory and must exist. Parts moved to another partition "leave" entry in /var/updatecache and therefore exist.

There is still a problem with the definition files. If you edit them on Windows they can use <CR><LF> as end of line marker. This is not corrected by the definitions compiler. Strings may end with <CR>, which gives faulty REs and types.



BeBiMa Profile: Posts: 2842, Joined: July 30th, 2011, 12:55 pm, Location: Mannheim

Re: Update Accelerator revisited

by RedneckMother » April 21st, 2015, 4:45 pm

Howdy! Thanks for your efforts. Please verify the u-g ownership and u+g+o permissions of the files... I found at least one discrepancy (for /srv/web/ipfire/cgi-bin/updatedxlrator.cgi) in v1.0-3 today.


Thank you, Paul

RedneckMother Profile: Posts: 94, Joined: June 21st, 2014, 1:34 am, Location: USA

Re: Update Accelerator revisited

by BeBiMa » April 21st, 2015, 5:34 pm

Will be fixed. Just checking against the "standard". EDIT: fixed after "return to standard" by installing Core Update 89




BeBiMa Profile: Posts: 2842, Joined: July 30th, 2011, 12:55 pm, Location: Mannheim

Re: Update Accelerator revisited

by BeBiMa » April 25th, 2015, 12:54 pm

Some update about the current state. In Bugzilla an issue with the statistics in the WUI is reported. Investigations showed problems if the number of files in the repository is great. State/statistics data are fold in separate files for each update file. Sampling these data needs much time (and memory space). Storing data in a data base saves much of this time. (Example for my system: Repository with 10714 files, implemented way: 130 seconds, with a data base: 6 seconds). Changing the Update Accelerator to the new concept will need some time. Errors found in the meantime will be fixed.

- Bernhard



BeBiMa Profile: Posts: 2842, Joined: July 30th, 2011, 12:55 pm, Location: Mannheim

Re: Update Accelerator revisited

by ckleuser » May 5th, 2015, 10:12 pm

How to clear the update accelerator cache without having to delete each file or waiting at least a week Any command line?

ckleuser Profile: Posts: 17, Joined: January 15th, 2015, 10:12 pm, Location: Florianópolis - SC / Brazil, Contact: [icon]

Re: Update Accelerator revisited

by zcworld » May 5th, 2015, 7:35 am

hi, first of all Love the new changes to the Update Accelerator its kicker but i love to see (not sure if just me or anyone else feels like this) like an que system so it like for updates when i did an windows update , there was 30 updates and it try to download all of them at once and well it slow the net access i down to a snail for a while to all of them where downloaded

and one qus, i cant work out how to add new sites to save download files from i need some help, i'm a bit lost , if you could send me a msg , maybe i can get my head around how it all works

zcworld Profile: Posts: 15, Joined: May 4th, 2015, 10:31 am

Re: Update Accelerator revisited

by BeBiMa » May 5th, 2015, 10:08 am

zcworld wrote: hi, first of all Love the new changes to the Update Accelerator its kicker

Nice to hear. Thx. zcworld wrote: but i love to see (not sure if just me or anyone else feels like this) like an que system so it like for updates

Don't understand, what you mean. zcworld wrote: when i did an windows update , there was 30 updates and it try to download all of them at once and well it slow the net access i down to a snail for a while to all of them where downloaded

That's normal behaviour. If a file is requested, that fits the defined "sources", a copy is downloaded to the repository. The download of the requesting client takes place in parallel. The requestor must not wait until the repository is filled. But you can limit the bandwidth/active time of download to the repository by Performance options "Lower CPU priority for downloads" and "Max. external download rate" in the WUI.

zcworld wrote: and one qus, i cant work out how to add new sites to save download files from

To add new "sources", follow these steps 1. sample the URLs, which are necessary to download the file(s) 2. build a regular expression for this set of URLs. With the new version you can even just use the pure set. 3. Chose an unambiguous "vendorid" 4. generate a new file in /var/ipfire/updatedxlrator/sources named "Vendorid" (first letter uppercase. Put in your list of regular expressions or URLs (syntax is described in the top post.) 5. Determine the type for this vendorid. "Unique" means the file is identified by the whole URL, "mirror" means file is identified by it's namet several URLs download the same file). Put the type into your vendor definition file. 6. run /var/ipfire/updatedxlrator/bin/sources.pl to update the definitions 7. restart Update Accelerator.

Bernhard



BeBiMa Profile: Posts: 2842, Joined: July 30th, 2011, 12:55 pm, Location: Mannheim

Update Accelerator revisited


Post Reply
Search this topic...
91 posts

Re: Update Accelerator revisited 66

by **BeBiMa** » May 5th, 2015, 10:11 am

kleuser wrote:
How to clear the update accelerator cache without having to delete each file or waiting at least a week 😊 Any command line?

At the moment the only way is to delete the contents of the directory /var/updatescache/<vendorid>. Maybe there are any ideas to solve this in more "elegant" way.



Re: Update Accelerator revisited 66

by **N0man** » May 5th, 2015, 10:58 am

zcworld wrote:
like an que system so it like for updates

I think he wants a queue system where each update is downloaded one-at-a-time in-order.

Currently the update accelerator downloads all updates at once. This can be an issue when 30 updates are trying to download at once.

Re: Update Accelerator revisited 66


by **BeBiMa** » May 5th, 2015, 11:38 am

Thanks for the explanation. 😊

Yes, you are right. I had this idea yet. At the moment each UpdxAcc process forks a download process for a new file. And WindowsUpdate may be very "aggressive" in requesting files simultaneously. Seems like the strategy "Download as much as we can get, then look what we need". A queue for download would help to reduce e.g. 60 downloads (30 to repository, 30 to client(s)) to 31 (1 to repository, 30 to clients).

I'll think of an implementation for this feature.

- Bernhard



New version 1.1 66

by **BeBiMa** » May 12th, 2015, 10:07 pm

After some struggle between intended and written source code 😊, I happily release version 1.1 of the new Update Accelerator for test.


What's new:

- the sorting of vendorids by usage is done at two places now. Each Updxlr process maintains it's own private copy. A server task maintains the overall list, getting information about chosen vendorids from the Updxlr processes through a pipe. At (re)start each process reads this global list.
- The meta data for each file in the update repository is now held in a data base. Thus the sampling of information for statistics and maintenance in the WUI should be faster (and hopefully more memory efficient).
- The installer was rewritten to use a more system conform structure. Install generates the new data base. Uninstall writes the information in the data base back to the file base meta data.

Further the naming of vendorids is more strict: Vendorids in the definition directories must start with an uppercase letter, allowed characters are letters, digits and '_'. Vendorids in the repository are all lowercase. Thus defined vendorids must be unambiguous after converting to lower case.

Next steps will be:

- More strict definition of repository contents. Only directories which are the lower case equivalent of a defined vendorid are part of the repository and can be managed by the WUI. Vendorids in directory sources_disabled are defined, but not used. This should speed up the repository file walk.
- Tidy up the sources to eliminate obsolete and double functions.
- Speed up the maintance page of the WUI



Re: Update Accelerator revisited 66

by **N0man** » May 13th, 2015, 2:00 am

I tried the new update.

The update accelerator doesn't seem to start.

When I run:
/usr/sbin/updxlrator

It gives:
Can't call method "as_string" on an undefined value at /var/ipfire/updatescache/bin/sources.pl line 182.

Re: Update Accelerator revisited 66

by **cardinszu** » May 13th, 2015, 6:11 am

same i got the same errors

Re: Update Accelerator revisited 66

by **BeBiMa** » May 13th, 2015, 10:45 am

Are there any messages, when you start /var/ipfire/updatescache/bin/?

What is the contents of /var/ipfire/updatescache/sources.pl?

The files in the installer archive are identical with my running version. And I've got no issues so far. Maybe I've got some file permissions wrong again. I'll check this next.



Re: Update Accelerator revisited 66

by **BeBiMa** » May 13th, 2015, 4:07 pm

😊 I think I've found the bug. 😊 The source definitions were not installed, due to a typo in the installer. New file is in the top post.



Re: Update Accelerator revisited 66

by **N0man** » May 14th, 2015, 11:43 am

I tried the new version 1-1-0

This is the error I get when I install it on a fresh IPFire Virtual Machine:

```
CODE: SELECT ALL
./updx_new install
copying files

converting the metadata to new format


generate cache of sources
Can't call method "as_string" on an undefined value at /var/ipfire/updatescache/bin/sources.pl line 182.

updatescache is installed.
You can disable update sources by moving definition files from
/var/ipfire/updatescache/sources to /var/ipfire/updatescache/sources_disabled.
(Don't forget to run /var/ipfire/updatescache/bin/sources.pl again!)
```

Re: Update Accelerator revisited 66

by **BeBiMa** » May 14th, 2015, 12:32 pm

@N0man: You've got a PM.



Re: Update Accelerator revisited 66

by **binky** » May 14th, 2015, 6:57 pm

I'm seeing exactly the same as N0man:

```
CODE: SELECT ALL
[root@ipfire updx_new.1.1-0]# ./updx_new install
copying files

converting the metadata to new format

generate cache of sources
Can't call method "as_string" on an undefined value at /var/ipfire/updatescache/bin/sources.pl line 182.

updatescache is installed.
You can disable update sources by moving definition files from
/var/ipfire/updatescache/sources to /var/ipfire/updatescache/sources_disabled.

(Don't forget to run /var/ipfire/updatescache/bin/sources.pl again!)
[root@ipfire updx_new.1.1-0]# ./var/ipfire/updatescache/bin/sources.pl
```

I only found ipfire today and I've got it all working as I'd hoped for my home network apart from the Update Accelerator not working, so I thought I'd try your 'revisited' version as you're asking for testers in a different thread and I'm only using it at home so absolutely no issue if I kill it 😊

Any suggestions?

Thanks,
Binky

N0man wrote:
I tried the new version 1-1-0


This is the error I get when I install it on a fresh IPFire Virtual Machine:

```
CODE: SELECT ALL
./updx_new install
copying files

converting the metadata to new format

generate cache of sources
Can't call method "as_string" on an undefined value at /var/ipfire/updatescache/bin/sources.pl line 182.

updatescache is installed.
You can disable update sources by moving definition files from
/var/ipfire/updatescache/sources to /var/ipfire/updatescache/sources_disabled.
(Don't forget to run /var/ipfire/updatescache/bin/sources.pl again!)
```




Re: Update Accelerator revisited 66

by **binky** » May 14th, 2015, 6:57 pm

I only found ipfire today and I've got it all working as I'd hoped for my home network apart from the Update Accelerator not working, so I thought I'd try your 'revisited' version as you're asking for testers in a different thread and I'm only using it at home so absolutely no issue if I kill it 😊

That's interesting. Did you have a hint, why the normal Update Accelerator does work. My version is just a enhancement of the existing framework.



Re: Update Accelerator revisited 66

by **BeBiMa** » May 14th, 2015, 9:17 pm


BeBiMa wrote:
Did you have a hint, why the normal Update Accelerator does work. My version is just a enhancement of the existing framework.

Sorry, no idea of why it wasn't working, but I do now have a solution as I've used the install.tar.gz from N0man from here:
<http://forum.ipfire.org/viewtopic.php?f=50&t=12868>

I'm not a stranger to Update Accelerator, in fact I used it in IPCop for many many months and even extended the provided version to have more sets of data to be added to the cache and added in icons for other sources, so I'm going to review what has been written for ipfire to see if I can extend it. My IPCOP additions are here: <https://github.com/pbinks/updx-accelerator>

I don't suppose you've got a Git repo for the current rewrite of the Update Accelerator that I could perhaps contribute to, or are you just passing updates between yourself and N0man?

Thanks,
Binky



Re: Update Accelerator revisited 66

by **BeBiMa** » May 14th, 2015, 9:37 pm

No, I haven't saved my files to the IPFire Git, yet. This will be one of the next steps, if installation and thus tests by others work.

At the moment I do the additions of new vendors by extracting them from N0man's archives. But this isn't really up to date. If someone runs my new version, it is very easy to add new definitions. I'll collect those to add them to the 'official' revision. Only the definition files will be needed.

I would be glad, if I would know where the install doesn't work properly.

I've installed to a separate directory and checked the files copied there. As far as I can see all files are transferred and are identical to my running version.



Re: Update Accelerator revisited 66

by **BeBiMa** » May 14th, 2015, 9:49 pm

For a check of integrity, I've added the MD5sum to the latest archive

EDIT new version 1.1-0a with correction of severe typos is available. The directories with the vendor definitions should be copied now. Next version will check for a empty definition dir.

Update Accelerator revisited

Post Reply Search this topic... 91 posts

Re: Update Accelerator revisited

by BeBiMa May 15th, 2015, 12:51 pm New version 1.1-1 available.

There was an error with vendors, when used the first time (what is normal for a clean repository). Due the changed calculation of disk usage and free space, these directories must exist to start the download. (Thx to N0man for giving me the hint).

- Just two tips for usage:
- Definition of new vendors
- Transfer of repositories from existing systems

```
Code block with shell commands for updating IPFire 2.25.
```

I also noticed an entry in the system log:

```
Code block showing system log output: 11:54:02 ipfire: UPDXL: no vendors defined
```

Re: Update Accelerator revisited by BeBiMa May 15th, 2015, 5:34 pm

To solve the main problem, edit the file updx_new (the installer) please. Line 32ff

```
Code block showing diff for updx_new file.
```

should read

```
Code block showing corrected code for updx_new file.
```

Don't know how this typo resisted all editing. The installer presumes a definitions directory, but it is not there because of the false cp command.

Thanks for the hint.

Re: Update Accelerator revisited by RedneckMother May 15th, 2015, 5:19 pm

Results for 1.1-1 install:

```
Code block showing terminal output of update process.
```

Re: Update Accelerator revisited by BeBiMa May 15th, 2015, 5:34 pm

Thanks for the hint.

No, no... thanks for the quick fix, and your diligence!

Re: Update Accelerator revisited by RedneckMother May 15th, 2015, 6:01 pm

Thanks for the hint.

No, no... thanks for the quick fix, and your diligence!

Re: Update Accelerator revisited by BeBiMa May 15th, 2015, 8:13 pm

RedneckMother wrote: No, no... thanks for the quick fix, and your diligence!

That wasn't really a quick fix. I knew this typo very well and had corrected it several times since I found it. Perhaps it is a problem of release time. Many corrections were made at times, when "my main tester" N0man was active. He lives also in the US and that means his 'activity time' is after midnight here in Germany.

BTW: Does it work now?

Re: Update Accelerator revisited by RedneckMother May 15th, 2015, 8:30 pm

BTW: Does it work now?

The uninstall / install sequence worked after the edit.

I'm having an issue with linux updates, though. I haven't had any luck tracing what's going on, but during a fedora update (after the 1.1-1 fix), I could see various rpm & drpm downloads in progress on the updx web interface, but they never show up in the stats or maintenance views.

I'll try to troubleshoot it as best I can.

I'm in the US, too (sorry 'bout that). Get some sleep!

Thanks again!

Re: Update Accelerator revisited by RedneckMother May 18th, 2015, 4:44 pm

A couple of issues:

I downloaded updx_new.1.1-1.tgz with the md5sum of 7487adec4e3e4701af27cbe5361a7de and extracted it.

I changed to the created directory, and ran './updx_new uninstall', which resulted in the following messages:

```
Code block showing terminal output of 'updx_new uninstall'.
```

I ran './updx_new install', which generated:

```
Code block showing terminal output of 'updx_new install'.
```

I then rm'd './usr/sbin/updxrator.std' and reran './updx_new install', which generated:

```
Code block showing terminal output of 'updx_new install' after rm.
```

After a './etc/init.d/squid restart', I attempted a few client software updates. The stats and maintenance views in the WUI do not reflect the new downloads.

What must I do to assist in troubleshooting?

Thank you, Paul

Edit: Also, saw 'UPDXL: no vendors defined' in system logs. In the maintenance WUI, existing entries won't delete... possible file/directory owner/permission problems?

D.

Re: Update Accelerator revisited by BeBiMa May 18th, 2015, 9:40 pm

I just checked the files of version 1.1-1 with the given MD5sum. I can't see a reason for the your issues. The error 'no sources defined' should printed to stdout also, if it is logged to the system log. Your output of the install doesn't show this error.

Also the error of 'uninstall' not deleting the .std files was there in a earlier version, but not in the latest.

To check for owner/rights problems do a

```
Code block showing 'ls -lR files' command.
```

from the installer directory. All files a copied with 'cp -p'. Therefore the attributes of the files in the archive should be propagated to the destination directories.

Thx in advance Bernhard

Re: Update Accelerator revisited by RedneckMother May 20th, 2015, 5:03 pm

Hello Bernhard.

I discovered that various files for update accelerator were corrupted, most likely from earlier attempts to update from N0man's thread and previous iterations of this thread. I recovered the files from an install ISO, and will try 1.1-1 again.

Sorry for the confusion.

Paul

Re: Update Accelerator revisited by BeBiMa May 20th, 2015, 8:45 pm

Paul,

this may be truely a reason for your issues. I know this situation from development, but in most cases I know what I've changed.

It is true, the installer grew up to a more or less stable version, that should allow a clean switch between the old and the new version. I am sorry for bothering my 'earlybirds'.

I'll try to check the installer once more deeply for the release of the next version.

Bernhard

Re: Update Accelerator revisited by vkykam May 28th, 2015, 2:43 am

Hi Bernhard,

Thanks for the great work.

Some feedback, and I don't know whether it's a bug or of my own doing... I've tested an earlier version (probably 1.0-1), and now 1.1-1 as well. If you look at the attached .jpg, you'll see that it appears the icon files are either missing for my install (or maybe I uploaded with the wrong permissions?), as this happened I believe after testing 1.0-1. There's also a snippet of what is a very long list (probably in the 1000's or 10000's) of Microsoft .cab updates that are miniature, 10kB or less, which only started happening after I tested 1.0-1. I may also have messed with the configuration files prior to testing 1.0-1, so I don't know whether it's of my own doing, or whether this is expected behavior, as this is my only IPFire install.

The other thing of note, and I believe this is also true of the old Update Accelerator, is that max external download rate is per update and not total aggregate of Update Accelerator, based on the observations. If my observations are correct, then I think the GUI should be changed to indicate that to be per update, as it's currently ambiguous. In an ideal world, it should be total aggregate, but I don't know if that's technically feasible. We are using Update Accelerator in a BYOD environment in a very rural part of the world, with connection speeds of 0.25Mbps shared among 20 or 30 people at times. The limit becomes useless, at least to us, when there are large quantity of update files to be downloaded, and the number of files multiplied by the speed limit will often saturate the entire bandwidth available, but we can't set the max update speed so low (like 5kpbs) that the updates will take a week or more to download.

```
Code block showing network speed test results.
```

Re: Update Accelerator revisited by BeBiMa May 28th, 2015, 10:06 am

Thanks for testing.

The first issue is generated by wrong permissions for the image files, I suppose. The images are located in /srv/web/ipfire/html/images/updboost. In my system they all have root/root 0644.

The second issue is a true bug. Download is limited for each file. You are right, this doesn't make sense for multiple concurrent downloads. To solve this problem, we need a UAdownload demon. I'll try to implement this in one of the next releases.

Thanks for the hint. Bernhard

Re: Update Accelerator revisited by vkykam May 29th, 2015, 2:40 am

I checked that already, and confirmed that it's root/root 0644, so I don't know why the image icons aren't being served.

BeBiMa wrote: Thanks for testing. The first issue is generated by wrong permissions for the image files, I suppose. The images are located in /srv/web/ipfire/html/images/updboost. In my system they all have root/root 0644.

Re: Update Accelerator revisited by BeBiMa May 29th, 2015, 2:27 pm

I fear, there were and are some problems with rights and ownership. I'm checking this at the moment after a fresh install of new UA because of IPFire upgrade, which restored the standard UA.

I hope, I can release a new version tomorrow.

Re: Update Accelerator revisited by RDaemon June 29th, 2015, 9:56 am

The limit becomes useless, at least to us, when there are large quantity of update files to be downloaded, and the number of files multiplied by the speed limit will often saturate the entire bandwidth available, but we can't set the max update speed so low (like 5kpbs) that the updates will take a week or more to download.

I faced the same situation and use "quick and dirty" fix for now:

1) find following code in "download" script:

```
Code block showing original download script code.
```

2) place this right before string "\$_ = system("\$cmd");"

```
Code block showing modified download script code.
```

Resulting code must look like this:

```
Code block showing final modified download script code.
```

This limits number of concurrent WGET tasks to 4, total speed limit is "max update speed"*4, another download tasks remains inactive, checking with 10-minute interval when tasks number of speed limit instances is less than 4 and waits, if more.

Post Reply Search this topic... 91 posts

Update Accelerator revisited

Post Reply Search this topic... 91 posts

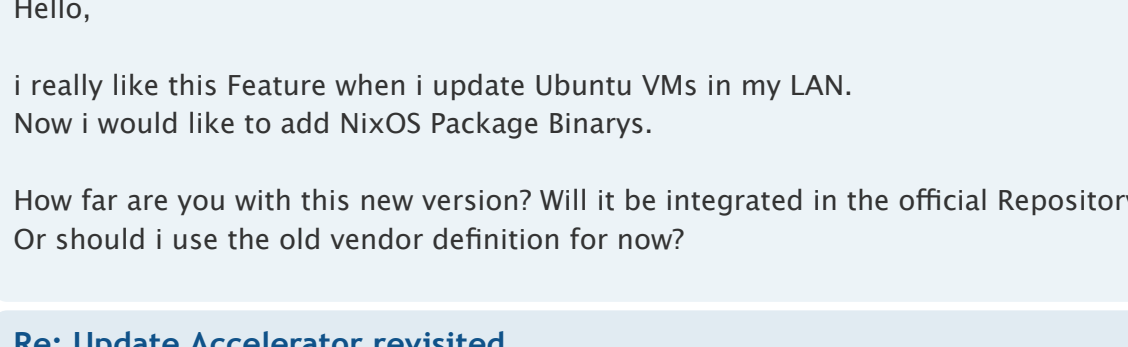
Re: Update Accelerator revisited

by BeBiMa » July 3rd, 2015, 9:41 am

Yes, that's just a "workaround". I'll try to implement this behaviour by semaphores. With a limit set, only one download process will start a wget process. A possible extension could be to limit the number of wgets to the number of UA processes in case of "unlimited" download.

But this will last some time. 😞 I'm struggling with some kind of influenza at the moment, which makes it easy to implement new features in a safe and effective way. 😊 Sorry.

Bernhard



Unitymedia Cable Internet (32MBit)

Re: Update Accelerator revisited

by davidak » September 8th, 2015, 12:05 pm

Hello,

I really like this Feature when I update Ubuntu VMs in my LAN. Now I would like to add NixOS Package Binaries.

How far are you with this new version? Will it be integrated in the official Repository? Or should I use the old vendor definition for now?

BeBiMa Posts: 2842
Joined: July 30th, 2011, 12:55 pm
Location: Mannheim

Re: Update Accelerator revisited

by BeBiMa » September 11th, 2015, 4:58 pm

davidak wrote:
Hello,
Now I would like to add NixOS Package Binaries.

If you know the URLs you can add them to your UA database.

IPFire 2.25 - Core Update 153 on i586 AMD Geode @ 498MHz - 242MB

Unitymedia Cable Internet (32MBit)

Re: Update Accelerator revisited

by EdoFedeb3 » October 13th, 2015, 10:20 pm

vkykam wrote:
... is a very long list (probably in the 1000's or 10000's) of Microsoft .cab updates that are miniature, 10k8 or less

To avoid saving that garbage in the cache, you can edit the microsoft section as explained here: [viewtopic.php?p=84401#p84401](#)

Great work Bernhard,
When I have a little time, I will try your Update Accelerator version on my system. Is there any possibility to have it integrated in the next official ipfire Core Update? 😊 😊

Bye,
Edoardo.

IPFire - An open source firewall solution

Re: Update Accelerator revisited

by xPIZit_xs » October 18th, 2015, 3:17 pm

Hi,

I was checking the updatextir file and noticed something.

in the Microsoft section is that a mistake or OK that near the end a "s" is missing before the "@" sign?

```
&& ($source_url |&#x2D; m@^http://([^\]*\.microsoft\.com)/:(/autoupd|selfupdate/).*\.cab@i)
should be this : ?
&& ($source_url |&#x2D; m@^http://([^\]*\.microsoft\.com)/:(/autoupd|selfupdate/).*\.cab&#x2D;@i)
```

let me know please 😊

xPIZit_xs



Re: Update Accelerator revisited

by BeBiMa » October 18th, 2015, 4:03 pm

Certainly, the 's' is necessary!
Otherwise a URL <http://website.microsoft.com/path/another...ubbish.exe> would be matched.

IPFire 2.25 - Core Update 153 on i586 AMD Geode @ 498MHz - 242MB

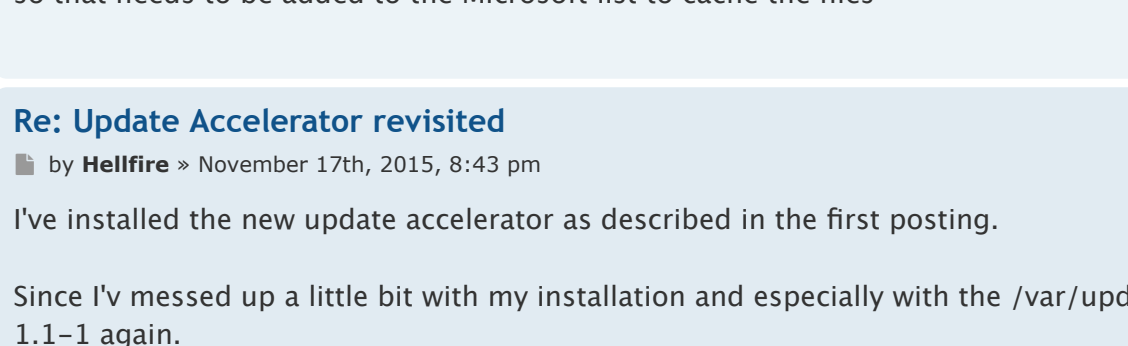
Unitymedia Cable Internet (32MBit)

Re: Update Accelerator revisited

by xPIZit_xs » October 18th, 2015, 5:41 pm

Ok, nice.
Guess you will update your package then.

xPIZit_xs



Re: Update Accelerator revisited

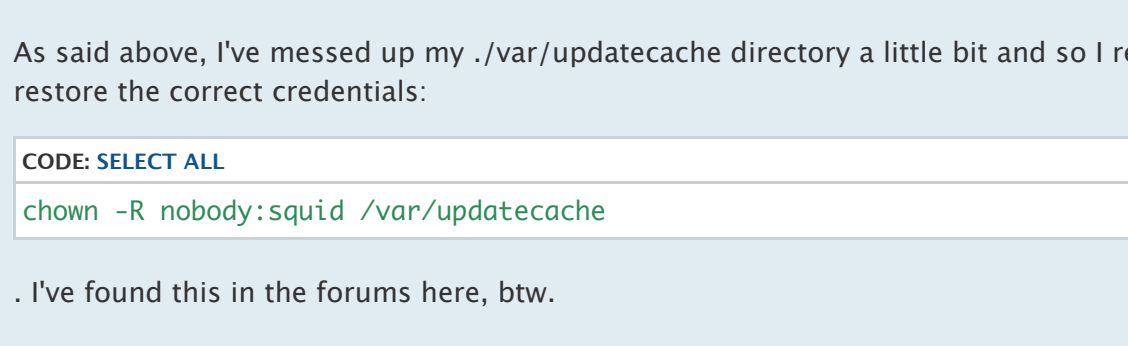
by Hellfire » November 12th, 2015, 8:03 pm

Hi,

is this update ready for the latest IPFire version?

I mean I will set up a new installation the next days and would like to use the new Update Accelerator from this thread. Can I still use the UI to lookup the statistics and more after the update?

Michael



Re: Update Accelerator revisited

by Cybermaze » November 13th, 2015, 12:38 pm

Hi Hellfire,

You can safely download and install the newest version of the Update Accelerator on page 1 of this thread. This work-in-progress works fine with the IPfire WUI, it will even work with and list statistics about the new sources.

Cybermaze Posts: 81
Joined: February 1st, 2014, 12:25 pm
Location: Siagelse, Denmark

Re: Update Accelerator revisited

by zcworld » November 13th, 2015, 9:30 pm

some of the windows 10 updates come down in esd file format so that needs to be added to the Microsoft list to cache the files

zcworld Posts: 15
Joined: May 4th, 2015, 10:31 am

Re: Update Accelerator revisited

by Hellfire » November 17th, 2015, 8:43 pm

I've installed the new update accelerator as described in the first posting.

Since I've messed up a little bit with my installation and especially with the /var/updates directory, I tried to install the latest version 1.1-1 again.

The installer told me to uninstall first but I did not had any clue how to do so I tried by any chance with running the command

```
CODE: SELECT ALL
./updx_new uninstall
```

which worked fine 😊 as it seems.

Afterwards, I tried to re-install version 1.1-1 but the installer told me that

new updatexlrator is installed yet.
Uninstall first, please.

So, how to proceed? Unfortunately after a positive first test running IPFire in a VM the update accelerator put some files in the above cache. Now on my fresh IPFire, running on a dedicated server, not a single file finds its way into the cache directory.

As said above, I've messed up my /var/updates directory a little bit and so I recreated it and fired the following command to restore the correct credentials:

```
CODE: SELECT ALL
chown -R nobody:squid /var/updates
```

.. I've found this in the forums here, btw.

Any hints?
Michael

IPFire 2.25 - Core Update 153 on x86_64 AMD GX-412TC SOC @ 966MHz x2 - 4GB

Re: Update Accelerator revisited

by Hellfire » November 17th, 2015, 9:13 pm

So, after removing the obviously wrong file updxlrator.std like others did in an earlier posting here, I re-run the installer to install version 1.1-1 again. No problems occurred.

Anyway, I'm currently downloading about 175 update from a IPfire server but none of those files is shown in the statistics nor in the maintenance section of the Update Accelerator WUI when refreshing the web-page from time to time.

Those files are still missing in directory /var/updates. Btw, I've set up the proxy as transparent and enabled on green if this does matter and of course the accelerator is enabled in the proxy settings. No firewall rules exist so far.

The file /var/log/updatexlrator does have current entries pointing to those Microsoft updates, but as already said, none of those files exist in the updatescache folder. Is this an async process - I mean is the cache filled after a certain period of time and not while the actual download runs?

Michael



Re: Update Accelerator revisited

by Hellfire » November 18th, 2015, 6:59 am

Hi,

since I can't solve this myself for obvious reasons, is there a way to get the previous update accelerator back?
Which files are needed?

The uninstall script does not work anymore.

And for my knowledge, which prerequisites are needed anyway for the update accelerator to work properly? Any of those proxy settings?

Edit: OK I got the original files back from a IPfire installation running in a VM, I uninstalled 1.1-1, restarted squid and tested some Windows updates. Now the cache is re-built correctly and the Microsoft files are present in the appropriate directories. Don't know for sure if my previous issues had to do with the beta here or anything else. If I find some time again I will return and install the beta once again.

Michael



Re: Update Accelerator revisited

by Hellfire » November 28th, 2015, 1:12 pm

Still alive?

I notice Win10 does update its store apps using file type .AppxBundle and may be .Appx, too.

So I've added both to the Microsoft source file, updating this line:

```
CODE: SELECT ALL
p,http://([^\]*\.microsoft\.com)/.*\.(exe|psf|msi|msu|cab|esd|AppxBundle|Appx)
```

Afterwards I've recompiled the source with the command given in the first posting and started the update process in Win10 Store again.

After looking into the proxy logs I see that some files were downloaded, the appropriate URL, however, is not pointing the file itself but rather to a streamed content. At least this is my guess:
<http://tlu.dl.delivery.mp.microsoft.com...l2ACp6Y%3d>
Anyway this URL results in a file with extension .AppxBundle.

I thought when adding the file's extension to the list of downloadable file types those will be cached either. Obviously this did not happen.

Any idea why not?

Michael



Re: Update Accelerator revisited

by Hellfire » December 3rd, 2015, 11:58 am

Is the revisited update accelerator a dead cow now?
After some good enhancements now great silence in this thread 😊

IPFire 2.25 - Core Update 153 on x86_64 AMD GX-412TC SOC @ 966MHz x2 - 4GB

Post Reply 91 posts

Update Accelerator revisited

Post Reply Search this topic... 91 posts

Re: Update Accelerator revisited


by BeBiMa February 27th, 2016, 1:32 pm

The cow isn't dead!

It only had no time to check and try the handling for real management of the download bandwidth. (A solution isn't implemented already).

But I hope to release a new version with updated source lists (and the code running for months now in my installation here), soon.

Greetings, Bernhard



Unitymedia Cable Internet (32MBit)

BeBiMa Posts: 2842

Re: Update Accelerator revisited

by tt2tt August 24th, 2016, 11:04 am

Hi, thank you for your work --) Is this an official addon which I can download in pakfire?

tt2tt Posts: 4

Re: Update Accelerator revisited

by deswong September 25th, 2016, 10:21 am

Just checking in to see if there is an ETA on a newer version? I have been having issues with my current install - and would like to try a later version rather than clean install and re-setup.

deswong Posts: 3


Re: Update Accelerator revisited

by JDSponky November 27th, 2016, 9:41 pm

Hello!

What's going on with this Tool? Will it soon replace the existing Updater?

JDSponky Posts: 5



Re: Update Accelerator revisited

by Mahagon July 17th, 2017, 8:03 am

Just found this Project

BeBiMa hasn't been active since Mon Aug 22, 2016. I just sent him an pm hoping he gets a mail notification ^^

Uploading it to GitHub or sth like that would be a solution. Maybe someone could continue his work.

Mahagon Posts: 5

Re: Update Accelerator revisited

by robertgpeterston October 9th, 2017, 9:01 pm

Is this rewrite reasonably ready for production use? Is there an uninstall/revert if necessary?

Update Accelerator has been a distinctive and key feature of IPFire for our context in Africa on a limited bandwidth satellite connection. It is a game-changer in comparison to any other firewall solution out there. I hope that it continues to be an integral and high priority part of IPFire's development. Our live version has 180+ GB of updates cached and used to average around 18:1 ratio. Windows 10 change to cumulative updates doesn't function as well with it and the ratio is closer to 12:1 now.

DNS filtering is another aspect that would be so helpful at the GUI level. I'm using external DNS filtering for the moment with selective URL filtering as backup.

Rob

robertgpeterston Posts: 33

Re: Update Accelerator revisited

by jlgtx December 2nd, 2017, 5:27 pm

One of the most significant problems with Update Accelerator is that it doesn't manage its cached files well. Even the manual maintenance tool is incomplete, e.g., there's no way to delete cached updates older than a certain date or smaller than a certain size, clean up duplicate updates, etc. If hard drive space were infinite, this would be no problem...but I'm running IPFire on a mini-PC with a fairly small CF filesystem, and Update Accelerator eats it up space pretty quickly. If you don't mind doing a bit of command-line stuff, you can use this shell script and crontab entry to clean up the "big two" problems--old files and duplicates.

Login to your IPFire host as root (replace 1.2.3.4 with your IPFire host IP):

```
# ssh root@1.2.3.4 -p 222
```

Now create the script...

```
# vi cleanupUA.sh
```

```
CODE: SELECT ALL
#!/bin/bash
# change this to the number of days you want to keep!
age=30
#
# This only works on the default cache directory. Things will break if you've moved it to a deeper or shallower path.
cachedir="/var/updates/cache"
# Touch the base update directories to avoid deleting them!
touch $cachedir/*
echo "Removing old updates..."
For upd in $(find $cachedir -type f -mtime +$age -print |cut -f1-5 -d'/');
do
  echo "$ Supd"
  rm -rf $upd
done
```

Make it executable:

```
# chmod 700 cleanupUA.sh
```

Run it to see it do its thing:

```
# ./cleanupUA.sh
```

To run the script automatically on a monthly basis, run "fcrontab -e" and insert the following rows at the end of the file:

```
CODE: SELECT ALL
# Clean up the Update Accelerator cache once a month.
#monthly * * * /root/cleanupUA.sh >/dev/null 2>&1
```

jlgtx Posts: 31

Re: Update Accelerator revisited

by robertgpeterston December 4th, 2017, 7:48 pm

Thank you. Appreciate the very clear instructions. Will run this soon.

robertgpeterston Posts: 33

Re: Update Accelerator revisited

by jluth@mail.com December 6th, 2017, 8:52 am

robertgpeterston wrote: ↑ Is this rewrite reasonably ready for production use?

History has shown (in both ipcop and ipfire) that you can't wait for pending updates to updxlrator to actually arrive. If you are capable of fixing any problems by yourself, then just move forward on your own.

An in regards to additional changes needed in updxlrator, many microsoft patches now ought to be matched on filename (\$mirror) rather than url (Sunique); specifically the files with an SHA1 hash in their name. That's really hitting me hard with Windows 10 updates being pulled from different mirrors.

Of course, this means that the existing cache files won't match, but some scripting can rename that. I've created a bug report for this at https://bugzilla.ipfire.org/show_bug.cgi?id=11558

jluth@mail.com Posts: 9

Re: Update Accelerator revisited

by robertgpeterston December 6th, 2017, 8:52 am

I'm not a programmer but this sounds very promising. I wonder if our situation is particularly disfunctional with being in Africa, the system can never decide which mirror to use and seems to pull from Australia, Europe and the US...sometimes all at the same time.

I think I could manage changing the code but the symbolic link section would require line-by-line instructions for me to try it.

robertgpeterston Posts: 33

Re: Update Accelerator revisited

by jluth@mail.com January 9th, 2018, 6:54 am

The next release of ipfire will include fixes for updxlrator. 1.) With Microsoft's new way of updating, it can easily make multiple requests for the same URL within one second. This was exposing a flaw that allowed multiple downloads of the same file to occur. Have you ever seen a download climb to greater than 100% of the file size? This is likely the reason - that multiple downloads were dumping on top of the same filename. Obviously in the end this is corrupt anyway, so that file, downloaded multiple times, ends up being useless, wasting a lot of bandwidth. That has been fixed. https://bugzilla.ipfire.org/show_bug.cgi?id=11567

2.) Updatxlrator has two modes - a mirror mode which requires that the filename be unique regardless of the server that it downloads from, or URL mode, where the entire URL must be matched exactly. ALL Microsoft downloads were based on URL mode. With the huge (5+GB) windows 10 updates that come each month, they could be downloaded multiple times, depending on which mirrors the client requested it from. Sometimes a single client will request from multiple mirrors for the same file. Most of Microsoft's updates now include an SHA1 number in their filename, indicating that the file is unique. This allows updatxlrator to just use the filename as the identifiable part, instead of the entire URL. So for Microsoft updates, updxlrator now notices when an SHA1 number is present. If so, it first looks to see if the file is already cached using the old method (the entire URL). If so, it uses that cache. If not, it will switch over to filename mode, check if the download exists in that cache, and if not it will download and save as filename mode. https://bugzilla.ipfire.org/show_bug.cgi?id=11558

This means that for any NEW updates that come in, they will only ever be downloaded one time. For 2017 updates, it is possible that you STILL might download an update once more IF a client requests it from a mirror that wasn't accessed previously. (Of course, that would also have been true before this update...) So there isn't really any need for the end user to do anything, but for any programming type of users who REALLY want to avoid a potential extra download, the bug report has a script that can be modified to rename the existing URL-based caches into filename-based caches. (It needs to be modified so that it renames the folder instead of creating a link.)

3.) filenames containing a + or a ~ were never being retrieved from the cache. That has been fixed. https://bugzilla.ipfire.org/show_bug.cgi?id=10504

@everyone who tinkers with updxlrator: Make sure you have a backup of your updxlrator files if you have modified them (to add custom caches etc.) I'm quite sure that the upgrade process will replace your custom version, wiping out your customizations.

@Robert: I'm also in Africa (South Sudan) and these issues over satellite really hit us hard, so I dug into it until I figured out what the problems / solutions were. The latency in getting a download started opens up a multi-second window for simultaneous downloads (problem #1). So the sites that need updxlrator the most were also the most affected by the bug. These updates should come automatically in Core 118, so no need for you to do any scripting any more.

@jlgtx: deleting duplicates is not a good idea, since the cache (for Microsoft updates) is not based on filename, but on the URL. The deleted duplicates would have downloaded again if a client asks for that particular URL. (But of course, that has now changed for SHA1-named files).

jluth@mail.com Posts: 9

Re: Update Accelerator revisited

by dni February 24th, 2018, 4:37 am

Hello,

I'm glad to see this is still being used and updated. Is somebody officially maintaining it now?

I was hoping to ask if it would be possible to add a feature. I'd like to split the functionality in to multiple files so that I can maintain my own "source_url" sections without them being overwritten every time the script is updated.

I have a network where I only want to cache downloads for one specific operating system. This is because there is at most 2 systems of any other operating system and the Update Accelerator results in two complete downloads anyway. The default settings for the Update Accelerator can result in less efficient usage of the download quota than not using it at all.

So I was wondering if it was possible to make the script modular. It could have 3 major parts:


- 1. The main perl script, which includes a reference to
2. a configuration file, listing all of the files which contain source_urls (today these are in "Section" heading) in which the main while loop
3. and split today's "Section" headings with source_urls in to multiple files. Let's call them "Sources".

The config file would allow a user to turn on and off individual modules ("Section" files containing source_urls). If necessary a user (like me) could add reference to their own custom "Source" file.

The idea being that the main perl script and the source_url files (modules) could be updated by the maintainer, while the config file can be left alone after initial install.

I don't know Perl. Would that be complex to do?

Thank you!



dni Posts: 375

Re: Update Accelerator revisited

by dni February 24th, 2018, 4:37 am

PS: I am NOT expecting a web interface for this. While that would be nice the update accelerator is incomplete in other important ways, like automatically pruning cache files (see jlgtx's post above).



dni Posts: 375

Re: Update Accelerator revisited

by scareycrow April 4th, 2018, 12:57 pm

Should we be expecting all Windows 10 updates to cache properly now?

I'm seeing pretty much everything except the Cumulative Updates (which are the largest) cached. Example KB4088776 is not caching.



scareycrow Posts: 8

Re: Update Accelerator revisited

by jluth@mail.com April 4th, 2018, 1:08 pm

Yes, I would expect it to be cached. Do you see it in the updatxlrator page being downloaded? What makes you think that it isn't being cached?

My only guess would be that your connection isn't stable enough to complete a very long download, so it aborts before finishing.

jluth@mail.com Posts: 9



Update Accelerator revisited

Post Reply



91 posts



Re: Update Accelerator revisited

by **jlgtx** » June 7th, 2019, 8:09 pm

So...trying to come up with a command-line equivalent of the cache maintenance panel. It seems to me that this should work, but it doesn't:

```
curl -k -X POST -u <admin_username>:<admin_pw> -H 'Content-Type: multipart/form-data' -F 'ACTION=Purge' -F 'REMOVE_OBSOLETE=on' -F 'NOT_ACCESSED_LAST=month3' -F 'REMOVE_NOSOURCE=on' -F 'REMOVE_OUTDATED=on'
https://localhost:444/cgi-bin/updatexlrator.cgi
```

The HTML returned by this command is just the Update Accelerator config page, and the maintenance form submission doesn't get processed. I've tried this with and without the -H component.

Anybody see what I'm missing?



jlgtx

Posts: 31
Joined: August 19th, 2016, 4:23 am

Post Reply



91 posts



[Return to "Development"](#)