

Servers - Bug #531

Suggestion to mirror hosts: sync from hosts other than repo.parabola.nu to avoid overload

2014-05-30 08:09 PM - icarious

Status: open	% Done: 0%
Priority: wish	
Assignee:	
Category:	
Description With the increasing number of mirrors for Parabola and most of them running rsync from repo.parabolagnulinux.org; it would be great if the hosts considered using other mirrors to run rsync from to reduce the load on the main repo server. A "Repo Server (repo.parabolagnulinux.org) -> Mirror 1 -> Mirror 2" distributed approach could be optimal. And individual users can hence connect to the individual mirrors. Would be great if every mirror host kept the increasing load in mind as mirroring from the same source will eventually congest it and slow it down. rsync://mirror.yandex.ru/mirrors/parabola/ might be considered, as they might have high bandwidth and run 24/7 and subsequently the other mirrors can run rsync from each other.	
Related issues:	
Related to Servers - Bug #611: [rsyncd] Reenable, properly limit bandwidth	fixed 2014-12-11

History

#1 - 2014-06-08 03:54 AM - lukeshu

We're having growth problems :)

How Arch handles this is: They have 3 tiers of servers; 0, 1, and 2. 0 is ftp.archlinux.org (repo.parabolagnulinux.org). Tier 1 servers rsync with ftp.archlinux.org, rsync to tier 0 is actually blocked for all non-approved tier 1 IPs. Tier 2 rsyncs to the tier 1 servers. Anyone can create a tier 2 server. Once a tier 2 server has a could enough reputation for syncing frequently and good uptime, they start talking about promoting it to a tier 1 server.

Also, ParabolaWeb has kinda-awesome mirror management facilities that we aren't using (<https://parabolagnulinux.org/mirrors>). It does statistics on uptime and lagtime and all kinds of cool stuff. Example: <https://www.archlinux.org/mirrors/gitlib.gatech.edu/>

#2 - 2014-06-10 08:28 PM - xylon

But if we have two tiers, and they sync every 24 hours, wouldn't that mean mirrors could be 48 hours out of date? Maybe we just need a stronger repo server.

#3 - 2014-06-10 08:30 PM - xylon

Or everyone could use pacman2pacman and we can make it use arch mirrors as webseeds and then we won't need to mirror any arch packages, only the parabola replacements. See [#536](#).

#4 - 2014-06-17 07:53 PM - lukeshu

The Arch mirrors sync way more frequently than 24 hours.

#5 - 2016-02-09 11:50 AM - alfplayer

- Subject changed from Suggestion to mirror hosts to Suggestion to mirror hosts: sync from hosts other than repo.parabola.nu to avoid overload

#6 - 2016-02-09 11:51 AM - alfplayer

- Related to Bug #611: [rsyncd] Reenable, properly limit bandwidth added