



LEALTERNATIVE



UPLOAD



Search

[ABOUT](#)[BLOG](#)[PROJECTS](#)[HELP](#)[DONATE](#) ❤️[CONTACT](#)[JOBS](#)[VOLUNTEER](#)[PEOPLE](#)

Contributors

Antony Antony
Bill "TTK" Moyer
Braam Aangstrom
Brad Tofel
Brewster Kahle
Bruce Baumgart
CR Saikley
Dan Aronson
Diane Barrera
Gavin Dapper
Guy Lauten
Jim Shankland
Joerg Bashir
John Berry
Joshua Coates
Marcel van der Peijl
Mario Murphy
Matt Laue
Ralf Muehlen
Rop Gonggrip
Sam Mortazavi
Sean Fagan
Shaun Fogarty
Tim Johnson
Tom Jones
Tracey Jaquith
Trevor von Stein
Will Mooney

Top Level Contributors

Petabox


Petabox

A few highlights from the Petabox storage system:

- No Air Conditioning, instead use excess heat to help heat the building.
- Raw Numbers as of December 2021:
 - 4 data centers, 745 nodes, 28,000 spinning disks
 - Wayback Machine: 57 PetaBytes
 - Books/Music/Video Collections: 42 PetaBytes
 - Unique data: 99 PetaBytes
 - Total used storage: 212 PetaBytes



The Previous Version



OVERVIEW

The PetaBox(tm), custom-designed by Internet Archive staff, was originally created to safely store and process one petabyte (a million gigabytes) of information. The goals and design points were:

- * Low power: 6kW per rack, 60kW for the entire storage cluster
- * High density: 100+ TB/rack
- * Local computing to process the data (800 low-end PC's)
- * Multi-OS possible, linux standard
- * Co-location friendly
- * Shipping container friendly: Able to be run in a 20' by 8' by 8' shipping container.
- * Easy Maintenance: One system administrator per petabyte
- * Software to automate full mirroring
- * Easy to scale
- * Inexpensive design
- * Inexpensive storage

The Internet Archive data center now houses ~3PB of PetaBox storage technology and is expanding steadily.

PETABOX HISTORY 6/2004

* The first 100TB Rack operational in Amsterdam as of June 2004.

* The second 80TB rack is operational in San Francisco

* Internet Archives spins off PetaBox production to newly-formed Capricorn Technologies.

2004 - 2007

* Capricorn replicates the Internet Archive's successful deployment of the PetaBox for major academic institutions, digital preservationists, government agencies, HPC and major research sites, medical imaging providers, digital image repositories, storage outsourcing sites, and other enterprises around the globe.

Petabox   RSS

New Post

Subject	Poster	Replies	Date
Web archive is not available for all the time	jsuzana320	1	Oct 11, 2023 5:02am
Re: Web archive is not available for all the time	Reclezon	0	Feb 18, 2024 2:56pm
Problems with connection	margotfq	1	Aug 25, 2023 4:42am
Re: Problems with connection	jsuzana320	0	Oct 11, 2023 5:04am
Throttled again.	ratherendit	2	Jul 19, 2023 7:00am
Re: Throttled again.	rilysh	0	Oct 16, 2023 1:50am
Re: Throttled again.	Texasluva	0	Dec 19, 2023 7:17pm
503ing for the very last time...	ratherendit	1	Jul 4, 2023 5:31am
Re: 503ing for the very last time...	Jeff Kaplan	1	Jul 4, 2023 10:32am
Re: 503ing for the very last time...	ratherendit	0	Jul 5, 2023 11:20pm
St. Louis Vacation Rentals	JZ Vacation Rentals	0	Jun 13, 2023 1:29am
Issue in url	Isaiah C. Allen	2	Mar 9, 2022 11:50pm
Ads	Ramazan Sancar	1	Sep 2, 2022 5:52am
Slow Loading Speeds	andersoncop85	0	Dec 16, 2023 2:57am
Re: Issue in url	jsuzana320	1	Oct 11, 2023 5:01am
Re: Issue in url	andersoncop85	0	Dec 16, 2023 3:05am
PetaBox	TimCA	0	Oct 7, 2021 6:57pm
Odd storage formating question	Reclezon	2	Jul 18, 2021 1:03pm
Re: Odd storage formating question	Jeff Kaplan	2	Jul 18, 2021 1:27pm
Re: Odd storage formating question	SeaMac	0	Dec 27, 2023 11:00am
about how much storage is ready for growth	twisted whiskers	2	Apr 9, 2021 8:02pm
Re: about how much storage is ready for growth	Richard Wayne530	1	Jun 5, 2021 6:56am

Bittorrent tracker	Evergreen_street	0	Dec 7, 2020 3:17am
"Welcome to US Petabox"	michaelreuter1	1	Jun 29, 2020 1:01pm
Re: 'Welcome to US Petabox'	varenhizzle	1	Dec 1, 2021 4:25pm
Silly question	Teutamos	1	May 18, 2020 12:57am
Re: Silly question	Richard Wayne530	1	Jun 5, 2021 6:52am
Yes	VUTHANHTUNG	0	May 5, 2020 9:16am
Internet Archive	https://www.HomelandInsecurity.net	0	Apr 28, 2020 12:28am
petabox os	openprotegent project	0	Feb 28, 2020 8:53am
Petabox software specifications?	crd6	0	May 21, 2019 8:21am
How much archive.org monthly operating cost?	tamim17	2	Nov 8, 2017 4:53am
Re: How much archive.org monthly operating cost?	spartos	1	Feb 4, 2018 4:24pm
not much	covidbloggermailnes	0	Apr 2, 2020 2:50am
Periodic updates	icebob99	2	Jan 9, 2017 4:10pm
Re: Periodic updates	Jeff Kaplan	0	Jan 10, 2017 6:48am
Re: Periodic updates	emijrp	0	Apr 30, 2018 2:30pm
Amazing	JayKae	0	Aug 13, 2016 5:02am
Wow...	makanal	1	Sep 21, 2015 9:52pm
Re: Wow...	Jeff Kaplan	0	Sep 21, 2015 10:22pm
2014 ... 1 petabyte in 8U	Javik	0	Dec 16, 2014 2:19pm
Design update?	ihitoit	3	Oct 9, 2012 7:24pm
Re: Design update?	Coderjo	0	Oct 10, 2012 11:10pm
Re: Design update?	Coderjo	0	Oct 10, 2012 11:18pm
Re: Design update?	GridEngine	0	Dec 24, 2012 10:00am
How long does the data last?	Seaware	1	Oct 3, 2012 9:49pm
Re: How long does the data last?	Coderjo	1	Oct 8, 2012 1:49pm
Re: How long does the data last?	Seaware	1	Oct 9, 2012 12:47am
Re: How long does the data last?	Coderjo	0	Oct 10, 2012 11:04pm
http://www.capricorn-tech.com/ is now dead ... new home?	Tyler	1	Dec 23, 2011 9:24am
Re: http://www.capricorn-tech.com/ is now dead ... new home?	Hydriz	1	May 1, 2012 2:01am
Re: http://www.capricorn-tech.com/ is now dead ... new home?	Coderjo	1	Jun 15, 2012 4:28pm
Re: http://www.capricorn-tech.com/ is now dead ... new home?	منصة شليلة	0	Nov 24, 2019 8:33am

We have 2TB of data to upload	Sheikh Naveed	2	Nov 13, 2011 6:00am
Re: We have 2TB of data to upload	fL0per	0	Sep 4, 2020 3:51pm
Re: We have 2TB of data to upload	Stright	1	Jan 15, 2021 7:54pm
6th pair of racks go into service: over 2PB of data space used	brewster	3	Mar 26, 2011 12:48pm
Re: 6th pair of racks go into service: over 2PB of data space used	baetis	1	Apr 20, 2011 1:09pm
Re: 6th pair of racks go into service: over 2PB of data space used	Sean_0000001	0	Apr 23, 2011 12:17pm
Re: 6th pair of racks go into service: over 2PB of data space used	Tyler	0	Dec 2, 2011 2:42pm
New pictures of the new petabox	brewster	3	Mar 23, 2011 10:44am
Re: New pictures of the new petabox	Tyler	0	Mar 25, 2011 9:12pm
Re: New pictures of the new petabox	Coderjo	1	Oct 10, 2011 3:09pm
Re: New pictures of the new petabox	GridEngine	2	Dec 24, 2012 10:19am
Re: New pictures of the new petabox	Coderjo	0	Jan 7, 2013 11:25am
Re: New pictures of the new petabox	emijrp	0	Dec 6, 2011 3:17pm
Current Raw Numbers	Sean_0000001	1	Mar 9, 2011 11:55pm
Re: Current Raw Numbers	brewster	0	Mar 21, 2011 7:34pm
could MapReduce be applied?	broffmann	1	Apr 20, 2010 7:55am
Re: could MapReduce be applied?	Coderjo	0	Apr 28, 2010 12:06pm
CAD files for Petabox?	deerpig	2	Mar 24, 2010 1:38pm
Re: CAD files for Petabox?	thattommyhall	1	Apr 6, 2010 7:34pm
Re: CAD files for Petabox?	deerpig	1	Apr 6, 2010 8:09pm
Re: CAD files for Petabox?	deerpig	0	Apr 6, 2010 8:14pm
Re: CAD files for Petabox?	baetis	0	Aug 5, 2010 2:26pm
Is the Archive site no longer archiving	maxadds	2	Sep 12, 2009 9:22am
Re: Is the Archive site no longer archiving	direwolf0701	0	Sep 12, 2009 9:27am
Re: Is the Archive site no longer archiving	brewster	0	Oct 13, 2009 9:09pm
I assume this forum is now obsolete	Javik	3	Jul 10, 2009 10:33pm
Re: I assume this forum is now obsolete	Coderjo	2	Jul 15, 2009 4:17am
Re: I assume this forum is now obsolete	tracey pooh	0	Oct 16, 2009 8:57pm
Re: I assume this forum is now obsolete	jack liu	0	Nov 24, 2009 1:49am
Visiting the installation	glenn	0	Dec 10, 2008 10:32pm
Intel Atom based motherboards?	TimSmall	1	Jul 22, 2008 2:29am

Re: Intel Atom based motherboards?	NoiseCollector	0	Jul 22, 2008 8:24am
VIA NAS 7800-15LST	dunno	0	Jun 28, 2008 5:26pm
SuperMicro Cascaded 128/256/384-drive SAS/SATA chassis	Javik	1	Nov 28, 2006 4:14pm
Understanding all that technical blather :-)	Javik	1	Nov 28, 2006 3:54pm
Re: Understanding all that technical blather :-)	elhoim	0	Jan 11, 2007 6:39am
FYI, don't need more PCI slots, just SATA-II or SAS	Javik	2	Nov 23, 2006 9:59pm
24-drives on one SAS channel, from SuperMicro	Javik	2	Nov 25, 2006 6:15pm
Oh look, the first ever 2.5 inch SAS hard drive	Javik	0	Nov 25, 2006 6:14pm
Re: Make that --144-- 2.5" drives on one SAS channel	Javik	0	Nov 29, 2006 12:34pm
Re: FYI, don't need more PCI slots, just SATA-II or SAS	Gobba	0	Apr 14, 2011 4:48am
Huge storage = huge challenges?	cchipmunks	0	Nov 14, 2006 3:32pm
hi	karatekid	1	Sep 22, 2006 1:53pm
Re: hi	cosmicharlie	0	Sep 22, 2006 2:02pm
As many pci slots as possible?	AWP	1	Jul 3, 2006 6:53am
Re: As many pci slots as possible?	elhoim	0	Jul 6, 2006 4:49am
all archive storage is on petabox now!	tracey pooh	1	Jun 2, 2006 9:26am
Re: all archive storage is on petabox now!	protobadger	0	Jun 4, 2006 9:22pm
powernode rack	brewster	1	Jan 28, 2006 2:18pm
Re: powernode rack	Coderjo	1	Jan 27, 2006 1:16am
Re: powernode rack	brewster	1	Jan 27, 2006 1:33am
Re: powernode rack	Coderjo	1	Jan 27, 2006 8:22am
Re: powernode rack	brewster	0	Jan 27, 2006 8:26am
current configuration?	fgleason	1	Jan 4, 2006 7:23am
Re: current configuration?	fgleason	1	Jan 12, 2006 12:06am
Re: current configuration?	jkrauska	0	Jan 17, 2006 1:59pm
Still sept / oct for LMA > petabox move?	Tyler	0	Aug 30, 2005 2:21am
software specs	jer0mey	1	Aug 18, 2005 4:27am
Re: software specs	billmoyer	1	Jul 29, 2006 10:30pm
Re: software specs	singod	0	Jun 11, 2007 10:31am
Hitachi 500GB drives	ngngokkiu	1	Jul 31, 2005 5:20am
Re: Hitachi 500GB drives	Nethacker	0	Jun 3, 2006 11:02am

hey. I built something like that	sxpert	1	Jun 21, 2005 7:27pm
Re: hey. I built something like that	JohnSebastianBach	0	Feb 23, 2007 8:47pm
1PB and Counting	CR Saikley	0	Jun 21, 2005 2:50pm
You guys still on track for LMA transfer to Petabox.. end of summer 05?	Tyler	0	May 2, 2005 3:01am
like my movie?	snipey_2	0	Apr 24, 2005 2:53am
18 PCI slots, 72 SATA drives per computer?	Javik	2	Mar 22, 2005 11:27am
Re: 18 PCI slots, 72 SATA drives per computer?	caiser	1	Sep 3, 2005 8:31am
Re: 18 PCI slots, 72 SATA drives per computer?	Javik	1	Apr 8, 2005 5:42pm
Re: 18 PCI slots, 72 SATA drives per computer?	elhoim	1	Apr 23, 2005 4:00am
Drive spin-down? Low-RPM drives? Power saving?	Javik	1	Mar 22, 2005 9:25am
Re: Drive spin-down? Low-RPM drives? Power saving?	CR Saikley	1	Mar 22, 2005 2:58pm
Re: Drive spin-down? Low-RPM drives? Power saving?	HDRJapan	0	Oct 3, 2007 6:14pm
Hitachi Case Study	Brad Leblanc	0	Mar 10, 2005 10:45am
higher speed node motherboard	brewster	0	Feb 11, 2005 9:21am
petabox case design	zippy	1	Jan 20, 2005 4:45am
Re: petabox case design	CR Saikley	0	Jan 20, 2005 6:39pm
via board being used	brewster	0	Jan 16, 2005 11:01pm
Massive network storage 2	Vladovlado	0	Dec 3, 2004 6:27am
Massive network storage	Vladovlado	1	Dec 3, 2004 6:11am
Re: Massive network storage	caiser	0	Apr 5, 2005 4:16am
about 400TB of this design shipped to the Archive	brewster	5	Oct 24, 2004 2:06am
Re: about 400TB of this design shipped to the Archive	foundation	0	Nov 10, 2004 11:58pm
Re: about 400TB of this design shipped to the Archive	Curator at the Security Digest Archives	1	Feb 17, 2005 8:02am
Re: about 400TB of this design shipped to the Archive	Brak	1	Feb 17, 2005 1:23pm
Re: about 400TB of this design shipped to the Archive	dunno	1	Mar 3, 2005 6:43am
Re: about 400TB of this design shipped to the Archive	brewster	1	Mar 3, 2005 1:06pm
Re: about 400TB of this design shipped to the Archive	jko	1	Mar 31, 2005 1:41pm
Re: about 400TB of this design shipped to the Archive	James Day	3	Apr 1, 2005 6:47pm
Re: about 400TB of this design shipped to the Archive	dunno	1	Jun 11, 2005 7:45am
Re: about 400TB of this design shipped to the Archive	James Day	0	Jun 11, 2005 12:27pm
Re: about 400TB of this design shipped to the Archive	indianews	0	Jan 21, 2010 7:27am

Re: about 400TB of this design shipped to the Archive	viswiss	0	Jan 23, 2010 7:55am
Re: about 400TB of this design shipped to the Archive	matt-genesi-usa	0	Jun 23, 2005 10:08am
Re: about 400TB of this design shipped to the Archive	indianews	0	Jan 21, 2010 7:28am
Serial console cables?	garan	1	Oct 21, 2004 4:27pm
Re: Serial console cables?	brewster	2	Oct 24, 2004 2:13am
Re: Serial console cables?	garan	0	Oct 24, 2004 2:55am
Re: Serial console cables?	Casandro	1	Mar 3, 2005 9:31pm
Re: Serial console cables?	sxpert	0	Jun 21, 2005 7:54pm
Possible solution: Disk Array Rackmounts	KirbyMeist	1	Aug 3, 2004 11:10am
Re: Possible solution: Disk Array Rackmounts	brewster	1	Oct 24, 2004 2:20am
Re: Possible solution: Disk Array Rackmounts	foundation	1	Nov 10, 2004 11:47pm
Re: Possible solution: Disk Array Rackmounts	HDRJapan	0	Oct 3, 2007 9:32pm
Filesystem	KirbyMeist	2	Aug 2, 2004 2:08pm
Re: Filesystem	brewster	4	Aug 2, 2004 3:11pm
Re: Filesystem	grignak	0	Oct 4, 2004 12:33am
Re: Filesystem	foundation	1	Nov 10, 2004 11:50pm
Re: Filesystem	brewster	1	Nov 11, 2004 12:39am
Re: Filesystem	foundation	1	Nov 11, 2004 1:18am
Re: Filesystem	brewster	1	Nov 11, 2004 1:40am
Re: Filesystem	dunno	1	Jun 16, 2005 1:55pm
Re: Filesystem	foundation	1	Jul 15, 2005 5:08am
Re: Filesystem	Plasma 000	0	Aug 2, 2018 9:58pm
my shot at massive network storage.	dunno	3	May 31, 2004 10:26pm
Re: my shot at massive network storage.	dunno	1	Jun 2, 2004 10:33pm
Re: my shot at massive network storage.	andyj	1	Jun 13, 2004 10:55pm
Re: my shot at massive network storage.	foundation	1	Nov 10, 2004 11:52pm
Re: my shot at massive network storage.	angelbassmuffin	0	Dec 21, 2005 1:11pm
Bootstrapping from USB	Youssef Eldakar	2	May 24, 2004 11:19pm
Re: Bootstrapping from USB	billmoyer	1	May 28, 2004 9:11am
1U power supplies	Jeff Anderson-Lee	1	May 20, 2004 2:17am

[View more forum posts](#)