

Petabyte Box for Internet Archive.

<http://www.Baumgart.org/petabytebox.pdf>

<http://www.Planetten.biz/MDC/mdc.ppt>

Bruce Baumgart and Matt Laue

pro bono for Internet Archive

8 November 2003

Abstract

Brewster Kahle, founder of the Internet Archive, suggested that a petabyte of data capacity can now fit inside a standard 40' x 8' x 8.5' cargo shipping container. This white-paper presents floor plans, cost estimates, electrical power and cooling requirements as well as the rationale for running large commodity storage systems in movable cargo containers. *The poster version has only some of the text.*

1 Innovative thinking inside the box.

Think about running a copy of the **Internet Archive** inside a cargo container. In late 2003, a full copy of the Internet Archive, space for 1024 terabytes can fit in one container. Although special mobile computer rooms exist for military and oil industry users, computer systems manufactured in room scale enclosures have not yet appeared in the general commercial market. In high rent areas the expensive buildings are best used for people, inexpensive containers can shelter the server farm computer equipment. The whitepape will be present details about implementing an outdoor petabyte NAS JBOD box; that is Network Attached Storage, Just a Bunch of Disks in a standard cargo shipping container.

2 Technical Details.



- 2.1 Raw data capacity \approx Secure data payload.
- 2.2 Cost preview: Dime/GB/month.
- 2.3 Space and Weight: 320 sq feet and 15 to 20 tons.
- 2.4 Electrical Power: 100 to 200 kilowatts.
- 2.5 Cooling Requirement: 20 to 40 tons.
- 2.6 Bandwidth: 1Gb by fiber and 7Gb by sea.

A gigabit fiber can, in theory, move 10 terabytes per day.

$$1 \text{ Gbit/sec} \times 1 \text{ day} \equiv \left(\frac{1 \text{ Gbit/sec} \times 24\text{hr/day} \times 60\text{min/hr} \times 60\text{sec/min} \times 2^{30}\text{bits/Gbit}}{8 \text{ bits/Byte} \times 2^{40} \text{ bytes/TeraByte}} \right)$$

$$\approx 10 \text{ TeraBytes per day.}$$

Shipping a container by sea from San Francisco to Sydney takes two weeks. A petabyte per fortnight is nearly 7 Gigabits per second, which is a little less than an OC192 at 10 Gigabits per second.

3 Rationale.

3.1 Archive Clones.

Although the initial reason for the Internet Archive to consider placing a lot of computer disks in a cargo container was a co-location lease dispute; the sustaining fundamental reason is that cargo containers now look like the easiest way to replicate copies of the whole Internet Archive for starting new sites around the world.

3.2 Romance of books on ships.

The ancient first library at Alexandria in Egypt copied books from the container ships of that day to build a collection of all human knowledge. Once such an archive exists, the flow of information needs to become duplex. Now copies of the archive are to be made and placed on ships to lower latency, to raise literacy, to lower the risk of losing old knowledge, and to help in the creation new knowledge.

3.3 Market size daydreaming.

Ten petabyte boxes in 2004. The Internet Archive hopes to build eight copies of itself around the world over the next three years pursuant of the International Internet Preservation Consortium work. At present the consortium members are: the Internet Archive, the United States Library of Congress, the British National Library, France, Italy, Australia, and the five Scandinavian countries. Additional Internet Archive sites have been discussed with people in Japan, India and China.

One hundred boxes in 2005. There are fifty major cities in the United States, with one major university or a digital library each. Many entities need to convert large collections of film, microfiche, analog video tape to near line digital media.

One thousand boxes in 2006. Next there is the S.E.C. reference data mandated archives for the Fortune 1000 companies. The reference data market might require a couple of hundred petabyte box units.

One further neat thing is that a petabyte box shrinks from forty feet in 2004, to twenty feet in 2005 and to a ten foot cube in 2006.

4 Answering objections.

4.1 Where do you park them ?

4.2 Security ?

4.3 What is the risk of damage during a move ?

4.4 Maintenance costs are unknown.

4.5 Why hasn't this been done already ?

4.6 It doesn't make a big difference.

4.7 It will look ugly.

4.8 The proposal lacks details.

4.9 It's only a niche market.

5 Node parameters.

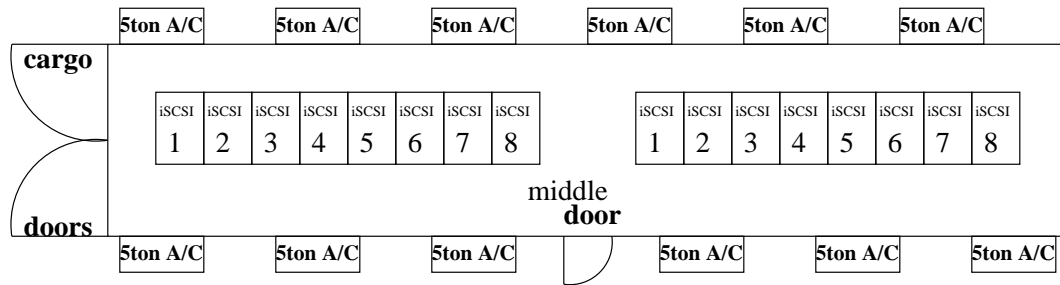
Table 1: Node parameters

Enclosure node,box,brick	# disks	size GB	total GB	volts	amps	watts	height	width	depth	volume	weight
IA HP	4x	160	640	120 v	1.70 A	200 w	24 in	8 in	30 in	5760 in^3	? lb
IA MidTower	8x	160	1280	120 v	? A	? w	17 in	8 in	17 in	1156 in^3	? lb
IA CasEdge	4x	250	1000	120 v	0.95 A	125 w	14 $\frac{1}{2}$ in	5 $\frac{1}{4}$ in	16 in	1218 in^3	20.5 lb
IBM Grid Acase	1x	40	40	120 v	? A	95? w	4 in	13 in	16 in	832 in^3	? lb
ITX low power	2x	250	500	120 v	? A	50? w	4 in	11 $\frac{1}{4}$ in	13 $\frac{1}{4}$ in	608 in^3	? lb
ITX custom case	4x	250	1000	120 v	? A	70? w	4 in	11 $\frac{1}{4}$ in	13 $\frac{1}{4}$ in	608 in^3	? lb
ITX Shuttle box	3x	250	750	120 v	? A	90? w	7 $\frac{1}{2}$ in	8 in	11 $\frac{1}{4}$ in	675 in^3	? lb
WF 4U whitebox	8x	160	1280	220 v	? A	? w	4 U	17 in	26 in	3094 in^3	47 lb
WF 1U xSeries 335	2x	36	72	220 v	0.83 A	180 w	1 U	17 in	? in	? in^3	28 lb
WF iSCSI engine	6x	73	438	220 v	1.00 A	220 w	4 U	17 in	? in	? in^3	? lb
WF iSCSI extension	14x	73	1022	220 v	0.92 A	200 w	3 U	17 in	? in	? in^3	? lb
WF iSCSI full	34x	73	2482	220 v	2.84 A	625 w	10 U	17 in	? in	? in^3	? lb


6 Appendix - Floor plans.

Standard 19" equipment racks, wide shelves and narrow shelves for current IA EdgCase boxes.


Web Fountain 16 racks of iSCSI floor plan.



Internet Archive one aisle, 30" shelf, floor plan.



Internet Archive 3 aisles, 18" shelves, floor plan.



7 Dime per gigabyte per month details.

Table presents an estimate of the TCO, Total Cost of Ownership, of one petabyte for three years.

Table 2: Example TCO for IA a cancelled IBM innovation demonstration.

Component	Who	Quantity	Unit Cost	Extended Cost
customized container	IBM	1	\$23,000.	\$23,000.
air conditioners 5 ton	IBM	6	\$5,000.	\$30,000.
whitebox 4-slot nodes	IA	1000	\$300.	\$300,000.
Maxtor 250GB disks	IA	4000	\$250.	\$1,000,000.
network switches	IA	30	\$500.	\$15,000.
open software	IA	many KLOC	\$0.	\$0.
proprietary software	IBM	0	\$?.	\$0.
3 years operations labor	IA	2 days/week	\$50/hr	\$124,800.
3 years electricity	IA	1,200 Mwh/year	\$150/Mwh	\$549,000.
3 years T1 communication	IA	36 mo	\$500/month	\$18,000.
3 years parking fees	IA	36 mo	\$250/month	\$9,000.
decommission/salvage	IA	1	\$guess?.	net \$5,000.
Total Cost Ownership	both	3 years 1 petabyte	estimate	\$ 2,073,800.

NOTE: This is 6 cents per gigabyte per month, compared with the IBM AIS charge to research of \$40 per gigabyte per month. NOTE: The electricity estimate is based on 120 watts per 4 disk node, 30 tons average cooling, and 15 cents per kilowatt hour. Your mileage will vary. NOTE: At an archive expansion rate of 80T per month it will take about one year to populate the box.

Table 3: cargo container customization cost estimate

Item	Cost estimate
Container Steel Re-Con 40' HC	\$3,000.
Paneling and Insulation	\$3,000.
Electrical package 800 amps 240v	\$10,000.
Extra personnel door	\$1,000.
Shelving	\$3,000.
Extra ducting and ventilation fans	\$1,000.
Fluorescent lights, locksmith, logo, misc	\$2,000.
Total cost estimate	\$23,000.