



Environmental efficiency of tape storage at CERN

Vladimír Bahyl

IT-SD-TAB

[Newsroom](#) / [Information Technology](#) / [Press Release](#)

Gartner Predicts Power Shortages Will Restrict 40% of AI Data Centers By 2027

Sydney, Australia, November 12, 2024


Rapid Growth in Energy Consumption For GenAI Will Exceed Power Utilities' Capacity

AI and **generative AI** (GenAI) are driving rapid increases in electricity consumption, with data center forecasts over the next two years reaching as high as 160% growth, according to Gartner, Inc. As a result, Gartner predicts 40% of existing AI data centers will be operationally constrained by power availability by 2027.


“The explosive growth of new hyperscale data centers to implement **GenAI** is creating an insatiable demand for power that will exceed the ability of utility providers to expand their capacity fast enough,” said **Bob Johnson**, VP Analyst at Gartner. “In turn, this threatens to disrupt energy availability and lead to shortages, which will limit the growth of new data centers for GenAI and other uses from 2026.”

Gartner estimates the power required for **data centers** to run incremental AI-optimized servers will reach 500 terawatt-hours (TWh) per year in 2027, which is 2.6 times the level in 2023 (see Figure 1).

CERN electricity consumption structure



CERN Data Centre electricity consumption structure




Source: M. Schulz: CERN-IT Sustainability

... and WLCG computing


Unprecedented data rates from the experiments to Tier-0 and WLCG during pp and heavy-ion runs.

In 2024 the experiments collected ~ 220 PB of raw data (to be compared with ~100 PB in 2023 and ~ 40 PB in 2022). Average of ~ 1 PB/day, peaking at 2 PB/day in September.

Tier-0 and WLCG services supported the computing needs of the experiments with high performance and reliability.




The average weekly WLCG data transfer rate since 2021 (Data includes raw, reconstructed, Monte Carlo)



Amount of LHC data stored

Data grouped by VO or experiment ⓘ



Tape Infrastructure

(February 2025)



- Archive of the physics data
- Provisioned capacity: ~1.25 EB
- Libraries:
 - 4 x IBM TS4500
 - 3 x Spectra Logic TFinity
- Drives:
 - 40 x IBM TS1170, 46 x IBM TS1160
 - 88 x LTO-9, 10 x LTO-8
- Media:
 - 250 PB on 3592JF, 150 PB on 3592JE, 227 PB on 3592JD
 - 551 PB on LTO-9, 17 PB on LTO-8, 59 PB on LTO-7M



- Backup of the business data
- Licensed capacity: ~15 PB
- Libraries:
 - 1 x IBM TS4500 (partitioned)
 - 1 x Spectra Logic TFinity (partitioned)
- Drives:
 - 10 x LTO-9
 - 10 x LTO-8
- Media:
 - 12 PB on LTO-8
 - 11 PB on LTO-7M

IBM CO₂e Report

50.0
metric tons
CO₂e


1.4
metric tons
CO₂e

Physical Tape
CO₂e Report


Breakdown of the CO₂e and Other Positive Sustainability Impacts of IBM Physical Tape Products
The Product Life Cycle

Weight	2.9kg
Recycle Offset	3.2kg
Energy Usage (Average) Life Cycle	.031 kWh
	6.85Yrs

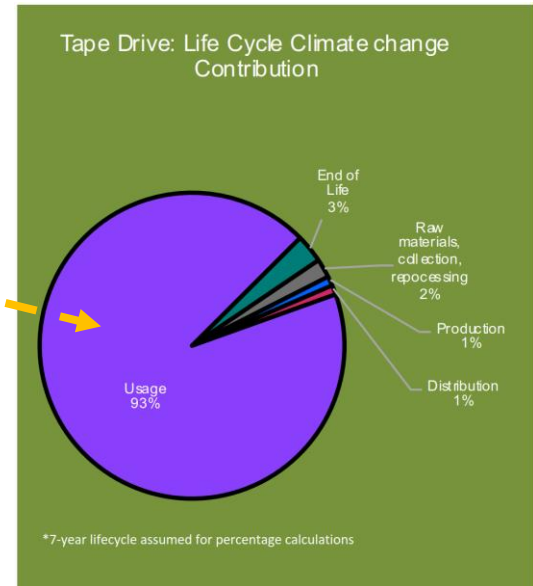
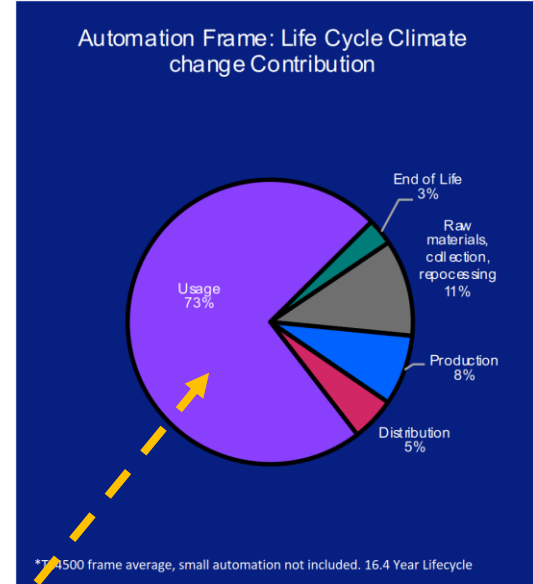
Production	41.3kg
MFG Distribution	.4kg
Operational Energy	836.4kg
Cooling Energy	347.1kg
End-of-Life	4.94kg



- Aluminum
- CU
- Plastics
- Silicon Card
- Silicon Processor
- Tin



- FE
- Aluminum
- CU
- Plastics
- Silicon Card
- Silicon Processor
- Tin



Shawn O. Brume Sc.D.
IBM Report Analyst
Shawn.Brume@ibm.com
January 2022



Environmental footprint of 1.2 EB – best case

- Could comfortably fit into two IBM TS4500 tape libraries, each with 15000 slots (IBM TS1170 drives and 3592JF cartridges)

Disclaimer: First version of the sustainability calculator tool

Tape solution being chosen:	TS4500
Geographic region of deployment:	Europe
TS4500 Library frame	1
Number of TS4500 Drive frames:	2
Number of TS4500 S-frames:	15
Number of TS1100 tape drives:	48
Number of LTO Full High tape drives:	0
Enter number of Cartridges:	15000
Enter type of media:	3592 Media
The total CO ₂ e contribution for a TS4500 configuration as listed, with all shipping to Europe is:	
	530 mt
Based on the TS4500 having a 9-year life cycle.	

$$2 \times 530 \text{ mt} = 1060 \text{ mt} / 9 \text{ years} = 118 \text{ mt per year}$$

$$118 \text{ mt} / 1200 \text{ PB} = \sim 0.1 \text{ mt of CO}_2\text{e per PB per year}$$

$$118 \text{ mt of CO}_2\text{e} = \sim 30 \text{ petrol cars driving 15000 km per year}$$

Source: <https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator>

- Reality – more hardware is needed:
 - Throughput requirement – at least 40 GB/s
 - Cost optimisation = using cheaper media
 - Multiple generations of tape drive technologies
 - Protection against library failures / outages
 - Avoiding vendor lock-in
 - Configuration not constant / Variable growth

Environmental footprint of 1.2 EB – reality – 1/2

IBMLIB1		IBMLIB2		IBMLIB3		IBMLIB4	
Tape solution being chosen:	TS4500	Tape solution being chosen:	TS4500	Tape solution being chosen:	TS4500	Tape solution being chosen:	TS4500
Geographic region of deployment:	Europe	Geographic region of deployment:	Europe	Geographic region of deployment:	Europe	Geographic region of deployment:	Europe
TS4500 Library frame	1	TS4500 Library frame	1	TS4500 Library frame	1	TS4500 Library frame	1
Number of TS4500 Drive frames:	2	Number of TS4500 Drive frames:	3	Number of TS4500 Drive frames:	2	Number of TS4500 Drive frames:	2
Number of TS4500 S-frames:	15	Number of TS4500 S-frames:	14	Number of TS4500 S-frames:	15	Number of TS4500 S-frames:	15
Number of TS1100 tape drives:	0	Number of TS1100 tape drives:	0	Number of TS1100 tape drives:	48	Number of TS1100 tape drives:	48
Number of LTO Full High tape drives:	38	Number of LTO Full High tape drives:	10	Number of LTO Full High tape drives:	0	Number of LTO Full High tape drives:	0
Enter number of Cartridges:	17500	Enter number of Cartridges:	1014	Enter number of Cartridges:	15394	Enter number of Cartridges:	13639
Enter type of media:	LTO Media	Enter type of media:	LTO Media	Enter type of media:	3592 Media	Enter type of media:	3592 Media
The total CO2E contribution for a TS4500 configuration as listed, with all shipping to Europe is:	450 mt	The total CO2E contribution for a TS4500 configuration as listed, with all shipping to Europe is:	296 mt	The total CO2E contribution for a TS4500 configuration as listed, with all shipping to Europe is:	534 mt	The total CO2E contribution for a TS4500 configuration as listed, with all shipping to Europe is:	520 mt
Based on the TS4500 having a 9 year life cycle.		Based on the TS4500 having a 9 year life cycle.		Based on the TS4500 having a 9 year life cycle.		Based on the TS4500 having a 9 year life cycle.	

Disclaimer:
Extrapolated calculation, using IBM methodology on Spectra Logic TFinity product

SPECTRALIB1		SPECTRALIB2		SPECTRALIB3	
Tape solution being chosen:	TS4500	Tape solution being chosen:	TS4500	Tape solution being chosen:	TS4500
Geographic region of deployment:	Europe	Geographic region of deployment:	Europe	Geographic region of deployment:	Europe
TS4500 Library frame	1	TS4500 Library frame	1	TS4500 Library frame	1
Number of TS4500 Drive frames:	2	Number of TS4500 Drive frames:	2	Number of TS4500 Drive frames:	1
Number of TS4500 S-frames:	12	Number of TS4500 S-frames:	12	Number of TS4500 S-frames:	12
Number of LTO Full High tape drives:	28	Number of LTO Full High tape drives:	26	Number of LTO Full High tape drives:	6
Enter number of Cartridges:	10571	Enter number of Cartridges:	11295	Enter number of Cartridges:	810
Enter type of media:	LTO Media	Enter type of media:	LTO Media	Enter type of media:	LTO Media
The total CO2E contribution for a TS4500 configuration as listed, with all shipping to Europe is:	332	The total CO2E contribution for a TS4500 configuration as listed, with all shipping to Europe is:	334 mt	The total CO2E contribution for a TS4500 configuration as listed, with all shipping to Europe is:	232
Based on the TS4500 having a 9 year life cycle.		Based on the TS4500 having a 9 year life cycle.		Based on the TS4500 having a 9 year life cycle.	

Environmental footprint of 1.2 EB – reality – 2/2

- Summary:
 - IBM: $250 + 296 + 534 + 520 = 1800$ mt / 9 years
 - Spectra Logic: $332 + 334 + 232 = 898$ mt / 9 years
 - Total: 2698 mt / 9 years = **300 mt / year**
- 300 mt / 1200 PB = **0.25 mt of CO₂e per PB per year**
- 300 mt of CO₂e = ~80 petrol cars driving 15000 km per year
Source: <https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator>

- Comparison: *The Dirty Secret of SSDs: Embodied Carbon* (<https://arxiv.org/pdf/2207.10793>)
 - Table 1 – 1 TB HDD generates 100 Kg of CO₂e per 5 years = 20 Kg per year
 - This includes server hardware
 - With ~100000 HDDs at CERN, that would be around 2000 tons of CO₂e per year

Energy consumption of 1.2 EB tape archive – 1/2

- IBM

- Considering maximum values from: <https://www.ibm.com/docs/en/ts4500-tape-library?topic=planning-power-consumption-cooling-requirements>

- IBMLIB1 : 1925 W
 - IBMLIB2 : 939 W
 - IBMLIB3 : 3543 W
 - IBMLIB4 : 3543 W


- **Total IBM (maximum continuous): 9950 W**

Maximum continuous power is consumed when the drives are actively reading or writing to the tape and the cooling fan rotates at normal speed.


	IBMLIB3		
Frame model, feature codes, and drives	Quantity	Per unit	Max total
L25	1,00	130,00	130,00
FC 1460 (redundant accessor power/network)	1,00	5,00	5,00
FC 1531, 1532, 1533, and 1534 (one of each FC per frame with drives)	12,00	3,00	36,00
D25	2,00	11,00	22,00
FC 1450 (TS4500 FCA including one LCC and two ac/dc power supplies)	2,00	38,00	76,00
S25	15,00	11,00	165,00
FC 1442 (HA kit with second accessor; for Dx5 or Sx5)	1,00	85,00	85,00
LTO drives	0,00	37,00	0,00
TS1160 drives	48,00	63,00	3024,00
Total library power consumption (watts)			3543,00

Energy consumption of 1.2 EB tape archive – 2/2


- Spectra Logic
 - Considering maximum values as reported by the libraries:
 - SPECTRALIB1: 3200 W
 - SPECTRALIB2: 3200 W
 - SPECTRALIB3: 1600 W
 - Total Spectra Logic (estimation): 8000 W
- TOTAL: **~18 kW per 1.2 EB = ~15 W per PB**
- Comparison: According to PDU monitoring, one rack at CERN usually consumes ~3 kW (example PDU-SJ06-01)



33 PB in 4 weeks = ~1 PB/day



18 PB in 9 days = ~2 PB/day



Conclusion

- Quantify environmental and sustainability impact of any (IT) equipment is a complex exercise
- The carbon footprint is dominated by the usage pattern and the capacity (number of tape cartridges)
- Calculations are based on IBM report (Spectra Logic extrapolated (as drives and cartridges are the same))
- CERN Tape Archive:
 - Generates around 300 mt CO₂e per year (80 petrol cars driving 15000 km per year)
 - Consumes around 18 kW (around 15 W per PB)
- Spikey usage demands larger configuration and has bigger environmental footprint
- Adding more:
 - Capacity will generate more CO₂e, but will not increase power consumption
 - Throughput (assuming same drive generations) will generate more CO₂e as well as increase power consumption
- Future work:
 - Include server hardware
 - Improved IBM tool (allow for different energy sources; consider growth of the archive)