

# AI Spending and ESG Pressures Drive Demand for TPM




CAPRAE CAPITAL

# “Millions of GPU data centers are coming”

-Jensen Huang (CEO, NVIDIA)

What does this mean for you? Well, it reflects the fastest growth of an industry that will change human life forever. This is as revolutionary as the invention of the automobile, allowing humankind to travel further than ever before, but with the side effects of devastating pollution.

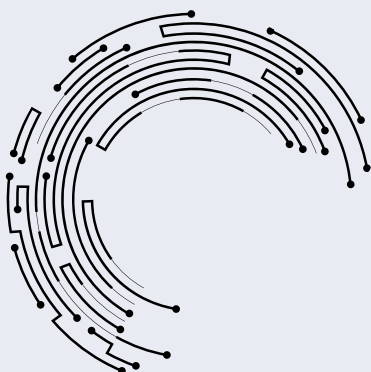
Similarly, AI will allow us to travel to a new reality and reach a higher plane of understanding. It is common knowledge that GPUs demand an overwhelming amount of power. Even worse, without extended maintenance, companies would dispose of their IT assets, often unsafely leaking toxic chemicals into the environment. In fact, only 22% of companies dispose of their IT assets safely. Additionally, the massive spending on high-priced GPUs and data centers will force companies to slash their capital expenditures, preventing them from refreshing their existing IT assets. This could leave systems outdated and vulnerable to critical failures and security breaches. Thus, the GPU revolution could make the world unlivable and cause data infrastructure to crumble.



Yet OEMs stand as a formidable adversary to financially and environmentally compelling server extensions, doing their utmost to prevent customers from extending the useful lives of servers. How, then, will technology companies simultaneously act for the noble aims of financial competitiveness and environmental sustainability, saving us from the grip of these giants? The solution lies with third-party maintenance companies.

## THESIS

As the demand for GPUs/AI increases, companies will need to extend their refresh cycles beyond the limit imposed by OEMs. This compels companies to move towards TPM. Additionally, as major technology companies experience pressure to act on their ESG goals, the TPM imperative will only strengthen.



Specifically, the growth rate of GPUs was projected at 34.6% and is likely even higher at 60% per year, thus reaching a predicted market size of \$120 bn by 2032. In parallel, the AI market growth rate was 28%. The accelerated growth by both markets is unlike anything we’ve seen in this era, and could be compared to the rapid rise of internet users in the early 2000’s. Already, the bulk of capex has been displaced towards the GPU market and away from buying current servers, as reflected in the extended refresh cycles.

Additionally, there is a strong environmental case to be made for switching to TPM. Without TPMs, companies would be forced replace their servers every 3 – 5 years, creating a catastrophic increase in E-waste, which is already up 82% since 2010, resulting in a long trail of environmental damage. The Earth is struggling to bear the weight of this mounting electronic debris, threatening our environment and health with irreversible damage. On top of that, the new servers that companies must then buy require an unfathomable amount of carbon emissions: in fact, 50% of the emissions over the entire server life comes from manufacturing. This is where TPM comes in, an indispensable part of the quest for Green IT.

As TPM is far cheaper and environmentally sustainable than buying servers more often, companies will be under strong pressure to move towards TPM. We expect TPM to grow by 20% per year over the next few years, especially with the additional focus on TPM and ESG. Investing in TPM is a tremendous opportunity to offset the reduction in the ITAD market due to these very same factors.

**MARKET LANDSCAPE/OVERVIEW**

**Third-Party Maintenance Market Overview:**

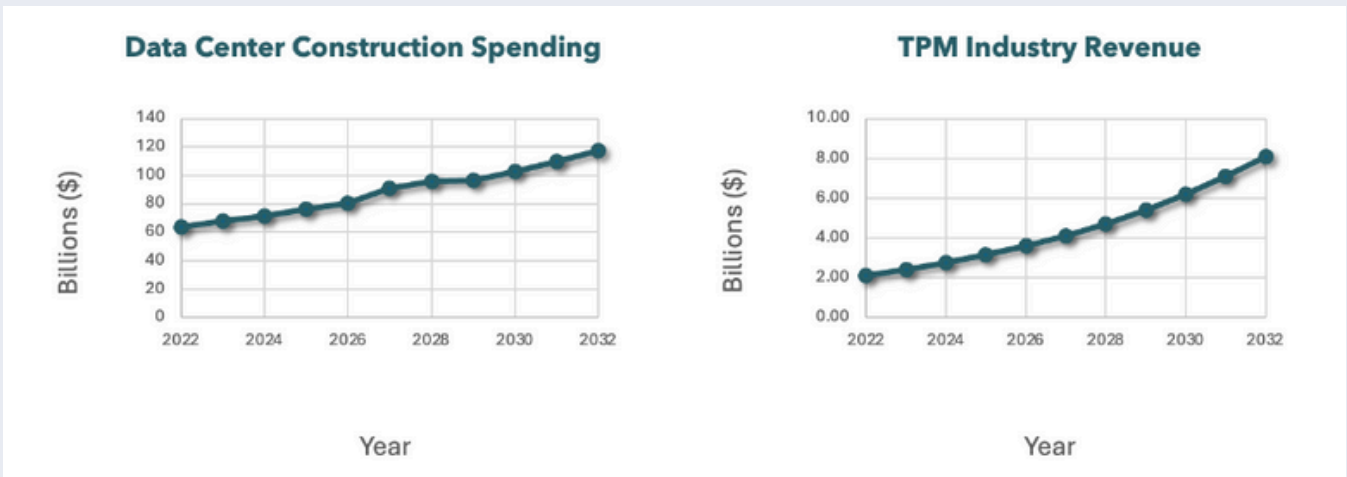
Third-Party Maintenance refers to IT maintenance by a company other than the OEM (original equipment manufacturer) – a server manufacturer like Dell, Supermicro, or Lenovo. Typical candidates for TPM include server, storage, and network. The Third-Party Maintenance (TPM) Industry was sized at \$2.1 billion in 2022 with a projected growth rate of 14.43 – 14.52%.



This implies that it has reached \$2.75 billion now and will reach \$6.2 billion by 2030. However, with the current extended refresh cycles, TPM may grow at a higher rate going forward. According to Gartner, over 10 million data center and network devices were under a TPM plan, and 71% of large companies used a TPM for support in 2016.




## Comparison



There is a clear correlation between the construction of data centers and the growth of TPM in data centers. As data center construction has sped up, with the United States constructing over 2,000 data centers in the past two years, so has TPM. We believe that the direct causation is actually lagged as TPM comes in 3 – 6 years after a data center is constructed. As companies have invested heavily into data centers and hardware, the demand to maintain the hardware will increase accordingly.

## Main Players and Consolidation



Market Shares of Key Players <sup>1920</sup>		
TPM Company	TPM Revenue	Market Share
Park Place	\$674,000,000.00	28.04%
Evernex	\$296,900,000.00	12.35%
Service Express	\$290,000,000.00	12.06%
Cxtec	\$215,700,000.00	8.97%
Maintech	\$108,200,000	4.50%
InKnowTech	\$92,400,000.00	3.84%

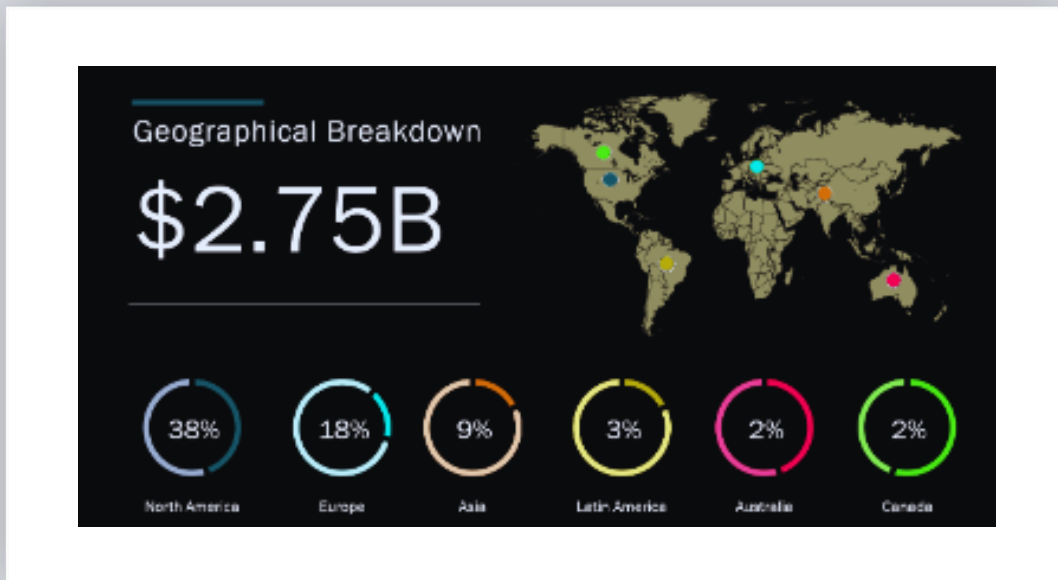
In 2023, Park Place Technologies, Service Express, CxTec, and Evernex dominated the TPM industry, holding over 60% of the market. Park Place, controlling nearly 30% through 18 acquisitions (including Curvature, NorthSmart, and Riverstone), is the leader—earning \$674M in revenue, the highest among its peers. Backed by \$2B from Blackstone in April, Park Place plans to continue acquiring companies to maintain its dominance.

The next largest player is Evernex, which controls about 12.4% of the market with nearly \$300,000,000 of revenue. Evernex has also pursued acquisitions such as EmconIT, Technogroup, and A Systems. Evernex has also received funding from major private equity firms. It has received a \$451 million LBO investment from the 3i Group in 2019, financed with an undisclosed amount of Debt from CIC Private, Eurazeo, and Tikehau Capital Debt on October 29, 2019. Previously, the owner was another major investment firm, the Carlyle Group.




**SERVICE AND GEOGRAPHICAL SEGMENTS**

Across the board, the Third-Party Maintenance market consists of Storage Maintenance, Server Maintenance, and Network Maintenance. Many companies conduct all three main types of maintenance. North America is expected to contribute 33% of the growth.



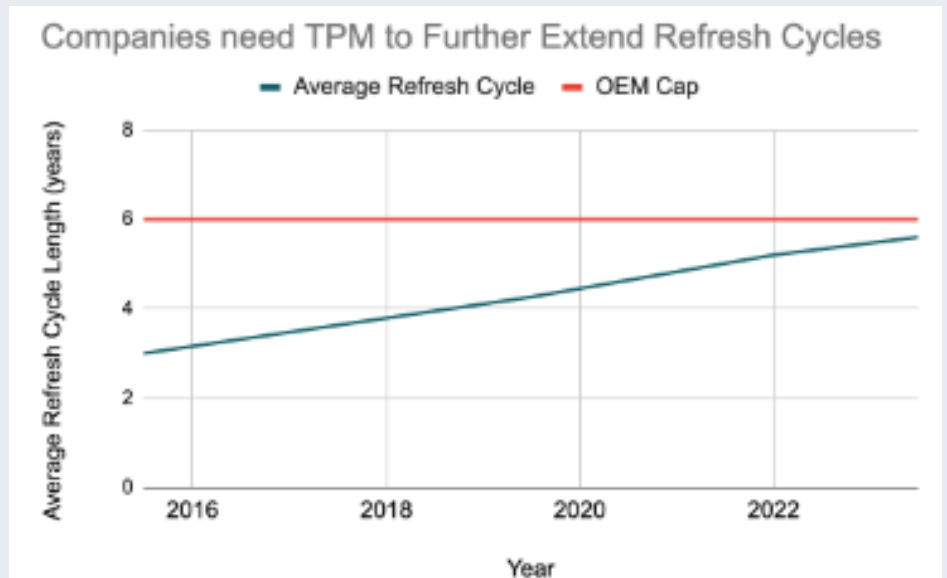
According to the number of data centers, approximately 37.79% of the market is in the USA, 17.68% in Europe, 8.47% in Asia, 2.72% in LatAm, 2.36% in Canada, and 2.15% in Australia. Asia will grow the fastest at 38% alongside the data center growth in the region. This is because growth in the TPM industry essentially lags the growth in data centers.



**OPERATIONS AND BUSINESS MODELS**

Companies may have different models for engineering and parts procurement. Some handle all support, whereas others use subcontractors to handle parts and hire engineers. All Third-Party Maintenance companies require capital expenditures on repair and replacement materials and the warehouses to store them in. It is not necessary to work completely with either an OEM or TPM, but it is possible to blend the two.


For certain IT maintenance services on specific devices a company might choose OEM, but for others it could choose TPM. For the first 3 years, there is usually full warranty from the OEM, but once a device gets to 3 - 6 years, End-of-Sale announcements come out, and maintenance becomes more expensive. After 6 years, OEMs put out End-of-Serviceable Life announcements, withdrawing service.



In effect, OEMs have a tactic similar to Apple or car manufacturers, where when the old model passes a certain date, they refuse to maintain it any longer. This forces customers to buy a new product. However, the refresh cycle reached as high as 5.6 years as of 2023, pushing the upper limit of OEM support, and some companies plan on using servers for as long as 10 years. When OEM maintenance is not practical, companies can switch to TPM – allowing them maintain servers for up to 12 years. to save money and make a better impression from an environmental standpoint.

Note that the average refresh cycle does not mean that all companies have that refresh cycle. It is an average across all companies and there is a distribution, with the refresh cycle differing for each company. Thus, many companies have >6-year refresh cycles. Furthermore, within each company, different types and specific items of IT could have different ages.

**6 Years is the Maximum Timeline Until Servers Have to be Replaced**



Typical hardware consists of servers, storage, racks, network equipment, etc. This timeline depicts how OEMs provide a warranty for only the first three years, hoping the customer will replace the hardware after the third year. However, they also offer support (post warranty) for maintenance for years 3-9 but confidently expect their old hardware after 9 years, which they

consider the "End of Service Life" for hardware. We see that servers have a maximum refresh timeline, but companies may refresh their hardware before the 6 – 9 year-maximum.

## Typical 3-year Refresh Cycle

### OEM Refresh

60 member PowerEdge x86 Cluster

1/3 Refreshes Annually

6-Year Total Cost  
Hardware \$900,000

All Hardware refreshed at warranty expiration

## Typical 6-year Refresh Cycle

### OEM Maintenance/Refresh

60 member PowerEdge x86 Cluster

1/6 Refreshes Annually

6-Year Total Cost  
Hardware \$450,000  
OEM Maintenance \$270,000

**Total Cost \$720,000**

### TPM Maintenance/Refresh

60 member PowerEdge x86 Cluster

1/6 Refreshes Annually

6-Year Total Cost  
Hardware \$450,000  
TPM Maintenance \$120,000

**Total Cost \$570,000**

For 60 individual PowerEdge servers, all based on the x86 architecture. 3-year refresh cycles would cost \$900,000 over 6 years. In contrast, OEM maintenance to reach a 6-year refresh cycle would reduce the cost to \$720,000 and TPM would reduce it further to \$570,000. TPM maintenance allows companies to save 21% compared to an OEM-enabled 6-year refresh cycle and 37% compared to a 3-year refresh cycle. As GPU prices surge, companies are scrambling to find efficient ways to save money on other IT.

### PAIN POINTS THAT IT MAINTENANCE COMPANIES ADDRESS

#### Server Extensions:



As reported in our other papers, technology companies need to extend server refresh cycles to buy GPUs for AI. Technology companies can extend the actual useful lives of IT through IT maintenance.

#### Security protection:



IT maintenance can help reduce cybersecurity risk and protect companies from the legal and reputational ramifications. 79% of organizations have suffered a detected ransomware attack, so this protection is crucial.

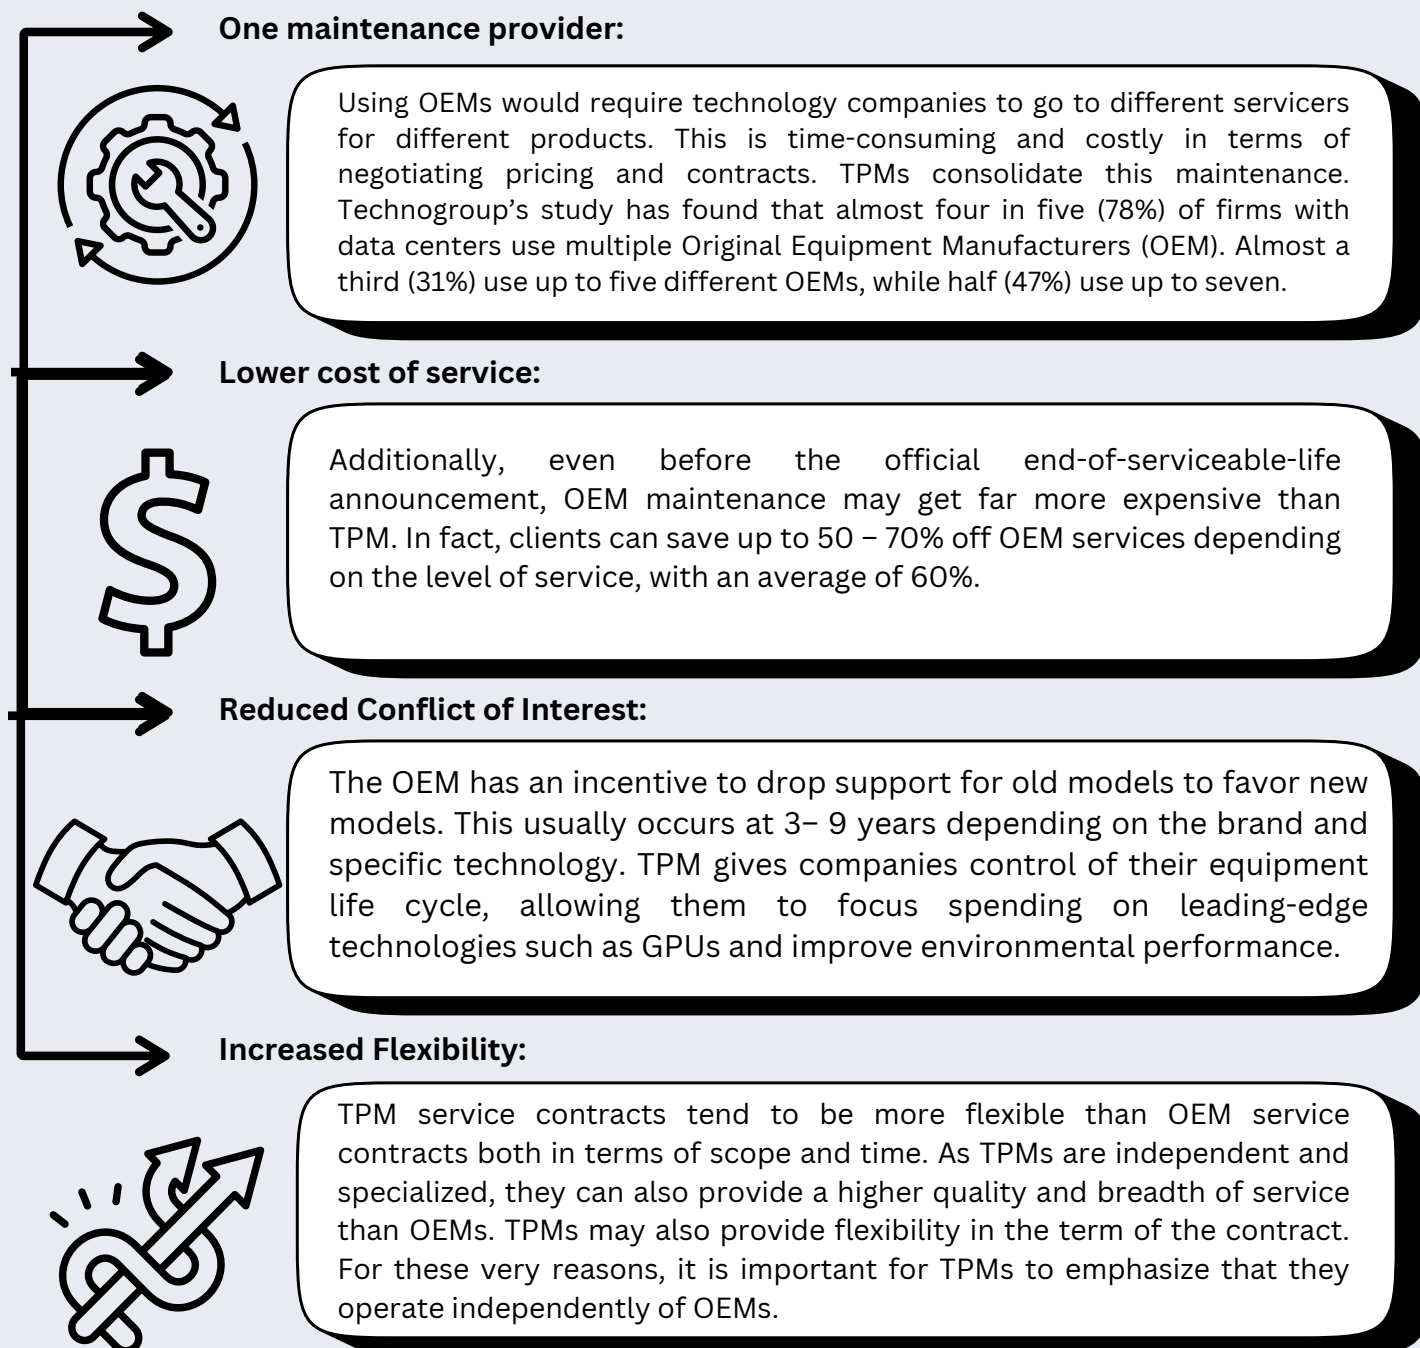
#### Need for scale:



Since IT maintenance is inherently a troubleshooting area, deep expertise, time flexibility, and variety of expertise make a strong case for economies of scale. Therefore, technology companies should outsource IT maintenance. In fact, outsourcing maintenance can lead to 60% savings. Outsourcing may also allow organizations to focus on their core areas.



## WHY TPM SPECIFICALLY?



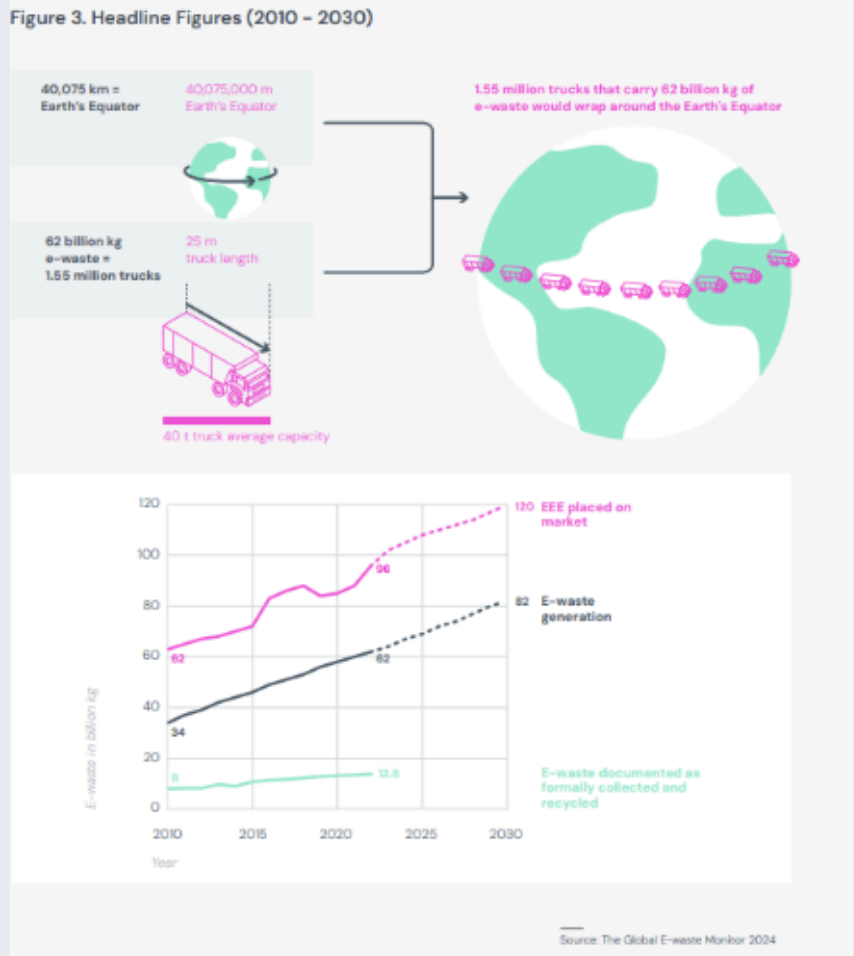
With so many reasons to choose third-party maintenance over OEM maintenance, it is no wonder that 70% of large companies now use TPM services.

## Amount of E-waste Generated

Electronic waste (e-waste), is a byproduct of data center refresh activity and is harming the environment significantly. The world generated 20–50 metric tons of E-Waste each year as of 2018, sending toxins such as heavy metals straight into landfills.

An EPA report states that e-Waste accounted for a staggering 70% of all toxic waste worldwide. The 62 million tonnes of e-waste generated in 2022 would fill 1.55 million 40-tonne trucks, roughly enough trucks to form a bumper-to-bumper line encircling the equator, according to the report from ITU and UNITAR.

If the unsustainable status quo continued, e-Waste would grow 8% each year, doubling every 9 years. The volume of global eWaste would reach 74.7– 82 metric tons by 2030. Specifically, the UN estimated that the electronic waste is rising five times faster than documented e-waste recycling.



Just as harmful, replacing machines and equipment leads to emissions. The carbon contribution for an average 1 MW data center includes 33,000lbs of emissions from its plastic, 73,000lbs from aluminum and 377,000lbs from steel.

**81 countries** have adopted e-waste policy, legislation or regulation.

**67 countries** have legal provisions on EPR for e-waste.

**36 countries** have provisions on e-waste recycling rate targets.

**46 countries** have provisions on e-waste collection rate targets.

Source: The Global E-waste Monitor 2024

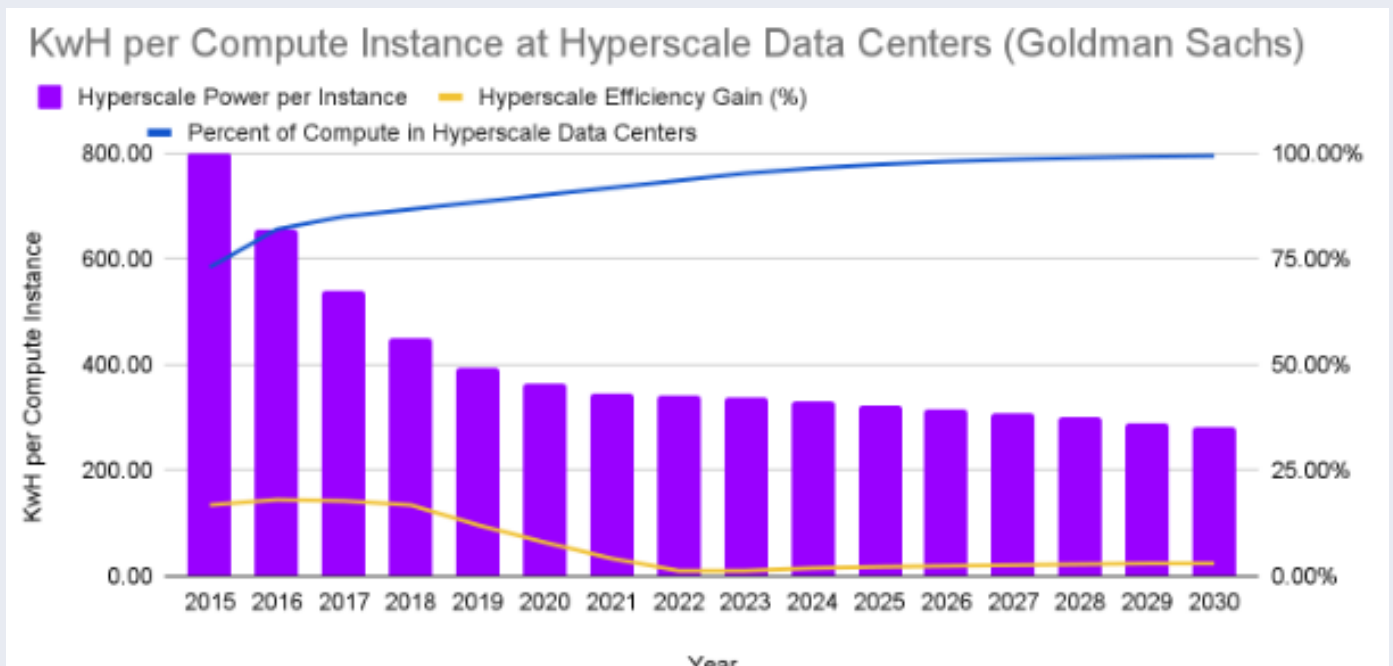
This highlights the severe environmental impact and the pressing need for sustainable solutions. E-waste reduction is thus an imperative on a global scale, but who is to blame? On average, 47% of data centers were refreshing their systems every 1-3 years and another 28% every 4-5 years and large companies with over \$1B+ in market cap were twice as likely to refresh systems each year compared to small businesses. Hyperscale data centers fit the bill perfectly, and this is exactly where TPM comes in.

TPM is by far the most efficient solution to the E-Waste problem. With having the ability to extend the life of hardware to 9-12+ years, it reduces the number of hardware being thrown out every year. If more companies adopted TPM for the majority of their hardware, we would see a major reduction in the waste problem. As mentioned before, Only 22% of this e-waste was disposed of safely, and even in Europe and Oceania, the best performing regions, only 42% of E-Waste was documented. This is where the true hero comes in, ITAD companies, which are known to securely and safely dispose of all E-waste with environmentally safe measures. We see companies recognizing this within the next 5 years because of the hard push for ESG and strict environment regulatory measures. Allowing ITAD companies to dispose of the harmful E-Waste on a massive scale would increase the percentage of e-waste disposed safely significantly.

Third-party IT support not only extends the working life of your equipment with data center service and parts, it also maximizes your initial investment: you can opt to keep your hardworking storage library for months or years without paying premium rates or risking downtime.

## Energy and Emissions

According to a report from Thinkstep, the manufacturing phase causes 50% of the total global warming potential across the server lifecycle. This means that replacing a server not only generates waste from disposal, but also leads to extensive Scope 3 emissions from the new server. Additionally, a report by Goldman Sachs projects the power use efficiency in hyperscale data centers to improve by much less than it had in the prior decade. This would mean that the energy use improvements from new hardware is not worth the emissions.



```
[6] import pandas as pd
from statsmodels.formula.api import ols
data = pd.read_csv("final_reg.csv")
model = ols("refresh ~ eff_gain", data = data).fit()
print(model.summary())
```

```
[7] OLS Regression Results
=====
Dep. Variable: refresh R-squared: 0.944
Model: OLS Adj. R-squared: 0.933
Method: Least Squares F-statistic: 84.26
Date: Sat, 15 Jun 2024 Prob (F-statistic): 0.000257
Time: 18:11:08 Log-Likelihood: 2.9986
No. Observations: 7 AIC: -1.997
Df Residuals: 5 BIC: -2.105
Df Model: 1
Covariance Type: nonrobust
=====
coef std err t P>|t| [0.025 0.975]
-----
Intercept 5.5546 0.125 44.338 0.000 5.233 5.877
eff_gain -11.2629 1.227 -9.179 0.000 -14.417 -8.109
=====
Omnibus: nan Durbin-Watson: 1.694
Prob(Omnibus): nan Jarque-Bera (JB): 0.784
Skew: -0.326 Prob(JB): 0.703
Kurtosis: 1.590 Cond. No. 17.5
=====
```

The source provides the historical hyperscale power per “compute instance”, the historical and forecasted hyperscale efficiency gains, and the historical percentage of compute in hyperscale data centers. A compute instance refers to a virtual machine, and in this diagram is standardized in size. GS defined the efficiency gain as the percentage decrease in power consumption per compute instance. Since over 90% of the world’s computing power is now located at hyperscale data centers, the effect of power efficiency gains at traditional data centers is insignificant. Therefore, we used the percentage hyperscale efficiency gain from the source to deduce the future hyperscale power per instance according to a multiplicative formula based on the definition of percentage change.

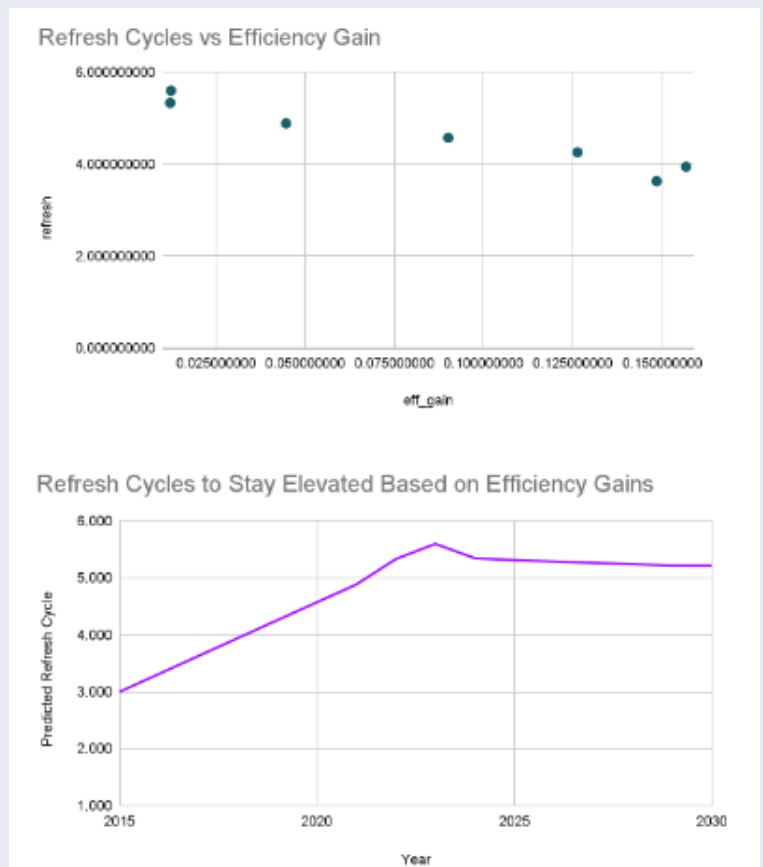
$$E(\text{kwh/cit}) = E(\text{kwh/cit-1}) * [1 - E(\text{powereffgaint-1})]$$



The chart is thus a combination of the Goldman Sachs forecasts and their implications for future energy consumption. Notice how the energy efficiency per instance is still improving, but much more slowly than before.

Our research question from this data was whether the reduced efficiency gain at servers contributed to the lengthened refresh cycles. We ran a regression analysis of the average refresh cycle length interpolated from the Refresh Cycles spreadsheet combined with the Goldman Sachs historical data. There was not enough evidence to state at conventional statistical confidence levels that either the number of GPUs or the GPU vs CPU market size ratio affects the refresh cycle with the limited data available.

Therefore, we only analyzed the efficiency gain, and found that as power use per instance decreases more slowly, the refresh cycles tend to lengthen.



Note that the statistical fit measures may not be very accurate due to limited data and interpolation: Only four of these data points were actually available and the rest were approximated. However, given refresh cycles likely stay continuous, the regression could be interpreted as if it had approximately 1.4x the standard error and  $R^2$  of ~70%, which is still a very good fit. Furthermore, the GPU boom has only begun, meaning that the effect has not yet factored into the data. Thus there is still room for GPU capex in the full explanation. However, using power efficiency gain alone, we can interpret the regression using prior-year data. According to the coefficients, a 1% decrease in the power efficiency gain in any given year, the refresh cycle tends to increase by 0.1126 years or 1.35 months. Since the power efficiency gain has been rapidly decreasing, we expect the average refresh cycles to stay very high based on this factor alone – above 5 years up to 2030. When refresh cycles are this high, TPM usage will likely be high as well. However, this association also suggests that a major risk is the adoption of newer, more efficient chips such as AMD EPYC 4th Gen Chips. Still, these efficient chips can be quite expensive even for their 96 cores, with one AMD Epyc 9654 costing \$11,800+ – as much as its GPUs. In other words, more efficient and powerful yet more expensive CPUs could become an indirect substitute for TPM.

## Environmental Regulations

Recent regulations around the world are starting to force OEMs to move towards a more circular, green IT lifecycle. Since 2020, UK manufacturers have been required to provide product and assembly information alongside replacement components, allowing TPMS to work more easily. The Right to Repair legislation would also reduce the power of OEMs to enforce their desired refresh cycle; a common provision is requiring OEMs to provide the information and tools necessary to repair machines. This means that the quality of TPM will increase and the costs related to TPM will decrease. Unfortunately, regulations that require OEMs to design more sustainable servers to begin with may work against the TPM industry. Given the emissions potential of the manufacturing process, however, we do not believe this impact will be severe. As the impacts of global warming and pollution become more evident, we predict that these regulations will become more common and stringent.

## Empirical Increase in Companies using TPM with Projection

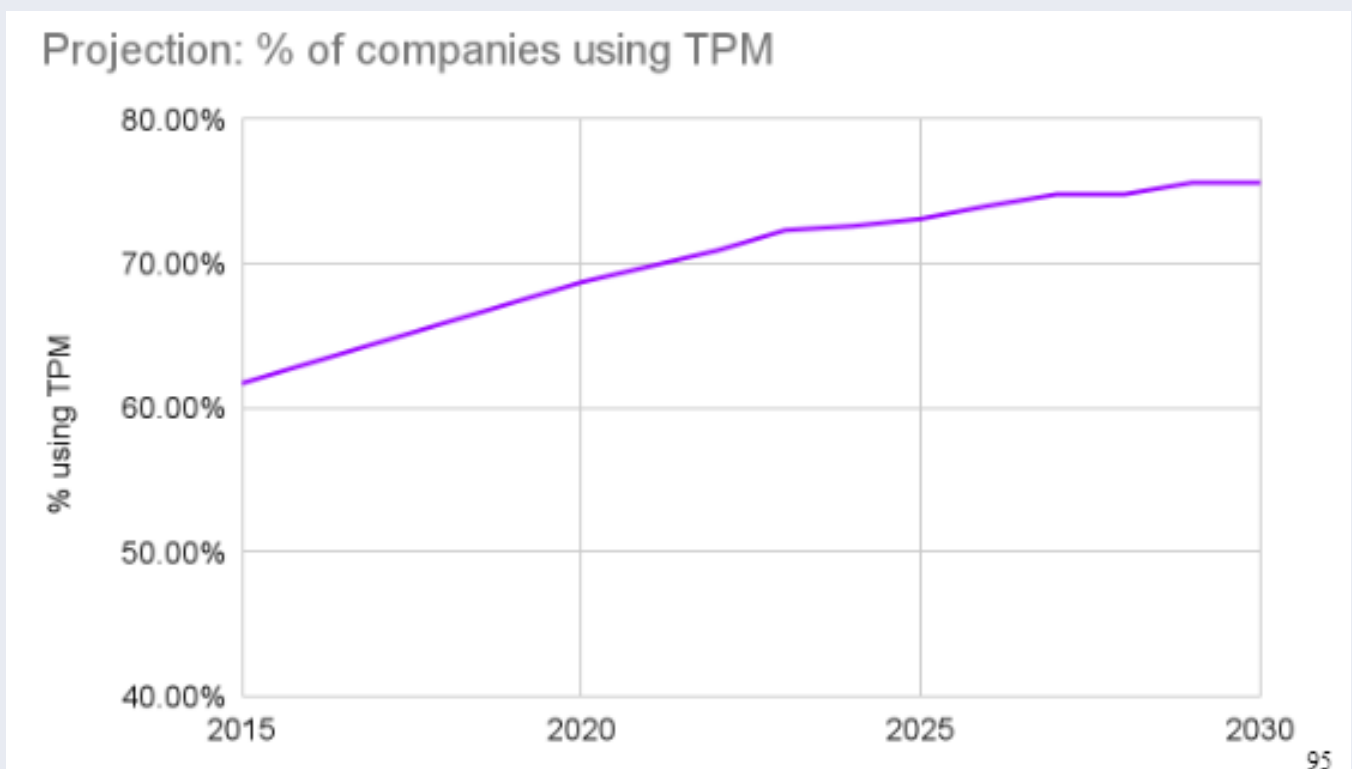
We conducted research on the past numbers of servers with a certain number of years. According to Horizon Technology, the distribution of average refresh cycles among companies was as follows:

	1-2 yr	3 yr	4 yr	5yr	>5yr
2015	9%	37%	20%	20%	14%
2020	9%	26%	15%	31%	19%
2022	7%	25%	15%	33%	19%

We constructed a sum-product model, which indicates that 70% of companies used TPM in 2022, matching the percentage previously cited, and 75% of companies will do so by 2030. Below we have attached an interpolation/projection graph of the percentage of companies using TPM.

1-2 yr	3 yr	4 yr	5 yr	>5 yr
10%	40%	70%	90%	100%

We constructed a sum-product model, which indicates that 70% of companies used TPM in 2022, matching the percentage previously cited, and 75% of companies will do so by 2030. Below we have attached an interpolation/projection graph of the percentage of companies using TPM.



# The Future of TPM



CAPRAE CAPITAL



## Arguments For/Against

# FOR

**Environmental responsibility:** Customers want to see technology companies becoming more environmentally sustainable. In fact, only 21% of people said that technology companies have done enough to minimize environmental impact, and 20% said that social media companies have done enough. This represents a dismal record only above financial services and airlines. Combined with the necessity of reducing e-waste and emissions highlighted above, we can see how important it is for companies to cut back on server replacement. Beyond a certain point, the only option to do so is TPM.

**Allows companies to save for innovation:** As illustrated above, using a TPM allows companies to reduce their capital expenditures on existing IT, allowing them to buy the GPUs critical to AI and the additional servers and storage for data.

**Favorable regulations:** Recent regulations aim to force OEMs to provide TPMs with the “tools and information” to repair their devices. More and more legislatures around the world are enacting these laws with the environment in mind.

**Space to grow:** The TPM market was only \$1.2 billion in 2019 and \$2.1 billion by 2022, whereas total data center and network maintenance was over \$65 billion already in 2019, so there is plenty of room to grow.

**Diversification:** If refresh cycles increase, TPM market growth increases and ITAD market growth decreases and vice versa. TPM thus serves as a good buffer to the impact of refresh cycles.

# AGAINST

OEMs may oppose TPMs as they reduce demand for new devices and cut into revenue. OEMs may consider TPM “unauthorized maintenance” as they are direct competitors. To discourage usage of TPMs, they may impose damages and/or recertification fees for returning to OEM maintenance. However, not every OEM has equally stringent anti-TPM policies.



Companies may be worried to use TPM for fear of not being able to access “microcode upgrades”. However, according to Park Place Technologies’ company Curvature, almost all such upgrades occur in the first three years.

Businesses in TPM extensive capital expenditures to set up warehouses for tools and replacement parts. This serves as a challenge to smaller players and reinforces economies of scale.

The TPM industry is extremely concentrated. We would be up against companies like Park Place and Service Express that have strong economies of scale in knowledge and the number of brands they are able to serve.

Breakthroughs in energy efficiency like in AMD Genoa/4th Generation Epyc chips could offer a viable substitute and challenge the TPM market – though the TPM market will remain competitive.

## Conclusion

The future of Third-Party Maintenance (TPM) is promising with an expected 14.5% yearly growth until 2030 and is poised to become increasingly mission-critical for several reasons. The surge in demand for GPUs and AI technologies compels companies to save funds for the requisite capital expenditures. To stay competitive, they must find a way to reduce spending on other IT infrastructure, requiring them to extend IT refresh cycles beyond the limits imposed by OEMs. Consequently, TPM services will become indispensable in ensuring the longevity and reliability of IT assets, thereby reducing environmental impact while empowering the rise of AI.



Titans such as Park Place Technologies stand to monopolize the TPM industry for multiple decades. With enormous capital expenditures needed to forge vast networks, the path to dominance seems ominously clear. Park Place Technologies, with its 18 acquisitions, and Blackstone's hefty investment in the company, cast long, unsettling shadows over the market with its 28% market share. Blackstone also owns QTS, which has 36 data center properties scattered across the United States and Europe. With their \$2 billion investment in Park Place Technologies one can easily infer that Blackstone is trying to create an all-in-one data center lifecycle (IT Installation, IT Maintenance, and IT Deinstallation). Blackstone's investment in Park Place Technologies, alongside Carlyle and 3i's backing of Evernex, underlines their recognition of the critical role TPM services play in supporting the technological and environmental ambitions of modern enterprises. We see a growth in the number of financial institutions building large players to wring value out of this crucial industry going forward.

---

As the world wakes up to the nightmarish reality of e-waste and the staggering carbon footprint of manufacturing new hardware, TPM services can save us from this impending catastrophe by lengthening the lifespan of existing equipment. Yet the impending oligopoly in the industry foreshadows high prices, which could hinder the transition to a sustainable economy. Thus, regulations like Right to Repair will gain traction, forcing OEMs to arm third-party maintainers with the tools and knowledge they need. As a side product, this will tear down barriers for smaller TPM providers to challenge the reign of Park Place and Evernex.

TPM will become mission-critical and environmentally vital in the coming years. As the world moves from a dispose-and-refresh model to a continuous maintenance model, TPM will rapidly become a leading substitute for IT Installation and Deinstallation. Thus, TPM is the market to hedge the risks of ITAD investment in an age of longer refresh cycles. Finally, investing in smaller TPM companies will not only deliver hefty returns, but also liberate technology companies from the scourge of unbearable prices as they seek to transition into this new era.