Hewlett Packard Enterprise

Technical white paper

Check if the document is available in the language of your choice.

# EFFICIENCY BENEFITS AND CHALLENGES OF INFORMATION TECHNOLOGY LIFECYCLE MANAGEMENT

# CONTENTS

An information technology (IT) lifecycle management strategy can yield measurable benefits in power savings, space reduction, cooling efficiency, IT staff productivity, and innovation opportunities. In the paper "Beyond PUE: Tackling IT's Wasted Terawatts," the Uptime Institute identifies the optimization of the server refresh lifecycle as one of the few, significant energy saving opportunities in IT, along with increasing server utilization, right-sizing redundancy requirements, and consolidating the infrastructure.[1] Yet, multiple disruptions often interfere with the initial lifecycle management strategy: acquisitions, divestitures, reorganizations, digital transformations, cloud migrations, and the emergence of new technologies.

Professionals implementing lifecycle strategies should understand the benefits, while acknowledging the challenges of operationalizing technology refresh plans. IT lifecycle strategies should also embrace the inherent environmental benefits included by considering the entire product lifecycle.

## BENEFITS OF LIFECYCLE MANAGEMENT

Various IT efficiency sources, including manufacturers, efficiency organizations, and consulting organizations, have published resources for considering the total lifecycle costs of IT products to define optimal lifecycles and asset management practices. These resources often propose utilizing Total Cost of Ownership (TCO) as a primary metric. However, the definition of TCO is often inconsistent, as most sources do not consider implications for the entire product lifecycle, such as performance, staff implications, and environmental impacts, or do not discuss barriers to adoption. While financial cost is a relevant metric, focusing primarily on TCO optimization fails to consider the benefits and opportunities associated with a strategic lifecycle management approach.

There are four primary benefits of optimizing technology refresh lifecycles: power and space reduction; operational efficiencies; residual value return; and reduced risk. Each of these benefits also provides an associated cost benefit.

### Performance, power, and space benefits

#### Compute

The most apparent benefits of server technology refresh:

- Improved application performance, requiring fewer servers, due to associated higher clock speeds, higher memory densities, and increased core counts
- Reduced heat generation and associated cooling costs from fewer servers
- Reduced data center space and electrical requirements from smaller cooling, electrical, and UPS capacity demanded by fewer servers
- Lowered environmental impacts associated with fewer servers, reduced power demand, and less support equipment

Server performance per watt, known as energy effectiveness, increases with each new generation of processors from both AMD and Intel®. This performance improvement is clear when viewing results of the SPEC Power® performance benchmark, which was developed by an industry organization known as the Standard Performance Evaluation Corporation.[2] The results demonstrate application performance and power effectiveness gains over time, as shown in Figure 1. On a 4-year window, power effectiveness gains, while slowing down, still provide substantial power savings opportunities.

[1] uptimeinstitute.com/beyond-pue-tackling-it%E2%80%99s-wasted-terawatts
[2] SPEC, spec.org/

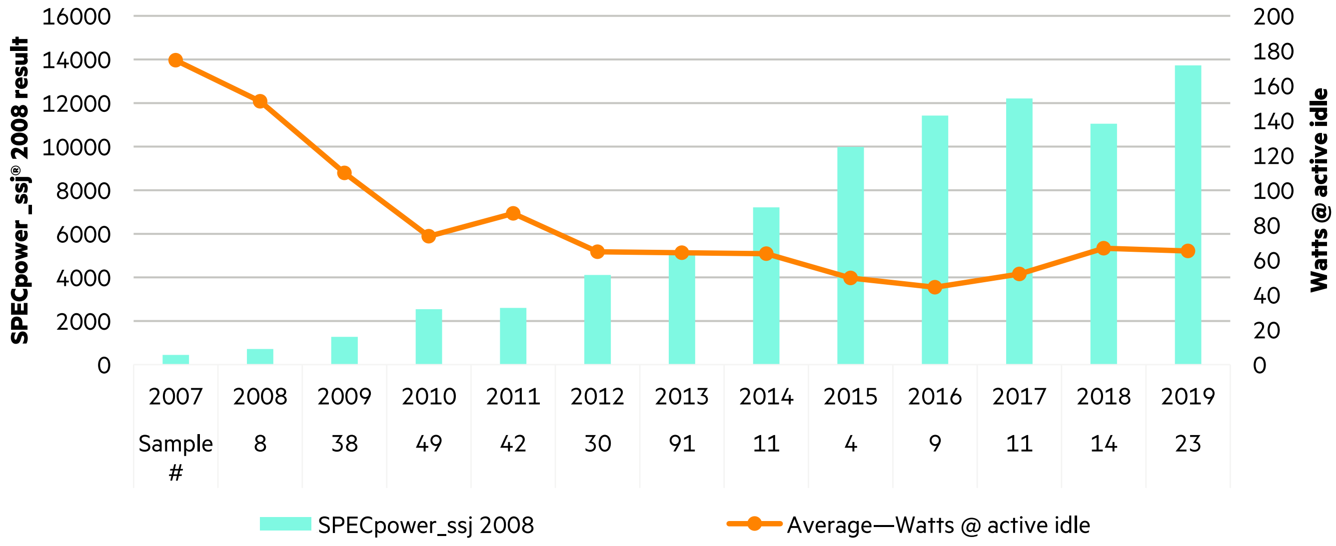


**FIGURE 1.** SPECpower_ssj 2008 public results over time

The SPEC Power committee also developed the SERT™ Suite, a next-generation toolset for measuring and evaluating the energy efficiency of servers. SERT Suite has been adopted by the U.S. Environmental Protection Agency (EPA) for assessing eligibility for ENERGY STAR® server program. These SPEC tools test power efficiency over a variety of load conditions to simulate actual use cases.

**Storage**

The quantity of data generated by people and "things" is growing exponentially while applications are increasingly deployed on hybrid cloud architectures, raising new challenges for enterprise data strategies. In order to define a storage lifecycle management strategy that will yield measurable efficiency benefits, it's critical to first define a data lifecycle strategy. This first step will better inform the technology refresh cycle for the storage infrastructure.

Storage technology refreshes can yield performance, power, and space benefits:

- Improved application performance and power effectiveness resulting from the adoption of flash arrays. Opting for SSDs instead of HDDs reduces and yields absolute power consumption and yield several orders of magnitude better power effectiveness both in capacity per watt and IOPS per watt[3, 4]
- Improved operational and power efficiency from a continuum of tape, hybrid and flash technology aligned with availability tiers and the data's lifecycle stage
- Reduced rack and floor space required, resulting from improved data reduction technology performance and higher density media

Data center storage and associated cooling is estimated to consume over 25% of the power in an average data center, yet there is no industry standard benchmark providing power efficiency insights for storage arrays.[5] The ENERGY STAR label, which relies on a standard test suite developed by the Storage Networking Industry Association (SNIA), partially addresses this gap, but does not disclose detailed benchmark results. This void requires IT departments make decisions based on vendors' dimensioning tools or the advice of storage consultants. When assessing storage efficiency, evaluators should consider unique workload requirements such as the size and capacity of the data footprint, the energy performance, and data redundancy objectives. Storage footprint and capacity for a given rack space varies widely depending on the RAID level of storage, compression, and deduplication technologies in use; and capacity and type of media used.

At the disk level, power efficiency data is available from manufacturers' data sheets. Some manufacturers offer disk-level power-throttling options to optimize power consumption in large deployments.

[3] researchgate.net/publication/320025946_A_Comparative_Study_of_HDD_and_SSD_RAIDs'_Impact_on_Server_Energy_Consumption
[4] bswd.com/FMS09/FMS09-101-Baral.pdf
[5] Kim and Rotem, "Energy Proportionality for Disk Storage Using Replication" faculty.tamuc.edu/jkim/documents/LBNL3936E.pdf

Energy effectiveness should be calculated considering the power requirements of all components in the storage system including drives, storage array controllers, network components, required cooling, and any other associated equipment. After total power demand is determined, power effectiveness can be determined based on metrics such as capacity per watt, IOPS per watt, or "activity" per watt. Capacity per watt will be highest when the array is fully or close to fully populated and being operated at the designed capacity level while avoiding over-dimensioning.

### Network

Networks are estimated to consume about 2% of the data center power envelope,[6] and enterprise networks consume about 8% of the overall ICT power consumption.[7] Similar to storage, there are no standardized third-party published test results to equip IT managers in helping assess power effectiveness of various generations of networking equipment. More than power effectiveness, the need to cope with data growth, new standards, or new industry regulations are the driving factors behind technology refreshes of network elements.

New Ethernet and Wi-Fi standards bring performance improvements in speed, security, and wireless density. **Figure 2** illustrates the progression of the Wi-Fi standard.

| **802.11n (2008)**<br>**Wi-Fi 4** | **802.11ac (2012)**<br>**Wi-Fi 5** | **802.11ax (2018)**<br>**Wi-Fi 6, 6E** | **802.11be (est. 2023)**<br>**Wi-Fi 7** |
| --- | --- | --- | --- |
| • 2.4 GHz and 5 GHz supported<br>• Wider channels (40 MHz)<br>• Better modulation (64-QAM)<br>• Additional streams (Up to 4)<br>• Backward compatibility with 11a/b/g<br>• Standard support up to 600 Mbps | • 5 GHz only<br>• Wider channels (80, 160 MHz)<br>• Better modulation (256-QAM)<br>• Additional streams (Up to 8, implemented up to 4)<br>• Backward compatibility with 11a/b/g/n<br>• Standard support up to 7 Gbps | • 2.4 GHz and 5 GHz supported<br>• Wider channels (80, 160 MHz)<br>• Better modulation (1024-QAM)<br>• Additional streams (Up to 8, implemented)<br>• Backward compatibility with 11a/b/g/n/ac<br>• Standard support up to 9.6 Gbps | • 2.4 GHz, 5 GHz, and 6 GHz supported<br>• Wider channels (40, 80, 160, 320 MHz)<br>• Better modulation (4096-QAM)<br>• Additional streams (Up to 16)<br>• Backward compatibility with 11a/b/g/n/ac/ax<br>• Standard target throughput of 40 Gbps |

**FIGURE 2.** IEEE 802.11 Wi-Fi standard evolution

Favoring wireless networks in remote, branch offices, or campus is an established trend that delivers a step change in efficiency by removing the need to lay cables and offering flexible role-based access control options. The more recent wireless networking standards, such as Wi-Fi 6, also address power management efficiency gains opportunities in both the network and connected devices.

Network upgrades can be disruptive and require downtime, leading to infrastructure management occurring at night or during weekends. Network operations support systems are complex and have not traditionally benefited from a widespread adoption of software automation. These combined factors make it operationally challenging to refresh network elements.

Having a proactive approach to network lifecycle management enables IT departments to selectively upgrade parts of their networks to take advantage of the new functionalities where they have the highest impact, while minimizing disruption. The emergence of live software upgrade options on select network elements helps minimize the disruption of technology refresh initiatives.

[6] iea.org/reports/tracking-buildings/data-centres-and-data-transmission-networks
[7] mdpi.com/2071-1050/10/9/3027

### Edge

Edge servers and gateways benefit from the same performance, power, and space improvements as data center servers and storage. These similarities are the result of components being leveraged for edge uses in more ruggedized form factors. However, additional considerations must be included when determining the optimal lifecycle of edge infrastructure:

- **Data gravity:** As compute power moves from the data center to where the data is being generated, planning the edge infrastructure refresh cycle requires an understanding of the data strategy and how the data will be used at the edge.
- **Operational technology equipment lifecycle:** Edge infrastructure has been deployed as part of dedicated technology systems historically. As these systems get modernized, opportunities will emerge to address new use cases, consolidate, and increase utilization of equipment at the edge.
- **Internet of Things:** Components in the edge infrastructure are sometimes specific to the use cases they are serving and are driven by compute storage and network bandwidth requirements, sensor technology evolution, or updated networking standards.
- **Asset management:** Edge infrastructure often also presents unique asset management complexity associated with the logistics involved in physically reaching them.

As a result, the refresh strategy for edge technology supporting preventive maintenance of manufacturing equipment, retail video analytics, or autonomous driving would each have to be approached differently. As edge use cases rapidly evolve, and are required to support a broader data strategy, they inherently carry the risk of deploying infrastructure which may no longer be optimal or adequate much sooner than expected. Assuming a shorter lifecycle at the onset and keeping options open to refresh early will mitigate the technology risk—for both the edge devices and the network elements that connect and enable them.

## Operational and cost efficiency benefits

In addition to the benefits noted above, a formal technology refresh strategy can also have significant operational and cost benefits.

Performance improvements result in a smaller number of compute and storage systems to manage, monitor, patch, and support. Smaller infrastructures allow IT staff to spend less time managing the infrastructure and more time on innovation and development opportunities. IDC research found that a three-year refresh cycle for servers cut staff management time and costs by 59%. Each refreshed server resulted in operational cost savings of $76,000 over three years compared to continuing to run the older server.[8]

Shorter refresh cycles also allow hardware-based improvements to be implemented into the infrastructure more quickly. Remote monitoring, measurement, and administration hardware, such as HPE Integrated Lights Out (iLO), dramatically simplify management and increase device security. Hardware-based security, such as HPE's Silicon Root of Trust, provides a series of trusted handshakes from lowest level firmware to BIOS and software to ensure a known good state.

Shorter refresh cycles also allow organizations to take advantage of recent processor architectural changes, allowing staff to control the number of cores or sockets being used. This change may sometimes offer significant cost advantages for software that is licensed on a "per-core" or "per-socket" basis. For example, one report suggested savings of $40K to $180K per processor in an Oracle environment and the potential for an almost $5M savings per processor for a business intelligence application.[9]

Finally, shorter refresh cycles allow organizations to roll out and reap the benefits of software innovations faster. For example, further effectiveness can be gained by the adoption of:

- **Anticipatory and self-correction technology tools:** Tools like HPE InfoSight helps organizations achieve 100% availability[10] on a single unit, minimizing the number of redundant systems deployed. The combination of self-diagnosis and correction, along with effective maintenance servicing, drives a very efficient use of deployed resources.
- **Performance monitoring tools:** These inform pertinent technology updates or upgrades. For example, HPE OneView manages device configurations and monitor operational status, making administrators aware of system status, utilization, and available firmware updates. Aruba NetInsight makes recommendations for network configuration tuning and anonymously benchmarks against peers.

## Residual value return

Another benefit of shortening refresh cycles, particularly for servers, is that newer generation models may have residual value at their end of use. Savvy IT organizations have identified this as an additional opportunity associated with proactive server lifecycle management; not only do organizations benefit from the power, space, and cooling efficiencies noted in the previous section, IT refreshes also maximize residual value at end of use.[11] However, operationalizing this aspect of the technology refresh strategy requires the coordination between IT, finance, facilities management, and procurement organizations to most effectively realize this value.

[8] techrepublic.com/resource-library/whitepapers/idc-accelerate-business-agility-with-faster-server-refresh-cycles/
[9] moorinsightsstrategy.com/wp-content/uploads/2017/11/AMD-EPYC-Presents-Opportunity-To-Save-On-Software-Licensing-Costs-By-Moor-Insights-And-Strategy-v2.pdf
[10] h20195.www2.hpe.com/v2/getmobile.aspx?docname=a00074521enw
[11] Various TCO calculators are at hpe.com/us/en/solutions/tco-calculators.html

### Technology risk mitigation

Because information technology evolves quickly, any technology solution carries an inherent risk of becoming obsolete sooner than expected. Technology refresh cycles are also a good opportunity to reassess requirements and make any adjustments to the type and quantity of equipment being procured. Composable systems, system-on-chip technology, hyperconverged infrastructure, liquid-cooled systems, and edge compute devices have matured into general use. Assessing the merits of these various options is beyond the scope of this paper but would nonetheless be an important element to consider as part of the wider benefits of technology refresh. Each of these technologies provides an opportunity to drive innovation as well as power, cooling, space, and IT staff efficiencies. Proactive lifecycle management is particularly relevant as a risk mitigation strategy for edge and edge-network components.

## BUSINESS REQUIREMENTS AND TECHNOLOGIES DRIVING SHORTER REFRESH CYCLES

Many organizations have already adopted lifecycle management strategies with shorter refresh cycles, particularly where IT systems are directly linked to business outcomes, such as research teams in fossil fuel exploration or high-tech manufacturing, computational sciences, and high-performance data analytics. These business functions recognize the value that a technology refresh approach can bring when the most efficient and highest performing IT assets are used.

Several other technology trends have increased the movement to shorter refresh cycles including:

- The widespread adoption of hybrid cloud architectures with containerized workloads “detached” from the infrastructure—thereby removing the risk of application migration between hardware versions
- The appeal of a step change in productivity granted by software-driven composable architectures
- The ability to upgrade equipment sub-assemblies, particularly on composable frames and storage systems, making it easier for IT staff to consider refresh initiatives rather than riskier migration projects
- Adoption of IT technologies with AI-driven self-diagnosis and healing capabilities which dramatically accelerate incident troubleshooting and resolution time and therefore reduce operating expenses
- Data availability commitments and “live” upgrade capabilities. These features make it easier for IT staff to consider storage arrays and network elements refresh initiatives which historically have required the undertaking of fairly complex migration projects

## BARRIERS TO ADOPTION OF SHORTER REFRESH CYCLES

Despite all these benefits, the reality is that server and storage infrastructure is not being refreshed at a rate that drives the highest efficiency benefits for organizations. IDC found that the average server is maintained for 5.15 years.[12] Research by The Uptime Institute found servers older than 3 years make up 40% of the installed base, yet perform less than 7%[13] of the work while consuming 66% of the power.[14]

There are notable barriers to technology refreshes:

- Technical debt in the application layer makes the asset migration complex with many legacy applications having legacy architecture requirements for specific OS and hardware versions. A hardware refresh would require an application modernization, requiring investments at a different scale.
- Depreciation policies for IT assets and software are not necessarily tied to technology cycles and are often rooted in larger financial considerations governed by the industry the enterprise operates in.
- For CAPEX-budget-driven organizations, technology refresh tends to be deprioritized in the face of more urgent projects and only address technology refresh needs once it has already become a major issue.
- IT organizations tend to be measured on application availability and do not have any incentive to reduce power consumption. Power-efficiency metrics usually are the responsibility of the facilities management team, making the benefits of refresh less clear to individual organizations.
- IT organizations often underestimate the benefits that power, performance, and operational improvements can achieve.

[12] ibm.com/downloads/cas/AV1ZWNZM
[13] uptimeinstitute.com/beyond-pue-tackling-it%E2%80%99s-wasted-terawatts
[14] intercompbusiness.com/wp-content/uploads/2019/06/Why-Faster-Refresh-Cycles-And-Modern-Infrastructure-Management-Are-Critical-To-Business-Success.pdf

It is also important to acknowledge that newer IT infrastructure may put pressure on the data center infrastructure by driving higher power density requirements. High-performance compute and supercomputer systems are at the forefront of power density requirements with current systems requiring as much as 250 kW and line of sight to 400 kW systems.[15] Newer generations of equipment will increase power density requirements and may push legacy data centers to their limits, requiring operators to rethink the design of their data centers. While a data center redesign or modular design update may have initial cost implications, it is also certainly an opportunity to drive higher standards of energy efficiency which will reduce costs and environmental impact.

## ASSET END-OF-USE STRATEGIES

A critical part of any lifecycle management strategy is the appropriate disposition of end-of-use assets being replaced. Circular economy principles provide a helpful framework to inform end-of-use decisions to minimize the environmental impact. Technology refresh strategies need to consider three primary areas of risk:

- **Security and compliance:** Are the assets being disposed following cybersecurity best practices and applicable government privacy regulations?
- **Environmental:** Is the equipment processed following labor, chemical substances, e-waste, and environmental regulations?
- **Economic:** Is the economic model transparent, scalable, and replicable?

At end of use, assets can either be recycled or upcycled, in whole or in parts. Without examining the specifics of recycling versus upcycling benefits, as a rule, upcycling is less energy intensive than recycling. Upcycling opportunities for the system, or its parts, are correlated to the asset age, configuration, and associated secondary market demand.

When IT assets cannot be refurbished and reused due to age or disrepair, they should be recycled in accordance with applicable regulations and best practices. The appropriate processing of e-waste is a growing challenge that has been acknowledged by the United Nations and many governments. The World Health Organization (WHO) states that “primitive recycling techniques, such as burning cables for copper recovery, expose both adult and child workers, as well as their families, to a range of hazardous substances.” Recycling or refurbishing assets through organization without adequate certifications, asset tracking, and assurance programs throughout the commodity lifecycle poses material, reputational risks for organizations.

## LIFECYCLE STRATEGY RECOMMENDATION

As part of a broader IT strategy, lifecycle management allows IT departments to be more intentional in the way they drive and execute their IT lifecycle decisions, taking advantage of the opportunities presented by performance improvements and newer technologies featured in current generation equipment.

### Environmental impact

Enterprise IT equipment has a significant environmental impact stemming from the materials, manufacturing, and logistics required to move the equipment to a customer, known as embodied impact. However, the greatest environmental impact results from the use phase, or when enterprises are using the asset. According to an internal HPE server Lifecycle Cycle Analysis (LCA), over 80% of the environmental impact of compute is caused during the use phase. To receive the greatest benefit from the embodied impact while minimizing the use-phase impact, the total environmental impacts of a product lifecycle should be considered when evaluating refresh cycles and end-of-use options for assets.

Enterprises can evaluate currently available referenceable data such as data center electricity consumption, mix of electricity sources, IT product lifecycle analysis, product material composition, recyclability of the various components, and carbon footprint to make informed lifecycle decisions. Tracking these data points are simple steps that organizations can take today to start measuring and reporting on the environmental impact of their IT equipment, supporting the organization’s larger business and environmental objectives.

### Determining the optimal lifecycle

The industry is converging to a 2- to 4-year compute refresh cycle that balances power efficiency gains, environmental impact, and operational constraints. Edge servers should broadly follow the compute lifecycle, and include additional considerations associated with the specific use case, as noted earlier. To effectively mitigate the technology risk associated with some rapidly evolving use cases, information technology and operations technology departments should assume refreshing equipment after 2 years and reconsider the assumption at that time, based on the evolution of sensors, standards, edge applications, bandwidth, and processing intensity requirements of the data handled by the edge infrastructure.

[15] insidehpc.com/2020/03/video-the-cray-shasta-architecture-for-the-exascale-era/

It is harder to draw a generic conclusion for storage. Nevertheless, considering the constant pressure on storage infrastructure associated with data growth, IT departments should monitor storage technology innovation developments and reassess every 3 to 4 years. This assessment should include an evaluation of the data and associated storage tiering they have implemented and if this level is still optimal. Storage lifecycle must align to the data lifecycle strategy it supports. Often, newer technology offers efficiency gains for specific data segments. Therefore, storage lifecycle management is less about a traditional technology refresh approach and more about augmenting existing capabilities with newer and more efficient capabilities for select sets of data.

The optimal lifecycle of a network device is a function of which part of the network the device belongs to: core, data center, campus, or edge. Data growth, number of devices, and density of connected end points each impacts network elements technology refresh cycles, particularly at the edge. For example, the adoption of Power over Ethernet (PoE) technology to power edge devices presents opportunities and challenges. PoE is an opportunity to reduce cabling and allow technology to be deployed at locations that were designed without power outlets. However planning and dimensioning PoE requirements is a challenge, as edge use cases, and the type and quantity of connected devices are still evolving rapidly. Power consumption and supported features requirements may outpace original plans. The recommendation for edge and campus network elements lifecycle management is therefore to consider the use cases that the network elements are supporting. As for other edge elements, assuming a shorter lifecycle at the onset and keeping options open to refresh early will mitigate the technology risk.

### Operationalizing the decision

Once an optimal lifecycle has been defined, it should be implemented with other business and operational processes. The refresh process should be determined up front and integrated with either the equipment acquisition process or the equipment disposal process.

### Refresh options

**From the start:** Defining asset lifecycle at the time of equipment acquisition is an approach that allows IT departments to derive up-front economic benefit from the improved residual value of the equipment at the time of return. Aligning refresh decisions with the procurement process allows organizations to leverage a very well defined and documented set of governance and compliance processes associated with lifecycle management of assets. It is also the most flexible option, as it allows organizations to delay the decision of the actual technology refresh options until a predefined time period at the end of use.

**After the fact:** Some organizations prefer an “after the fact” approach where assets due to be refreshed are systematically sold for upcycling or recycled if too old for upcycling. This approach may be harder to implement and monitor as initially there may be less governance attached to the end of use than to the procurement of assets. An “after the fact” approach may also negate the economic benefit if assets have little residual value and result in financial complexities, such as accounting for revenue from fully depreciated IT assets. Because the risks associated with IT asset disposal are material, an “after the fact” approach to asset upcycling requires the same robust governance processes as those attached to the asset procurement phase, including transparent reporting, governance, and compliance.

**As-a-service:** The emergence of “as-a-service” models, such as HPE GreenLake, embed technology refresh cycles which include tools and metrics to evaluate equipment effectiveness, along with the necessary framework to evaluate optimal lifecycles and effectively manage end-of-use operations with the least cost and environmental impacts.

## CONCLUSION

A comprehensive IT lifecycle strategy that considers all phases of the lifecycle and incorporates business, technology, and environmental implications, will yield significant positive efficiency, cost, and environmental benefits. Organizations that seek input from technology partners, impacted internal organizations, and trusted advisors are best positioned to overcome common barriers to adoption. As the rate of technology improvement accelerates, effective IT infrastructures which include effective lifecycle management are a business imperative. Thankfully, the information and resources needed to develop and implement effective lifecycle strategies are readily available, along with numerous examples of success from companies that have implemented these strategies.

## LEARN MORE AT

hpe.com/us/en/living-progress.html





a50002031ENW, July 2020