

EOSL Strategy

for Server, Storage & Network.

End-of-Service-Life is a contractual and commercial milestone — not a technical expiry date. When extending pays off, when security forces replacement, and what a defensible, documented decision looks like.

June 2026 · TechCare Research

TechCare Solutions GmbH · Hahnstätten · techcaresolutions.de

Contents

01	Management Summary
02	EOL, EOS, EOSL — the terms cleanly separated
03	The vendor lifecycle (Cisco · HPE · Dell)
04	How long does hardware really last?
05	The hyperscaler proof: six-year useful life
06	Security & compliance: NIS2, BSI OPS.1.1.3, CVE/EOL risk
07	The EOSL risk score (6 dimensions)
08	Cost & budget: extend vs. replace
09	Sustainability: embodied carbon, CSRD/ESRS E1, e-waste
10	Decision: extend, replace, monitor
11	A 6-step approach
12	Decision checklist
13	Conclusion
14	About TechCare & Sources

CHAPTER 01

Management Summary

The core thesis of this whitepaper is simple: **End-of-Service-Life (EOSL) is a commercial and contractual milestone, not a wear-out date.** The day a vendor finally ends support for a product says nothing about whether the hardware will fail the next morning. Yet many organisations equate EOSL reflexively with "must be replaced" — tying up budget that would deliver more value elsewhere.

A sound EOSL decision is therefore not a yes/no answer pinned to a calendar date, but a **scored, documented keep/extend/replace decision**. It weighs the actual reliability of the hardware, the security and compliance position, spare-part and support availability, as well as cost and sustainability. Three evidence strands underpin this approach.

First — the hardware reality. Failure data from large disk fleets (Backblaze, Google) show the classic "bathtub curve": high infant mortality, then a long, flat and stable useful-life phase, and only afterwards the wear-out rise. That flat plateau typically extends past the vendors' five-year standard support window.

Second — the hyperscaler proof. Microsoft, Google/Alphabet, Amazon and Meta have, in audited financial filings, raised the accounting useful life of their servers from four or five to **six years** — with billions in depreciation effects. This is not marketing but an auditor-attested statement about how long modern server hardware is economically used.

Third — the security gate. Extension is subject to a clear proviso: hardware that is no longer securely operable — missing security patches, open exploited vulnerabilities, regulatory obligations from NIS2 and the BSI baseline — must be replaced or decommissioned. Security is not a weighting factor; it is a veto. This inverts the usual logic: it is not the end of support that forces replacement, but the security position.

What this means for the operator

In practice this approach shifts the EOSL question from a one-off procurement decision to an ongoing risk process. Instead of "when do we buy new?" the question becomes "which assets may keep running, under what conditions, for how long — and when do we review that again?". The vendor thus remains an important but no longer the sole voice: its LDoS date enters the assessment as *one* input alongside failure data, security position and cost. The result is not only cheaper and ecologically cleaner but also audit-proof — exactly the documented risk trade-off that NIS2, the BSI baseline and CSRD demand of management.

Key takeaway

EOSL terminates a contract, not a device. The correct response is a scored, documented decision — backed by failure data and hyperscaler evidence, gated by a hard security and compliance threshold.

CHAPTER 02

EOL, EOS, EOSL — the terms cleanly separated

Four terms are routinely conflated in everyday use, yet describe clearly distinct stations in a product's life. Separating them cleanly is the foundation of any defensible lifecycle plan.

End-of-Life (EOL) — the announcement

EOL is the vendor's official **announcement** that a product is reaching the end of its lifecycle. It is the starting gun for a series of staggered deadlines — not the end of operation. From here, the orderly retreat of sales and support begins.

End-of-Sale (EoS) — the sales stop

From End-of-Sale, the product is no longer sold new, but remains supported under existing contracts. EoS is usually the anchor from which the downstream support windows are counted.

End-of-Software-Support — no more updates

From this point, the vendor delivers no further software or firmware updates — in particular no security patches. This milestone is security-critical: the attack surface now ages unchecked, because new vulnerabilities are no longer closed.

End-of-Service-Life (EOSL) / Last-Date-of-Support (LDoS)

EOSL — often termed **Last-Date-of-Support (LDoS)** by the vendor — is the definitive end of all vendor support: no support, no spare parts, no RMA, no bug fixes. From this date the organisation alone owns operation and risk (optionally covered by third-party maintenance).

The most common misconception

"EOSL = hardware fails." Wrong. EOSL ends the *contractual* relationship with the vendor. The physical condition of the hardware is unaffected — it follows the failure curve (Chapter 04), not the calendar.

The practical consequence: in the lifecycle plan, each asset carries at least three dates — the end of software support (security-relevant), the LDoS/EOSL (contractual-commercial) and the estimated wear-out window (physical). Only these three together enable a decision that neither burns budget too early nor leaves a security risk standing too late.

For the operator the key takeaway is: the security-relevant date is the end of software support, not the later EOSL. Anyone aligning their risk clock to LDoS may run gear without patch supply for months while believing they are still "in support"; conversely, the early EoS date tempts premature procurement — both errors spring from the same single, wrongly chosen date.

CHAPTER 03

The vendor lifecycle (Cisco · HPE · Dell)

The major infrastructure vendors follow a similar phase model but differ in deadlines and terminology. The figures below reflect each published policy; in any specific case the product-level announcement governs.

Cisco

Per Cisco's "End-of-Life Policy", the End-of-Sale is typically announced around six months in advance. After EoS, contract customers receive staggered support: typically about one year of hardware failure analysis, around two years of bug fixes, and up to five years of critical security fixes. From EoS to the Last-Date-of-Support is therefore roughly five years.

HPE

HPE generally grants support for about five years from End-of-Sale per its published EOS/EOSL documentation. HPE reserves the right to shorten or extend these windows by product.

Dell

Dell offers up to five years of standard support from ship date; with the **Lifecycle Extension** under ProSupport up to six years, and in Post-Standard-Support a total of roughly seven years — subject to spare- part availability.

The pattern

Across vendors, a **standard support window of roughly five years from End-of-Sale** emerges, with vendor extensions to six or seven years. This is precisely where the strategic question begins: the end of support does not coincide with the end of the hardware's useful life.

Typical phases and windows (policy-based, product-dependent)

Vendor	Standard support	Extension	Anchor
Cisco	~5 years to LDoS	critical security fixes up to ~5 yrs	from End-of-Sale
HPE	~5 years	product-dependent shorter/longer	from End-of-Sale
Dell	up to 5 years	up to 6 yrs (Lifecycle Ext.), ~7 yrs total	from ship date

The lifecycle phases from the operator's view

Behind the vendor terms sits a recurring phase model. The table below translates it into what actually happens to running operations at each station — based on the Cisco, HPE and Dell policies above. The windows are typical orders of magnitude; the product-specific announcement governs.

Phase	Vendor milestone	What changes for the operator	Typical window
Active	Product on sale, full support	Normal operation, full RMA & updates	until EoS
End-of-Sale	Sales stop announced (Cisco: ~6 mo. notice)	No new purchase; support countdown starts here (anchor for Cisco/HPE)	EoS = day 0
Full contract support	Bug fixes + HW failure analysis (Cisco: ~1–2 yrs)	Updates and RMA as usual; risk low	~0–2 yrs after EoS
Security fixes only	Cisco: critical security fixes up to ~5 yrs	Feature updates end; security patches only — watch patch status	~2–5 yrs after EoS
Extended support	Dell Lifecycle Extension (~6 yrs) / Post-Standard (~7 yrs)	Support only by arrangement, spares scarcer	~5–7 yrs (from ship date for Dell)
LDoS / EOSL	Definitive end of support	No OEM support, no RMA — third-party maintenance or decommission	from LDoS

The table makes two things visible. First: the transition from "full support" to "security fixes only" is often more important to the operator than the distant LDoS — from here no more feature and stability updates arrive, only patches. Second: the extension options (Cisco security fixes, Dell Lifecycle Extension) show that the vendors themselves do not treat the five-year window as a hard end — they extend it commercially. This is exactly the logic this whitepaper picks up and carries past the LDoS.

CHAPTER 04

How long does hardware really last?

If the end of support is not a wear-out date — how long does hardware actually last? The best publicly available evidence comes from large disk fleets, because rotating magnetic storage is among the most heavily stressed components and its failure behaviour is well documented.

The bathtub curve

Failure behaviour follows a **bathtub curve**: high "infant mortality" in the first weeks, then a long, flat and stable useful-life phase with a low failure rate, and only afterwards the wear-out rise. Crucially, the flat phase typically extends beyond the five-year OEM support window.

The three phases deserve a closer look, because each carries its own operational lesson. **Infant mortality**: in the first weeks and months an above-average number of devices fail — material and manufacturing defects that surface early under load. In practice: a device that survives its first six months has its riskiest phase behind it, not ahead. **Useful-life plateau**: then comes the long, flat middle with a low, largely constant failure rate. Here the hardware is at its most reliable — and this is exactly where, statistically, the fifth and often the sixth year of operation fail. **Wear-out rise**: only afterwards does the rate climb noticeably, as mechanical and thermal ageing accumulate. The art of the EOSL decision is to recognise the end of the plateau — not to confuse the end of support with the onset of wear-out.

Evidence

The **Backblaze Drive Stats 2024** (a fleet of about 301,000 drives) report an annualized failure rate (AFR) of roughly 1.57% across all drives (lifetime AFR around 1.31%). Drives aged two to four years are "low and stable"; failure rates "noticeably increase as drives exceed five years" — some 5-plus-year models rose to around 3% AFR.

The older but methodologically foundational **Google study (Pinheiro et al., USENIX FAST '07)** across more than 100,000 drives found age-dependent AFR ranging from roughly 1.7% to 8.6%. An important corollary: SMART values alone are a weak predictor of individual failure — age and population say more than a single sensor reading. (The data is from 2007 and is used here for curve shape, not absolute figures.)

Both datasets show the same picture across two decades: the 2024 Backblaze figures — an AFR of around 1.57% over a modern fleet of well over 300,000 drives — and the 2007 Google data jointly trace a flat plateau whose rise begins only *past* five years. Precisely where the standard OEM support window ends, the hardware is empirically not yet in its wear-out phase. This temporal gap between contract end and physical risk rise is the quantitative core of this whitepaper: it is the room in which an extension is not a gamble but a statistically grounded decision. The caveat remains important — this is fleet statistics, not a guarantee for the single device; that is why observed failure history enters as its own scoring dimension in Chapter 07.

The conclusion

Hardware does not fail at the end of support. The low, stable useful-life phase routinely extends beyond the five-year OEM window. Empirically the rise begins only *past* five years — creating the room in which an extension becomes commercially sensible.

CHAPTER 05

The hyperscaler proof: six-year useful life

The most convincing evidence comes not from lab tests but from audited financial filings. The world's largest server operators have raised their accounting useful life — an auditor-reviewed statement about how long servers are economically used.

Extended accounting useful life (audited disclosures)

Operator	Change	Effective	Balance-sheet effect (order of magnitude)
Microsoft	server & network 4 → 6 years	FY2023 (ann. Jul 2022)	around \$3–4bn benefit FY2023
Alphabet/Google	servers 4 → 6 yrs; some network 5 → 6 yrs	2023	depreciation ~\$3.9bn lower; net income ~+\$3.0bn (\$0.24/share)
AWS / Amazon	servers 5 → 6 years	from Jan 1, 2024	~+\$900m net income Q1 2024
Meta	servers toward 5 years (non-AI)	2022	~\$1.5bn benefit 2022

The common denominator is unambiguous: for **general-purpose server hardware**, a six-year consensus has emerged. Operators running hardware at maximum utilisation and under maximum cost pressure do not stretch useful life beyond what is technically defensible — the change is therefore a conservative, market-wide signal.

The honest nuance: general-purpose vs. AI/GPU

In 2025, AWS partially reverted server useful life from six back to five years, and Microsoft also trimmed some refresh cycles. **The driver, however, was not wear-out but AI/GPU refresh pressure:** AI accelerators become functionally obsolete quickly because new generations deliver dramatically more performance per watt. For classic, general-purpose server hardware — the class relevant to the mid-market — the six-year consensus holds.

Why is this evidence so much stronger than a vendor or consultant promise? Because an accounting useful life is not marketing but an attested assumption with a direct impact on profit and share price. That no fewer than four of the largest operators took this step is therefore a deliberately conservative signal — and for the mid-market it grows even stronger: hyperscalers run their servers under near-constant load, a typical mid-sized server sits in a far milder load profile. If even the hardest-worked hardware in the world holds for an attested six years, the reserve under moderate load is more likely larger than smaller.

Relevance to the mid-market

The general-purpose hardware relevant to mid-sized data centres is depreciated over six years, with audit attestation, by the world's largest operators. The five-year OEM support window is therefore not a technical limit but a commercial marker.

CHAPTER 06

Security & compliance: NIS2, BSI OPS.1.1.3, CVE/EOL risk

Extension is subject to a non-negotiable proviso: hardware must remain **securely operable**. Three regulatory strands define what that means in practice.

NIS2 — Directive (EU) 2022/2555

The NIS2 Directive had to be transposed into national law by 17 October 2024 (applicable from 18 October 2024). It markedly widens the scope and obligations of cybersecurity and makes EOSL hardware an explicitly regulated risk. Three strands of the directive bear directly on the lifecycle decision.

Article 21 — risk-management measures. Art. 21 requires a proportionate catalogue of technical and organisational measures — expressly including **supply-chain security** and the **handling of vulnerabilities**. This is where EOSL hardware becomes tangible: a device without patch supply can no longer be properly treated against vulnerabilities in the sense of Art. 21 — unless compensating controls (segmentation, virtual patching, heightened monitoring) demonstrably take its place.

The reporting cascade. NIS2 couples a three-stage reporting duty with hard deadlines to security incidents: an **early warning within 24 hours** of awareness, an **incident notification within 72 hours** with a first assessment, and a **final report within one month**. This cadence is only achievable in daily operations if the inventory, dependencies and patch status of assets are already documented — a further reason for the structured EOSL process.

Fines. For essential entities, fines reach up to **€10m or 2% of global annual turnover** (the higher value applies). The order of magnitude alone moves the EOSL question out of pure IT cost accounting and into entrepreneurial risk consideration.

Germany — NIS2UmsuCG (BSIG amendment)

Germany missed the October 2024 deadline; the NIS2 implementation act was enacted at the end of 2025 (promulgated in December 2025). Its scope covers around 30,000 German companies across 18 sectors (roughly 8,250 "particularly important" and around 21,600 "important" entities). The key BSIG paragraphs map clearly:

BSIG paragraphs (NIS2UmsuCG) — mapping to the EOSL decision

§	Subject	EOSL relevance
§28	Scope (who is covered)	Clarifies whether the asset is in NIS2 scope at all (dim. 6)
§30	Risk-management measures	Counterpart to Art. 21 — EOSL gear needs compensating controls
§32	Reporting duties	24h/72h/1-month cascade presupposes a documented inventory
§38	Management liability	Leadership is personally accountable — lifts the decision to board level
§65	Fine provisions	Sanction framework for breaches of §30/§32

The decisive paragraph for this whitepaper is **§38**: management must approve the risk-management measures and oversee their implementation — and is personally liable for doing so. With this, unassessed EOSL hardware is no longer a pure IT matter but a management liability risk. A documented, score-based keep/extend/replace decision is exactly the evidence §38 demands of the leadership level.

BSI IT-Grundschatz — module OPS.1.1.3

The BSI module OPS.1.1.3 "Patch and Change Management" (Edition 2023) is the decisive hinge here: if products are used that are no longer supported by their manufacturers, it **MUST** be verified whether they can still be operated securely; **if not, such products MUST NOT be used**. This is a binding **MUST** requirement — exactly the security gate that every extension decision must pass.

Regulators treat EOL gear as a live threat

That this is no theoretical concern is shown by CISA Binding Operational Directive 26-02 (February 2026): a US federal order to inventory and remove end-of-support edge devices on a staged timeline (inventory within 3 months, listed devices within 12 months, others within 18 months). Rationale: threat actors increasingly exploit unsupported edge devices. CISA maintains the Known Exploited Vulnerabilities (KEV) Catalog for this purpose. The message is clear: EOL hardware is an active, not a latent, risk to regulators.

The security gate

"Securely operable — or decommission." Hardware beyond software support, with open exploited vulnerabilities and no compensating controls, fails this gate. Then the rule is: replace, regardless of the score.

CHAPTER 07

The EOSL risk score (6 dimensions)

To turn the preceding chapters into a reproducible decision, the EOSL risk score rates each asset along **six dimensions**, 1 to 5 points each (drawing on ISO 31000 and common asset-valuation / criticality models). One dimension carries veto power.

The six scoring dimensions (1–5 points each)

#	Dimension	Guiding questions
1	Business criticality	What damage on failure? Single point of failure?
2 	Security & patch status (VETO)	Past LDoS? Open exploited CVEs? Securely operable per BSI OPS.1.1.3?
3	Spare-part & support availability	OEM or third-party maintenance available? Spares on the market?
4	Reliability / age vs. curve	Where on the bathtub curve? Observed failure history?
5	Performance / capacity headroom	Enough headroom? Bottleneck foreseeable?
6	Compliance / audit exposure	NIS2 scope? CSRD Scope-3? Documented risk acceptance?

The veto in dimension 2

Dimension 2 is **not a weighted summand but a hard veto**. An asset that is not securely operable — beyond software support, with open exploited vulnerabilities and no effective compensating controls — is a replace case regardless of all other dimensions. No score in criticality, cost or headroom, however good, can offset a breached security gate.

Decision bands

Band	Profile
EXTEND	moderate criticality, securely operable, spares available, flat failure curve, sufficient headroom
MONITOR	middle band — instrument, compensating controls, re-score quarterly
REPLACE	unmitigable security risk, BSI gate breached, no spares or hard performance blocker — especially when business-critical

A worked example

The following example shows the score in action on three typical asset types. **The point values are chosen illustratively** to demonstrate the mechanics — they are not a real assessment result and not a statement about specific devices. Higher points mean higher risk or more urgent action; dimension 2 (security) acts as a veto.

Asset (example)	D1 Crit.	D2 Sec. ▴	D3 Parts	D4 Age	D5 Headr.	D6 Compl.	Recommendation
Core switch, 6 yrs old	4	2	2	2	2	3	EXTEND
Business-critical (D1=4), but security gate passed (D2=2, patchable), spares available, flat plateau, enough headroom → extend in a controlled way with a spare pool and a quarterly review.							
Backup storage, 7 yrs old	3	3	4	4	3	2	MONITOR
Secondary system (D1=3), security borderline but mitigable (D2=3), spares growing scarce (D3=4), past the plateau (D4=4) → instrument, secure spares, budget the replacement, re-score quarterly.							
Edge firewall, past LDoS	5	5	4	3	3	5	REPLACE
Exposed edge device past LDoS without patches (D2=5 → veto), business-critical (D1=5), NIS2/KEV-exposed (D6=5) → the breached security gate forces replacement regardless of all other values.							

The third row shows the veto in pure form: although age (D4) and headroom (D5) are middling, the security value of 5 overrides every commercial consideration. The first row shows the opposite case — high criticality alone forces no replacement as long as the security gate is passed. This very asymmetry — security as a veto, everything else as a trade-off — is the core of the model.

CHAPTER 08

Cost & budget: extend vs. replace

The commercial question is not "new or old" but: does the effort of **one more operating period** outweigh the cost and risk of replacement? The two options differ fundamentally in their cost structure.

Capex vs. opex

Replacing is capex — a lumpy investment depreciated over five to six years, plus migration effort. **Extending is opex** — predictable annual costs for (third-party) maintenance, spare parts and energy. The decision therefore weighs capex plus migration labour plus the written-off embodied carbon of the replacement against the cost and risk of one more operating period.

A quantified anchor instead of invented percentages

Neutral "€-per-server" savings benchmarks from vendor-independent sources are scarce — this should be stated honestly. What is robust, by contrast, are the audited hyperscaler depreciation deltas from Chapter 05: extending from five to six years cut annual depreciation per asset class by roughly the **order of magnitude of about 20%** (derived estimate). This figure is the most credible available proxy for the commercial lever of an extension.

TCO rule of thumb

A widely cited TCO rule of thumb holds that the pure hardware purchase makes up **less than 20% of lifecycle TCO**; around 80% falls in the period after — support, energy, labour. Looking only at the purchase price optimises the smallest lever.

The commercial effect of an extension draws on two mechanisms: the **time-value effect** (every safely bridged operating period pushes a large capex item back by a year) and the **smoothing effect** (predictable smaller opex items instead of a lumpy investment block). Both are real but site-specific — which is why this whitepaper deliberately quantifies no blanket saving.

Illustrative model — no customer commitment

Illustratives Modell / Illustrative model. Suppose a depreciated server, kept in extended operation, costs only third-party maintenance and energy per year, whereas replacement triggers immediate capex plus several person-days of migration: then every safely bridged operating period defers the investment and smooths the budget. Concrete amounts are strictly site-specific and deliberately not quantified here — this is not a promise of savings but a structural model.

CHAPTER 09

Sustainability: embodied carbon, CSRD/ESRS E1, e-waste

In an EOSL context, sustainability is no longer a footnote but increasingly a reporting obligation — and it frequently overturns the gut feeling that "new is greener".

Embodied carbon vs. operational share

A server's carbon footprint splits into two parts. The **embodied carbon** arises *once* in manufacturing, raw-material extraction and transport — it is already "paid" at purchase, no matter how long the device then runs. The **operational share** arises *continuously* from power consumption in operation. Across studies, embodied carbon accounts for roughly **20–30%** of the lifecycle footprint (range about 15–43%), the operational share correspondingly around 70–80%.

This split is the key to the sustainability argument. A worked example from the **Tech Carbon Standard** reports roughly **5,657 kg CO₂e** lifecycle for a server, of which about **1,726 kg (30.5%)** is embodied. These 1,726 kg accrue regardless of useful life — they are amortised the better, the longer the device runs. A premature replacement throws away the not-yet-"lived-off" part of this embodied carbon and immediately triggers the full embodied carbon of the new device. This is exactly where the gut feeling "new is greener" flips: as long as the old device runs efficiently, continued operation is frequently the lower-carbon option.

E-waste

The second ecological dimension is the waste stream. The UN Global E-waste Monitor 2024 puts global e-waste in 2022 at **62 million tonnes**, on track for **82 Mt by 2030**. Only **22.3%** was documented as collected and recycled in 2022 — Europe leads at 42.8%, in many regions the rate lies far below. Every avoided premature retirement directly reduces this stream and at the same time avoids the embodied carbon of a replacement device. Life-extension therefore acts twice: less new emission, less new waste.

CSRD / ESRS E1

Under the CSRD, Scope-3 reporting is mandatory where material, governed by ESRS E1. Purchased IT hardware is Scope 3 (Category 2 Capital Goods / Category 1). Extending useful life amortises the already-incurred manufacturing emissions and defers the replacement emissions — a **reportable Scope-3 reduction lever**.

An honest caveat (Uptime Institute)

Replacement does *not* automatically cut CO₂. Unless new hardware is highly utilised, efficiency gains may not recover the embodied carbon of the replacement. The Uptime Institute favours consolidation over blanket refresh. An extension only wins on carbon if the gear stays efficient and utilised.

CHAPTER 10

Decision: extend, replace, monitor

The score and bands from Chapter 07 resolve into one of three recommendations. The order of evaluation is decisive: the security veto comes first, the weighted trade-off only afterwards.

Extend

Recommended for moderate business criticality, a passed security gate, available spare parts (OEM or third-party), a flat position on the failure curve, and sufficient performance headroom. The extension is agreed for a defined period and tied to a quarterly re-score.

Replace

Mandatory for unmitigable security risk, a breached BSI gate, missing spare parts or a hard performance blocker — especially when the asset is business-critical. Here the veto beats every commercial consideration.

Monitor / evaluate

The middle band: assets that still carry today but whose risk is rising. They are instrumented (monitoring, spare pool, compensating security controls) and **re-scored quarterly**. This turns a one-off decision into an ongoing process — the only form that does justice to the realities of failure curves and the threat landscape.

Order of evaluation

- 1) Security veto (dim. 2) — breached? → Replace.
- 2) Otherwise form the weighted score.
- 3) Derive the band.
- 4) Document the decision and set a review date.

Importantly, none of the three recommendations is a terminal state. "Extend" carries a review within it, "monitor" is by definition an interim status, and even a "replace" results in a new asset that sits in the lifecycle plan again from day one — three states in a cycle, not an endpoint.

CHAPTER 11

A 6-step approach

The following sequence turns the score into a repeatable, audit-proof process.

- 1 Inventory.** Build a complete asset list with model, age, LDoS date, function and dependencies. Without a reliable inventory, every EOSL decision is piecemeal.
- 2 Score.** Rate each asset along the six dimensions, 1–5 points each — with dimension 2 explicitly flagged as a veto.
- 3 Security gate.** Before any commercial trade-off, check: securely operable per BSI OPS.1.1.3? Open exploited CVEs (KEV)? No → straight into the replace path.
- 4 Cost & carbon model.** Compare capex/opex, migration effort, embodied carbon and Scope-3 effect per option — with the hyperscaler deltas as an anchor and honestly labelled assumptions.
- 5 Decide.** Derive the band (extend / replace / monitor) and record the recommendation with its rationale.
- 6 Document & re-score quarterly.** Store the decision, assumptions and risk acceptance in an audit-proof manner (NIS2/CSRD-relevant) and update it every quarter.

Why quarterly?

The threat landscape and the failure curve are dynamic: a newly exploited vulnerability or a cluster of failures can flip an "extend" into a "replace" overnight. Re-scoring keeps the decision alive — and incidentally produces the documentation that §38 BSIG and ESRS E1 demand.

CHAPTER 12

Decision checklist

This list condenses the decision into checkable points. It is intended as a per-asset working template.

- ☐ Complete inventory with model, age and LDoS date available
- ☐ Business criticality and dependencies assessed per asset (dim. 1)
- ☐ Security & patch status checked: past LDoS? open exploited CVEs? (dim. 2 / veto)
- ☐ Secure operability per BSI OPS.1.1.3 documented and confirmed
- ☐ Spare-part and support availability (OEM or third-party) clarified (dim. 3)
- ☐ Position on the failure curve and failure history assessed (dim. 4)
- ☐ Performance and capacity headroom checked (dim. 5)
- ☐ Compliance exposure clarified: NIS2 scope, CSRD Scope-3 (dim. 6)
- ☐ Cost and carbon model computed per option (extend/replace)
- ☐ Decision band derived and recommendation justified
- ☐ Risk acceptance and rationale documented in an audit-proof manner
- ☐ Review set for quarterly re-scoring

CHAPTER 13

Conclusion

EOSL is a date in a contract, not a state of the hardware. Equating the two ties up budget without need — or, worse, runs insecure gear in misplaced trust in the calendar page.

The evidence is consistent: failure curves show a long, stable useful-life phase beyond the five-year OEM window, and the world's largest server operators depreciate general-purpose hardware over six years, with audit attestation. A measured extension is therefore often the better choice both commercially and ecologically — **provided the security and compliance gate is passed.**

The right answer to EOSL is thus not a reflex but a procedure: a scored, documented keep/extend/replace decision with a hard security veto and a quarterly review. This protects budget, satisfies NIS2/CSRD, and makes the decision audit-proof — right up to the management that is personally liable under §38 BSIG.

In one sentence

Score rather than replace by reflex — and measure every extension against the security gate.

CHAPTER 14

About TechCare & Sources



About TechCare Solutions

TechCare Solutions GmbH, based in Hahnstätten, supports mid-sized organisations in operating, maintaining and lifecycle-managing their server, storage and network infrastructure. We combine vendor-independent maintenance with a structured EOSL evaluation process — so that lifecycle decisions are commercially defensible, security-compliant and documented.

Contact: techcaresolutions.de · 24/7 service desk (DE/EN) · Hahnstätten

Methodology & transparency

All concrete figures in this whitepaper are drawn from the publicly available primary and secondary sources listed below (vendor policies, audited financial filings, regulatory orders, research and industry studies). Where sources contain ranges or unverified data, this is flagged in the language ("around", "approximately", "order of magnitude", "per financial filings"). Cost statements are presented as an illustrative model and constitute no customer commitment. Vendor deadlines follow the respective published policies; in any specific case the product-level announcement governs.

Sources

- [1] Cisco — End-of-Life Policy. cisco.com/c/en/us/products/eos-eol-policy.html
- [2] HPE — EOS/EOSL Support. support.hpe.com/hpesc/public/docDisplay?docId=a00125046en_us
- [3] Dell — Lifecycle Extension / Networking Warranty. dell.com/en-us/lp/networking-warranty
- [4] Backblaze — Drive Stats 2024. backblaze.com/blog/backblaze-drive-stats-for-2024/
- [5] Pinheiro/Barroso (Google) — USENIX FAST '07. research.google/pubs/failure-trends-in-a-large-disk-drive-population/
- [6] The Register — Amazon six-year useful life. theregister.com/2024/02/02/amazon_q4_2023/
- [7] The Register — Microsoft server life extension. theregister.com/2022/08/02/microsoft_server_life_extension/
- [8] Datacenter Dynamics — Google server life. datacenterdynamics.com/en/news/google-increases-server-life-to-six-years-will-save-billions-of-dollars/
- [9] Computer Weekly — Google savings. computerweekly.com/news/366557152/Google-saves-almost-3bn-by-running-servers-for-six-years
- [10] Data Center Frontier — Meta. datacenterfrontier.com
- [11] European Commission — NIS2 Directive. digital-strategy.ec.europa.eu/en/policies/nis2-directive
- [12] OpenKRITIS — NIS2UmsuCG / BSIG. openkritis.de/it-sicherheitsgesetz/nis2-umsetzung-gesetz-cybersicherheit.html
- [13] BSI — IT-Grundschutz, module OPS.1.1.3. bsi.bund.de
- [14] CISA — Known Exploited Vulnerabilities (KEV) Catalog. cisa.gov/known-exploited-vulnerabilities-catalog
- [15] UN — Global E-waste Monitor 2024. ewastemonitor.info/the-global-e-waste-monitor-2024/
- [16] Uptime Institute — Embedded Carbon of IT. journal.uptimeinstitute.com/what-does-embedded-carbon-of-it-really-represent/
- [17] Tech Carbon Standard. techcarbonstandard.org
- [18] ISACA Journal — IT Asset Valuation & Risk Model. isaca.org
- [19] ESRS E1 (CSRD) — EU delegated regulation. eur-lex.europa.eu