Spot GPU Training Break-Even & Interruption Calculator

GPUSmith · GPU Smith · 2026-09-19 · spot gpu training break-even calculator · sagemaker spot training · checkpoint interval · gpu interruption cost · ml.p4d.24xlarge · aws spot pricing · finops · gpu training cost

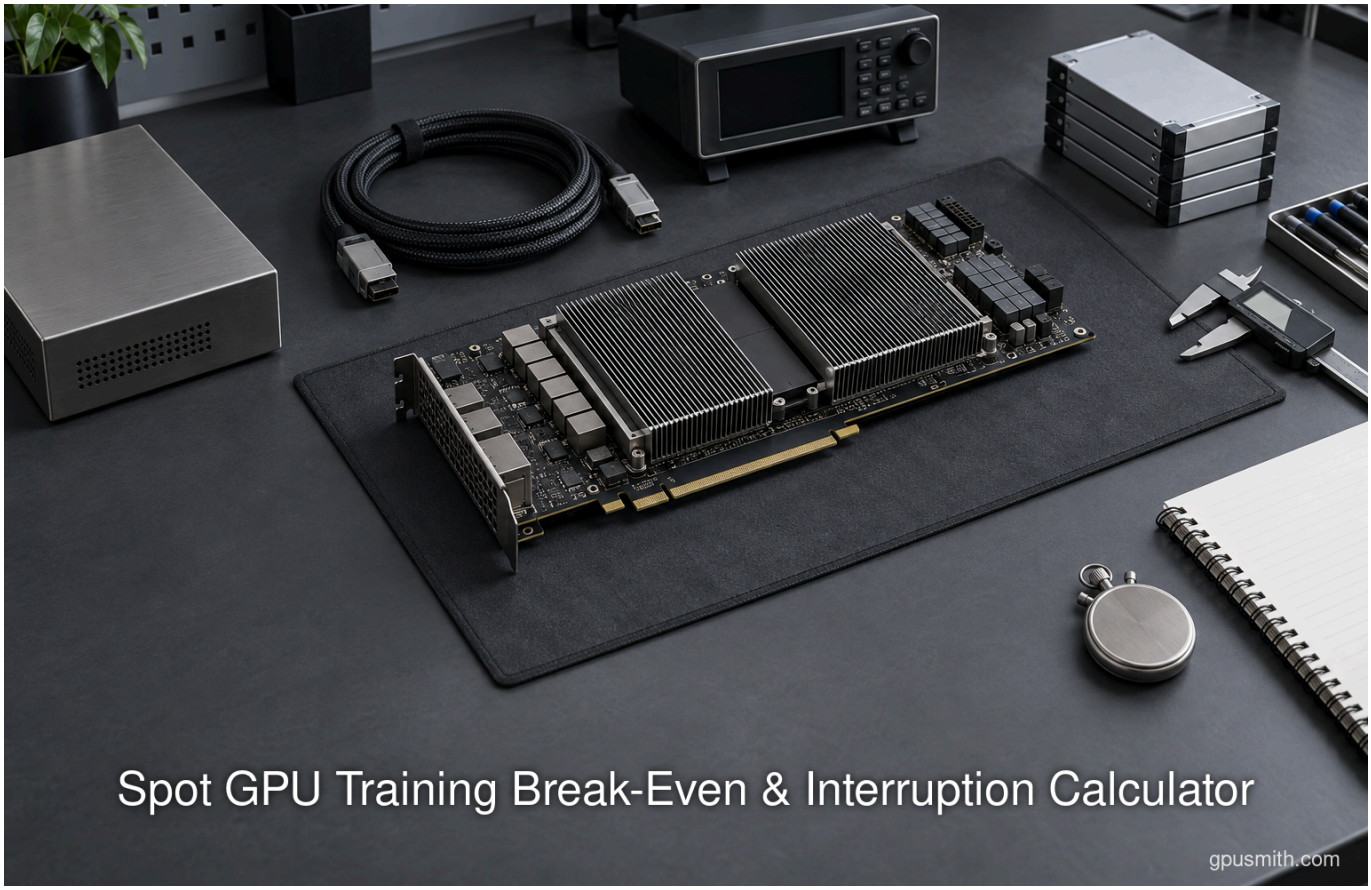

Executive Summary

An interruptible **SageMaker GPU training job is cheaper only when its workload-specific all-in cost stays below the On-Demand baseline**, and its completion time still satisfies the buyer's deadline. AWS's claim of up to **90%** savings is a ceiling, not a planning discount (docs.aws.amazon.com). As of September 19, 2026, the AWS price list quotes **ml.p4d.24xlarge training in US East (N. Virginia) at $25.2512883 per instance-hour** (pricing.us-east-1.amazonaws.com). AWS also describes On-Demand as having no minimum fee or upfront commitment (aws.amazon.com). The instance has eight 40 GB A100 GPUs (aws.amazon.com). Actual Managed Spot capacity, quota, and prices must still be verified in the buyer's account and chosen Availability Zone.

The calculator separates **billed compute**, **elapsed waiting**, **checkpoint storage and requests**, and **deadline penalty**. For a hypothetical 50-hour continued-pretraining pilot with 6-minute writes, 8-minute restores, 5-minute initialization, 30-minute reacquisition waits, and a buyer-supplied $8.00 Spot scenario rate, the 60-minute checkpoint policy costs **$470.17** after three interruptions, versus **$1,388.40** for the defined On-Demand baseline.

Yet it takes **60.13 hours**, so it fails a 58-hour deadline. Adding a $1,000 missed-deadline penalty makes its decision cost **$1,470.17**, more than On-Demand. These are reproducible scenario outputs, not public AWS benchmark results.

The model is:

```
C_spot = r_spot × H_billable + C_storage + C_requests + C_other + C_deadline
```

```
H_elapsed = H_billable + H_nonbilled_wait
```

For interruption `j`, incremental billable time is `lost_work_j + restore_j + initialization_j`; waiting for capacity remains in elapsed time unless the job record shows it was billed. AWS defines billable time as absolute wall-clock time once training is running (docs.aws.amazon.com), while `MaxWaitTimeInSeconds` includes both capacity wait and run time (docs.aws.amazon.com). The purchase decision therefore needs two outputs, not one blended savings percentage.

Checkpoint intervals should be chosen from a **sensitivity grid**, not from a universal interruption rate. Classical interval analysis starts from a total-lost-time objective (arxiv.org), while broader models incorporate failure rate, checkpoint cost, failure detection, and restart cost (arxiv.org). The practical calculator instead evaluates zero, half-interval, and full-interval lost work under buyer-supplied interruption counts. Approval should require a pilot, raw `DescribeTrainingJob` fields, CloudWatch timestamps, a checkpoint-object inventory, a dated Spot extract, restart trials, and final-model equivalence evidence. GPU Smith's stated method centers workload modeling and cloud baselines (gpusmith.com), which is the appropriate adjacent-advisor posture for this decision rather than a vendor comparison row.

Introduction and Background

The phrase **spot GPU training break-even calculator** sounds like a price comparison. It is actually a recovery-engineering model. An On-Demand run pays a stable published rate for billable runtime. A Managed Spot run adds uncertain interruption timing, replayed work, checkpoint I/O, restore and initialization, capacity waits, and possibly a business cost when a delivery window is missed. Therefore, a percentage discount alone cannot answer the procurement question.

This report uses **one ml.p4d.24xlarge in us-east-1** for a hypothetical continued-pretraining job. P4d supplies 400 Gbps networking (aws.amazon.com) and 8 TB of local NVMe storage (aws.amazon.com). Its processors total 96 vCPUs (aws.amazon.com). Those specifications identify the workload boundary, but they do not prove Managed Spot capacity. AWS says account quotas may differ from defaults (docs.aws.amazon.com). The approval pack must show the selected instance is requestable for SageMaker training in the account and Region.

The worked inputs are intentionally labeled **hypothetical pilot measurements**. Readers should replace them with timestamps and bytes from their own run. Public price history does not supply workload interruption probability. The EC2 API returns prices for a selected time range (docs.aws.amazon.com) and can filter by Availability Zone (docs.aws.amazon.com). It is not a forecast of interruptions or future capacity.

The resulting decision framework has two audiences. Training engineers establish useful-work rate, durable checkpoint timing, and correct restoration. FinOps practitioners attach dated rates and ancillary charges. Platform leads define the deadline, fallback, and acceptable evidence. Keeping those roles explicit prevents an apparent compute saving from masking a wall-clock failure or an irreproducible recovery path.

Key Changes

Change 1: Define the decision boundary before collecting prices

The calculator begins with a reproducible job specification. Without it, teams can unintentionally compare an eight-GPU training job with a different GPU count, Region, checkpoint policy, or completion objective.

- **Useful-work target:** 100,000 optimizer steps, measured at 1.8 seconds per step, or 50.0 useful GPU-instance hours.
- **Procurement unit:** one **ml.p4d.24xlarge** SageMaker training instance in **us-east-1**.
- **Model state:** weights, optimizer, scheduler, scaler, dataloader position, and random-number-generator state.
- **Checkpoint policy:** evaluate 30, 60, and 120 minutes of useful work between durable checkpoints.
- **Recovery objective:** restart from the last complete S3 checkpoint and verify no step discontinuity.
- **Completion window:** 58 hours from job submission to durable final artifact.
- **Fallback:** submit On-Demand after a buyer-set maximum wait or repeated interruption threshold.
- **Price timestamp:** preserve the On-Demand price file and Spot extract used for approval.

The state list matters. PyTorch says optimizer state should also be saved for a resumable training checkpoint (docs.pytorch.org). Hugging Face Trainer restores optimizer, scheduler, and random-number-generator state (huggingface.co). Without optimizer files, the optimizer restarts from scratch (huggingface.co). NeMo separately identifies dataloader, RNG, and step-scheduler state as necessary for accurate resumption (docs.nvidia.com), including Python, NumPy, Torch, and CUDA generators (docs.nvidia.com). A weights-only save can produce a loadable model without reproducing the original training trajectory.

Change 2: Build an input ledger from quoted and measured evidence

Table 1 separates facts quoted from AWS from values that a pilot must measure. Confidence is high only when the exact workload, container, object layout, and Region match the approval run.

Input	Worked value	Unit	Timestamp/source	Evidence class and confidence
On-Demand rate	**$25.2512883**	instance-hour	AWS price list, 2026-09-01 (pricing.us-east-1.amazonaws.com)	Quoted, high for listed SKU, recheck at submission
Spot scenario rate	**$8.00**	instance-hour	Buyer-supplied hypothetical input	Not a market quote, replace with dated AZ extract
Useful training	**50.0**	hours	Hypothetical pilot: 100,000 × 1.8 seconds	Measured, medium until repeated
Checkpoint size	**220**	GB	Hypothetical object inventory	Measured, medium

Input	Worked value	Unit	Timestamp/source	Evidence class and confidence
Exposed checkpoint write and sync	**6**	minutes	Hypothetical application and S3 timestamps	Measured, medium
Restore from durable checkpoint	**8**	minutes	Hypothetical restart trial	Measured, medium
Container and framework initialization	**5**	minutes	Hypothetical CloudWatch timestamps	Measured, medium
Capacity wait per restart scenario	**0, 30, 120**	minutes	Buyer scenario bands	Assumed, low, retain separately
S3 Standard storage	**$0.023**	GB-month, first 50 TB	AWS price list, 2026-09-01 (pricing.us-east-1.amazonaws.com)	Quoted, high, confirm storage class
S3 write requests	**$0.005**	1,000 PUT, COPY, POST, or LIST	AWS price list, 2026-09-01 (pricing.us-east-1.amazonaws.com)	Quoted, high, count multipart operations
Deadline penalty	**$0 or $1,000**	per missed run	Buyer scenario	Business input, not cloud spend

The table's most important distinction is **measured versus quoted**. Object size cannot reveal CPU staging, asynchronous buffer pressure, multipart requests, or synchronization stalls. PyTorch asynchronous checkpointing first copies state into CPU buffers (docs.pytorch.org), and its interface exposes separate staging and upload completion signals (docs.pytorch.org). The same documentation recommends one outstanding asynchronous request at a time (docs.pytorch.org). Distributed Checkpoint exposes staging and upload completion separately (docs.pytorch.org). Measure staging, device pause, upload, and memory pressure separately.

The hardware path matters too. AWS lists 16 GB/s local NVMe read throughput for P4d (aws.amazon.com), but that specification is not an S3 checkpoint benchmark. The instance also exposes up to 600 GB/s bidirectional NVSwitch bandwidth (aws.amazon.com). Serialization, CPU staging, network routing, object operations, and synchronization can each dominate, so only the pilot's end-to-end timestamps belong in the cost model.

Change 3: Model SageMaker's lifecycle rather than an abstract VM

SageMaker synchronizes checkpoints from a local path to S3 (docs.aws.amazon.com). On restart, it copies S3 data back into the local path. The default local directory is `/opt/ml/checkpoints` (docs.aws.amazon.com). The training code still has to discover the checkpoint and restore the complete state.

The relevant sequence is:

1. **Starting:** container, image, and input preparation begin.
2. **Training:** useful work and checkpoint I/O occur during billable runtime.
3. **Interrupted:** the current attempt stops and work after the last durable checkpoint is lost.
4. **Waiting:** the service seeks Spot capacity, increasing elapsed time.
5. **Downloading:** the next attempt receives data and checkpoints.

6. **Restoring:** the script loads state and reconstructs the training loop.
7. **Training:** useful work resumes from the durable recovery point.
8. **Uploading:** the final artifact is persisted before completion.

AWS documents the interrupted-to-starting lifecycle for a successfully resumed job (docs.aws.amazon.com). These records, rather than an assumed hazard rate, provide the evidence for a pilot-specific interruption audit. CloudWatch itself is another metered input: AWS's us-east-1 example prices standard log ingestion above the first 5 GB at $0.50 per GB (aws.amazon.com). The calculator should use the actual billing export rather than assuming observability is free.

Change 4: Keep cost, elapsed time, and deadline value separate

For the defined On-Demand baseline with 50 useful hours, 49 six-minute checkpoints, and one five-minute initialization:

```
H_on_demand = 50 + 49 × 0.1 + 5/60 = 54.9833 hours
```

```
C_on_demand = 54.9833 × \$25.2512883 = \$1,388.40
```

For a Spot scenario with checkpoint interval `I`, interruption count `N`, observed lost work `L_j`, restore `R_j`, initialization `A_j`, and non-billed wait `W_j`:

```
H_spot_billable = H_useful + H_checkpoint + H_initial + Σ(L_j + R_j + A_j)
```

```
H_spot_elapsed = H_spot_billable + ΣW_j
```

```
C_spot = r_spot × H_spot_billable + C_S3 + C_requests + C_other + C_deadline
```

For distributed training, AWS directs users to multiply `BillableTimeInSeconds` by `InstanceCount` (docs.aws.amazon.com). That means the calculator must scale every billed restart minute by the node count. Waiting belongs in wall-clock time unless billing evidence says otherwise. A network appliance can also add charges: AWS notes that services in the transfer path may add data-processing cost (aws.amazon.com).

The break-even inequality is:

```
r_spot < (C_on_demand - C_storage - C_requests - C_other - C_deadline) /
H_spot_billable
```

This inequality is more useful than an advertised discount. It produces the **maximum tolerable realized Spot rate** for each scenario, and it makes a deadline penalty explicit rather than hiding it in a narrative caveat.

Change 5: Treat the checkpoint interval as a sensitivity variable

Young-style analysis balances checkpoint overhead against re-execution. A broader model incorporates checkpoint interval, failure rate, checkpoint cost, detection, and restart cost (arxiv.org). Lost time depends on checkpoint interval, save time, and average failure time (arxiv.org). But the classical optimum depends on distributional assumptions. An observed interruption count from one workload is not a provider-wide benchmark.

The calculator should therefore compute three lost-work bands for every interruption:

- **Best case:** zero useful work lost because interruption follows a completed checkpoint.
- **Midpoint case:** `I/2`, used only when the buyer explicitly accepts a uniform-location assumption.
- **Conservative case:** the full interval `I`, suitable for deadline stress testing.
- **Observed case:** actual steps replayed, derived from logs and checkpoint metadata.

The checkpoint overhead fraction is `total checkpoint write time / useful training time`. In this worked model, it is **19.8%** at 30 minutes, **9.8%** at 60 minutes, and **4.8%** at 120 minutes. Longer intervals lower planned I/O but increase the maximum replay per interruption. One alternative objective is utilization, defined as the fraction of total time available for useful work (arxiv.org). The first-order model also assumes error-detection latency is small relative to the interval (arxiv.org). No single value is optimal across all interruption and deadline scenarios.

Implementation Considerations and Process Changes

Make checkpoint durability observable

The service's S3 synchronization is necessary, but application-level completeness is equally important. TensorFlow checkpoints preserve parameter values but not the computation definition (tensorflow.org). TensorFlow optimizer slots are preserved when both the variable and optimizer are saved (tensorflow.org). In multi-worker saves, each worker writes its own section (tensorflow.org), and all workers need a common visible filesystem for merging (tensorflow.org). Checkpoints capture exact parameter values (tensorflow.org), but framework-specific restore tests must still be part of the pilot.

A durable checkpoint protocol should record:

- **Checkpoint identifier:** globally unique job, attempt, worker, and step keys.
- **Manifest:** expected objects, byte counts, hashes, and framework version.
- **Completion marker:** created only after all required shards are durable.
- **Training position:** optimizer step, samples or tokens consumed, and dataloader cursor.
- **State coverage:** model, optimizer, scheduler, scaler, RNG, and any custom sampler state.
- **Write timing:** first local write, last local write, first object, last object, and marker time.
- **Restore timing:** S3 copy, deserialization, state application, and first resumed step.
- **Retention:** rotating complete checkpoints without deleting the only valid recovery point.

Orbax uses an atomic rename to signal completion (orbax.readthedocs.io), and its staged approach leaves the previous checkpoint untouched if interrupted before finalization (orbax.readthedocs.io). The precise mechanism differs by storage and framework, but the decision principle is constant: restore only a checkpoint that has a verifiable completion boundary.

Hugging Face applies the same broad principle through an incomplete-checkpoint sentinel that is removed only after writing finishes (huggingface.co). Orbax explicitly stages changes while preserving the prior checkpoint before finalization (orbax.readthedocs.io), while its default protocol uses an atomic rename for completion (orbax.readthedocs.io). These patterns support a calculator input called **durable completion time**, which may be later than the application's local save return.

Prevent distributed workers from colliding

SageMaker's high-level configuration points to one S3 location without per-instance suffixes (docs.aws.amazon.com). A distributed script must create unique paths or coordinated shards. PyTorch Distributed Checkpoint produces multiple files, normally at least one per rank (docs.pytorch.org). DeepSpeed requires every process to participate in the save (deepspeed.ai). Hugging Face warns that per-node saves on shared storage reuse the same names (huggingface.co). NeMo assigns each data-parallel rank its own optimizer-state shard (docs.nvidia.com).

Use one of two patterns:

- **Rank-sharded:** each rank writes to `job/attempt/step/rank`, followed by a coordinator manifest.
- **Coordinated consolidation:** the framework safely creates a single artifact after collective completion.
- **Never shared filenames:** do not let workers independently overwrite a common object.
- **Topology test:** restore on the exact intended topology, then test permitted resharding separately.

PyTorch supports load-time resharding between cluster topologies (docs.pytorch.org). DeepSpeed can preserve model, optimizer, and learning-rate scheduler state (deepspeed.ai) and can consolidate ZeRO 2 or 3 checkpoints (deepspeed.readthedocs.io). These documented capabilities still require a container-version-pinned restore trial.

Set stop conditions and fallback rules before launch

`MaxWaitTimeInSeconds` covers capacity waiting plus running and must be at least `MaxRuntimeInSeconds`. A practical runbook should define:

- **Maximum pending time:** how long the business accepts before falling back.
- **Maximum total wait:** enough to include expected runtime plus scenario capacity delays.
- **Interruption trigger:** an explicit count or elapsed-time point for On-Demand fallback.
- **Deadline trigger:** latest time at which a resumed Spot run could still finish.
- **Budget trigger:** stop when accumulated billed cost plus expected completion cost exceeds baseline.
- **Artifact trigger:** do not restart from a checkpoint lacking the completion marker.
- **Quota check:** confirm the account can request the chosen training instance.
- **Equivalence gate:** compare loss curve, final metrics, and artifact hashes where meaningful.

Quota and capacity checks belong before economic approval, not after a calculator has declared a theoretical saving. AWS's public pricing posture also states no minimum fees or upfront commitments for On-Demand (aws.amazon.com). That commercial flexibility is part of the fallback comparison, even though it does not remove quota requirements.

Preserve an evidence pack

The evidence pack should include raw machine-readable artifacts:

- **Job description:** full `DescribeTrainingJob` response for every terminal attempt.
- **Status history:** all secondary status transitions with start and end timestamps.
- **Logs:** CloudWatch stream identifiers and timestamp export.
- **Metrics:** step time, useful tokens or samples, and checkpoint-duration series.

- **Objects:** S3 inventory with keys, versions, bytes, modification times, and request counts.
- **Prices:** dated On-Demand file and Spot history extract for exact instance, product, and AZ.
- **Retries:** interruption index, last durable step, resumed step, and work replayed.
- **Equivalence:** fixed evaluation set, seed policy, software digest, and comparison result.

AWS notes that some timing and billing attributes may be absent when a job fails (docs.aws.amazon.com). The audit design needs fallbacks to CloudWatch and the buyer's own timestamps rather than assuming every API field exists.

Data Analysis and Evidence

Table 2 crosses checkpoint interval with interruption count using the conservative full-interval replay assumption, 30 minutes of non-billed wait per interruption, the hypothetical $8.00 rate, and a flat $1.10 placeholder for S3 storage and requests. It is a calculator demonstration, not an AWS performance claim.

Useful interval	Interruptions	Checkpoint overhead	Work replayed	Billable hours	Elapsed hours	Spot scenario cost	58-hour deadline	Break-even Spot rate
30 min	0	9.90 h	0.00 h	59.98	59.98	**$480.97**	Fail	$23.13/h
30 min	1	9.90 h	0.50 h	60.70	61.20	**$486.70**	Fail	$22.86/h
30 min	3	9.90 h	1.50 h	62.13	63.63	**$498.17**	Fail	$22.33/h
60 min	0	4.90 h	0.00 h	54.98	54.98	**$440.97**	Pass	$25.23/h
60 min	1	4.90 h	1.00 h	56.20	56.70	**$450.70**	Pass	$24.69/h
60 min	3	4.90 h	3.00 h	58.63	60.13	**$470.17**	Fail	$23.66/h
120 min	0	2.40 h	0.00 h	52.48	52.48	**$420.97**	Pass	$26.43/h
120 min	1	2.40 h	2.00 h	54.70	55.20	**$438.70**	Pass	$25.36/h
120 min	3	2.40 h	6.00 h	59.13	60.63	**$474.17**	Fail	$23.46/h

The table shows why the **cheapest compute scenario is not automatically acceptable**. Thirty-minute checkpoints miss the deadline without interruption because I/O consumes 9.90 hours. At three interruptions, every policy misses. A $1,000 penalty makes the 60-minute decision **$1,470.17**, or **$81.77** above On-Demand.

Replace full-interval loss with observed replay retrospectively, but retain it for a conservative approval band. Hugging Face describes exact continuation as potentially slower (huggingface.co), while omitting optimizer data can restart optimization from scratch (huggingface.co). DeepSpeed can automatically preserve model, optimizer, and learning-rate scheduler state (deepspeed.ai). NeMo's distributed layout assigns each data-parallel rank an optimizer shard (docs.nvidia.com). Both state coverage and resume behavior belong in measured restart time and equivalence evidence.

Table 3 is the interruption audit the buyer should populate. The sample rows are labeled hypothetical and use the 60-minute scenario.

Interruption	Last durable checkpoint	Interrupted step	Steps replayed	Restore plus init	Capacity wait	Rate applied	Evidence artifact
1, hypothetical	20,000	22,000	2,000	13 min	30 min	$8.00/h scenario	Job JSON, log export, S3 manifest
2, hypothetical	48,000	50,000	2,000	13 min	30 min	$8.00/h scenario	Job JSON, log export, S3 manifest
3, hypothetical	76,000	78,000	2,000	13 min	30 min	$8.00/h scenario	Job JSON, log export, S3 manifest

This audit converts interruption cost into inspectable evidence. A true run should use exact step timestamps, not evenly spaced illustrative rows. PyTorch warns that complete reproducibility is not guaranteed across releases (docs.pytorch.org), and deterministic operations can be slower (docs.pytorch.org). PyTorch's distributed checkpoint API can reshard at load time (docs.pytorch.org), but that capability does not guarantee identical numerics. TensorFlow likewise distinguishes saved parameter values from the computation definition (tensorflow.org). Version pinning and an agreed equivalence tolerance are therefore part of the economic model.

Implications and Future Directions

The largest improvement over a discount calculator is **decision traceability**. Every cost term has a source, timestamp, unit, and evidence class. Every interruption can be reconciled to a recovery point. Every wait is reported in wall-clock time without being automatically charged as compute. This structure lets FinOps, platform engineering, and model owners review the same run without redefining savings.

The model also clarifies where engineering investment pays back:

- **Faster writes** reduce planned overhead at every checkpoint.
- **Faster restores** reduce the cost of each realized interruption.
- **Smaller state** can reduce storage, requests, network use, and recovery time.
- **Asynchronous saves** may hide I/O but must be measured for CPU memory and staging cost.
- **Unique rank paths** prevent a nominal checkpoint from becoming unusable.
- **On-Demand fallback** caps deadline exposure when Spot wait grows.
- **Scenario bands** remain useful when interruption evidence is sparse.
- **Pilot repetition** gives a workload distribution without claiming a provider-wide rate.

NVIDIA NeMo, for example, documents checkpointed dataloader, RNG, and step-scheduler state for accurate resumption (docs.nvidia.com). Its RNG state spans Python, NumPy, Torch, and CUDA generators (docs.nvidia.com). TensorFlow multi-worker checkpoints split sections across workers (tensorflow.org), and DeepSpeed requires collective participation (deepspeed.ai). These details illustrate why framework support should be verified at the exact container version.

GPU Smith is an **adjacent independent engineering advisor**, not a cloud or Spot-capacity provider. Its site states that the firm has no cloud of its own (gpusmith.com) and describes costed comparisons against cloud baselines (gpusmith.com). Its published assessment scope includes workload characterization and tokens per day

(gpusmith.com). The appropriate role is to make assumptions and acceptance criteria inspectable, not to insert the firm into a procurement table as though it sells SageMaker capacity.

Frequently Asked Questions (FAQs)

What is a spot instance checkpoint interval calculator?

It evaluates checkpoint overhead against work replay under explicit interruption scenarios. Enter useful runtime, write and restore duration, initialization, checkpoint interval, interruption count, lost work, wait, and rates. It should report **cost and elapsed time separately**. Classical formulas can be a starting point, but their assumptions must be disclosed.

How should GPU training interruption cost be calculated?

For each interruption, add observed work since the last durable checkpoint, restore time, and initialization time to billable compute. Add capacity wait to elapsed time. Then add storage, requests, other metered services, and any buyer-defined deadline penalty. Do not silently value every interruption at half an interval.

What is the optimal checkpoint interval for Spot instances?

There is no universal optimum. The answer changes with measured checkpoint duration, restore time, useful step rate, interruption scenarios, and deadline. DeepSpeed can consolidate ZeRO checkpoints to one full-precision state dictionary (deepspeed.readthedocs.io), but consolidation time and storage behavior still need measurement.

How should teams compare Spot versus On-Demand GPU training cost?

First calculate the uninterrupted On-Demand baseline from the dated rate and billable runtime. Then compare it with each Spot scenario's billed compute and ancillary charges. Report wall-clock completion beside cost and apply the same final-artifact and model-equivalence criteria to both modes.

What is the spot GPU interruption break-even point?

It is the maximum realized Spot rate, interruption count, or deadline penalty at which the all-in Spot scenario still costs less than On-Demand. In the worked 60-minute, three-interruption case without a deadline penalty, the rate ceiling is **$23.66 per billable hour**. With a $1,000 deadline penalty, that scenario is already above the baseline at the hypothetical $8.00 rate. The baseline rate is anchored to the dated AWS price file (pricing.us-east-1.amazonaws.com).

How is machine-learning checkpoint overhead calculated?

Use `total exposed checkpoint write time / useful training time`. Measure local serialization, staging, sync, and durable completion separately. PyTorch recommends limiting outstanding asynchronous saves to one request (docs.pytorch.org), which underscores that asynchronous I/O still consumes finite resources.

What is the expected cost of interrupted GPU training?

When a buyer has a defensible interruption distribution, weight each scenario cost by its probability. When evidence is sparse, publish scenario bands instead of an invented expected value. The EC2 price API can reach back up to 90 days (docs.aws.amazon.com), but price history is not interruption probability.

Do Spot savings with checkpointing include waiting time?

AWS's documented savings formula uses billable time relative to training time (docs.aws.amazon.com). A buyer should still show capacity waiting in elapsed time and deadline evaluation. The procurement answer can be “lower cloud spend but unacceptable completion risk.”

Conclusion

A defensible **spot GPU training break-even calculator** is a ledger of observed work, recovery behavior, prices, waits, and business constraints. For the defined ml.p4d.24xlarge example, the current On-Demand baseline is **$1,388.40**. The hypothetical 60-minute policy with three conservative interruptions is **$470.17** at a buyer-supplied $8.00 rate, but it misses the 58-hour deadline. A $1,000 penalty reverses the procurement result.

The approval decision should therefore require five things: a dated exact-SKU price record, a real Spot history extract for the selected scope, checkpoint and restore measurements from the actual model state, an interruption-by-interruption audit, and a deadline-aware fallback policy. The calculation must never infer an interruption rate from a vendor savings ceiling or from one workload's history.

The practical go or no-go rule is simple: approve Managed Spot only when every accepted scenario satisfies the all-in cost inequality, the completion window, and the artifact-equivalence gate. Otherwise, shorten the checkpoint path, change the interval, relax the deadline with explicit business approval, or choose On-Demand. That conclusion is reproducible because every input can be replaced with the buyer's own pilot evidence.

The output should be retained with the job artifacts and price snapshot. That record lets reviewers reproduce the approval, distinguish assumptions from measurements, and update only the inputs that changed before the next run.
