



WHITEPAPER · ZCP OBJECT STORAGE

Object storage by industry

Six buying paths, the S3 workflows each one runs, and the questions to settle before moving data.

PUBLISHED BY
ZSoftly Technologies Inc.

DATE
September 2026

VERSION
1.0

AUDIENCE
Engineering, data and
platform leads

CONTENTS

Executive summary	Page 2
AI and machine learning · SaaS and software vendors	Page 3
Healthcare and life sciences · Finance, fintech and insurance	Page 4
Media and gaming · Public sector and research	Page 5
What ZCP Object Storage supports today	Page 6
Choosing a storage model	Page 7
Migration checklist	Page 8
How ZSoftly can help	Page 9
Sources	Page 10

Product facts were checked against the ZCP website and documentation on 29 September 2026. Prices change. The current rates are at zcp.zsoftly.ca/pricing.

EXECUTIVE SUMMARY

The same S3 API, six different sets of questions

Teams in every industry store objects: datasets, uploads, exports, media, records and backups. The S3 API is the same. The data, the people who touch it and the rules around it change. This paper walks through six buying paths and lists, for each one, the workflows teams run and the questions to answer before any data moves.

The problem

- Transfer charges make large datasets and frequent downloads expensive to move and hard to forecast.
- Teams want to keep the S3 tools and SDKs already built into their pipelines and applications.
- Buyers and their customers ask where the data lives, and need a clear answer.
- Regulated teams often start with a product decision when they should start with a data classification.

What ZCP Object Storage is

ZCP Object Storage is the self-service, S3-compatible object store on the ZSoftly Cloud Platform. It runs on Ceph RGW and works with standard S3 clients and SDKs. You place each bucket in the Montreal or Ottawa region you choose. There is no charge for data transfer in or out, and no minimum commitment. Storage and API requests are billed, at the rates published on the pricing page.

For teams that need dedicated hardware, root access, or block and file storage next to object storage, ZSoftly Cloud Storage (ZCS) is a separate offer: a dedicated, single-tenant Ceph cluster. Page 7 helps you choose between the two.

S3-compatible

Standard clients and SDKs, on Ceph RGW.

Your region

Montreal or Ottawa, chosen per bucket.

No transfer charges

Data in and out is not metered.

Six paths

Workflows and questions for each industry.

How to read this paper

Each industry section has the same four parts: the buyer's problem, the common S3 workflows, what ZCP supports today, and the questions to settle first. We describe workloads, not certifications. ZCP Object Storage makes no compliance claim for any industry. Where rules apply to your data, they decide the design, and we review that design with you.

The six paths at a glance, pages 3 to 5

AI and machine learning

Datasets, checkpoints, RAG sources

SaaS and software vendors

Tenant uploads, exports, backups

Healthcare and life sciences

Imaging archives, study data

Finance, fintech and insurance

Statements, claim files, exports

Media, streaming and gaming

Source media, game builds

Public sector and research

Records, open datasets, archives

Who should read this

Engineering and data leads planning an S3 move, **platform teams** choosing shared or dedicated storage, and **security and risk teams** who confirm the design before data moves.

BUYING PATHS 1 AND 2 OF 6

AI and machine learning · SaaS and software vendors



AI and machine learning data platforms

Machine learning teams need a durable store for datasets, checkpoints, evaluation output and the source documents behind retrieval-augmented generation (RAG). It has to work with the S3 tools already in the pipeline.

COMMON S3 WORKFLOWS

- Upload raw and curated datasets
- Write checkpoints and training logs
- Keep evaluation artifacts
- Store RAG source documents before chunking and embedding
- Export artifacts to other S3-compatible tools

WHAT ZCP SUPPORTS TODAY

An S3-compatible API for the SDKs and CLIs your jobs already use. Buckets in the region you choose. No charge to move data in or out.

Dataset versioning stays in your data tooling. For a dedicated AI data plane or a custom Ceph layout, talk to us about ZCS.

SETTLE FIRST

- Where should the datasets live?
- Is shared storage enough, or do you need a dedicated cluster?
- Which tools version your datasets and models today?



SaaS and software vendors

SaaS teams need application object storage that keeps their S3 integration, keeps transfer costs predictable, and gives enterprise buyers a clear answer about where data lives.

COMMON S3 WORKFLOWS

- Tenant file uploads
- Generated reports and exports
- Avatars and attachments
- Import staging
- Build artifacts and application logs
- Database backup objects

WHAT ZCP SUPPORTS TODAY

Existing S3 clients and SDKs, pointed at a new endpoint. A region per bucket. Published transfer terms with no charge for data in or out.

Tenant separation belongs in your application design. We review it with you.

SETTLE FIRST

- How are uploads, exports and tenant access designed today?
- Which parts of your application change in an S3 migration?
- Which S3 features does your code call, such as presigned URLs, bucket policies or lifecycle rules?

BUYING PATHS 3 AND 4 OF 6

Healthcare and life sciences · Finance, fintech and insurance



Healthcare and life sciences

Providers, imaging vendors, research teams and life-science companies handle imaging, study data, documents and exports that need defined access and location controls.

COMMON S3 WORKFLOWS

- DICOM object archives from PACS
- Pathology images
- Clinical-trial data exports
- Genomics files and research datasets
- Application attachments
- Backup copies


WHAT ZCP SUPPORTS TODAY

S3-compatible storage in the region you choose, with an architecture review of your workload.

ZCP makes no compliance claim for health data. Your classification sets the controls. ZCS offers dedicated storage.

SETTLE FIRST

- How is the data classified?
- Which controls do you need, and who confirms them?
- Should it run on a dedicated deployment?



Finance, fintech and insurance

Banks, fintechs and insurers keep documents and large data collections while managing third-party risk, portability, access control and operating cost.

COMMON S3 WORKFLOWS

- Statements and policy documents
- Claim attachments
- Regulatory report exports
- Market and risk data inputs
- Data-lake and data-science outputs
- Backup copies

WHAT ZCP SUPPORTS TODAY

The S3 API keeps workflows portable, and free transfer makes large moves practical.

Object Lock on a versioned bucket stores objects write-once-read-many. It does not prove compliance (page 6).

SETTLE FIRST

- What are your residency and retention requirements?
- If OSFI B-10 applies, how will you meet its cloud and portability expectations [5]?
- How will you show portability?
- What does your third-party risk process need from a provider?

BUYING PATHS 5 AND 6 OF 6

Media and gaming · Public sector and research

**Media, streaming and gaming**

Studios, publishers, platforms and game teams need low-cost storage for large asset libraries, user uploads, distribution files and production archives.

COMMON S3 WORKFLOWS

- Source media files
- Transcoding inputs and outputs
- User-generated content
- Press assets
- Game builds, patches and downloadable content
- Production archives

WHAT ZCP SUPPORTS TODAY

S3-compatible tools for object workflows, no charge for data transfer in or out, and a region per bucket.

ZCP does not document a CDN or a streaming service. Plan delivery to end users as a separate layer.

SETTLE FIRST

- Does delivery need a CDN or streaming layer?
- Who downloads the files, and from where?
- Has the delivery design been scoped before storage is chosen?

**Public sector, research and higher education**

Government teams, universities and researchers manage records, research data, open datasets, course materials and project archives under procurement and retention rules.

COMMON S3 WORKFLOWS

- Records and public datasets
- Research data repositories
- Collaboration exports
- Course material and recorded lectures
- Project archives
- Backup copies

WHAT ZCP SUPPORTS TODAY

The S3-compatible API, a region per bucket and published transfer terms. ZCS for a dedicated storage deployment.

We make no claim of government approval or suitability for classified data. Tell us your requirements and we review the design against them.

SETTLE FIRST

- What classification and authorization requirements apply?
- What retention and accessibility rules apply?
- What does procurement need from the provider?

ACROSS ALL SIX PATHS

What ZCP Object Storage supports today

The six paths share one product. The table separates what ZCP documents today from what you should settle with us in an architecture review before you depend on it.

Topic	Documented today	Settle in an architecture review
API	S3-compatible API on Ceph RGW. Standard S3 clients and SDKs [1].	Each S3 feature your code calls, such as bucket policies, presigned URLs, lifecycle rules or event notifications
Location	Each bucket lives in the Montreal or Ottawa region you choose [1].	A copy in a second region, backup design, recovery targets
Transfer	No charge for data transfer in or out [1].	Delivery to end users, including any CDN or streaming layer
Retention	Bucket versioning, and Object Lock on a versioned bucket for write-once-read-many storage [3].	Retention periods, and whether that setting meets your own rules
Security	Access with S3 credentials through standard clients.	Encryption, key management and access policy for your data classification
Pricing	Storage and API requests are billed, with no minimum commitment. Rates are published in CAD [2].	An estimate from your stored volume, object count and request rate

Table 1. Product scope for ZCP Object Storage, checked 29 September 2026.

What Object Lock does

A versioned bucket with Object Lock enabled stores objects in a write-once-read-many model. Object Lock requires bucket versioning, and object versions count toward your storage costs [3]. Choose both settings when you create the bucket. Plan the extra storage for every version you keep.

Two details that save a failed first run

- **Bucket names carry a suffix.** ZCP appends a numeric suffix to the name you request. A bucket requested as `app-backups` is created as a name like `app-backups-001024`. Use the full name in endpoint URLs, SDK calls and `aws s3` commands. `zcp object-storage bucket list` shows it [3].
- **Change configuration, not code.** Most S3 clients need only a new endpoint URL and new credentials. Keep both in configuration, then test every S3 call your application makes before switching traffic.

Where the price lives

This paper does not repeat rates, because they change. The current storage and request prices for each region are on zcp.zsoftly.ca/pricing [2]. Build your estimate from three numbers: the data you will store, the number of objects, and the requests per month.

DECISION

Choosing a storage model

Pick the storage model before you pick the bucket. Most S3 workloads fit ZCP Object Storage. A smaller set needs its own cluster.




Figure 1. Two questions decide most cases. Anything unusual goes to an architecture conversation.

Question	ZCP Object Storage	ZSoftly Cloud Storage (ZCS)
Tenancy	Shared, self-service platform	Dedicated, single-tenant Ceph cluster [4]
Your level of access	Buckets and objects, through the portal, CLI and S3 API	Root-level access to the cluster [4]
Storage types	Object (S3)	Object (S3), block (RBD) and file (CephFS) [4]
Where it runs	The Montreal or Ottawa region you choose	A ZSoftly region or your own hardware [4]
Commitment	No minimum commitment [1]	Sized and scoped with you
How to start	Create a bucket in the portal	An architecture conversation
Choose it when	You run ordinary S3 workloads and want to start today	You need dedicated hardware, root access, a custom layout, or object, block and file on one cluster

Table 2. The two storage models side by side.

Start on ZCP Object Storage and move to ZCS later if you need to. Both speak the S3 API, so the application change is an endpoint and a set of credentials.

When to start a ZCS conversation

- An AI team wants a dedicated data plane or a custom Ceph layout.
- A health or research team needs dedicated storage control after its data classification.
- A public-sector buyer asks for a dedicated storage deployment.
- An application needs block or file storage next to its objects.
- Your operations team needs root access to the storage cluster.
- The hardware has to be yours, in your own facility [4].

BEFORE, DURING AND AFTER

Migration checklist

Move one workload at a time. Keep the source intact until the new copy has been checked and has served real traffic.

Before you move data

- 1 List every bucket in scope with its size, object count and monthly request volume.
- 2 List the S3 calls each application makes. Mark any feature you have not confirmed on ZCP, such as bucket policies, presigned URLs or lifecycle rules.
- 3 Classify the data and write down who must approve the move.
- 4 Choose the storage model on page 7 and the region for each bucket.
- 5 Estimate the monthly cost from the pricing page and your own volumes.

While you move it

- 1 Create the bucket. Decide on versioning and Object Lock at this step.
- 2 Record the full bucket name, including the suffix ZCP adds.
- 3 Copy the data with the S3 tool your team already uses, such as `aws s3 sync` with a custom endpoint.
- 4 Check the copy. Compare object counts, total size and checksums for a sample of objects.
- 5 Change the endpoint and credentials in each application's configuration, one environment at a time, starting with `dev`.

After the switch

- 1 Watch error rates and latency in the application for an agreed window.
- 2 Keep the source bucket read-only until that window closes, so you have a way back.
- 3 Compare the first bill with your estimate.
- 4 Retire the source copy and record the new bucket, region and owner.

Mistakes that stall a first migration

Mistake	What happens	Prevention
Using the requested bucket name	Calls fail with bucket not found	Use the full name with the suffix ZCP adds [3]
Deciding retention late	Data lands in a bucket without the settings you need	Decide versioning and Object Lock when you create the bucket [3]
Assuming every S3 feature	A policy, presigned URL or lifecycle call behaves differently	Test each S3 call your code makes before switching traffic
Forgetting version storage	The bill is higher than the estimate	Count kept object versions in the estimate [3]

No charge for data in or out. Copying into ZCP, reading back to verify, and switching back if needed carry no ZCP transfer charge [1]. Check whether your previous provider charges you to read data out, and read its terms before you plan the copy.

WORKING WITH US

How ZSoftly can help

ZSoftly Technologies Inc. builds and runs the ZSoftly Cloud Platform. Our engineers also design and move workloads for clients in Canada, the United States and the United Kingdom. The engineers who plan your move are the ones who carry it out.

Assess

Inventory your buckets, S3 calls and data classes. Estimate the monthly cost.

Design

Choose ZCP Object Storage or ZCS, the region and the bucket settings, in writing.

Build

Create the buckets, copy the data and check every object count.

Operate

Switch applications over one environment at a time, with a way back.

Hand over

Runbooks, credentials and bucket records transfer to your team.

Where to start

- 1 **Try it yourself.** Create a bucket in the ZCP portal and point one S3 client at it.
- 2 **Book a discovery call.** Your first hour with our engineers is free. Bring one workload and its questions from this paper.
- 3 **Architecture review.** For regulated data, dedicated storage or large moves, we write up the design, the open questions and the cost estimate before any data moves.

What you keep

- Your data, in buckets in your own account, reachable with any S3 client.
- The migration plan, scripts and runbooks we write for you.
- The freedom to move again later, because the API is S3 and transfer out carries no ZCP charge.

Bring these to the discovery call

- The buckets in scope, with size, object count and monthly requests.
- The S3 calls each application makes.
- The data classification, and who must approve the move.
- The region you want for each bucket, and whether you need a second copy.
- Your retention rules, and whether they need Object Lock.
- What your current provider charges to read data out.

What the architecture review covers

The review works through the right-hand column of Table 1 on page 6 for your workload: the S3 features your code depends on, a second copy and recovery targets, delivery to end users, retention periods, and encryption, key management and access policy for your data classification. You receive it in writing with a cost estimate built from your own volumes.

Talk to an engineer. zcp.zsoftly.ca/contact or +1 (647) 771-1065. Product details: zcp.zsoftly.ca/services/object-storage and zcp.zsoftly.ca/services/cloud-storage.

REFERENCES

Sources

All sources were checked on 29 September 2026.

1. ZCP. Object Storage service page.
<https://zcp.zsoftly.ca/services/object-storage/>
2. ZCP. Pricing.
<https://zcp.zsoftly.ca/pricing>
3. ZCP documentation. Create a bucket (versioning, Object Lock, bucket naming).
<https://docs.zcp.zsoftly.ca/public-cloud/storage/object-storage/create-bucket/>
4. ZSoftly Cloud Storage, dedicated Ceph clusters.
<https://zcp.zsoftly.ca/services/cloud-storage/>
5. OSFI. Guideline B-10, Third-Party Risk Management.
<https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/third-party-risk-management-guideline>

Notices

This paper is for general information. It describes workloads and product features. It is not legal, regulatory or compliance advice, and ZSoftly Technologies Inc. makes no claim that ZCP Object Storage or ZSoftly Cloud Storage meets any regulation, certification or standard for your data. Product features and prices change. Check the linked pages before you plan. Amazon S3 is a trademark of Amazon.com, Inc. or its affiliates. Other names are trademarks of their owners.

© 2026 ZSoftly Technologies Inc. Version 1.0, September 2026.