Journal Platform — Hosting Evaluation & Production Quotation

Summary and recommendation

The platform is one codebase deployed as six isolated stacks (five journals + a shared platform), not a 3-service app, so the quote now prices that system in two phases and two options. Cloudflare stays as the edge in both options.

Monthly, incl. 18% GST	Launch (IJMAT + platform, 11 Dec 2026)	All 5 journals
Option A — per SoW: one Hostinger VPS (Mumbai) + Cloudflare + S3 backups	~₹8,900	~₹23,400
Option B — high availability: AWS Mumbai, multi-AZ, Cloudflare in front	~₹54,900	~₹1,06,600

- **Option A** is cheapest; a server failure means downtime and a manual restore of 2–6 hours.
- **Option B** absorbs a server or zone failure with zero to two minutes of impact and is the configuration the client's high-availability request describes. It is not a single-server setup.
- **DevOps proposal:** go live on Option B if the client approves by 7 Nov 2026. Otherwise launch IJMAT on Option A on 11 Dec and move to Option B before J2 goes live; containers, Cloudflare, S3 and backups are the same in both, so the move is an origin switch plus a database migration.
- All data stays in India: Mumbai primary, Hyderabad backup copies.

File storage adds about ₹2,700/month per extra 500 GB in either option. Figures are list prices as of 23 Oct 2026 at ₹97/USD; the AWS Pricing Calculator export for both phases is still to be attached.

What the system runs

One Node 20 / NestJS 11 codebase (Prisma 6, PostgreSQL 16) is deployed as six isolated stacks: one per journal plus a shared platform. Isolation comes from separate containers, databases, queues and storage, not separate machines; `JOURNAL_CODE` and per-journal env select the journal at runtime.

Component	What it does	Size per copy (est.)	At launch (IJMAT)	All 5 journals
journal-api	Public site API, author and admin routes for one journal	0.5 vCPU, 0.5–1 GB	1	5
journal-worker	Email, SMS, WhatsApp, PDFs (Chromium), Word conversion, DOI deposits, Scholar checks, scheduled jobs	1 vCPU, 1–2 GB, bursty	1	5
platform-api	Staff login, roles per journal, reporting	0.25–0.5	1	1

Component	What it does	Size per copy (est.)	At launch (IJMAT)	All 5 journals
		vCPU, 0.5 GB		
Journal websites	Public sites; article pages server-rendered for Google Scholar	~0.5 vCPU, 0.5 GB if SSR	1	5 domains (one multi-site app proposed)
Admin dashboard	Staff UI, static React build	static	1	1
PostgreSQL databases	j1_db…j5_db + platform_db on one server	—	2	6
Redis	BullMQ queues and cache, one key prefix per journal (`j:IJMAT:`)	1–2 GB, persistent	1	1
ClamAV	Virus scan of every upload	0.5 vCPU, 1.5 GB	1 per node	1 per node
File storage	Manuscripts, revisions, PDFs, figures, certificates, invoices	—	1 bucket	5 + platform

Every backend service runs two copies in production, so the full system is about 11 services × 2 plus websites and ClamAV.

Workload used for sizing (estimates, to confirm)

Input	Value used
Readers	50k–200k page views/month per journal; peak 100–300 concurrent users across journals
Submissions	20–100 per day per journal; uploads ~20 MB per manuscript, ~50 MB per submission
Staff	10–30 concurrent dashboard users
Bursts	Email batches 08:30, 08:45, 09:00 IST; reminders 13:00 and 19:00; issue publication days; Scholar crawling
Database	20–100 reads/s, 5–20 writes/s, up to ~250 connections; under 5 GB per journal DB in year 1
Files	200 GB placeholder for the existing archive (client to confirm); +500 GB/year as a safe upper bound; 10–25 MB per published paper
Bandwidth	Mostly PDFs through Cloudflare, under 1 TB/month

Input	Value used
Exchange rate and tax	₹97/USD , 18% GST on all cloud charges

Third-party services (ZeptoMail, MSG91, Razorpay, Stripe, Crossref, SerpAPI, reCAPTCHA) run on the client's own accounts and are listed separately in the cost section.

Option A (per SoW) vs Option B (high availability)

Option A costs about a sixth of Option B but accepts hours of downtime on a server failure; Option B removes that risk at roughly ₹46,000 more per month at launch. Both use Cloudflare, S3 in Mumbai for files and backups, and the same containers.

	Option A — per SoW	Option B — high availability
Setup	One Hostinger KVM VPS in Mumbai running Docker (all stacks, PostgreSQL, Redis, ClamAV); Cloudflare in front; files and backups in S3 Mumbai	AWS Mumbai: ECS on 2–3 nodes across 2 AZs, RDS PostgreSQL Multi-AZ, ElastiCache Multi-AZ, S3, backups copied to Hyderabad; Cloudflare in front
Monthly at launch, incl. GST	~₹8,900 (KVM 4)	~₹54,900
Monthly, 5 journals, incl. GST	~₹23,400 (KVM 8)	~₹1,06,600
Server or hardware failure	Site, submissions, payments and scheduled jobs stop until the VM is recovered or rebuilt	Absorbed by the other zone; no visible downtime
Recovery time	2–6 h: new VPS, deploy from Git, restore DB and Redis	Automatic, 0–2 min
Data loss on failure	Minutes, with WAL archiving to S3 (up to 24 h if only daily dumps)	None for server or zone failure
OS and host maintenance	Short downtime per update	Rolling, no downtime
Capacity ceiling	8 vCPU / 32 GB on the largest plan	Add or resize nodes, DB and Redis without migration
One-time setup effort	6–8 engineer-days	16–20 engineer-days
11 Dec 2026 go-live	Low risk	Feasible if approved by 7 Nov 2026; tighter

	Option A — per SoW	Option B — high availability
Commercials	Inside the fixed-price SoW	Change request: the SoW quotes the AWS move separately

Hostinger lists KVM 4 at ₹1,099/month and KVM 8 at ₹2,199/month on long prepaid terms, excluding GST; renewal is roughly double .Hostinger offers a Mumbai data centre, which keeps Option A inside India

DevOps proposal: Option B for go-live if approved by 7 Nov. If not, Option A for 11 Dec, then move to Option B before J2 launches; the move is a database migration and a Cloudflare origin change, with a 30–60 minute read-only window.

Edge: Cloudflare or CloudFront

Proposal: keep Cloudflare (Option 1). It is already in use, edtechpublishers.com has moved to the client's account, the app already reads `CF-Connecting-IP`, and it works unchanged for both hosting options. CloudFront is slightly cheaper on AWS but costs a migration. The two are never stacked.

	Option 1 — keep Cloudflare (proposed)	Option 2 — move to CloudFront + Route 53
What it covers	DNS, WAF, CDN, DDoS, TLS for all journal domains	Same, on AWS; Cloudflare retired
Plan	Pro per domain: $20/month on annual billing ($25 monthly)	CloudFront Pro flat-rate: $15/month per distribution, no overage
Monthly at launch, incl. GST	~₹5,700 (2 domains + ~$10 AWS egress on cache misses)	~₹3,400 (2 distributions; origin transfer waived)
Monthly, 5 journals, incl. GST	~₹16,600 (6 domains + ~$25 egress)	~₹10,300 (6 distributions)
Extra effort	None	2–3 days: move DNS for every domain, rebuild WAF and cache rules, certificates, switch real-IP header
Works with Option A (Hostinger)	Yes	Yes, but adds AWS to an otherwise non-AWS setup
Upload body limit	100 MB per request on Free and Pro, above the ~50 MB submission maximum	To verify against CloudFront quotas

Cloudflare Free ($0) also covers DNS, CDN, DDoS and basic custom WAF rules; Pro adds the managed WAF rulesets and is advised on domains that take submissions and payments. Cloudflare bills per

domain. CloudFront plans are per distribution. Edge caches hold only public article pages and PDFs; authenticated and personal-data routes send `Cache-Control: no-store`.

Provider comparison

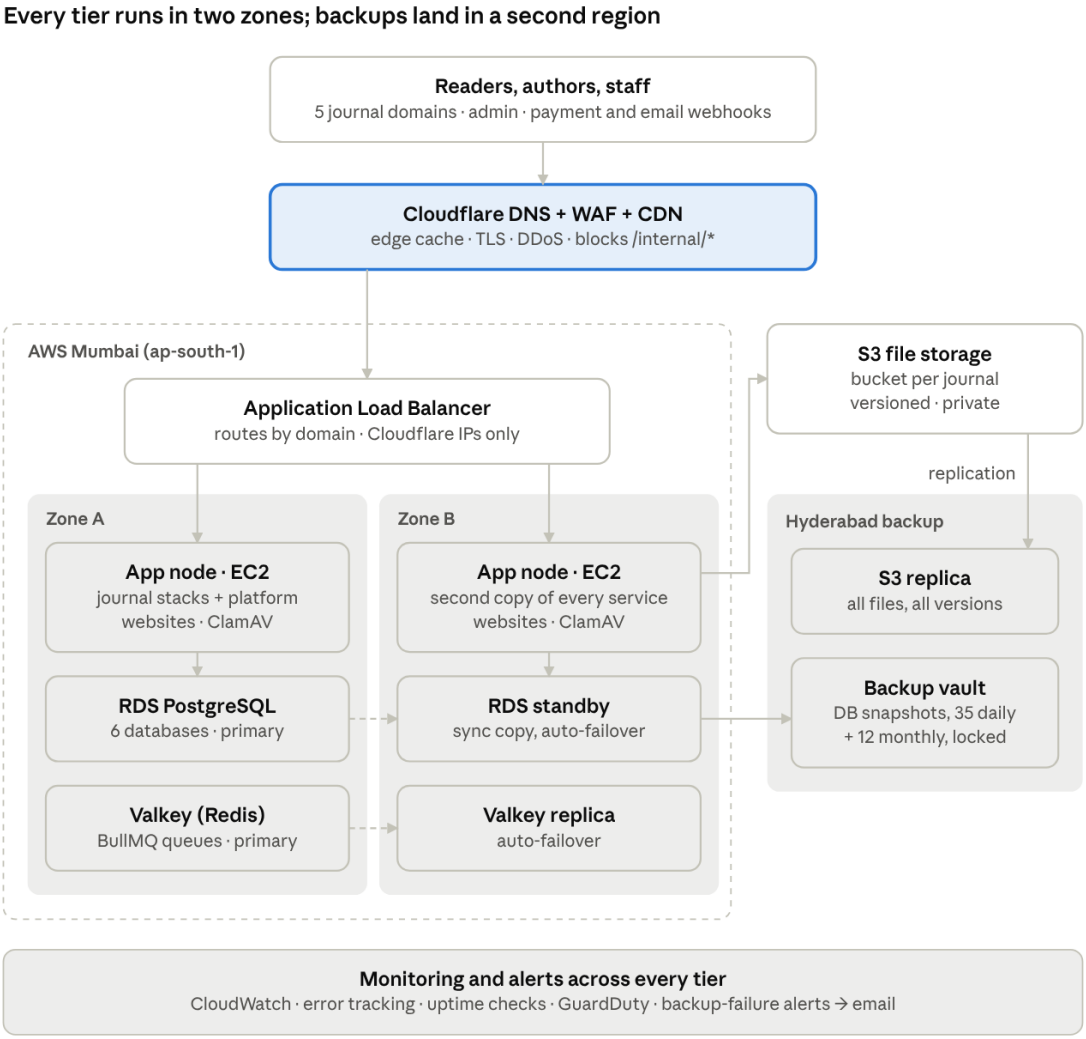
For the high-availability option, AWS and DigitalOcean both qualify; AWS is proposed for its managed Multi-AZ database and cache, locked cross-region backups and two Indian regions. Hostinger fits Option A only.

Requirement	AWS (Mumbai + Hyderabad)	DigitalOcean (Bangalore)	Hostinger VPS (Mumbai)
Managed PostgreSQL with standby	RDS Multi-AZ, automatic failover	Yes; new clusters need Advanced Edition (from $130/month per node, standby billed the same)	None; self-managed on the VPS
Managed Redis with failover	ElastiCache Multi-AZ	Managed caching offered; not priced here	None; self-managed
Second region in India for backups	Yes (Hyderabad)	No second Indian region assumed; to verify	No
Deletion-locked backups	AWS Backup Vault Lock	Not evaluated	Off-site copies to S3
Object storage in India	S3 Mumbai	Spaces; India availability to verify	None; use S3 Mumbai
Fit	**Option B (proposed)**	Viable alternative for Option B	**Option A (per SoW)**

- DigitalOcean's Standard Edition databases stop offering standby nodes for new clusters from 30 Nov 2026, so HA requires Advanced Edition.
- Hostinger VPS prices are promotional on long prepaid terms and renew substantially higher.
- Azure and Google Cloud offer the same building blocks as AWS at similar or higher cost and were not priced in detail.

Option B architecture (AWS, high availability)

Every service runs two copies on app nodes in separate Mumbai zones, behind a load balancer that routes each journal domain to its own stack. PostgreSQL and Valkey each have a standby in the other zone, Cloudflare stays at the edge, and every backup is copied to Hyderabad.



production architecture · Mumbai primary, Hyderabad backups

Layer	Launch (IJMAT + platform)	All 5 journals	Notes
Edge	Cloudflare Pro on ijmat domain + platform domain	Cloudflare Pro on 6 domains	Load balancer accepts traffic from Cloudflare IP ranges only
Routing	ALB host rules: journal domain → its journal-api and the web app; admin and platform hostnames → dashboard and platform-api; `/internal/*` → 403	Same, one rule set per journal	Health checks on `/api/v1/health`

Layer	Launch (IJMAT + platform)	All 5 journals	Notes
App nodes	ECS on 2× EC2 t4g.large (2 vCPU, 8 GB), one per zone	3× t4g.xlarge (4 vCPU, 16 GB)	Memory reserved ~11 GB at launch, ~35–40 GB at 5 journals; CPU shared so bursty workers borrow idle CPU. Confirm with a load test; move to m7g if CPU stays high
Services	2 copies each of journal-api, journal-worker, platform-api, web app; ClamAV on every node	11 backend services × 2	Each journal's services, DB user, bucket and Redis user are separate
Websites	One Next.js multi-site app, server-rendered, host-based	Same app, more domains	New journal by configuration, no rebuild
Admin dashboard	Static build in a small nginx container	Same	Buckets stay private
Database	RDS PostgreSQL 16 db.t4g.medium Multi-AZ, 2 databases	db.t4g.large Multi-AZ, 6 databases	~250 connections at full scale fits; pool sizes set per service
Queues	ElastiCache Valkey cache.t4g.small, primary + replica	cache.t4g.medium	`noeviction` policy for BullMQ; ACL user per journal prefix
Files	S3 bucket per journal + platform, private, versioned	6 buckets	PDFs streamed by journal-api on the journal's own domain (`/pdf/…`) and cached by Cloudflare
Outbound	2 NAT Gateways with fixed Elastic IPs; S3 gateway endpoint (free)	Same	The two IPs are whitelisted at MSG91 and payment gateways
Staging	Single-AZ, small sizes, same Terraform	Same	Tests payments, email and DLT templates before go-live
Delivery	Terraform + GitHub Actions, ARM64 images, rolling deploys	Same	Includes ClamAV and Chromium images

Requirements from the application

The backend already depends on these behaviours, so the hosting must provide them in either option.

Requirement	Option B (AWS)	Option A (Hostinger)
Isolation: one journal stopping leaves the other four running; platform-api or dashboard stopping leaves sites, submissions, payments and jobs running	Separate ECS services, health checks, DB users, buckets and Redis ACL users per journal; platform-api is its own service	Separate containers and DB users on one host; the host itself is shared
Domain routing	ALB host rules per journal domain, plus admin and platform hostnames	Reverse proxy (nginx/Traefik) host rules
Inbound webhooks (Razorpay, Stripe, ZeptoMail, MSG91) reach `/api/v1/webhooks/*`	Cloudflare rule skips bot challenge and rate limits on that path	Same
`/internal/*` never public	Blocked by a Cloudflare WAF rule and an ALB 403 rule; services call it on the private network	Cloudflare rule + proxy deny
Real client IP for rate limits and upload-session binding	`CF-Connecting-IP` passed through; ALB accepts Cloudflare IPs only; `TRUST_PROXY` kept	Same, proxy accepts Cloudflare IPs only
Upload size above the app's maximum (~50 MB per submission, to confirm)	Cloudflare Free/Pro allow 100 MB per request; ALB has no lower limit	Cloudflare 100 MB; nginx `client_max_body_size` set above the app limit
Static outbound IP for provider whitelisting	2 NAT Gateway Elastic IPs (both whitelisted)	The VPS IP
Scheduled jobs run exactly once, including the 08:45 batch, during a node or zone failure	BullMQ on Multi-AZ Valkey; no cron on nodes. ElastiCache uses a replica plus snapshots rather than AOF; Amazon MemoryDB is the drop-in if append-only durability is required	Redis with AOF on the VPS; jobs stop while the VPS is down
Health checks	`/api/v1/health` on every API for ALB, ECS and uptime checks	Same, for uptime checks and container restart
Secrets per journal (DB URL, Redis prefix, ZeptoMail token, domain, author-token secret)	Secrets Manager, one secret per service, injected by ECS	Per-stack env files, root-only, backed up encrypted

Requirement	Option B (AWS)	Option A (Hostinger)
ARM64 images	Required on Graviton: every image, including ClamAV and Chromium, built for arm64	Not required (x86 VPS)
New journal by configuration	Terraform journal list adds services, DB, bucket, secrets, ALB rule and Cloudflare zone	Compose file and proxy rule per journal
Data in India (DPDP)	Mumbai primary, Hyderabad backups	Mumbai VPS, S3 Mumbai backups

Monthly cost in INR

Option B costs ~₹54,900/month at launch and ~₹1,06,600 with all five journals, incl. GST; Option A costs ~₹8,900 and ~₹23,400. On-demand list prices, Mumbai, ₹97/USD; files at the 200 GB placeholder.

Option B — AWS high availability

Service	Launch configuration	Launch ₹	5-journal configuration	5 journals ₹
Compute	2× t4g.large	6,344	3× t4g.xlarge	19,034
Server disks	2× 30 GB gp3	534	3× 30 GB gp3	800
Load balancer	ALB, host rules	2,328	ALB, more rules	2,716
Outbound + static IPs	2 NAT Gateways	8,245	2 NAT Gateways	8,536
Database	RDS db.t4g.medium Multi-AZ	11,834	RDS db.t4g.large Multi-AZ	23,668
Database storage	50 GB gp3 Multi-AZ	1,261	same	1,261
DB backup storage	PITR beyond free allowance	485	same	485
Backup copies	AWS Backup to Hyderabad, vault-locked	1,261	same, 6 DBs	1,455
Queues	Valkey cache.t4g.small × 2 (est.)	3,880	Valkey cache.t4g.medium × 2 (est.)	7,760
Files	S3, 200 GB versioned, bucket per journal	776	same	776

Service	Launch configuration	Launch ₹	5-journal configuration	5 journals ₹
File replica	S3 replica in Hyderabad	485	same	485
Data transfer	AWS → Cloudflare on cache misses (~$10)	970	~$25	2,425
Container registry	ECR	194	ECR	194
Monitoring	CloudWatch logs, metrics, alarms	2,425	more services and logs	4,365
Error tracking	Sentry free plan	0	Sentry Team	2,522
Secrets and keys	Secrets Manager + KMS	485	per-journal secrets	776
Threat detection	GuardDuty + CloudTrail	1,164	same	1,455
Edge	Cloudflare Pro × 2 domains	3,880	Cloudflare Pro × 6 domains	11,640
Subtotal	$480	**46,551**	$931	**90,353**
GST 18%		8,379		16,264
Total incl. GST		**54,930**		**1,06,617**
Budget with 10% buffer		60,423		1,17,279

Option A — per SoW (Hostinger)

Service	Launch	Launch ₹	5 journals	5 journals ₹
VPS (Mumbai)	KVM 4: 4 vCPU, 16 GB	1,099	KVM 8: 8 vCPU, 32 GB	2,199
Files + off-site backups	S3 Mumbai, versioned; DB WAL + daily backups; copy to Hyderabad (~$16)	1,552	same	1,552
Data transfer	S3 to VPS and internet (~$10 / ~$20)	970		1,940
Edge	Cloudflare Pro × 2 domains	3,880	Cloudflare Pro × 6 domains	11,640
Error tracking +	Free tiers	0	Sentry Team	2,522

Service	Launch	Launch ₹	5 journals	5 journals ₹
uptime				
Subtotal		**7,501**		**19,853**
GST 18%		1,350		3,574
Total incl. GST		**8,851**		**23,427**

VPS prices are first-term rates on long prepaid plans; at renewal they roughly double, adding ~₹1,500 (launch) or ~₹2,600 (5 journals) per month incl. GST. On Cloudflare Free instead of Pro, both options drop by ₹4,600 (launch) or ₹13,700 (5 journals) incl. GST.

Staging, storage growth and other charges

- **Staging:** Option B single-AZ copy (1× t4g.large, RDS db.t4g.small, small Valkey, ALB) ~₹12,000/month incl. GST, roughly halved if stopped outside working hours. Option A: a KVM 2 VPS, ~₹950/month incl. GST.
- **Files:** each extra 500 GB adds ~₹2,700/month incl. GST (primary, replica, replication). End of year 3 at +500 GB/year: ~+₹8,200/month.
- **Commitments (Option B, from month 2–3):** 1-year RDS reserved instance and EC2 Savings Plan cut ~₹6,000 (launch) to ~₹14,000 (5 journals) per month incl. GST.
- **Optional:** cross-region DB read replica for minutes-level RPO on a Mumbai outage, ~₹7,000–₹14,000/month.
-
- Third-party running costs (client's own accounts)

These are billed by each provider at its current price list and are the same in both options.

Service	Used for	How it is charged
ZeptoMail (Zoho)	Transactional email	Prepaid email credits
MSG91	SMS and WhatsApp	Per message; WhatsApp also carries Meta conversation charges; DLT registration for SMS templates
Razorpay	Domestic payments	Percentage per transaction
Stripe	International payments	Percentage + fixed fee per transaction
Crossref	DOI registration	Annual membership fee + per-DOI registration fee
SerpAPI	Google Scholar checks	Monthly plan by search volume
reCAPTCHA	Bot protection on forms	Free within Google's quota

Backups, point-in-time recovery and deletion protection

Every data store has an automatic daily backup, a copy in a second region, and a guard against accidental or malicious deletion.

Data	Backup method (Option B)	Retention	Deletion protection	Option A equivalent
PostgreSQL, all 6 databases	RDS automated backups with point-in-time recovery (~5 min behind live)	14 days (up to 35)	RDS deletion protection; final snapshot forced	WAL archiving + nightly base backup to S3 Mumbai (pgBackRest), PITR
PostgreSQL	AWS Backup daily snapshot at 02:00 IST + monthly, copied to Hyderabad	35 daily, 12 monthly	Backup Vault Lock (compliance mode)	Backup bucket copied to Hyderabad, Object Lock
Redis (queues, scheduled jobs)	ElastiCache daily snapshot + live replica	7 days	Snapshot copies kept on delete	AOF on disk + nightly RDB copy to S3
Files, per journal bucket	S3 versioning + replication to Hyderabad	Old versions 90 days; replica indefinite	Versioning undoes deletes; replica writable only by replication role	Same S3 setup
App configuration	Terraform + Git; images in ECR	Git history; last 20 images	Branch protection; immutable images	Compose files in Git
Secrets, per journal	Secrets Manager	7–30 day recovery window	Scheduled deletion only	Encrypted copy in the backup bucket

App servers hold no data, so they are rebuilt from code rather than backed up. Every backup job failure raises an alert by email (and Slack, if used).

Restoration procedures

1. **One or more deleted/overwritten files:** list object versions in S3, restore the previous version. Minutes, no downtime.
2. **Bad data in the database (wrong update, deleted rows):** restore RDS to a new instance at the minute before the incident; copy the affected rows back, or switch the app's DB endpoint to the restored instance. 30–90 minutes for a 50 GB database.
3. **Database instance lost or corrupted:** restore the latest snapshot or PITR to a new Multi-AZ instance, update the secret holding the endpoint, restart services. Under 1 hour.

4. **Mumbai region unavailable:** apply the same Terraform in Hyderabad, restore the latest copied DB snapshot, point the app to the replica bucket, switch the Cloudflare origin. 4–8 hours.

A restore drill runs every quarter into a scratch environment and its time is recorded against these targets.

Expected recovery times

Server and zone failures are absorbed automatically with zero to two minutes of impact; only human error and full-region loss need manual recovery. RTO is time to restore service, RPO is how much recent data can be lost.

Failure	What happens (Option B)	Recovery time (RTO)	Data loss (RPO)	Action needed
One container crashes	ALB stops routing to it within ~30 s; ECS starts a replacement	None visible	None	Automatic
One journal's stack fails	Other journals and the platform keep serving	That journal: minutes	None	Automatic restart, then investigate
App node (EC2) fails	Other zone's node keeps serving; autoscaling replaces the node	None visible; full capacity in ~5 min	None	Automatic
Bad deployment	Health checks fail, deploy rolls back	5–10 min	None	Automatic or one click
Database primary fails	RDS promotes the standby	1–2 min	None (synchronous standby)	Automatic
Valkey primary fails	Replica promoted; BullMQ reconnects; scheduled jobs continue and run once	~1 min	Jobs written in the last moments before failover, to reconcile from DB job rows	Automatic
Whole availability zone fails	Traffic shifts to the other zone; DB and Valkey fail over; second NAT keeps outbound working	1–2 min	None	Automatic
Accidental file deletion	Restore previous S3 version	Minutes	None	Manual

Failure	What happens (Option B)	Recovery time (RTO)	Data loss (RPO)	Action needed
Accidental DB change or deletion	Point-in-time restore to just before the mistake	30–90 min	Up to ~5 min	Manual
Mumbai region outage	Rebuild in Hyderabad from code and copied backups	4–8 hours	DB up to 24 h; files minutes	Manual (DR runbook)
Option A: VPS or host failure	Everything stops; new VPS, deploy, restore DB and Redis	2–6 hours	Minutes with WAL archiving	Manual

The region-outage RPO can be cut to minutes later by adding a cross-region DB read replica (~₹7,000/month incl. GST).

Security and monitoring

All costs for these controls are already in the monthly figure.

Area	Control
SSL/TLS	Cloudflare edge certificates; Full (strict) mode to the origin with a free AWS certificate on the ALB; TLS 1.2+; HSTS
Firewall / WAF	Cloudflare WAF managed rules, rate limiting and DDoS protection; webhook paths `/api/v1/webhooks/*` exempt from bot challenges and rate rules; `/internal/*` blocked at Cloudflare and at the ALB
Origin lock-down	ALB security group allows Cloudflare IP ranges only; real client IP from `CF-Connecting-IP`
Network	App nodes, database and Valkey in private subnets; DB and Valkey reachable only from app nodes; outbound through 2 NAT Gateways with fixed IPs
Server access	No SSH ports; admin access via AWS Systems Manager with audit log
Per-journal isolation	Separate IAM task role per journal (own bucket only), DB user per database, Valkey ACL user per key prefix
Uploads	Every upload scanned by ClamAV before it is stored as accepted
Data	Database, disks, snapshots and S3 encrypted at rest; buckets never public
Secrets	Per-journal secrets in Secrets Manager; app refuses placeholder secrets in production

Area	Control
Identity	Per-person IAM users with MFA; least-privilege roles for services and CI/CD
Threat detection	GuardDuty; CloudTrail on every API action; image vulnerability scans on push
Error tracking and logs	Sentry for journal-api and journal-worker; CloudWatch Logs per service
Uptime monitoring	External checks on each journal's `/api/v1/health` and homepage every minute
Resource monitoring	Dashboards for CPU, memory, 5xx rate, latency, DB connections and storage, queue depth, failed jobs
Alerts	Email (Slack too, if the client uses it): site down, 5xx spike, CPU/memory >80%, DB storage >80%, failover, failed-job spike
Backup alerts	Notification on any failed or missed backup, snapshot or replication
Cost alerts	AWS Budgets warning at 80% and 100% of monthly budget

Scaling path and one-time setup

In Option B every growth step is a configuration change; none needs a migration. Option A scales only up to the largest VPS plan, after which the move to Option B is the next step.

Trigger	Step (Option B)	Added cost, incl. GST
A new journal (J2–J5)	Add it to the Terraform journal list: services, database, bucket, secrets, Valkey user, ALB rule, Cloudflare domain. No image rebuild	~₹12,000–15,000 per journal (its share of the 5-journal total)
App nodes above ~70% memory or CPU	Autoscaling adds a node (max 4), or resize to the next size	~₹7,500 per t4g.xlarge
Sustained CPU on burstable nodes	Move nodes from t4g to m7g	Depends on load test
DB CPU or memory above 70%	Resize the RDS instance one size up	~₹14,000–28,000
Heavy read traffic or search	Add a read replica	~₹7,000–14,000
Queue memory above 70%	Resize Valkey one size up	~₹4,500

Trigger	Step (Option B)	Added cost, incl. GST
Need minutes-level RPO on a region outage	Cross-region DB replica + warm standby in Hyderabad	~₹15,000–25,000
File growth	Nothing to do; older files move to a cheaper class after 90 days	Lower than estimated

DevOps validation, decisions and next steps

Validated: Option B prices the full production system — six isolated stacks, two copies of every service across two zones, Multi-AZ PostgreSQL and Valkey, per-journal buckets, Cloudflare, locked backups in Hyderabad, staging and monitoring. It is not a single-server setup. Option A is priced as the SoW describes it.

Not yet a fixed price: totals rest on the workload estimates above. After the questions below are answered, both phases are rerun in the AWS Pricing Calculator and the export is attached.

Decisions for the client

Decision	DevOps proposal
Hosting: Option A (SoW) or Option B (HA)	Option B for go-live if approved by 7 Nov 2026; otherwise A for 11 Dec, B before J2
Edge: Cloudflare or CloudFront	Keep Cloudflare; never run both
Cloudflare plan	Pro on domains that take submissions and payments; Free elsewhere
Setup effort and who pays	A inside the fixed-price SoW; B as a change request

Open questions

- ☐ Size of the existing archive to migrate (files in GB, number of articles, DOI list)
- ☐ Confirmed maximum upload size per file and per submission
- ☐ Current Cloudflare plan on the client's account
- ☐ Does the client use Slack for alerts, or email only?
- ☐ Backend lead: can scheduled jobs be rebuilt from DB job rows after a Redis failover, or is MemoryDB needed?
- ☐ Backend lead: confirm ClamAV and Chromium images build for ARM64; otherwise use x86 nodes (~20% more compute cost)
- ☐ Backend lead: one multi-site Next.js app for all journal websites, server-rendered (today's site is client-rendered)

Next steps

1. Backend lead reviews this version before it goes to the client.
2. Rerun launch and 5-journal totals in the AWS Pricing Calculator; attach the export.
3. Recheck the ₹97/USD rate on the day this is sent.
4. Load-test one journal stack to confirm node size.