



# SpazeNode

## User Guide

Mock Frontend Operating Manual

Version 1.0 - May 2026

A complete walkthrough of the SpazeNode console,  
host provider portal, and marketing site.

# Table of Contents

- 1. What is SpazeNode? ..... 3
- 2. Quick Start ..... 3
- 3. User Roles ..... 4
- 4. Navigation Map ..... 5
- 5. Marketing Site (/) ..... 6
- 6. Authentication & Onboarding ..... 7
- 7. Onboarding Tour ..... 8
- 8. Dashboard ..... 8
- 9. Deploy Wizard (4 Steps) ..... 9
- 10. Clusters ..... 10
- 11. Deployments & Metrics ..... 11
- 12. Models & Marketplace ..... 12
- 13. Real-time Monitoring ..... 13
- 14. API Playground & Keys ..... 13
- 15. Billing & Payment Checkout ..... 14
- 16. Storage, Audit, Notifications ..... 15
- 17. Settings, Integrations, Team ..... 15
- 18. Host Provider Console ..... 16
- 19. Documentation Portal ..... 17
- 20. Downloadable Reports ..... 18
- 21. Troubleshooting & Support ..... 19

# 1. What is SpazeNode?

SpazeNode is an AI Infrastructure platform packaged as Data Center as a Service (DCaaS). We aggregate the world's idle GPUs — household gaming rigs, ex-mining hardware, and underutilized regional data centers — and sell them as on-demand AI infrastructure at 40-70% lower cost than hyperscalers like AWS, Azure, or GCP.

## WHAT THIS MOCK DEMONSTRATES

This frontend is a fully interactive demo of the entire SpazeNode product surface: a customer Console (buy compute, deploy models), a Host Console (sell idle GPU time), and a public Marketing site. All data is mocked; nothing connects to a real backend.

## Stack

- Next.js 14 (App Router) + TypeScript
- Tailwind CSS with the SpazeNode Glass design system
- Framer Motion for animations
- Recharts for interactive charts
- jsPDF + jspdf-autotable for downloadable reports

# 2. Quick Start

## Start the dev server

```
cd /Users/parth/Desktop/Spazencode/spazencode-app
npm run dev
```

Then open your browser to:

```
http://localhost:3000
```

## Demo flow (5 minutes)

- / — Marketing landing page
- Click "Start Deploying" ! /signup ! /onboarding/org ! /onboarding/invite ! /d
- The Onboarding Tour auto-starts on /dashboard
- Click 'Deploy Now' ! walk through the 4-step Deploy Wizard
- Click 'Pay invoice' on /billing ! complete the checkout to see receipt PDF
- Visit /host for the supply-side flow

## 3. User Roles

SpazeNode has five distinct user personas, each with different access and views:

Role	Surface	Key permissions
Anonymous Visitor	Marketing site	Browse landing, pricing, models, contact sales
Admin (Customer)	Console	Full access — deploy, billing, team, settings
Member (Customer)	Console	Deploy + view; cannot change billing or invite team
Billing (Customer)	Console (Billing only)	View invoices, manage payment methods
Host Provider	Host Console	Connect GPUs, manage availability, withdraw earnings

### HOW TO SWITCH ROLES IN THE DEMO

Customer console and Host console use separate sidebars. Use the 'Switch to Console' link in the Host sidebar, or navigate directly: `/dashboard` (customer) vs `/host/dashboard` (provider).

## 4. Navigation Map

### Marketing site (public)

- `/` — Landing page (hero, world map, GPU showcase, stats)
- `/marketing/how-it-works` — 3-layer architecture explainer
- `/marketing/pricing` — Tiered pricing with cost calculator
- `/marketing/models` — Public model catalog
- `/marketing/contact` — Enterprise sales lead form

### Customer Console (authenticated)

- `/dashboard` — Overview
- `/clusters`, `/clusters/[id]` — GPU clusters
- `/deployments`, `/deployments/[id]`, `/deployments/[id]/metrics`
- `/deploy` — 4-step deploy wizard
- `/deploy/progress`, `/deploy/success` — post-deploy flow
- `/models`, `/models/[id]`, `/models/upload` — Model catalog + BYO
- `/marketplace` — GPU SKU marketplace with real product photos
- `/monitoring` — Real-time GPU utilization
- `/playground` — Interactive API tester
- `/api-keys` — Manage API keys
- `/docs` — Documentation portal

- /billing, /billing/checkout, /billing/methods, /billing/alerts
- /storage — Artifacts & datasets
- /audit — Activity log
- /notifications, /integrations, /settings

### Host Console (provider)

- /host — Public host landing
- /host/signup, /host/onboarding/install, /host/onboarding/benchmark
- /host/dashboard — Provider overview
- /host/gpus, /host/jobs, /host/earnings, /host/controls, /host/benchmark, /host/settings

## 5. Marketing Site

URL: <http://localhost:3000/>

The landing page is the front door for anonymous visitors. It uses a hero with a data center background image, real H100/A100 product photos in the GPU showcase strip, and an interactive SVG world map with pulsing region pins.

### Sections

- Hero: 'AI is growing exponentially. Infrastructure isn't.' with primary CTA
- GPU showcase: 4 real GPU product cards with hourly pricing
- Global GPU network: world map + 6 region pins + live counter
- Deploy in seconds: 4-step preview card
- Stats strip: 672+ GPUs, \$2.4M+ paid, 99.99% uptime
- Enterprise security: SOC 2, E2E encryption, IAM
- CTA section + footer

## 6. Authentication & Onboarding

### Sign up flow

- /signup — Email + OAuth (Google, GitHub)
- /onboarding/org — Org name, logo, default region, use case
- /onboarding/invite — Invite teammates with role selector (skippable)
- /onboarding/welcome — Quick start with templates
- Lands at /dashboard with the onboarding tour auto-starting

### Login

- /login — Email/password + OAuth providers
- Forgot password / reset (UI only, demo)

## 7. Onboarding Tour

On first visit to /dashboard the tour auto-starts with a darkened backdrop and a spotlight on each key section. Six steps walk users through the entire dashboard layout.

### Tour controls

- Next / Back buttons advance between steps
- Skip tour dismisses and remembers via localStorage
- Click anywhere on backdrop to close
- Restart tour button in the Dashboard topr

### Tour sequence

- 1. Stat cards (clusters, GPUs, deployments, calls)
- 2. Deploy AI in 4 Steps wizard hero
- 3. Cluster status visualization
- 4. Recent deployments
- 5. Usage overview donut
- 6. Sidebar navigation

## 8. Dashboard

URL: /dashboard

### Top stats row

- Active Clusters with weekly delta
- GPUs Online by type
- Running Deployments
- Total Inference Calls with weekly delta

### Deploy AI in 4 Simple Steps hero

A glass-panel callout summarizing the wizard. Clicking 'Deploy Now' opens the full /deploy flow.

### Bento grid

- Cluster Status — 3 visual server racks with load percentages + uptime / load stats
- Recent Deployments — last 3 deployments with status pills, clickable
- Usage Overview — donut showing compute / storage / network breakdown

## 9. Deploy Wizard (4 Steps)

URL: /deploy

Each step is animated with Framer Motion: slide transitions, pulsing rings on the active step.

### Step 1 — Choose Compute

- Pick GPU type from 5 SKUs (H100 SXM/PCIe, A100 SXM/PCIe, RTX 4090)
- Slider: cluster size (1-16 GPUs)
- Dropdown: region

### Step 2 — Select Model

- Catalog tab: 6 pre-configured models
- Bring your own tab: HF repo URL or safetensors upload

### Step 3 — Configure

- Deployment name
- Scale-to-zero toggle
- Min/max replicas
- Environment variables (key/value pairs)

### Step 4 — Review & Deploy

- Configuration summary
- Live cost estimate (hourly + monthly)
- 52% savings vs AWS callout
- Deploy now button (shimmering gradient)

After clicking Deploy: `/deploy/progress` shows a live pipeline (Reserving ! Pulling !) with progress bar and log stream. Then `/deploy/success` reveals the endpoint URL, one-time API key, and curl/Python/JS code samples.

## 10. Clusters

URL: `/clusters`, `/clusters/[id]`

Lists all provisioned GPU clusters across regions. Each card shows region, GPU type/count, uptime, load, and hourly cost. Clicking opens the cluster detail page.

### Cluster detail

- Status header pills (Status, GPUs, Uptime, Load, Cost/hr)
- Nodes table with per-GPU health (status, utilization, uptime)
- Active deployments on this cluster
- Download cluster report PDF

## 11. Deployments & Metrics

URL: `/deployments`, `/deployments/[id]`, `/deployments/[id]/metrics`

### Deployments list

- Filterable table: deployment, model, cluster, status, req/hr, cost/hr
- Download full deployments report PDF

### Deployment detail

- 5-step pipeline (Reserving ! Pulling ! Loading ! Warming ! Serving)
- Pipeline visualization with progress for currently-running step
- Job Details (Job ID, GPUs, start time, ETA)
- Real-time monitoring tab embedded

### Deployment metrics page

- KPI cards: Throughput, P95 Latency, Error Rate, GPU Util
- Interactive throughput chart (Recharts area)
- Active replicas list
- Environment variables

- Live aggregated log stream

## 12. Models & Marketplace

### Models catalog (/models)

- Filter by type (Chat / Image / Code / Embedding)
- 6 pre-configured models with deploy counts
- Click any model! detail page

### Model detail (/models/[id])

- Description, license, recommended GPU
- 3 deployment tiers (Best / Balanced / Budget)
- Per-million token pricing
- Use cases checklist

### BYO model upload (/models/upload)

- From Hugging Face — repo URL + access token
- Direct upload — drag .safetensors files
- License & safety acknowledgment

### GPU Marketplace (/marketplace)

- Filters: GPU Type, Memory, Location
- Real product photos for H100 + A100 cards
- Hourly pricing with In Stock indicator
- Trust strip: verified hardware, instant provisioning, 40-70% cheaper

## 13. Real-time Monitoring

URL: /monitoring

- 4 KPI cards: Active GPUs, GPU Util, Memory Util, Running Jobs
- Interactive utilization chart (Recharts area, 24h)
- Donut by GPU type with legend
- Activity by region (US-EAST, EU-WEST, US-WEST, AP-SOUTH)
- Two report downloads: full utilization PDF + region report

## 14. API Playground & Keys

### Playground (/playground)

- Pick endpoint (inference / clusters)
- Edit request parameters with live curl preview
- Mock response with status code, latency, payload size

### API Keys (/api-keys)

- Table of active keys (Name, Prefix, Scope, Created, Last used)
- Quick start curl example below

## 15. Billing & Payment Checkout

### Billing overview (/billing)

- 3 KPI cards: MTD Spend, Projected Total, Last Invoice
- Spend-over-time area chart
- Spend by service (compute / inference / storage / egress)
- Invoices table with Download PDF per row
- Top action: Export usage report PDF

### Payment Checkout (/billing/checkout)

- Stage 1 — Card form with brand badges, billing address, save card toggle
- Stage 2 — Processing animation: validate ! reserve ! confirm ! receipt
- Stage 3 — Success: receipt download, transaction ID, status pill

### Other billing pages

- /billing/methods — Cards on file, invoice address
- /billing/alerts — Monthly cap, threshold alerts

## 16. Storage, Audit, Notifications

### Storage (/storage)

- Files inventory: model weights, datasets, log archives
- Per-file size, region, last updated
- Download storage report PDF

### Audit log (/audit)

- Every action with actor, target, timestamp
- Download full audit log PDF

### Notifications (/notifications)

- Email + in-app toggle matrix
- 9 event types covered

## 17. Settings, Integrations, Team

### Settings (/settings)

- Profile — name, email, password, photo
- Organization — name, logo, default region
- Team & roles — members table with role pills

### Integrations (/integrations)

- 8 integrations: HuggingFace, S3, GitHub, Slack, PagerDuty, Datadog, Stripe, Webhook
- Connected / Disconnected status per integration

## 18. Host Provider Console

The supply side. Hosts contribute GPU compute and earn money. Separate sidebar with 'Switch to Console' link to return to the customer view.

### Host onboarding flow

- /host — Landing with earnings calculator + 4-step explainer
- /host/signup — Provider signup
- /host/onboarding/install — Download agent (Linux / Windows / macOS) + registration token
- /host/onboarding/benchmark — Hardware verification with live benchmark console

### Host dashboard (/host/dashboard)

- Welcome banner
- 4 KPI cards (Earnings 24h, Active GPUs, Active Jobs, Uptime)
- 30-day earnings chart with best day / avg / projected
- Recent jobs list with earnings per job

### Other host pages

- /host/gpus — Per-GPU detail with real product photos
- /host/jobs — Job history table
- /host/earnings — Available balance, lifetime, payout history, withdraw
- /host/controls — Resource allocation, weekly schedule, automation rules
- /host/benchmark — Re-run benchmarks, view certificate
- /host/settings — Profile, payout method (Stripe Connect), notifications

## 19. Documentation Portal

URL: /docs

- Quick Start Guide — first cluster in 5 minutes
- API Reference — programmatic infrastructure control
- CLI Tools — local management and automation
- Deployment Patterns — high availability, model serving
- Security & Compliance
- Cost Optimization

## 20. Downloadable Reports

Every major surface has a Download PDF button. Files are named 'spazencode-{type}-{id}.pdf' and saved to your Downloads folder. Generated client-side using jsPDF + jsPDF-autotable.

Page	Button	Output
/billing/checkout (success)	Download receipt	Receipt with line items, tax, total, card
/billing	Export usage (PDF)	Full billing report with KPIs, services, invoices

Page	Button	Output
/billing (row)	Download PDF	Individual invoice with line items + status pill
/clusters	Fleet report	All clusters as table
/clusters/[id]	Download report	Cluster config + nodes + deployments
/deployments	Deployments report	All deployments with KPIs
/deployments/[id]/metrics	Metrics report	KPIs + config + throughput sample
/monitoring	Utilization (PDF)	GPU stats + utilization by type
/monitoring	Region report	Per-region GPU + load breakdown
/audit	Audit log	Every action with actor + target
/storage	Storage report	Artifacts inventory
/host/dashboard	Host report	Provider earnings + jobs
/host/earnings	Earnings statement	Balance + payout history
/docs	User guide	This document

### PDF DESIGN

Every report has the blue brand bar, SpazeNode wordmark, generated timestamp, title block, KPI cards, primary-blue table headers with alternating row fills, status pills, and footer with page numbers.

## 21. Troubleshooting & Support

### Common issues

- Dev server won't start — Run 'rm -rf .next' and 'npm run dev' again.
- Onboarding tour replays — Clear localStorage or click 'Restart tour'.
- PDF download blocked — Check browser pop-up blocker.
- Images not loading — Ensure /public/gpus/h100.png and /public/datacenter/hero.png exist.

### Restart the app

```
cd /Users/parth/Desktop/Spazenode/spazenode-app
pkill -f "next dev"
rm -rf .next
npm run dev
```

### Regenerate this guide

```
node scripts/generate-user-guide.mjs
```

### Contact

- Founder: parth@skatechain.org
- Issues: track in your issue tracker of choice (mock has no backend)

**THANK YOU**

This mock demonstrates the full SpazeNode customer experience end-to-end. Walk investors, designers, or future engineers through it to align on the product vision before shipping the real backend.