



AI Enabled Capacity Operations Platform

Phase 1 Production Implementation

AI-first capacity management — from the moment information arrives to the moment a customer sees it, in one connected working system.

Proposal by ThirdEye Data

Prepared for: SPYRE Consulting



This document is confidential and proprietary to ThirdEye Data and is submitted solely for evaluation by SPYRE Consulting. It shall not be reproduced, disclosed, or distributed to any third party without prior written consent of ThirdEye Data.



Table of Contents

Document Control	4
What Phase 1 delivers	4
Why this matters for SPYRE	4
What SPYRE has asked for in Phase 1	5
What we learned from the sample workbook SPYRE shared	5
How things work today	6
Today versus with the platform	6
What changes for the day-to-day user	6
4.1 Efficient Intake	7
How we'll know it's working	8
4.2 Data Structure	8
4.3 Data Query	8
4.4 Inventory / Capacity Management	9
4.5 Customer Output	9
What each role can do	10
A few everyday journeys	10
How TEDRA operates	11
From source to outcome, end to end	11
How the pieces fit together	11
The technical stack, layer by layer	12
How we'll get there	13
Excel after cutover — what's still allowed	13
Milestone gates	15
The delivery team	16
How we'll stay aligned	16
What counts as “done”, area by area	17
How we treat defects	17
How we protect the platform	18
Getting ready to run it day to day	18
What SPYRE receives	19
Try it yourself: the interactive prototype	19
What we're assuming	20
What's out of scope for Phase 1	20



Key risks and how we'll manage them	20
Principles guiding what comes after Phase 1	21
Proposed next steps	22
The decision we're asking for	22

Document Control

Item	Detail
Document	SPYRE Capacity Operations Platform — Phase 1 Production Implementation
Prepared for	SPYRE Consulting
Prepared by	ThirdEye Data
Version / date	1.0 / 28 August 2026
Status	Customer-ready scope proposal; commercials pending internal confirmation
Confidentiality	Confidential — for SPYRE Consulting and ThirdEye Data review only

NOTE

Scope baseline

This proposal covers a Phase 1 production implementation — not another round of proof-of-concept work. Pricing and commercial terms will follow, once approved internally.

What Phase 1 delivers

Priority	What it means in practice
Efficient Intake	AI-assisted intake for email, Excel, PDF, call transcripts, notes and manual forms, reviewed on one screen instead of many.
Data Structure	Operators, sites, schedules, customers, deals, activity, sources and audit history, all linked with permanent IDs.
Data Query	Plain-language questions and structured filters, both backed by TEDRA and grounded in real evidence.
Inventory / Capacity Management	Clear ownership, freshness tracking, controlled status changes, deal visibility and protection against double-booking.
Customer Output	Polished, redacted capacity reports and proposals in Excel, PDF and PowerPoint.

Why this matters for SPYRE

- Less time spent reading, re-typing, filtering spreadsheets and formatting documents by hand.

- One dependable view of capacity, deal activity and customer-facing output.
- Fewer risks from outdated information or two people reserving the same capacity.
- Relationship knowledge, sources and decisions are kept on record — not just in someone's inbox or memory.
- A solid foundation to build on later, when SPYRE is ready for CLM, a customer portal or Salesforce integration.

SECTION 2

Understanding of the Requirement

SPYRE is a data-center capacity brokerage and consulting business. It doesn't own the inventory it sells access to — its edge comes from the relationships that surface market information, how current that information stays, and how quickly the right options reach a customer.

What SPYRE has asked for in Phase 1

01	Move quickly into a working system The first phase should be a production build — centered on Efficient Intake, Data Structure, Data Query, Inventory / Capacity Management and Customer Output.
----	--

What we learned from the sample workbook SPYRE shared

What we saw	What it means for the build
78 records in the master Sites sheet, spread across 48 columns	Phase 1's data model has to keep all of the existing technical, geographic, operational and time-based capacity fields — nothing gets dropped.
Five extra custom / opportunity-specific capacity sheets	These become saved shortlists and views tied to a deal — not separate copies of inventory.
No reliable permanent ID across sheets	Migration has to assign stable IDs and reconcile duplicate or similar-looking records before go-live.
Dates, prices and locations written in mixed formats	We need normalization rules plus a human review step for anything unclear.
Status (available, sold, etc.) isn't shown consistently	A dedicated, controlled lifecycle status is needed, kept separate from facility type or data-quality flags.
27 supplier contact rows, several with gaps	Operators and contacts need proper structured records, with validation and room to fill gaps over time.

Put simply: this is not a quick Excel-to-database upload. It's a controlled project to reconcile real data and change how the day-to-day work gets done.

SECTION 3


Current State and Target Operating Flow

How things work today


- Inventory, deal activity and customer conversations live in separate, disconnected places.
- People repeatedly read, clean up, copy, filter, redact and reformat the same information.
- Custom tabs drift away from the master sheet, because nothing links them together permanently.
- Whether a site is available, reserved or sold depends heavily on someone following up and remembering correctly.

Today versus with the platform

TODAY (AS-IS)



WITH THE PLATFORM (TO-BE)



What changes for the day-to-day user

Today	With the Phase 1 platform
Read and re-type every incoming source by hand	AI drafts the structured record and only flags what's missing, conflicting or uncertain


Filter through many spreadsheet columns	Ask a plain-language question, or use structured filters
Copy sites into a customer-specific tab	Save a live link to the inventory record inside a deal shortlist
Track activity through email and memory	See every email, note, transcript, document and decision in one timeline
Manually update availability and status	Use a controlled lifecycle, with suggested actions, approvals and built-in conflict protection
Rebuild customer documents from scratch	Generate an approved, redacted, versioned output package

SECTION 4

Phase 1 Scope: Danielle's Five Priorities

Phase 1 is built around five connected priorities. Each one solves a specific daily problem, and together they form one continuous working flow.

Phase 1 moves through five connected priorities



4.1 Efficient Intake

Nobody should have to fill in a long form for information that already exists in an email, spreadsheet, PDF, call transcript or note. Phase 1 introduces one shared intake queue: AI prepares the update, and the user reviews only the parts that need judgement.

- Accepted sources: email or forwarded mail, Excel, PDF, call transcripts, pasted notes, and manual entry.
- What AI does first: works out what kind of source it is, pulls out the key fields, normalizes units/dates/currency, and suggests which operator or site it matches.
- Built-in checks: duplicate detection, side-by-side comparison of conflicting fields, a confidence score, flags for missing values, and a link back to the original source.
- Human sign-off: approve, correct or reject the whole proposed update from one review screen before anything changes the master record.
- Notifications: the right owner is notified directly, instead of having to search through mail and files.

How we'll know it's working

- Every uploaded source stays linked to whichever approved records or fields came from it.
- A low-confidence or conflicting value can never quietly overwrite approved inventory.
- A user can approve, correct or reject an entire proposed update from one screen.

4.2 Data Structure

The platform keeps one reusable operating model instead of copying workbook tabs. Capacity stays at the center and can connect to several customer opportunities at once, without spinning off disconnected copies.



Record type	What it's for
Operator and supplier contact	Relationship ownership, outreach, source and contact history.
Site and capacity schedule	Location, technical details, total capacity and time-based availability.
Customer and deal	Requirement, owner, stage, tasks, activity and linked documents.
Shortlist	Live references to the right inventory for a specific opportunity.
Source and activity	Emails, PDFs, spreadsheets, transcripts, notes and decisions, in order.
Audit and output version	Who changed, approved, generated or shared what, and when.

4.3 Data Query


Business users can ask questions in plain language as the fastest way in, while structured filters stay available whenever precision or verification is needed.

ASK	Example business question "Show available sites in Finland with at least 50 MW becoming available by Q4 2027."
------------	--

- Turns a plain-language request into an explicit structured filter behind the scenes.
- Leaves out anything unavailable, sold or in conflict, based on configured rules.
- Explains why each result matched, and shows its source and last-verified date.
- Saves a result as an opportunity shortlist, without duplicating the inventory record.
- Asks a clarifying question when location, capacity, timing or another key detail is missing.
- Never invents inventory facts that aren't supported — it flags the gap instead.

4.4 Inventory / Capacity Management

Proposed lifecycle below — final business definitions to be confirmed with SPYRE in Week 1.



- Ownership: assigned sales or operations owner, reassignment history and open tasks.
- Freshness: last-verified date, staleness thresholds and reminders to refresh.
- Lifecycle: controlled changes, role-based approval, and automatic removal from active search where appropriate.
- Reservation protection: a database-level rule that blocks two active reservations against the same capacity.
- Deal visibility: shows where a piece of capacity has been shortlisted, presented, put under LOI, or reserved.
- Activity timeline: emails, notes, transcripts, documents, tasks and status changes, all in one chronological view.

4.5 Customer Output

The platform builds customer-facing capacity reports or proposals straight from an approved shortlist, using one frozen snapshot of data and redaction settings.

Capability	Phase 1 behavior
Redaction	Configurable hiding or masking of operator identity, exact location, pricing and internal notes.

Content generation	Executive summary, requirement, site options, comparison and a narrative grounded in sources.
Formats	Downloadable Excel, PDF and PowerPoint.
Versioning	Keeps the records, field values, redaction rules, author, approver and timestamp for every output.
Approval	Manager review before release, where configured.
Branding	Applies SPYRE-approved colors, typography and templates.

SECTION 5

Users, Roles and Low-Click Journeys

What each role can do

Capability	Sales user	Manager / Operations approver
Inventory	View assigned and permitted records; propose additions or updates	View all permitted inventory; approve exceptions and protected changes
Deals	Create and manage assigned opportunities, activities and shortlists	View pipeline across the team; reassign ownership
Lifecycle	Request or perform permitted transitions	Approve reservation, sold and exception transitions
Outputs	Prepare drafts from approved shortlists	Approve release and review output history
Administration	No platform-wide configuration	Manage users, role policies, fields, redaction and controlled lists

A few everyday journeys

Journey	What AI does	What the user does
New supplier update	Reads, extracts, normalizes, matches and checks for conflicts	Reviews exceptions and approves

New customer requirement	Structures the requirement and proposes matching capacity	Confirms the requirement and shortlist
LOI received	Recognizes the document and suggests the lifecycle change	Confirms the protected transition
Customer proposal	Drafts, redacts, formats and versions the output	Reviews, approves and downloads

How TEDRA operates

TEDRA is ThirdEye Data's intelligent agent and automated research analyst — the same assistant introduced in Section 1. It works in four modes, depending on what's needed.



CONTROL

Fewer clicks doesn't mean less control

AI can prepare and navigate, but a person still has to confirm anything that materially changes inventory, moves a reservation to sold, resolves a conflict, or releases a customer output.

From source to outcome, end to end



SECTION 6 Functional and Technical Architecture


How the pieces fit together

Experience layer	Operational services	Control and intelligence
Sales workspace, Manager	Inventory and schedules,	AI extraction and



workspace, TEDRA assistant, guided intake, search and dashboards	operators and contacts, customers and deals, activity timeline, shortlists and outputs	normalization, duplicate / conflict rules, confidence and evidence, role-based access and approvals, audit and notifications
--	--	--

The technical stack, layer by layer



NOTE

Architecture note

The final cloud provider, identity provider, AI model, mailbox connector and deployment setup will be confirmed during the Week 1–2 design baseline. At this proposal stage, the solution stays cloud-neutral.

SECTION 7

Data Migration and Excel Retirement

Excel stops being the internal system of record only after the data has been reconciled, users have validated it, and a controlled cutover has happened. It stays a perfectly acceptable format for the outside world, since operators will keep sharing spreadsheets with SPYRE.

OUTCOME

Excel is retired as the internal system

Sales and Operations work inside a HubSpot-style capacity workspace. Excel stays supported for external uploads, migration and exports — but the approved platform data becomes the trusted operational record.

How we'll get there

1. Profile the master workbook, supplier contacts and every custom capacity sheet.
2. Agree the field dictionary, controlled values, unique-ID rules and lifecycle definitions.
3. Separate master inventory from opportunity-specific copies, and reconcile anything that doesn't match.
4. Normalize dates, locations, units, prices and quarterly capacity schedules.
5. Load proposed records into a staging area, with duplicate and conflict reports.
6. Get SPYRE's sign-off on the migration baseline.
7. Run a time-boxed parallel period to reconcile any differences, then make the platform the source of truth.
8. Keep the workbook as an archived source / export file, but stop maintaining it day to day.

Excel after cutover — what's still allowed


Use	Allowed after cutover?	How it's controlled
Operator-provided spreadsheet	Yes	Uploaded to staging; AI mapping and human approval before it's committed
Bulk migration or approved correction	Yes	Manager-controlled batch, with validation and audit
Customer / report export	Yes	Generated straight from the operational database
Internal master workbook maintenance	No	The platform is now the system of record
A separate customer inventory tab	No	Use a deal shortlist linked to the

		canonical inventory instead
GATE	Cutover condition Excel is only retired once migrated totals, record samples, opportunity links, permissions, lifecycle behavior and generated outputs all pass the agreed UAT baseline.	

SECTION 8

Delivery Plan and 12-Week Timeline

This is one Phase 1 production build, delivered in controlled, working increments — not a string of separate proof-of-concept exercises.



Weeks	Workstream	Key activities	What proves it's done
1-2	Baseline and design	Confirm workflow, field dictionary, lifecycle definitions, UX baseline, architecture, migration rules and UAT plan	Approved Phase 1 baseline and prioritized backlog
2-5	Operational core	Users/roles, operators/contacts, inventory, schedules, deals, activities, lifecycle, audit and	Core journeys demonstrated with

		initial dashboards	SPYRE's own sample data
3-7	Efficient, AI-assisted intake	Excel/manual intake, email/PDF/transcript extraction, staging, mapping, confidence, duplicate/conflict rules, source traceability	Representative intake evaluation passes agreed thresholds
5-8	Query, lifecycle and outputs	Grounded natural-language query, filters, saved shortlists, protected transitions, redaction, Excel/PDF/PowerPoint and versioning	Approved search, lifecycle and output workflows
7-9	Migration and reconciliation	Data cleansing, unique IDs, custom-list reconciliation, trial migration and business validation	Approved migration reconciliation report
9-10	Release hardening	Security, performance, monitoring, backup, end-to-end, negative-role and regression testing	Production release candidate
11	Business UAT and training	Role-based UAT, priority fixes, user guides, admin and user training	Business UAT sign-off
12	Production deployment and cutover	Final migration, production deployment, go-live, Excel retirement and start of hypercare	Production acceptance and handover baseline

Milestone gates

Gate	What gets approved
G1 — Scope baseline	Field dictionary, journeys, roles, lifecycle and acceptance measures
G2 — Core workflow	That the operational core reflects the real day-to-day process
G3 — AI and output validation	Intake/query quality, plus output and redaction behavior
G4 — Release readiness	Migration, security, performance, documentation and open-defect position
G5 — Production acceptance	Cutover, operational ownership and hypercare entry
NOTE	Timeline assumption This timeline assumes prompt access to subject-matter experts, representative source samples, output templates and the required technical environments. Delays

	here may push the schedule.
--	-----------------------------

SECTION 9

Team, Governance and Communication

The delivery team

Role	Indicative allocation	What they own
Project Manager / Business Analyst	1.0	Scope, process, backlog, decisions, plan, risks, UAT and stakeholder coordination
Solution Architect	0.3–0.5	Architecture, security, integration boundaries and technical governance
Full-Stack Engineers	2.0–3.0	Web experience, APIs, workflow services, data model and output services
Data / AI Engineers	1.0–1.5	Ingestion, extraction, normalization, matching, grounded query and evaluation
QA Engineer	1.0 from Week 3	Functional, integration, security, regression, performance and UAT support
UI/UX Designer	0.3 during design/build	Low-click journeys, everyday usability and SPYRE's visual style
DevOps / Security Engineer	0.3–0.5	Environments, CI/CD, security controls, monitoring, backup and deployment
SPYRE Business SME(s)	As required	Rules, data validation, prioritization, output review and acceptance

How we'll stay aligned

- A weekly delivery review covering progress, working software, decisions, dependencies and risks.
- Focused process and data sessions in Weeks 1–2, and whenever exceptions come up.
- Fortnightly demos of the working increment, starting Week 4.
- One shared log for decisions, dependencies, risks and changes — each with a named owner and due date.
- Formal sign-off at every milestone gate; any scope change is assessed for effort and timeline impact before it's accepted.

Unauthorized Distribution or Publication is Strictly Prohibited.

© 2026 Third Eye Data Analytics Services India Pvt. Ltd (India) | Third Eye Consulting Services and Solutions LLC. (USA)

SECTION 10

QA, UAT and Acceptance

Acceptance is judged on real, end-to-end business evidence — not on isolated claims about the AI model. Extraction and query quality will be measured against a jointly agreed representative test set, with human review and exception handling built in.

What counts as “done”, area by area

Area	Acceptance evidence
Efficient intake	Agreed source types are processed; extracted values, confidence, conflicts and source links are all visible; unapproved changes can't commit.
Data structure	Approved Phase 1 fields and relationships are represented; migrated records have unique IDs and traceable sources.
Data query	Agreed natural-language scenarios reconcile to equivalent structured filters; unsupported facts are never fabricated.
Inventory management	Ownership, freshness, lifecycle and deal usage behave as approved; conflicting active reservations are blocked.
Customer output	An approved shortlist produces Excel, PDF and PowerPoint from the same snapshot, redaction rules and version history.
Roles and audit	Sales and Manager permissions pass both positive and negative tests; protected actions and outputs are traceable to a person.
Migration	Approved record counts, totals and sample records reconcile to the signed migration baseline.
Usability	Named business users complete core journeys at normal browser zoom, without technical help, after training.
Performance	Agreed concurrent-user and response-time targets pass in the chosen production-sized environment.

How we treat defects

- No open critical or high-severity defect affecting data integrity, access control, reservation protection or output confidentiality at go-live.
- Lower-severity items can only be accepted with a documented workaround, an owner and a target release date.

- AI exceptions must fail safely into a review or clarification step — never silently invent an unsupported fact.

SECTION 11

Security, Deployment and Operations

This proposal stays deployment-neutral until SPYRE confirms its preferred cloud or managed environment, identity provider, mailbox platform, data-residency needs and security standards. Production sizing and ongoing infrastructure costs will be finalized alongside the commercial proposal.

How we protect the platform

Control area	Phase 1 approach
Identity and access	Role-based access for Sales, Manager/Operations and administration; final authentication method confirmed during design.
Data protection	Encryption in transit and at rest; restricted document access; secrets kept outside source code.
Auditability	Traceable inventory changes, approvals, lifecycle transitions, source access and generated outputs.
AI controls	Grounded retrieval, visible source evidence, confidence scoring, clarification, and human approval for protected actions.
Application security	Input validation, malware/file checks where relevant, least privilege, dependency scanning and secure CI/CD.
Monitoring	Application, job, error, security and AI-exception monitoring, with alerting and clear ownership.
Backup and recovery	Scheduled backup, tested restoration, and documented recovery objectives for the chosen environment.
Environment separation	Separate development/test and production controls, data and deployment permissions.

Getting ready to run it day to day

- Health checks, structured logs, error monitoring and a clear support escalation path.
- Backup, restoration and rollback procedures for deployments.
- An administrator guide, a user guide, a runbook and a known-issue register.
- A hypercare period to review adoption, data exceptions, output usage and priority enhancements.

SECTION 12

Deliverables, Training and Handover

What SPYRE receives

Category	Deliverables
Product	The responsive SPYRE Capacity Operations Platform, covering the approved Phase 1 scope
AI	TEDRA workflows for extraction, normalization, matching, grounded query, recommendations and controlled generation
Data	Approved schema, data dictionary, mapping rules, migration utilities, reconciliation report and archived source baseline
Outputs	SPYRE-approved Excel, PDF and PowerPoint customer-output templates, with redaction and versioning
Technical	Solution architecture, API/data design, deployment configuration, monitoring, backup and security documentation
Quality	Test plan, automated and manual test evidence, security checks, UAT pack, defect register and release notes
Enablement	Administrator guide, non-technical user guide, recorded/live walkthrough and role-based training
Handover	Source/configuration handover approach, runbook, environment inventory, support contacts and knowledge-transfer sessions

Try it yourself: the interactive prototype

This proposal comes with a working prototype that shows the intended experience, using representative records from SPYRE's own sample workbook.

[LINK](#)

Prototype walkthrough

A working prototype built from SPYRE's own sample data:

Open the SPYRE Capacity Operations Platform prototype ↗

1. Start the Guided Walkthrough from the top-right of the platform.
2. Process the sample source and review TEDRA's extracted fields, confidence score and recommendation.



3. Look at the connected operator, site, capacity schedule, opportunity and audit trail.
4. Run the Finland capacity query and review the evidence-backed matches.
5. Review the controlled lifecycle and how reservation protection works.
6. Generate the customer-output preview and try the redaction controls.
7. Open TEDRA from the bottom-right and try the suggested questions.

SECTION 13

Assumptions, Dependencies, Exclusions and Risks

What we're assuming

- SPYRE will name business and technical contacts who can make timely decisions.
- The shared workbook is a representative baseline; SPYRE will confirm the current master and opportunity-specific data for the real migration.
- SPYRE will provide representative email, PDF, spreadsheet, transcript/note and output examples, handled under the right confidentiality controls.
- Final lifecycle definitions, role permissions, redaction rules, staleness thresholds and customer-output templates get approved during Weeks 1–2.
- Email intake starts with controlled upload/forwarding; direct mailbox integration depends on the confirmed platform, permissions and security sign-off.
- AI-assisted extraction cuts down manual work, but doesn't remove human review for anything low-confidence, conflicting or sensitive.
- Infrastructure, model/API usage, licenses and ongoing support will be covered in the commercial submission.

What's out of scope for Phase 1

- Two-way Salesforce integration.
- CLM integration, or automated contract authoring/review.
- External customer portal access.
- Direct integrations with data-center operators.
- Public or unauthenticated inventory access.
- Autonomous reservations, deal execution, contract negotiation or customer release without configured human approval.
- External market-intelligence feeds or open-internet enrichment, unless separately approved and governed.

Key risks and how we'll manage them

Unauthorized Distribution or Publication is Strictly Prohibited.

© 2026 Third Eye Data Analytics Services India Pvt. Ltd (India) | Third Eye Consulting Services and Solutions LLC. (USA)

Risk	Impact if it happens	How we'll manage it
Inconsistent legacy records	Incorrect migration or duplicate capacity	Staged reconciliation, permanent IDs, duplicate/conflict reports and business sign-off
Limited representative source samples	AI extraction misses real-world variation	Early sample collection, an evaluation set, confidence thresholds and an exception workflow
Parallel Excel use after cutover	Two competing sources of truth	A time-boxed parallel period, a clear operating policy, access controls and adoption monitoring
Unclear lifecycle definitions	Inconsistent availability or reservation decisions	Week 1 definitions, a role matrix, transition rules and UAT scenarios
Sensitive operator information	Confidentiality breach in a customer output	Field-level redaction, role control, frozen versions and release approval
Over-automation expectations	An unsupported AI action, or reduced trust	Evidence, confidence scoring, clarification and confirmation before any material action

SECTION 14

Directional Roadmap

Phase 1 lays the inventory foundation, and should prove itself in production before SPYRE commits in detail to what comes next. The roadmap below is directional — it will be refined once we see how Phase 1 performs.



Principles guiding what comes after Phase 1



- Build the canonical inventory and activity model once, so later integrations can reuse it rather than rebuild it.
- Keep Salesforce and CLM as future integration systems — don't duplicate their job prematurely.
- Take on advanced capabilities selectively, based on real production adoption, data quality and demonstrated value.
- Avoid a second re-platforming exercise as SPYRE grows from a lean team into its planned scale.

SECTION 15

Commercials and Next Steps

COMMERCIALS

To be submitted after internal confirmation

The Phase 1 scope, timeline, resourcing plan, assumptions and acceptance baseline are presented here for business review. One-time implementation, infrastructure, recurring AI/cloud usage, support and payment milestones will follow ThirdEye Data's internal commercial approval.

Proposed next steps

1. SPYRE reviews the interactive prototype and this Phase 1 scope proposal.
2. Both teams hold one focused scope-confirmation session, covering lifecycle, roles, source channels, output templates and the production environment.
3. ThirdEye Data folds in the confirmed decisions and issues the commercial submission and final statement of work.
4. Both parties confirm the start date, named stakeholders and the Week 1–2 workshop schedule.

The decision we're asking for

CONFIRM

Decision requested

That this proposal reflects SPYRE's intended Phase 1 production outcome across Efficient Intake, Data Structure, Data Query, Inventory / Capacity Management and Customer Output — subject to the assumptions above and commercial confirmation.

Thank you

ThirdEye Data



AI Services, Solutions and Products