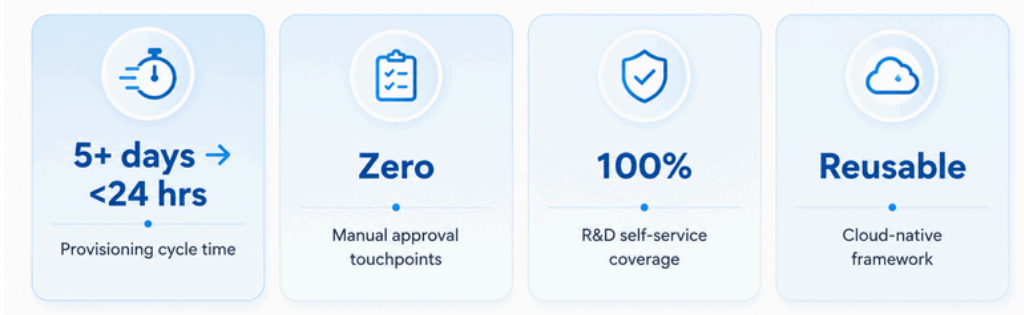


CASE STUDY

Compressing the R&D provisioning cycle from days to hours



An AI-enabled self-service provisioning platform for a global biopharma R&D organization



CLIENT SITUATION

The client is a global biopharmaceutical organization advancing the discovery, development, and delivery of innovative medicines and vaccines. Its R&D function operates a sizeable Databricks and Azure estate used by scientists worldwide for computationally intensive analyses across the drug discovery lifecycle. As experimentation volumes scaled, the existing process for standing up Databricks workspaces and supporting Azure resources had become a constraint on research velocity – manual, opaque, and dependent on a small group of platform engineers.

- Cycle time
- Manual effort
- Bottleneck
- Visibility

OUR APPROACH

Working alongside the client's platform engineering organization, the team designed a cloud-native, API-driven self-service portal for Databricks and associated Azure resources. The work was organized across four parallel workstreams.

- Guided intake
- AI-driven configuration
- Secure, automated provisioning
- Reusable platform framework

OUTCOMES

- **Speed:** Provisioning cycle time reduced from more than five days to under 24 hours.
- **Automation:** Approval and configuration handoffs eliminated from the standard request path.
- **Capacity:** Platform engineering capacity redirected from repetitive provisioning to higher-value work.
- **Quality:** Configuration errors and downstream rework materially reduced through AI-driven validation.
- **Researcher experience:** R&D users gained a direct, self-service channel with real-time visibility into request status.
- **Governance:** Role-based access, secret management, and audit logging maintained end-to-end across requests.

LOOKING AHEAD

With the core provisioning experience in production, the client is positioned to extend the same model across the wider R&D platform estate.

Near-term priorities include expanding the catalog of self-service resource types, deepening the AI agent's reasoning over platform policies, and applying the underlying framework to additional cloud services and research workloads – converting a single point of automation into a consistent self-service operating model for R&D.

