

Executive Guide

Built an AI App in Lovable? Here's What It Really Takes to Run It in Production

Over the last year, something remarkable has happened.



Prepared by

Vikram Katyani

Founder, IntelliMinds Digital

Executive Summary

Building AI has never been easier. Running it every day has never been harder. A practical guide to taking an AI prototype built in Lovable, Cursor or Bolt into a secure, monitored, production-ready business system.

Over the last year, something remarkable has happened.

Software development has become dramatically more accessible.

Tools such as Lovable, Cursor, Replit, Bolt and GitHub Copilot have made it possible for founders, consultants, operations teams and business specialists to build working applications in days rather than months. Tasks that once required a team of developers can now be achieved by a single person with a clear understanding of a business problem and the willingness to experiment.

Contents

1. Introduction
2. The Prototype Trap
3. Why AI Development Has Changed Forever
4. The Six Questions Every Organisation Should Ask Before Going Live
5. Production Isn't the Finish Line. It's the Beginning.
6. The Human-in-the-Loop Still Matters
7. The Three Mistakes We See Most Often
8. From AI Experiments to Business Capability
9. AI Is Becoming Infrastructure
10. A Simple Production Readiness Checklist
11. The Future Belongs to Organisations That Can Operate AI
12. Key takeaways
13. Where Should You Start?
14. Continue the Conversation

Introduction

Over the last year, something remarkable has happened.

Software development has become dramatically more accessible.

Tools such as Lovable, Cursor, Replit, Bolt and GitHub Copilot have made it possible for founders, consultants, operations teams and business specialists to build working applications in days rather than months. Tasks that once required a team of developers can now be achieved by a single person with a clear understanding of a business problem and the willingness to experiment.

From an innovation perspective, this is fantastic news.

Ideas can be tested faster. Business problems can be explored more cheaply. Organisations can validate concepts before committing significant investment.

I've spoken with several organisations recently that have used these tools to create genuinely impressive AI applications. Some had built internal knowledge assistants. Others had developed proposal-writing tools, customer support assistants or operational workflows powered by large language models.

Technically, many of these applications worked well. The demonstrations were impressive. Users were excited. The business could clearly see the potential.

Yet every conversation eventually reached exactly the same point.

"We're now trying to work out how to put this into production."

Interestingly, very few of the questions that followed were actually about AI.

They were about everything surrounding it.

- Where should it run?
- How do we secure it?
- What happens if OpenAI changes its pricing?
- Who monitors it?
- What happens if it stops working on a Monday morning?
- How do we update it safely?
- Who owns it once it's live?

These aren't development questions. They're operational questions.

And they're the questions that determine whether an AI application becomes an everyday business system—or simply another prototype that looked impressive in a demonstration.

The Prototype Trap

Building an AI application proves that an idea works. Running it in production proves that a business can rely on it.

Those two achievements are often confused.

A prototype has one purpose: Demonstrate possibility.

A production system has a very different purpose: Deliver consistent business value every single day.

The difference sounds subtle. In reality, it's enormous.

Consider an internal AI assistant built to answer questions from company documentation.

The prototype may work beautifully. It answers questions accurately. Staff are impressed. The demonstration goes well.

From a technical perspective, success.

But what happens next?

- Can fifty employees use it simultaneously?
- What happens when someone uploads confidential information?
- How are prompts logged?
- Where are API keys stored?
- Who receives an alert if the service becomes unavailable?
- How are software updates deployed?
- What happens if the underlying AI model changes behaviour?
- Can responses be audited?
- Can access be revoked immediately when someone leaves the organisation?

None of these questions make the prototype more intelligent. Yet every one of them determines whether the organisation can safely depend upon it.

This is where many AI initiatives begin to slow down.

Not because the technology isn't capable. Because production introduces an entirely different set of responsibilities.

Why AI Development Has Changed Forever

Historically, building software was the expensive part. Deployment was comparatively straightforward.

Today, AI has reversed that equation.

Generating working software has become dramatically easier. Keeping it reliable has become the difficult part.

Modern AI applications depend upon multiple external services.

Language models. Vector databases. Authentication providers. Cloud infrastructure. Monitoring platforms. Storage. Third-party APIs.

Every additional dependency introduces another operational consideration.

A prototype can tolerate occasional failure. A business system cannot.

This is one reason why organisations are increasingly asking different questions.

Twelve months ago, conversations usually started with: "Can AI do this?"

Today they often begin with: "How do we make sure this keeps working?"

That shift is significant. It marks the transition from AI experimentation to AI operations.

And, in my view, that's where the next phase of AI adoption will be won or lost.

The organisations that succeed won't necessarily be those building the cleverest AI applications. They'll be the organisations that make those applications dependable, secure, monitored and maintainable.

In other words, they'll treat AI applications like any other critical business system.

Because that's exactly what they become once people begin relying on them.

The Six Questions Every Organisation Should Ask Before Going Live

The encouraging news is that most AI prototypes don't need rebuilding. In many cases, the application itself is perfectly capable of delivering value.

What's missing is everything around it.

Production isn't about making the AI smarter. It's about making the entire system dependable.

Before any AI application becomes part of day-to-day business operations, I believe every organisation should be able to answer six straightforward questions.

During prototyping, deployment is often an afterthought. The application might run on a developer's laptop, inside a temporary cloud environment, or on infrastructure designed purely for experimentation.

That's perfectly reasonable while validating an idea. It's not appropriate once people begin relying on it.

A production environment should answer questions such as:

- Where is the application hosted?
- Which country is the data stored in?
- How is availability managed?
- What happens if a server fails?
- Can the environment scale as usage grows?
- How quickly can it be recovered following an outage?

For many organisations this naturally leads towards managed cloud platforms such as Microsoft Azure or Amazon Web Services, where infrastructure can be monitored, secured and maintained using established operational practices.

The important point isn't choosing Azure over AWS—or vice versa. The important point is recognising that hosting is part of the application, not something added afterwards.

Almost every AI application depends on credentials.

API keys. Database passwords. Authentication tokens. Storage accounts. Connection strings.

During development it's tempting to keep these inside configuration files or environment variables on a local machine. In production, that approach quickly becomes a security risk.

Credentials should be managed through dedicated secret management services with appropriate access controls, auditing and rotation policies.

Equally important is understanding the data flowing through the application. Many organisations are rightly asking:

- Are prompts stored?
- Is personal information retained?
- Who can access conversation history?
- Does this align with our GDPR obligations?
- Can sensitive information be removed if requested?

These questions are no longer just technical considerations. They're governance questions.

If an organisation can't confidently explain how its AI application handles sensitive information, it probably isn't ready for production.

One of the biggest differences between a prototype and a production system is visibility.

When you're demonstrating an application yourself, you immediately notice if something breaks. Production doesn't work like that.

Users don't necessarily report problems immediately. Sometimes they simply stop using the system.

That's why monitoring matters.

A production AI application should provide visibility into areas such as:

- application availability
- response times
- API failures
- infrastructure health
- usage trends
- unexpected cost increases
- failed integrations

Monitoring isn't just about detecting outages. It's about understanding whether the system continues to deliver value.

For example: An AI assistant may still be technically available while response quality quietly deteriorates because the underlying model has changed. Without monitoring, nobody notices until confidence has already been lost.

That's often far more damaging than a complete outage.

One characteristic of modern AI systems is that they're constantly evolving.

Models improve. Frameworks change. Security vulnerabilities are discovered. Dependencies receive updates. Your own business processes continue to evolve.

Production systems therefore require a repeatable deployment process.

Updates should be:

- tested before release
- deployed through controlled pipelines
- capable of being rolled back quickly if required
- documented
- monitored after deployment

Many organisations now use Continuous Integration and Continuous Deployment (CI/CD) pipelines to achieve this consistently. While these practices have existed for years in software engineering, they're becoming equally important for AI applications.

An AI system that cannot be updated safely will gradually become more expensive, less secure and increasingly difficult to maintain.

This is one of the most overlooked production risks.

Most modern AI applications rely on external providers.

OpenAI. Anthropic. Google. Microsoft.

These providers continually improve their services—which is usually beneficial. However, change also introduces uncertainty.

Pricing evolves. Models are retired. Capabilities improve. Rate limits change. APIs are updated.

A prototype built around today's assumptions may behave differently six months from now.

Production systems should therefore be designed with change in mind. Rather than tightly coupling every workflow to a single provider or model, organisations should think about flexibility.

Can models be changed without rewriting the entire application? Can costs be monitored? Can fallback strategies be introduced?

The objective isn't to predict every future change. It's to ensure the application can adapt when change inevitably arrives.

This is perhaps the most important question of all.

Many AI initiatives begin as innovation projects. A small team builds a prototype. Everyone celebrates. The application goes live. Then responsibility becomes unclear.

Who monitors it? Who applies security updates? Who responds if users encounter problems? Who approves model changes? Who reviews costs? Who decides what features are developed next?

Successful AI applications are never "finished." They're operational products.

Like any business-critical platform, they require ongoing ownership, governance and continuous improvement.

In many organisations, this responsibility naturally sits with internal technology teams. In others, it becomes part of a managed service delivered by an external partner.

The model itself matters less than having clear accountability.

Because when nobody owns the system... eventually nobody trusts it.

Production Isn't the Finish Line. It's the Beginning.

One misconception I encounter regularly is that "going live" marks the end of an AI project.

In reality, it's the point where the real work begins.

Think about any business system your organisation relies on today.

Your CRM. Your finance platform. Your learning management system. Your website.

None of them were built once and left untouched. They've evolved over years. They've received updates, security improvements, new integrations, performance enhancements and feature releases.

AI applications are no different. In fact, because they're built on rapidly evolving technologies, they often require even more attention.

The organisations seeing the greatest success with AI aren't those treating it as a one-off implementation. They're treating it as an operational capability that improves continuously over time.

The Human-in-the-Loop Still Matters

One of the biggest misconceptions surrounding AI is that the objective should always be complete automation. In practice, that's rarely the right answer.

The most successful AI systems I've seen don't replace people. They make people significantly better at what they already do.

Consider an AI-powered proposal assistant. Its role isn't necessarily to write the final bid. Its role might be to:

- prepare a strong first draft
- identify missing information
- suggest previous winning responses
- highlight risks
- improve consistency

The final decision still belongs to the bid manager.

Similarly, an AI assistant supporting customer service might draft responses, summarise conversations or retrieve relevant documentation. The support adviser remains responsible for the customer interaction.

This approach delivers two important benefits.

First, it improves trust. Users remain in control.

Second, it allows organisations to improve AI gradually rather than attempting complete automation from day one.

The most effective AI often works alongside people—not instead of them.

That's why we place so much emphasis on designing the entire workflow rather than simply introducing an AI model into an existing process.

AI should fit naturally into the way people work. Not force people to work around the AI.

The Three Mistakes We See Most Often

As organisations move from experimentation into production, the same patterns appear repeatedly. They're understandable. But they're also avoidable.

Mistake 1 — Treating the Prototype as the Final Product

A prototype answers one question: Can this work?

It doesn't answer: Can the business depend on this every day?

Skipping the operational work because the prototype already "works" is one of the quickest ways to lose confidence in an AI initiative.

Mistake 2 — Thinking AI Is a Technology Project

Technology is only one part of the solution. Successful AI adoption also depends on:

- business processes
- governance
- user adoption
- operational ownership
- change management
- security
- ongoing improvement

When organisations focus exclusively on the model, they often overlook everything that determines whether the solution succeeds commercially.

Mistake 3 — Assuming the Work Is Finished

Large language models continue to evolve. Business processes change. Employees discover better ways of working. New opportunities emerge.

The best AI systems improve continuously.

Treating deployment as the finish line usually means the application slowly loses relevance. Treating deployment as the beginning creates an environment where the solution becomes increasingly valuable over time.

From AI Experiments to Business Capability

One phrase I often hear is: "We're experimenting with AI."

Experimentation is important. It's how organisations learn.

But there comes a point where experimentation needs to evolve into capability.

That transition requires different thinking.

Instead of asking: Can AI do this?

The questions become:

- Should AI do this?
- Can people trust it?

- Can we support it?
- Can we scale it?
- Can we govern it?
- Does it improve the way our business operates?

Those are very different conversations.

They're also the conversations that separate organisations generating occasional AI demonstrations from those building sustainable competitive advantage.

AI Is Becoming Infrastructure

Twenty years ago, organisations debated whether they needed websites. Later, they debated cloud computing. Then mobile applications.

Today, AI is following a similar trajectory.

Over time, AI won't be viewed as a standalone innovation project. It will become another layer of business infrastructure.

Just like identity management. Just like cloud hosting. Just like cybersecurity.

That shift changes how organisations should think about investment.

You're no longer investing in "an AI project." You're investing in a long-term business capability.

Capabilities require planning. Capabilities require governance. Capabilities require ownership. Most importantly, capabilities require operational excellence.

That's why production matters. Not because it's technically impressive. Because it's the point where AI begins delivering measurable business value every single day.

A Simple Production Readiness Checklist

If you've recently built an AI application—or are close to deploying one—take five minutes and ask yourself the following questions.

Can you confidently answer "Yes" to each of them?

If several of those questions produced uncertainty, don't worry. You're in exactly the same position as many organisations adopting AI today.

Building the prototype was the easy part. Operationalising it is the next step.

The Future Belongs to Organisations That Can Operate AI

We're entering a fascinating stage of AI adoption.

The conversation is no longer centred on whether AI works. It clearly does.

The question is becoming: Can organisations operate AI reliably enough for people to depend on it every day?

That's a much more interesting challenge.

Because it isn't solved by choosing a different language model. Or writing another prompt. Or generating more code.

It's solved through good engineering. Good architecture. Good governance. Good operational processes. And, above all, understanding how AI fits into the wider business.

The organisations that succeed over the next few years won't necessarily be those experimenting with the newest models every week. They'll be the organisations that quietly build dependable AI capabilities into their everyday operations.

Their employees will trust the systems. Their customers will benefit from them. And AI will simply become another reliable part of how the business works.

That's the real destination.

Not a clever prototype. A trusted business system.

Key takeaways

- Building an AI prototype is no longer the difficult part.
- Production readiness is about security, governance, monitoring and ownership—not just AI models.
- AI applications should be treated like any other business-critical platform.
- Human oversight remains essential in most enterprise AI workflows.
- Organisations that invest in operational excellence will realise far greater long-term value from AI than those focused solely on experimentation.

Where Should You Start?

Every organisation's AI journey is different. Some already have working prototypes. Others are still identifying the right opportunities.

The common challenge is understanding where to focus first.

That's exactly why we created the AI Readiness Assessment .

In around five minutes, it provides:

- An overall AI Readiness Score
- A structured review across strategy, technology, data, governance and people
- The highest-value AI opportunities for your organisation
- Key implementation risks
- Practical recommendations for the next steps

Whether you're planning your first AI initiative or preparing to move a prototype into production, it provides a clear, structured starting point.

Continue the Conversation

If your organisation is exploring how to move AI beyond experimentation—and into secure, production-ready business systems—we'd be happy to talk.

At IntelliMinds Digital , we help organisations bridge the gap between proof of concept and production.

From AI strategy and readiness assessments to custom development , deployment and long-term managed AI operations , our focus is simple:

Building AI systems that people rely on every day.

Author Vikram Katyani — Founder, IntelliMinds Digital. Helping organisations move AI from experimentation into production through practical strategy, custom development and managed AI operations.

Continue the Conversation

IntelliMinds Digital helps organisations turn AI ambition into dependable business capability. If this guide has raised questions about your own AI programme, we'd welcome the conversation.

Next steps

- **Take the free AI Readiness Assessment** — a five-minute diagnostic with a personalised report.
- **Book a 30-minute discovery call** to discuss your priorities.
- **Explore the full library** of AI Executive Guides at intellimindsdigital.com/insights/



Read the latest version online

<https://intellimindsdigital.com/insights/built-an-ai-app-in-lovable-what-it-really-takes-to-run-it-in-production/>

Scan to open the live article, which is kept up to date as our thinking evolves.

IntelliMinds Digital

Founder: Vikram Katyani

vikram@intellimindsdigital.com

www.intellimindsdigital.com