



AWS Startup Resiliency Baseline

AWS Prescriptive Guidance



AWS Prescriptive Guidance: AWS Startup Resiliency Baseline

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Introduction	1
Scope of guidance	1
Shared responsibility model	2
Stage 1: Set objectives	3
Stage 2: Design and implement	5
Stage 3: Evaluate and test	7
Stage 4: Operate	8
Stage 5: Respond and learn	10
Next steps	11
Resources	12
AWS Well-Architected Framework	12
AWS Prescriptive Guidance	12
Other AWS resources	12
Contributors	13
Authoring	13
Reviewing	13
Technical writing	13
Document history	14

AWS Startup Resiliency Baseline

Amazon Web Services ([contributors](#))

March 2026 ([document history](#))

Every startup founder knows the tension: the team is racing to ship new features while customers expect every release to just work. You wake up to a notification about last night's outage, customer complaints are flowing in, and your roadmap is already packed with promised features. This is the daily reality for startups—balancing the pressure to innovate with the need to maintain stable services.

The instinct to prioritize new features is natural. Investors want growth, customers ask for more capabilities, and competitors aren't standing still. Yet experience shows that resilience problems can derail even the most innovative startups. A service disruption doesn't just mean lost revenue; it erodes the trust you've worked so hard to build with your customers. In fact, rock-solid reliability can become your silent differentiator in a crowded market.

Building resilience into your startup doesn't mean slowing down or massive infrastructure investments. It's about making smart choices early that prevent problems later. Think of it like building a house. It's much easier to put in a solid foundation at the start than to fix structural issues once everything is built. By weaving resilient practices into your initial architecture and operations, you create a competitive advantage. You're not just delivering a product; you're delivering a reliable experience that builds customer loyalty.

The AWS Startup Resiliency Baseline (AWS SRB) was created specifically for teams facing this challenge. It provides a practical path to build reliability into your systems while maintaining the speed that makes startups special. No heavy processes or enterprise-scale complexity, just practical patterns that grow with your business.

This guide describes how to implement the AWS SRB in a startup environment. It helps you identify what really matters for resilience, where to focus your limited resources, and how to build practices that scale as your company grows.

Scope of guidance

When exploring the practical steps of building startup resilience, it's important to understand where this guidance fits within your organization's broader growth trajectory. The AWS Startup

Resiliency Baseline is aligned with the [resilience lifecycle framework](#). It serves as an initial roadmap, designed to help you implement fundamental resilience measures with minimal development overhead as you deploy your first applications on AWS. Think of it as your starter kit for building reliable systems that can recover effectively from disruptions.

The [AWS Well-Architected Framework](#) and this resilience guidance complement each other to help you build secure and reliable infrastructure. This guide dives deeper into resilience practices while the AWS Well-Architected Framework provides a broader set of architectural best practices. You can assess your architecture against these practices using the [AWS Well-Architected Tool](#) in your AWS account.

Shared responsibility model

It's critical to clarify an important aspect of building resilient systems in the cloud: the partnership between AWS and your startup. Resilience in the cloud operates on a [shared responsibility model](#), with AWS and your team each playing distinct roles in ensuring system resilience.

AWS takes responsibility for the resilience of the underlying cloud infrastructure that powers their services. Think of this as the foundation upon which you build your applications. This includes maintaining the resilience of data centers, networking components, and the core AWS services you use in your architecture.

Your startup's responsibility focuses on the resilience of the systems you build and deploy on this foundation. The recommendations in this guide fall within your domain of responsibility. By implementing these practices thoughtfully, you can create applications that use reliable AWS infrastructure and incorporate robust recovery capabilities.

This model means you're never building resilience in isolation. You're building upon a proven, reliable foundation while focusing your efforts on the aspects that directly impact your customers and business operations. This guide explores how to make the most of this shared responsibility model, using reliable AWS infrastructure while implementing practical measures to make sure that your applications recover gracefully from disruptions.

Stage 1: Set objectives

Imagine that your team is in the final sprint before a major product launch. The new features are groundbreaking, and investor excitement is building. Then, during a routine deployment, your core service goes down. As customer complaints flood your email, two questions become painfully clear: How long can you afford to be offline? What data can you afford to lose?

Hoping that everything will work fine isn't a good strategy. You need a systematic way to decide where resilience matters most and where it does not. This is where a business impact analysis (BIA) becomes critical. It helps you make informed decisions about where to invest in resilience. A BIA helps you understand which parts of your system truly need rock-solid reliability and which can tolerate some flexibility.

Start by mapping out your core user journeys. For each one, ask yourself the following:

- What's the impact if this is disrupted?
- How quickly must we restore service?
- What data is critical to protect?

This isn't just a technical exercise; it helps you understand the business impact of reliability issues. Lost revenue is just the beginning. Consider how outages might erode customer trust, violate regulatory requirements, or give competitors an edge.

From this analysis, you'll derive two critical numbers for each user journey: recovery time objective (RTO) and recovery point objective (RPO). RTO defines how quickly you must restore that journey. RPO defines how much data loss your customers can tolerate. These business-driven targets then guide which components you choose and how you architect them, without over-engineering every part of the system.

The beauty of this approach is that it helps you focus your limited resources where they matter most. Perhaps your core transaction processing needs near-instant recovery and zero data loss, but your recommendation engine can tolerate longer downtime. By setting clear objectives, you create a framework that lets you continue rapid feature development while strategically building in resilience.

Document these objectives clearly. They're not just for your engineering team. When you're pitching to enterprise customers or going through technical due diligence with investors, this documentation demonstrates that you've thought critically about business continuity.

These targets evolve as your startup grows. The resilience needs of your first thousand users are different from those of your first enterprise client. Start with objectives that you can realistically meet today, but plan for how they'll tighten as you scale.

This guide explores how to implement resilience measures that meet these objectives. Setting these targets is your crucial first step. They're your compass for navigating the constant tension between innovation and stability, helping you build a system that dependably delivers value to your customers.

Stage 2: Design and implement

This section discusses turning your resilience objectives into reality. You've mapped out what matters most to your business, and now it's time to build it. How do you build in resilience without slowing down innovation?

Think of AWS managed services as a resilience shortcut. Instead of burning precious engineering hours maintaining infrastructure, use services that handle redundancy for you. For example, consider [Amazon Simple Storage Service \(Amazon S3\)](#). It automatically stores multiple copies of your data within an AWS Region for durability. It doesn't require extra code or late-night pager duties.

What about your core application components? Smart choices can multiply your team's impact. Consider a database that is the backbone of your service. Instead of building your own replication system, consider using [Amazon Aurora](#), which automatically handles failover. These features might cost more, but they shift your team's focus from maintaining infrastructure to solving business problems. This cost can be offset through faster feature delivery and avoided revenue loss during outages.

Sometimes startups need to build custom solutions. That's the nature of innovative startups. When you do, keep it simple but smart. Spread your application across multiple Availability Zones by using [Elastic Load Balancing](#) and [Amazon EC2 Auto Scaling groups](#). Set the Auto Scaling group minimum capacity to handle your baseline traffic even if one Availability Zone fails. This provides resilience against localized failures without complex architectural patterns. As your startup grows and customers demand higher resilience, you can evolve to more sophisticated approaches.

We recommend that you keep your production and development environments in separate AWS accounts. It's tempting to mix them when you're moving fast, but this boundary is your safety net. It prevents a well-meaning experiment from taking down your production service. Think of it as insurance for your "move fast and break things" development culture - break things in development, keep production stable.

If your application depends on third-party services, plan for their failures. When your payment processor has issues, can your system gracefully handle it? Build simple circuit breakers and fallback options. Maybe queue those transactions instead of showing error messages. Your customers will appreciate that you kept things working, even if not perfectly.

Document as you build, but keep it practical. Focus on recording the *why* behind key decisions, and create simple recovery playbooks. It's important to have these ready when incidents occur.

You're not building for perfect resilience; you're building for appropriate resilience. Every hour spent over-engineering resilience is an hour not spent on features that customers are asking for. Use AWS managed services as your foundation, add targeted resilience where it matters most, and create clear paths to scale up resilience as your business grows.

The next chapter discusses how to validate these design choices without burning through engineering resources. For startups, testing should be a reasonable lift and a smart investment in your application's resilience.

Stage 3: Evaluate and test

You've built a resilient foundation, but how do you know it actually works? Testing resilience might sound like a luxury when you're racing to prove product-market fit. However, there's a smart way to do this without derailing your feature development. This chapter describes lean, practical testing that fits a startup's pace.

Start with [AWS Resilience Hub](#), and think of it as an initial architecture assessment tool. It provides a helpful baseline review of an architecture's resilience foundations. It helps you evaluate if the basic infrastructure setup aligns with your recovery objectives by checking common configuration patterns and potential single points of failure. It can flag obvious gaps, such as missing multiple Availability Zone configurations or incomplete backup policies. Resilience Hub complements, but doesn't replace, thoughtful architecture reviews and targeted testing of critical paths.

To validate your documented recovery objectives, schedule monthly restore tests in [AWS Backup](#) in your development environment. Even though it requires engineering time, it might be cheaper than discovering your backups don't work during a real incident. Make it part of your regular development cycle, like running unit tests or code reviews. The goal isn't perfection; it's confidence that you can recover when you need to.

As your startup grows and customers start depending on you more heavily, gradually level up your testing game. When you're deploying new features, include basic resilience checks in your pipeline. Try simple chaos experiments by using [AWS Fault Injection Service](#). Start in your preproduction environment and start small. Test how your application handles a delayed API response in development before considering any production experiments. As your confidence grows, gradually expand these tests, but always validate in preproduction first. For a startup, breaking things in production is risky enough without doing it intentionally.

The key is balance. Every hour spent on testing is an hour not spent building new features. But a few strategic tests can prevent the kinds of outages that lose customer trust. Use automated tools provided by AWS to do the heavy lifting, and focus on the testing that matters most to your customers. This helps you build confidence in your application's resilience without slowing innovation.

The next chapter explores how to evolve this foundation as your startup scales.

Stage 4: Operate

You've built a resilient application and tested it. Now the daily reality is keeping it running. But in a startup, you can't watch all operations, and you shouldn't try to. The key is to stay alert to what matters without providing too many metrics or overburdening your team.

Start with the customer perspective. [Amazon CloudWatch Synthetics](#) canaries act as automated customers. They continuously test critical user journeys. Have them log in, simulate purchases using test accounts, or access key features, especially during your busiest hours. This helps you understand the customer experience and helps you catch issues before real users do. When a canary fails, you know immediately that something's wrong from a customer's perspective.

Build on this foundation with focused monitoring of the supporting infrastructure. What signals tell you there's trouble? [Amazon CloudWatch](#) helps you build dashboards that track these signs. Don't just monitor technical metrics; tie them to business impact. For example, high CPU usage matters, but that's because it might degrade the customer experience that you're tracking with canaries.

As a practical approach, map your monitoring to your customer journeys. If you're running a software as a service (SaaS) platform, you likely care about API response times, authentication success rates, and core feature availability. Set up alerts that tell you when these metrics drift. However, be selective. Every alert should demand action. If your team starts ignoring alerts because "it's probably nothing," you've set too many or are tracking the wrong metrics.

Route these alerts through tools that your team already uses. If your engineers live in a particular messaging application, send alerts there. The goal is quick awareness without creating a new process. When an alert fires, your team should know exactly what it means and what to do about it.

Keep your operational documentation lean and practical. Store runbooks with your code in version control, but remember that they're not novels. When something breaks, your team needs clear, actionable steps. Each alert should link to a corresponding runbook, and each runbook should answer three questions:

- What broke?
- Why does it matter?
- How do I fix it?

Implement a simple incident management process. You don't need complex frameworks, just clear definitions of what constitutes an incident and who to call when things escalate. Keep incident logs because they help you improve your application's resilience.

The key is finding the sweet spot between vigilance and overhead. Use AWS tools to automate what you can, focus on monitoring metrics that impact customers, and keep your processes light enough to evolve as you grow.

The next chapter explores how to foster a resilience mindset without sacrificing the speed and innovation that make startups special. At the end of the day, resilience is as much about people as it is about technology.

Stage 5: Respond and learn

When you're running a startup, complex post-mortem processes can slow down your team. This chapter explores how to learn from incidents without turning them into bureaucratic exercises.

Integrate incident learning into your existing rhythms. If your team already has regular meetings, use ten minutes to discuss recent incidents. Focus on practical questions, such as:

- Did the runbooks help?
- Did the alerts happen at the right time?
- Could AWS managed services have prevented this?

Stay focused on actions, not blame. In a startup, you're not building a perfect system; you're building one that gets better every time something goes wrong.

You can use your ticketing system to track incidents; there's no need for specialized tools. Create a simple template that includes the incident timeline, customer impact, recovery steps taken, and lessons learned. This can become institutional memory if you actively use it. Review past incidents during onboarding to bring new engineers up to speed. Reference them in architecture reviews when designing similar systems. Pull them into game days to create realistic failure scenarios based on actual events. The template captures what happened, and regular use transforms it into organizational learning.

As startups grow, patterns emerge. Maybe certain components fail more often, or perhaps particular types of changes cause problems. Use these patterns to guide resilience investments. If database failovers cause issues, consider improving your multiple Availability Zone setup. If third-party service disruptions are a common theme, consider improving circuit breakers.

The goal isn't to prevent every possible failure. That's impossible and would slow you down too much. The goal is to learn fast, adapt quickly, and keep the application reliable enough while you're growing rapidly. Use each incident as a chance to make your system a little more resilient, your team a little more knowledgeable, and your customers a little more confident in your service. For startups, speed and learning beat perfection. Create lightweight processes that help you learn from incidents without slowing innovation. The best resilience practices are the ones your team actually uses.

Next steps for startups

Remember that late-night deployment that made your heart race? Or that moment when a major customer asked about your resilience measures? A startup's resilience journey isn't about eliminating these moments—it's about facing them with confidence.

This guide describes a practical path to resilience, including: setting realistic recovery targets, using AWS managed services, implementing lean testing, monitoring what matters, and learning from incidents. This approach is powerful for startups because it grows with you. The same framework that helps your startup recover from issues today builds the foundation for serving enterprise customers tomorrow.

Think of resilience like a product roadmap; you don't build everything at once. Start with protecting your core services. Add monitoring for critical customer journeys. Implement basic testing that catches obvious issues. Document what you learn. Each step is manageable and practical, and most importantly, they don't stop you from shipping features.

As a startup grows, resilience needs evolve. Enterprise customers might ask tougher questions. Peak traffic can hit new highs. International expansion brings new challenges. Because you've built resilience into your foundation, you are ready to adapt.

The goal for startups is not perfect uptime or zero incidents. The goal is building a system reliable enough to earn customer trust while maintaining the speed that makes startups special. By following this approach, you can create something more valuable than creating a perfect system; you can create one that gets better every day by learning from each challenge while keeping pace with your startup's ambitions.

The resilience journey doesn't end. It evolves with every feature you ship, every customer you win, and every incident you handle. This guide helps you build a framework to help you make confident and practical steps forward. In the end, startup success isn't about choosing between resilience and speed. It's about finding smart ways to deliver both.

Resources

AWS Well-Architected Framework

- [Reliability Pillar](#)
 - [REL05-BP01 Implement graceful degradation to transform applicable hard dependencies into soft dependencies](#)
 - [REL06-BP01 Monitor all components for the workload](#)
 - [REL06-BP03 Send notifications \(Real-time processing and alarming\)](#)
 - [REL08-BP03 Integrate resiliency testing as part of your deployment](#)
- [Shared Responsibility Model for Resiliency](#)

AWS Prescriptive Guidance

- [Backup and recovery approaches on AWS](#)
- [Resilience lifecycle framework: A continuous approach to resilience improvement](#)

Other AWS resources

- [AWS Fault Isolation Boundaries](#) (AWS Whitepapers)
- [Establishing RPO and RTO Targets for Cloud Applications](#) (AWS blog post)
- [Overview of Deployment Options on AWS](#) (AWS Whitepapers)

Contributors

The following individuals contributed to this guide.

Authoring

- Dylan McCarroll, Associate Solutions Architect, AWS
- Igor Fil, Startup Solutions Architect, AWS
- Parth Shah, Senior Solutions Architect, AWS
- Vrajesh Prajapati, Solutions Architect, AWS

Reviewing

- Clark Richey, Principal Technologist, AWS
- Bruno Emer, Principal Technologist, AWS

Technical writing

- Lilly AbouHarb, Senior Technical Writer, AWS

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an [RSS feed](#).

Change	Description	Date
Significant updates	Updated recommendations and use of AWS services throughout.	March 6, 2026
Initial publication	—	January 24, 2024