

# Deploy in a Flexible Hosting Environment

Our services should be deployed on flexible infrastructure, where resources can be provisioned in real-time to meet spikes in traffic and user demand. Our digital services are crippled when we host them in data centers that market themselves as “cloud hosting” but require us to manage and maintain hardware directly. This outdated practice wastes time, weakens our disaster recovery plans, and results in significantly higher costs.

## Key Questions

1. Where is your service hosted?
2. What hardware does your service use to run?
3. What is the demand or usage pattern for your service?
4. What happens to your service when it experiences a surge in traffic or load?
5. How much capacity is available in your hosting environment?
6. How long does it take you to provision a new resource, like an application server?
7. How have you designed your service to scale based on demand?
8. How are you paying for your hosting infrastructure (e.g., by the minute, hourly, daily, monthly, fixed)?
9. Is your service hosted in multiple regions, availability zones, or data centers?
10. In the event of a catastrophic disaster to a datacenter, how long will it take to have the service operational?
11. What would be the impact of a prolonged downtime window?
12. What data redundancy do you have built into the system, and what would be the impact of a catastrophic data loss?

13. How often do you need to contact a person from your hosting provider to get resources or to fix an issue?

## Checklist

Resources are provisioned on demand

Resources scale based on real-time user demand

Resources are provisioned through an API

Resources are available in multiple regions

We only pay for resources we use

Static assets are served through a content delivery network

Application is hosted on commodity hardware

---

Revision #1

Created 26 February 2024 16:19:12 by Tom O'Malley

Updated 26 February 2024 16:19:46 by Tom O'Malley