Azure Resilience Review

Disaster Recovery | Scalability | Monitoring

Disaster Recovery planning for the cloud is different from Disaster Recovery planning for on-prem or co-lo. In the cloud, you are more at risk from individual services failing than from a whole data centre failing and you need an approach that is designed around that fact. NewOrbit has built systems in the Cloud for more than a decade, including several with five-nines uptime requirements.

Planning for Disaster Recovery for a cloud-hosted application can be daunting. Different cloud technologies have different SLAs and different abilities to deal with failures (see Azure Failover and Resilience 101 for an overview).

Different business and different systems have different resilience requirements; Some may require no downtime at all whereas others can cope with many hours of downtime. There are often also different contractual requirements in play for different systems. It is a well-known adage that for every 9 you add to the end of the SLA target, the cost goes up by an order of magnitude so it is very important to match your level of resilience with your business requirements to avoid spending more than you need to.

Our process goes through the following steps:

1. Understand your resilience requirements and your business constraints.
2. Understand your current setup.
3. Suggest a suitable setup based on your specific situation.
4. Help you to plan how to implement this.

Each step is outlined in more detail below, with example questions. We will ask you many more questions during the consultation. Do bear in mind that many of the questions are over the top for many scenarios; we will evaluate the appropriateness with you based on your specific context.

Requirements and Constraints

What level of resilience does your business really need?

Area	Example questions
Real world Impact	*What would be the real-world impact on your business, your customers and your users if the system was down for a minute? An hour? A day?* *What would be the real-world impact be on your business, your customers and your users if you lost one minute's worth of data? An hour's worth? A day's worth? All your data?*
RPO/RTO	*Do you have any externally or internally imposed Recovery Point Objective and/or Recovery Time Objectives?*
Standards compliance	*Do you need to meet external standards for recovery and resilience, i.e. FCA, ISO22301 etc? Would a data loss violate an individual's rights under the GDPR?*
Contracts	*What level of SLA are you expected to provide? What happens if you don't? Are there specific requirements in your contracts, for example a requirement to have a fail-over data centre or off-site backup etc?*
Constraints	*What time and financial constraints do you have?*

Current situation

What is your current situation?

Area	Example questions
What is your application architecture?	• What languages and frameworks do you use? • Monolithic or distributed? • Background jobs? • Do you use queues?
What does your infrastructure look like?	• How is the system hosted? • Where is data stored? • Do you have any fail-over in place?
Infrastructure and deployment	• How is your code deployed? Manual or CD? • Is your infrastructure setup scripted?
What monitoring do you have in place?	• Application trace logs? • An APM tool? • Uptime monitoring? • Dependency monitoring? • Response time, CPU and memory monitoring? • Any automated base lining, monitoring or alerts? • Do you monitor for the absence of activity? E.g., would you know if a background service stopped sending out registration emails?
What about your team?	• Do you have a team in place who can monitor "early signs"? • How quickly can your team mobilise if the system goes down?

Architecture and Plan

A plan to improve the resilience of a system of Azure usually requires activity in one or more of the following areas, depending on requirements;

- Deployment *and* code changes to facilitate hot fail over for different sub systems.
- Deployment changes to make it easier/faster to re-deploy in a secondary Data Centre.
- Code changes to make the system more resilient to failures in sub-systems.
- Hosting and possibly code changes to allow for "warm" failovers.
- Monitoring, especially early detection of impending failures.

One of the most important things to understand is that the biggest risk is not that an entire data centre fails but that individual services within a data centre fail. For example, SQL Azure could fail in the primary DC but the App Service might still be running. It is also important to have a sense of the likelihood of a particular sub system failing – something that goes beyond the official SLA numbers. A proper resilience plan for Azure will consider each of the individual services.

Note: The outcome from this review is a plan for how to improve the resilience of the system in accordance with the business requirements. Sometimes there is an external demand for a "disaster recovery plan" that has been "tested". These usually make a simplified assumption that the whole DC has failed and requires moving everything. It is our experience that these plans – and in particular the testing of them – takes several person-weeks of effort to put into place. If requested, NewOrbit can help you with this (at additional cost) but note that most of the effort will be from the people who actually operate the system on a daily basis.

Implement

If desired, NewOrbit can help you to implement parts or all of the plan;

- We can introduce your team to the selected tools and help you design the solution.
- We can second a Azure developer to your team to pair-program on the initial implementation of a particular technology.
- We can provide you with development and design capacity to help you build parts of the solution.
- We can be your Azure Cloud Solution Provider, providing you with Azure hosting and giving you access to Azure experts and support as needed.