Disaster Recovery Best Practices

OUTSCALE

About Outscale

- **Founded**: 2010
- **Executive Leadership**: Laurent Seror, Founder and Global CEO
  Robert Rosborough, US CEO
- **# of Employees**: 150
- **Global HQ**: Saint-Cloud, France
- **US HQ**: Ridgewood, NJ
- **R+D**: 15% of revenues reinvested in innovation / 30% of staff dedicated to R+D
- **Data Centers**: Ten globally, including four in the US. Outscale plans to add 12 additional data centers over the next three years
- **Key Partners**: Cisco, Intel, Nvidia, and NetApp

Regardless of the maturity of your business — startup to enterprise — a Disaster Recovery Plan (DRP) is a must. Why? Disasters of some level of magnitude must be anticipated. In order to ensure you're able to sustain your business in such a situation, you must have a disaster recovery (DR) plan in place. This will allow you to maintain business continuity and recover from the unexpected should a disaster impact operations

So, what exactly is a disaster? We define it as any event that can negatively impact your business, including hardware or software failure, network outage, power outage, physical impact to your building (such as a natural disaster like a tornado or hurricane, or even a flood or fire). In order to minimize the potential impact of a disaster, you must invest time and resources to plan and prepare for the unexpected. This plan must address people (policies, remote work, preparedness training, system access, etc.), process (what you will do, step-by-step, when disaster strikes) and systems (proactive monitoring, having redundant systems in place, etc.).

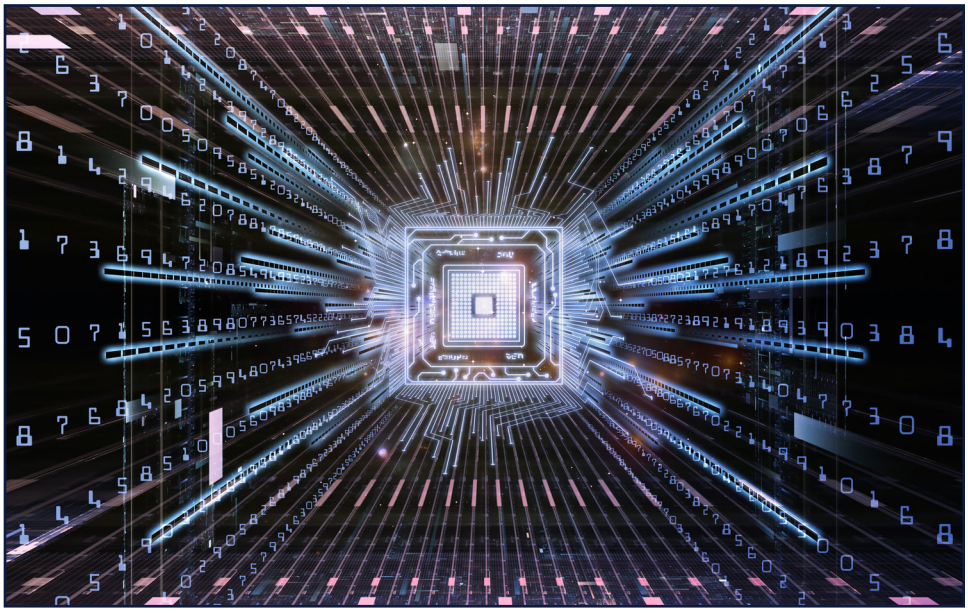

Outscale Disaster Recovery Services

Your investment in a disaster recovery plan and the associated redundancy required must take into account the potential impact on your business if, and when, disaster strikes. When it comes to your physical environment, we recommend that you duplicate your infrastructure and ensure you have available capacity available in advance.

With Outscale, you can scale up your infrastructure as you require it. You will have access to a scalable, reliable and highly secure environment, along with the ability to quickly modify and optimize resources, as required. Think of it as a DR insurance policy.

Key Considerations

Developing and implementing a DRP is not the equivalent of performing a simple backup. It is extremely important to weigh your risks (customer satisfaction, lost revenue, business continuity, lost data, etc.) as you define the DRP for your organization. How much data can you afford to lose, if any? How long can you afford to be down?

Key Definitions:

- **Downtime:** The time that it takes to bring files, applications, full servers, and a full site back into production after an outage and related costs associated with loss of productivity and revenue.
- **Capital Expenditures:** Cost of purchasing software, hardware, and implementing the solution.
- **Operational Expenses:** Cost for maintaining the solution, including time spent reviewing backup logs, ensuring successful completion of backup jobs, troubleshooting error messages, testing restores and running full DR tests.

- **RPO:** The Recovery Point Objective (RPO) is the age of files that must be recovered from backup storage for normal operations to resume if a computer, system or network goes down as a result of a hardware, program or communications failure. The RPO is expressed backwards in time from the instant at which the failure occurs, and can be specified in seconds, minutes, hours, or days.

- **RTO:** The Recovery Time Objective (RTO) is the maximum tolerable length of time that a computer, system, network or application can be down after a failure or disaster occurs.

Figure 1
All of the factors above are key considerations that will shape your technical solution.

Leverage the Outscale API to Manage Disaster Recovery through Load Balancers

For maximum flexibility, we provide an open API that allows our Cloud to interact with other Cloud providers. This is key to a seamless Disaster Recovery Plan or Multi-Cloud Strategy. Managing your infrastructure through code ensures the right behavior at the right moment. Based on metrics that you will have defined beforehand, you will be able to handle peaks of activity and detect when your infrastructure is unhealthy. If and when that happens, the Disaster Recovery Plan will be triggered and your system will keep running. That is where the defined RPO and RTO come into play.

Utilizing the tools at your company's disposal, you are able to plug this script with different APIs. The API behaves like a key, allowing you to access one or multiple cloud infrastructures (ex: US West or US East). Within the Cloud, the load balancers distribute the workload. For example, if something goes wrong in your infrastructure, the load balancer takes over and redirects everything to the Cloud.

Figure 2
In more technical terms, a load balancer distributes incoming network traffic between multiple Flexible Compute Unit (FCU) instances of the public Cloud or of a Virtual Private Cloud (VPC) to avoid overload and to improve the availability and reliability of your services.

Additional Reference Information

For additional information, please reference the links below, which are available on the Outscale Public Wiki:

APIs	docs.outscale.com
Load Balancers	wikioutscale.net/display/DOCU/About+Load+Balancers
Launching an Instance	wiki.outscale.net/display/DOCU/Launching+Instances
Adding/Removing Users	wikioutscale.net/display/DOCU/Adding+or+ Removing +a+User

Contact Us **www.contact-us@outscale.com**