

Be prepared when disaster strikes

Executive Summary ▶

Downtime is always unwelcome and costly but for some organisations it can prove disastrous. Gartner estimates that the average cost to the business of downtime for small or medium sized enterprises is approximately £27,000 per hour but for larger organisations, with IT as a crucial part of the core business, this figure can be much larger.


Fortunately there are now technology solutions available that can give some assurances against the impact of downtime such as clustering, replicated storage and virtualisation.

This whitepaper aims to uncover the techniques that can be used to test and pre-validate disaster recovery (DR) availability during normal daily operations to enable faster and more accurate decision-making



threeisquared
Infrastructure Insight Meets Innovation


www.threeisquared.com



2. Introduction ▶

With organisations of all types now dependent on IT continuity to keep operating, no matter what size of organisation, the need for faultless failover has never been greater.

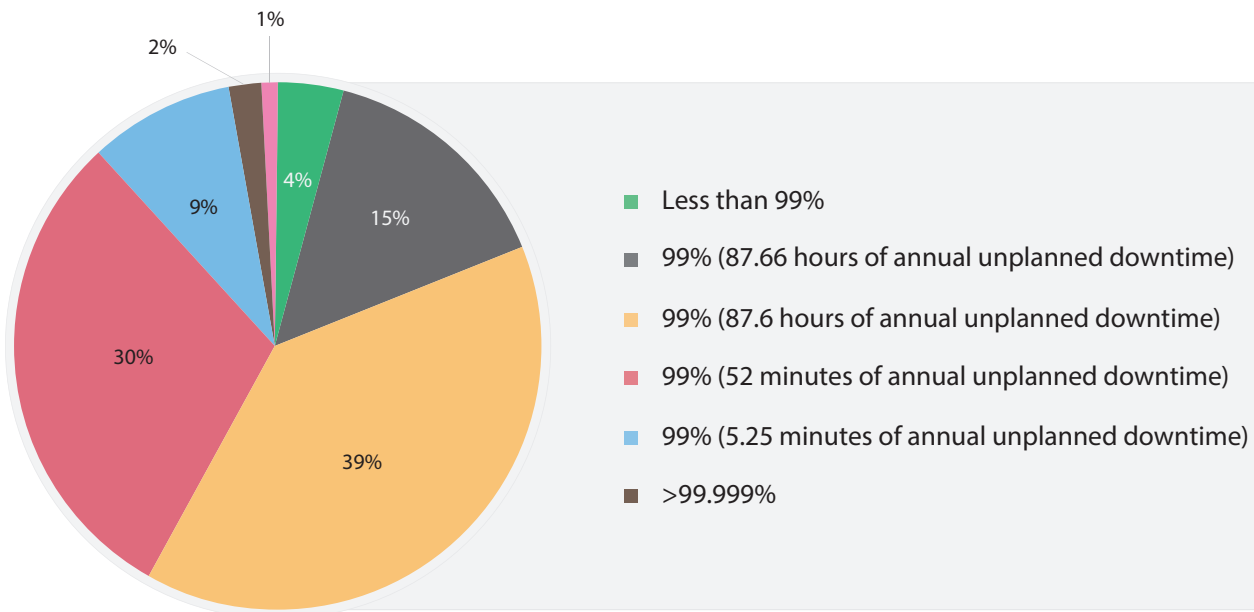
Downtime can be an organisation biggest IT expense. If an environment is down for a short while manual processes can be put into place to continue business processes but if the downtime is any longer than a few hours then it is likely that customer expectations will not be met.



In a competitive marketplace an organisations reputation can be it's most valuable asset and if an organisation can demonstrate that not only does it have disaster recovery in place for all critical applications and infrastructure but also that it has the ability to constantly test the critical service elements that are required for the application to perform faultlessly in a failover situation, customers will be immediately reassured.

“Gartner estimates that the average cost to the business of downtime for small or medium sized enterprises is approximately £27,000 per hour”

As can be seen from the pie chart below only 2% of companies surveyed by ITIC have anything approaching zero downtime annually. It is therefore vitally important that an organisation has confidence that disaster recovery works faultlessly first time.



Traditionally the ability to invoke DR whether on an application or an organisational level is only tested on a few occasions each year. These tests are generally pre-planned and heavily orchestrated Disaster Recovery tests, which require an enormous amount of resource to plan and implement. In general an organisation is only confident on the ability to invoke DR for a short period after a successful test.

Once change is introduced into the environment (for example cluster changes that may impact failover, or network routing changes that are not replicated correctly on the DR server) there is a real possibility that the ability to failover faultless is compromised.

3. A disaster! ▶

When Disaster strikes an organisation needs to have definitive and current information about its ability to failover either at an application or datacenter level. A lot of time is wasted during an incident trying to get enough information to decide whether to failover and resume service in an alternate datacenter or stay put and try to fix the problem in situ. This wasted time adds to the Mean Time To Recovery, impacts reputation, can cause regulatory issues and ultimately costs the business money.

4. How can threeisquared help? ▶

In our experience uncertainty around a number of elements can hamper the decision to move to DR or stay put and ride out the storm:

4.1 Failover host ▶

Problem:

There is often confusion over which server is the destination for an application environment. Shared or heavily consolidated environments with different failover targets often complicate this scenario.

Answer:

Depending on the underlying replication or failover technologies it is possible to automate the continual discovery of the failover server for each application environment. Once discovery for an environment is complete further checks can be completed.

For environments where automated discovery is not possible a manually identified relationship can be used.

4.2 Basic infrastructure checks ▶

Problem:

There is often uncertainty as to whether the production and DR infrastructure is the same or has the same configuration.

Answer:

There are a finite number of infrastructure differences that will cause application environment problems; most of these are relatively easy to check for in an automated fashion

4.3 Application environment checks ▶

Problem:

There is often uncertainty as to whether the application will actually work once failed over to the failover host

Answer:

Application requirements are generally known and checks can be automated to make sure these requirements are constantly checked.