



Resilience and redundancy

Mark Jones | October 2020



A while ago, we took some technology phrases and acronyms and made them clear. Like CDR (Committed Data Rate) and 95th percentile billing.

Now feels like a good time to explore some other phrases that can often cause confusion.

What's resilience?

What does redundancy mean when it comes to technology?

And what's the difference between the two, if there even is one?

Let's find out ...



What's resilience?

Resilience, sometimes resiliency, is the ability of components in your technology ecosystem to recover and bounce back from a fault, so that the overall service being relied upon is wholly dependable.

You'll often see networks described as resilient. This means that, in the event of an outage, your infrastructure is intelligent enough to not only detect the issue but also to speedily identify alternative live connections through which your operations can run.

Key to resiliency is a degree of network automation, as waiting for a human being to spot the problem and redirect traffic isn't really an option.


What's redundancy?

Redundancy is when you have a duplicate set of components running in parallel with the first, providing a dependable backup should your core system fail for any reason.

In short, it's redundant- not used- most of the time. Until it's critical.

You'll often hear about networks being redundant. An example could be a leased line, which gives you a fast connection with uncontended bandwidth; yours alone. Redundancy comes in when that leased line is backed up with another data connection, typically a broadband line.

This line runs quietly in the background until (and if) your leased line fails. Your business can still keep running by using the backup line until your main connection is restored.



Effectively, redundancy is what the resiliency within your system is looking for when it matters most.

Assuming you've built either into your infrastructure.

So, how do these terms relate to data centres?

In the context of a data centre, these terms take on additional, vitally important meaning.

Resilience still refers to the data centre's ability to switch to alternative systems when the primary one fails. But, in a data centre, that means a lot more intelligent software keeping multiple systems on high alert for issues and immediate solutions.

In terms of redundancy, network redundancy is unquestionably important in a data centre. But the notion of redundancy also comes in when you consider power to customers' servers or the smooth running of the environmental controls that have to be in place to keep rack upon rack of hardware running optimally.

Redundancy can take many forms. It could be two connectivity providers. It could be the use of two separate mains power sources, from two different providers. It could also mean a backup generator in case the main power supply fails or additional cooling circuits if the HVAC system runs into problems.

Ultimately, resilience and redundancy both need to be present to keep data centres functioning, businesses operating, people talking and customers buying.

What redundancy means to us

At Safe Hosts, we designed and built our data centre and manage it ourselves; you'd be hard pushed to find an organisation that knows its facility as well as we do.

We've built in multiple redundant elements into our data centre including:

- › Concurrently maintainable LV power grid with multiple redundant components
- › 2 generators and physically separate UPS
- › Carrier-diverse connectivity around our core network
- › N+1 redundancy on all cooling plant

And that's just a snapshot of the measures we have in place to keep your hardware always available.

Take a tour around our data centre; we'd love to welcome you and show you what Colocation means when you work with us.

**Simply get in touch
and we'll go from there.**