



Combell
Your host on the internet

Reduce your
downtime
to the
minimum
with a
multi-data
centre concept

White paper

Put your business-critical activities in good hands

If your income depends on the continuous availability of your servers, you should ask your hosting provider for a high availability solution. You may e.g. run a web store and lose money and customers every minute that your web server is down, or you may run an e-commerce platform for which every second counts. In such cases, you should go for a multi-data centre concept: two data centres that operate as one logical network, thereby optimising the availability of your website.

This system spread over two locations guarantees you the availability of your Internet service. If a failure occurs in the first data centre, your servers continue to operate in the second data centre. In this white paper, we will examine the technology that makes this possible and explain what you need to consider when choosing a multi-data centre solution.

A traditional multi-data centre concept based on the DNS

Traditionally, a multi-data centre concept operates using the DNS (Domain Name System), the network protocol that is used to convert computer names (domain names) into IP addresses. The domain name `www.example.com` is e.g. redirected to IP address 1 on a web server in the first data centre. If a failure occurs in this data centre, the system redirects the same domain name to IP address 2 on an identical web server in the second data centre. From there, all the customers' requests for the web server are redirected to the second data centre.

Since the DNS settings usually do not change very frequently, the DNS uses no cache: by doing so, a computer that requests the IP address behind a domain name also receives a TTL (time to live) value that determines how long that address is valid. This value varies from a few minutes to several days. If the computer requires the same address that is associated with that domain name in this time frame, it will retrieve that address through a local cache. This way, the DNS servers are not constantly overloaded with redundant requests.

If you want `www.example.com` to be redirected to the web server in the second data centre in the event that the first data centre is no longer available, you need to preset the TTL of `www.example.com` to e.g. five minutes. The people who visit your website will then have to request the IP address of your web server every five minutes. In the event that the first data centre is no longer available, you have your domain name redirect to IP address 2 of your identical web server in the second data centre. In theory, after five minutes, every DNS server worldwide has assigned this new IP address to the domain name `www.example.com`, which means that visitors arrive on the second web server.

DNS does not offer enough guarantees

A multi-data centre concept using the DNS will work well in most cases, but has some major downsides:

- **The approach is not reliable**

In theory, your web server should be available again after five minutes via the second data centre, but in practice some Internet providers ignore TTL values shorter than a couple of hours. This means that some visitors will be able to visit your website after five minutes via the second data centre, while others will be unable to do so. So, this is not a reliable multi-data centre solution.

- **You overload the DNS infrastructure**

With such a short TTL, you actually misuse the DNS protocol, because it was not designed for this sort of use. With a TTL of 5 minutes, many redundant requests are sent every day to the DNS servers, which is also the reason why some Internet providers do not approve this: because you overload the DNS infrastructure.

- **Five minutes downtime is too much**

A TTL shorter than five minutes is impossible for all kinds of reasons. So, with a multi-data centre concept with DNS, the best-case scenario is five minutes downtime, which is still too much for many applications.

Although many hosting providers still use DNS-based multi-data centre concepts, Combell Solutions stopped using such a concept in 2009. Instead, we went for a technology that makes it possible to switch faster between two data centres.

A multi-data centre based on BGP for a level of availability close to 100%

BGP (Border Gateway Protocol) is the protocol that is used to route network traffic between several providers: BGP routers communicate with each other to determine which IP address ranges are available through them. This makes it possible to automatically determine the fastest route available to the data centre for each visitor of your website, at any time.

Combined with a load balancer, the BGP protocol is a more logical solution to switch between two data centres than DNS. In each data centre, you need to run identical web servers behind a load balancer. If visitors want to reach your website via the IP address of your web server, they will enter the public IP address of the load balancer. This latter constantly checks which of the two web servers (that are set up redundantly) is available and redirects the network packets to the available server(s).

By setting up the load balancer and the BGP router correctly, there will always be a route to reach the server. If one of the data centres fails, the router will notice that the load balancer is no longer available via one of the routes and will choose the other route. Simultaneously, the load balancer stops routing the network traffic to the non-responsive web server to the underlying web server. As soon as the failure has been repaired, the load balancer will resume routing network traffic to the web server. This entire system is completely transparent

for visitors, who will never notice that the data centre failed. To them, it is as if the web server was working the whole time. By using this solution, your website continues to be 100% available for the public: downtime is almost nonexistent

What uptime guarantees are possible technologically?

Thanks to the BGP-based multi-data centre technology, you can achieve the highest level of availability for your servers. The kind of guarantees you can get depends on what happens in the data centres and the distance between both data centres.

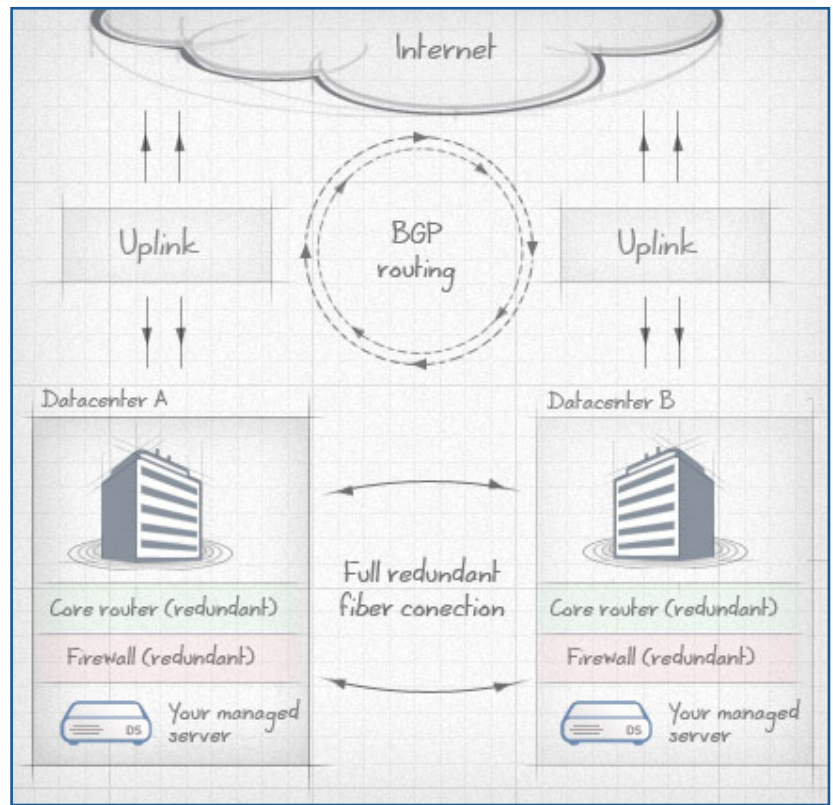
The benefit of Combell Solutions is that, for a relatively small additional fee, you can have your server infrastructure for one data centre converted into our multi-data centre infrastructure. And we will manage it for you, so that you can focus on your applications without having to worry about the uptime of the infrastructure.

- **What level of redundancy is offered by the data centre?**

A data centre can offer different levels of uptime, depending on its level of redundancy. If you are using just one data centre without taking any special precautions, your server will be down or unavailable in case of a failure. However, a smart interior architecture (multi-room) can help avoid this. The data centre can e.g. be composed of two rooms operating in a completely independent manner. This way, if a fire breaks out in one of the rooms, for instance, the other room will continue to operate and a server in this second room will take over. However, for a larger scale disaster that would affect the entire data centre, it is necessary to have a second data centre in a different location (a multi-data centre solution). Obviously, every data centre in a multi-data centre concept can also use a multi-room architecture.

- **Synchronous or asynchronous data mirroring?**

In the event of a disaster in one of the data centres, it is essential that your servers remain available, but also that you lose as little data as possible. This is why, for many applications, you need to synchronise the storage systems in each data centre (mirroring). Every time your web server in one of the data centres saves data to the storage system, the same data also have to be saved in the other data centre. This can be done in a synchronous or asynchronous manner.



When using synchronous mirroring, the data are saved simultaneously in the first data centre and sent to the second data centre. The application waits until the data in both data centres have been saved. The network connection to the second data centre causes the application to slow down, but data loss is excluded. When using asynchronous mirroring, the data are also sent to the second data centre, but the application does not wait until the operation is complete. So, the application runs faster, but you can always lose data in the event of a disaster in the first data centre.

The choice between both technologies is heavily dependent on the type of application that is to be spread over 2 data centres (how sensitive is the impact of the slowdown caused by synchronous replication) and the business requirements: how much data can you afford to lose and how fast should the 2nd data centre be available in case of a failure.

How much downtime can you afford?

With a BGP-based multi-data centre, you can achieve a level of availability close to 100% for your servers: in the event of a disaster in one of the data centres, your servers will run in the other data centre a few seconds later. However, what is possible technologically is not always possible financially. So, make sure to always start with a business impact analysis, in which you perform a risk analysis in order to determine what can go wrong and what the consequences are for your business. After that, you can determine the following criteria that the solution provided by your hosting provider must meet:

- **How long can your application be unavailable?**

The answer to this question is expressed in the RTO (recovery time objective), with which you indicate after how much time the application must be operational again in order not to compromise your business continuity. Be careful: make sure that you and your hosting provider have the same understanding of RTO. Imagine that a hosting provider with a DNS-based multi-data centre defines the RTO as the time that it takes to start a backup server in the second data centre. If that server is not available in that same time frame due to the unreliable nature of the DNS, you will not consider it operational: you will not agree with the fact that the server is running if the people who visit your website cannot actually access it.

- **How much data can you lose if a data centre fails?**

The answer to this question is expressed in the RPO (recovery point objective), with which you indicate the maximum time frame (before the failure) in which you can lose data without compromising your business continuity. Make sure you make unambiguous agreements with your hosting provider on when the clock starts ticking after a disaster.

Both the RTO and the RPO are expressed in time units. The achievable RTO and RPO depend on the technology that you use. Storage is one of the main determining factors. The following RTO and RPO can usually be achieved with synchronous and asynchronous mirroring:

	synchronous	asynchronous
RTO	a few seconds	15 - 90 minutes
RPO	0	15 - 30 minutes

These are theoretical values. In practice, you need to make certain technological choices to achieve them. You need to make a comparison between the costs incurred to achieve the RTO and RPO above with these choices and the loss that your business may suffer or the data you may lose during such a downtime period. Depending on your application, that loss may be easy to calculate. A web store, for instance, has figures on the average turnover per minute and can thus deduce what an acceptable level of downtime is.

About Combell Solutions

Combell has been the absolute market leader in hosting services for companies, IT integrators and software developers since 1999 and became the most reliable one-stop partner to host about any IT infrastructure, website or application. Combell Solutions is a division of Combell that handles advanced bespoke hosting projects.