

# *the* **Availability Digest**

[www.availabilitydigest.com](http://www.availabilitydigest.com)  
[@availabilitydig](https://twitter.com/availabilitydig)

## **Facebook Turns Off an Entire Data Center to Test Resiliency** November 2016

As a major social media network, Facebook is used by over a billion users to communicate with their friends and families. Facebook takes this responsibility very seriously. If one user can't be reached, other users may turn to alternative means for communication. The result could be a domino effect in which droves of users leave Facebook.



To prevent faults such as this, Facebook operates several data centers. Should a data center fail, its traffic is instantly routed to other data centers for processing. To ensure this continued resiliency, Facebook periodically takes down an entire data center to ensure that its load is transferred rapidly and correctly to other data centers.

### **Facebook Data Centers**

Facebook has data centers in Oregon, Iowa, North Carolina, and Sweden. It leases wholesale data center space in California and Virginia.

### **Facebook Can't Ever Crash**

It is imperative that Facebook remain online. A Facebook outage would be noticed immediately by millions of users. These users depend upon Facebook to communicate with friends and family, especially during disasters and other events. Facebook itself has 1.7 billion users. Its subsidiaries Messenger and Instagram have one billion and a half billion users, respectively.

Facebook's data centers must not only withstand power outages, but they must also be resilient to other failures that might take a data center down. In the unlikely event of a data center failure, traffic being handled by that data center must be immediately and reliably transferred to other surviving data centers.

In the early days of Facebook, an outage would imperil the company's survival. Today, with multiple data centers, that is no longer true. The primary concern of Facebook today is the impact an outage would have on its users

The data center resiliency required by Facebook requires extensive planning and practice runs. This is a very complex process since a data center processes tens of terabytes of data per second and consumes tens of megawatts of power driving thousands of servers supporting thousands of different software systems

### **The Germ of the Idea**

During Hurricane Sandy in 2012, several data centers were disrupted for days or weeks. The Facebook data centers survived, but some just barely. The storm hit its data centers in North Carolina and Virginia

with ferocity. Fortunately, these data centers were far enough away from the coast that they were able to weather the onslaught.

Facebook had built up a great deal of redundancy over the years. But following the storm, they asked themselves what would have happened if they lost a data center due to a storm such as this or due to some other event?

Facebook decided it was time to do some stress tests on its systems. It created a SWAT team called Project Storm to do just that. Facebook had procedures in place to analyze the causes of failures and to recover quickly, so the SWAT team was free to conduct whatever stress tests it thought appropriate.

The SWAT team decided that the proper resiliency test was to take down an entire data center. This might have such massive repercussions that the entire engineering team and major portions of the rest of the company were brought into the test.

## **Planning and Preparation**

A great deal of planning and preparation took place before the team could pull the plug. The SWAT team began with a series of mini-shutdown drills. The team had to make sure that the other data centers would pick up the load. Things didn't go well the first few times they tried, but preparedness ahead of time protected the Facebook users from being affected.

This preparation work paid off. After a series of drills, the mini-shutdowns became almost boring, according to one Facebook manager.

The SWAT team then worked up a 'massive-scale storm' drill that would simulate what would happen if some major outage should occur. They wanted to ensure that they had an environment that ensures that if one data center goes down, the unsupported traffic is immediately absorbed by other data centers.

## **Resiliency Testing**

After a successful execution of several mini-shutdown drills, the SWAT team finally pulled the plug on an entire data center. Not everything worked 100%, but the overall system persevered and all applications stayed up. No users were affected. The team put some improvements on the company's development roadmap.

A major lesson the team learned was that traffic balancing was really hard. The Facebook systems could shift traffic from a failed data center to other data centers quickly and reliably, but the traffic load on the other data centers varied from light to heavy. The Facebook engineers worked to build a better control system that now distributes the traffic fairly evenly across the surviving data centers.

## **Summary**

Facebook has added a whole new dimension to the idea of an infrastructure test. It periodically shuts down one of its data centers for a day to see if the safeguards it has put in place for such incidents perform in action. It makes this test on a regular basis to learn what improvements in the process can be made.

In addition to the load-balancing issue mentioned earlier, Facebook also found that it takes a long time to bring a data center back into operation. This is a very difficult and time-consuming operation. Facebook has invested significantly to make this process predictable and reliable.

## Acknowledgements

Information for this article was taken from the following sources:

Facebook Turned Off Entire Data Center to Test Resiliency, *Data Center Knowledge*; undated.

What Facebook Has Learned from Regularly Shutting Down Entire Data Centers, *Data Center Knowledge*; undated.

What Facebook Learned About Data Center Resiliency (and Why it Matters to You), *vXchnge*; undated.

Facebook and Natural Disasters: Extreme Scale Data Center Resiliency; *Enterprise Tech*; September 16, 2016.