

WHAT IS A CLOUD OUTAGE? CLOUD OUTAGES EXPLAINED



Organizations thrive on the availability of IT services powering their daily business operations. Most organizations leverage technology to deliver all kinds of services to end-users and customers, and would fail to do so without a functional IT infrastructure. With the proliferation of cloud computing, many organizations rely on third-party cloud vendors to operate and deliver the IT infrastructure services. While vendors promise adequate reliability of the service, such as Service Level Agreements (SLA) that guarantee availability for 99.999 percent of the time, outages of cloud services is a stark reality of modern enterprise IT industry. Even the largest cloud vendor and pioneer of the on-demand, subscription based networked IT infrastructure offering Amazon Web Services has had its fair share of cloud outages.

What is a Cloud Outage?

Cloud Outage simply refers to the duration when the cloud infrastructure service is unavailable for use. The unavailability may also refer to performance inadequacy of the service, as per the agreed SLA metrics. For instance, the incident during which an outage may have only partially impacted a data center may cause the vendor to perform the necessary maintenance and restoration measures. Until the service is fully restored as per the agreed SLA standards, it may be seen as a downtime for the end-user.

Common Causes of Cloud Outages

Cloud outages may result from a range of causes within and beyond the control of a cloud vendor. The following list briefly highlights the issues that cloud vendors take into consideration in order to ensure that the service always delivers on the SLAs with sufficient acceptability:

- **Power Outage:** One of the most common causes for the outage of a cloud service is the unavailability of the electric energy that powers the underlying datacenters. Cloud vendors inherently operate on a massive scale – a single datacenter may consume 10s to 100s of megawatts of power, for which they typically rely on the national grid or power plants independently operated by third parties. This makes the consistent availability of adequate electricity a challenge for datacenter companies, especially as rapid growth and scalable market demands require scalable power source, which is otherwise only available in limited quantity.
- **Cybersecurity:** Cyber-attacks such as the Distributed Denial of Service (DDoS) cause datacenters to overload with incoming traffic, preventing legitimate users from accessing the service via the same networking channels. Despite adequate protection systems in place, hackers tend to exploit hidden loopholes that either trigger protective mechanism isolating the services from legitimate users, leak data or shut the service altogether.
- **Human Error:** A single incorrect command can bring potentially bring the entire IT infrastructure service down, despite stringent protocols and systems in place to avoid such unforeseen issues. This can happen even with the largest of cloud vendors, as seen in 2017 when the global Internet suffered outage due to a [human error](#) at a AWS data center facility. While the systems were able to detect the anomalous behavior early enough, the infrastructure in many of the affected datacenters required a full restoration and restart.
- **Software and Technical Issues:** Cloud infrastructure comprises of a complex system of hardware and software technologies. Glitches and bugs are likely occurrences in enterprise-grade datacenters that power organizations of all sizes and verticals. These technical issues may be overlooked or remain under the radar until it translates into an actual service incident impacting end-users. When the solution to these issues is not apparent or applicable for immediate issue resolution, the service may remain in a state of outage.
- **Networking Issues:** Cloud vendors may partner with telecommunication service providers and government organizations operating the communication networks across long distances. Issues associated with networks beyond the organization, especially across borders may be well beyond the control of the service provider, especially in terms of resolving a connectivity issue. In this case, cloud vendors and customers rely on their telecommunication partners to ensure the service is restored. To address this limitation, most large-scale cloud vendors operate globally in multiple countries and are able to balance workloads dynamically across geographically disparate datacenters. This allows the company to continue delivering the service to end-users, even when resolving the networking issues is beyond their internal controls.
- **Maintenance:** Cloud vendors are responsible for the operations, maintenance and management of their IT infrastructure. End-users only pay for the services consumed, while vendors invest in service improvement on an ongoing basis. This includes both the scheduled and unscheduled maintenance and upgrades. The maintenance procedure may require service interruption, transfer of workloads across datacenters or general fixes that require a full system restart. During this period, the service may remain unavailable to end-users and is

regarded as a cloud outage.

Lessons for Customers

Cloud computing allows business organizations to invest resources on product development and innovation instead of keeping the infrastructure alive. The sheer scale of a modern cloud datacenter and the pressing internal and external threats make it virtually impossible to eliminate the possibility of a cloud outage. For this reason, customers of cloud computing must understand both the reality and unpredictability of the cloud outage. While it is possible to mitigate the issues leading to a cloud outage, the cloud system also suffers from the “unknown unknowns” – issues that the vendors don't know about what they don't know. This unpredictability has to be compensated by the acknowledgement that a cloud outage will happen, and corrective measures applied to reduce the impact. For some issues, it may be cost effective or less complex to suffer from an outage than to invest in mitigation efforts.

For instance, datacenter resources generate a deluge of service incident alerts pointing to possible technical issues in the future. Cloud vendors rely on advanced machine learning capabilities and automation technologies to identify the most impactful red flags and perform proactive maintenance on a small isolated section of the infrastructure related with the root cause. The alerting mechanism is set to a maximum threshold that may allow some risky alerts to go under the radar and only trigger action when the impact to the service is large enough. This tradeoff is optimized to reduce incident triggers that force frequent maintenance, while operating at an acceptable risk that may have a low probability of impact. In the real world however, this risk calculation may be inaccurate or complement other risks that lead to an eventual cloud outage.

Similar tradeoff should be considered by the customers of cloud services when investing in a cloud solution. If the impact of an outage for a certain duration is not acceptable for healthy business operations, it may be suitable to invest in high availability SLAs. Similarly, additional monitoring, visibility and control capabilities may be required on part of customers to ensure that a possible cloud outage is least impactful toward their business.