POWER OUTAGES AT PUBLIC CLOUD DATA CENTERS: HOW TO MITIGATE RISKS

Public cloud solutions allow businesses to access IT services without having to manage and operate the underlying infrastructure on-site. In effect, this statement also suggests that customers of public cloud have no control over infrastructure operations as they would for on-site deployments. The vendor is entirely responsible to deliver smooth operating environment as per the agreed Service Level Agreement (SLA), while customers can do little to ensure that downtime doesn't occur at periods of peak usage.

For instance, a massive cyclone that hit Loudoun County, Virginia caused power outages at the Equinix data center facilities and led to connectivity issues impacting several AWS customers earlier this year.

At the same time, businesses operating on-premise IT infrastructure left without power could have done little to bring the power back. Customers of public cloud infrastructure however, have several options to mitigate risks associated with power outages that impact public cloud data centers:

(This article is part of our Data Center Operations Guide. Use the right-hand menu to navigate.)

1. Understand the Risks of Public Cloud Outages

A few years ago, Gartner analysts suggested that IT outages pose a greater risk than security breaches in the cloud. Public cloud data centers are protected with several layers of sophisticated

security mechanisms to prevent security infringements and data leaks. Power outages occur more frequently and render customer data inaccessible during the service downtime. In some cases, the data is lost and irrecoverable.

Business organizations investing in cloud solutions should therefore understand the inherent risks of power outages in public cloud data centers to determine appropriate risk mitigation strategies. The risk may include geographical and zonal threats based on natural incidents such as cyclones and hurricanes. Organizations must also consider the true cost of losses they may incur due to IT service outages. These may include the actual business loss due to service disruption, impact on brand loyalty and reputation, lost revenue and business opportunities, as well as loss of workforce productivity.

2. Identify SLA Requirements and Know Your SLAs

Based on the potential cost of service downtime, organizations need to evaluate their SLA requirements for different workloads. Public cloud may not be the best option for security and availability sensitive mission-critical IT workloads and there may be specific SLA requirements to meet regulatory compliance standards. Yet, proving compliance may also not suffice if end-users and customers are not satisfied with the availability of data and services operating on outage-prone public cloud environments. The SLA figure should therefore be meaningful in terms of its impact on business operations, goals, revenue and profits, opportunities and other key business indicators.

Organizations also need to understand and know how the agreed SLA translates into service availability in the real-world. The impact of IT outage is significant during hours of peak service usage and for most organizations, it may not be possible to predict when a power outage to public cloud data center will occur.

It is important to identify critical processes and goals based on business impact. Organizations may monitor metrics that are most relevant and impactful to facilitate these processes and goals. If the SLA is designed to maintain desired performance standards of the most impactful metrics, the downtime will have minimal impact on customers and end-users.

3. Institute Redundancy and Multi-cloud Strategies

Cloud computing allows organizations to prevent the impact of power outages and downtime by introducing redundancy into their IT infrastructure strategies. Redundancy in cloud computing follows a simple approach: if one server instance fails or runs out of power, the workload can shift to another server instance. If the entire data center is impacted by the power outage, data replicated on data centers at distant geographic zones can take over to deliver the necessary IT service. The strength of redundancy overcomes the risk of power outages in public cloud data centers. This is further complemented by a multivendor cloud strategy, which involves pairing of services from multiple cloud providers. When a power outage impacts the primary cloud provider, the cloud service from a secondary vendor can serve as a failover solution to ensure business continuity. Additionally, a multi-cloud strategy also reduces the risk of vendor lock-in and allows organizations to optimize their public cloud investments by leveraging the best options available in the market in terms of features, support, reliability, price and other key factors that impact business decisions.

4. Test for Business Continuity and Disaster Recovery

Business continuity and disaster recovery plans can have a vastly different result during real disasters if the organization is not prepared to execute the plan during the real incidents. An effective business continuity and disaster recovery plan should be designed to identify the circumstances and address the limitations that may occur during a disaster. Execution of these risk mitigation plans is then a matter of following the documented checklist items and best practices. Without mock exercises, simulations and testing, organizations cannot understand the true circumstances and limitations that they may face during disaster incidents.

Some organizations may also be obliged by compliance regulations to conduct regular drills for business continuity during IT service outages. This approach ensures that the business continuity plan aligns with the evolving IT service availability needs of the organization in response to changing business and IT requirements.

5. Communicate with Stakeholders

Once the power outage has taken place and the organization executes their business continuity and disaster recovery plans, it is important to notify the end-users regarding the scope of impact and best practices for damage limitation. Proactive communication with the affected users may not eliminate the threat of a power outage but it will help lower the impact on customer trust and brand loyalty among users eagerly waiting for service uptime. Communication with business stakeholders, internal experts and external vendors may be required based on the documented disaster recovery plan designed to mitigate risks of power outages.

Power outages in public cloud data centers tend to occur without prior warning. It's important to not only understand the risks but also take the necessary steps to limit the damages. Organizations should follow a strategic approach in employing public cloud solutions and be well prepared for the threats that cause unannounced and unpredictable service downtime.