protiviti®
Face the Future with Confidence

Architecting your cloud infrastructure for failure (and resilience)

Companies aggressively moved to the public cloud in 2020, driven in part by the pandemic and the shift to a remote workforce. That pace is expected to accelerate in 2021, as more and more companies move applications and data to the cloud to achieve the benefits that have been touted for years, including cost savings, flexibility and on-demand scalability. It would be a challenge for an individual organization to match the ability of a cloud service provider in these areas. This is their business.

There is, however, one wrinkle to building on cloud services, and that is the shared responsibility model.

Cloud solutions cannot be simple "lift and shift," but must be intelligently designed to prevent, adapt and recover from an operational disruption in a manner that is aligned to and consistent with the importance of the service and data placed in the cloud.

Regardless of what cloud provider an organization is using, it is the responsibility of the organization to manage its own space within the cloud. The cloud provider will maintain the contracted digital space, but the organization must deploy and ensure diligence of use within the space. This starts with proper architecture within the cloud. And in addressing concerns of proper architecture, we are not initially contemplating the disruption event that could occur, but we assume the inevitability of a disruption event and we are architecting for the results of such an event. In other words, we are architecting for failure.

It is hard for most people to wrap their heads around cloud resiliency. "My data and my applications are somewhere 'in the cloud' and they are safe" is a hard concept to understand, mostly because the virtual, amorphous service doesn't have the same simplicity as a server in a firm's data center. However, in terms of resiliency and the processes that lead to disruption events, such as a security incident, cloud service providers are simply better than most. CSPs have architected their platforms with the intentional design for staying resilient regardless of component failure. And by using their platforms, organizations can embrace the same architectural patterns. By building on top of the provided cloud services, organizations can deploy in a manner architected for failure at a fraction

of the cost and administrative burden that achieving the same outcome would require with multiple physical data center environments.

Our supply chain will continue to grow more intertwined. Customers' expectations will continually advance. And regulatory bodies will continue to press firms on concepts of operational resilience, specifically designed to look at the ability of organizations to deliver goods and services within a stressed environment. This means ideas like "lift and shift" are outdated and may be detrimental to firms without appropriate risk analysis and consideration of re-architecting to meet the evolving threat and regulatory landscape.

We have noted below our five keys to architecting for failure:

Recognize prevention as a form of recovery

Every second matters in a recovery event, and minimizing the number of systems and files impacted is vital. This is particularly true when the event is a ransomware occurrence within the environment. Some prevention technologies can adjust based on observed behavior and can prevent the spread of ransomware after several files are encrypted and the behavior recognized. Further, well-segmented services can limit the scope of systems susceptible to an event and enable smaller contained areas to recover quickly. Think of such reduced segments as "blast zones" that are narrow in scope and purpose. This reduction of scope enables reduced complexity, and thus, quicker and more efficient recovery.

Risk-based controls

Implementation and operation of controls that can enable a near-zero Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are often both complicated and expensive. These controls should be implemented for only the most critical infrastructure and should be scaled back in applications to less critical systems, in alignment with the RTO and RPO objectives.

Replication is not the same as backup

Cloud service providers make systems highly available, more accessible to implement and more affordable than on-premise infrastructure. This **cloud service provider** capability for highly resilient systems in the face of a myriad of failures can provide **organizations with** a false sense of protection **within their enterprise** against ransomware or other forms of data corruption, such as an unexpected application error.

Since such an event impacts the integrity of data but not of system availability, backups of known good data remain critical, as the replicated data in a highly available system will also have corrupted data.

Whitelist, not blacklist

Critical systems should be under tight change control, which makes them well-positioned for the whitelisting of executables. This sometimes "extreme" form of system hardening is common on critical systems such as ATM and SCADA networks. By only allowing the software that has been explicitly permitted to execute, all other executables, whether benign or malicious, are prevented from running and impacting the system.

Identity is key

Much as system architecture can reduce the area of impact from a network and systems communication standpoint, identity architecture can be leveraged to reduce the extent of impact that any single identity can have on the environment. By reducing the impact of an event through both systems and identity architecture, an event on a given system is less likely to have a broad impact on the organization. This can enable a quicker, more contained recovery effort.

Cloud solutions cannot be simple "lift and shift," but must be intelligently designed to prevent, adapt and recover from an operational disruption in a manner that is aligned to and consistent with the importance of the service and data placed in the cloud.

Why it matters

Beyond technically architecting for failure, firms must recognize that not all services a firm moves to the cloud are of the same level of importance, nor do they require the same level of sophistication in their recovery effort. The business case for heightened disaster recovery is as integral in the process as the technical architecture, else firms will assuredly spend too much time, money and effort. In addition, with the complexity of running a dual cloud solution or CSP/on-premise solution, blanket solutions may enhance complexity, leading to heightened risk of failure.

While there is no set formula for prioritizing a business service, reviewing the tiering of applications and regulatory mandates on a business service, and integrating with a firm's overall resilience are clear paths to address the business side of architecting for failure. Disaster recovery ability and effort should easily align to regulatory obligations and application/service importance.

How Protiviti can help

Whether you have moved to the cloud, are in the process of moving, or are in the process of architecting a cloud solution, we can ensure the business and technical robustness of your solution. Protiviti's integrated solution combines our cloud, security and operational resilience practices to help you seamlessly move to the cloud while satisfying the needs of the business, IT and global regulatory obligations. Our services to help firms architect for resilience include:

- Cloud migration strategy
- Business Continuity and Disaster Recovery
- Firmwide Operational Resilience
- Security review, assessment and design

Contacts

Randy Armknecht
+1.630.281.0391
randy.armknecht@protiviti.com

Douglas Wilbert
+1.212.708.6399
douglas.wilbert@protiviti.com