Cloud Solutions Approach True Disaster Avoidance

Traditional Disaster Solutions

From your neighborhood small business to the global enterprise, today's companies depend heavily on a wide range of information technologies including file sharing, email, order and payment processing, customer interaction, inventory & logistics, coordination & collaboration, and countless custom applications. Technology has become such a central part of business operations that any downtime can lead to significant financial impact through lost productivity, lost opportunities, loss of time-sensitive data, and lack of accountability. As such, companies spend a growing amount of time, labor, and capital resources in an effort to replicate critical data and technologies so that in the event of a loss, failure, or natural disaster, the operation can be back up and running as quickly as possible. Traditional approaches to ensuring disaster recovery typically address several key vulnerabilities: data, hardware, and network.

Data Redundancy

Securing data is typically the first step in even the most basic disaster recovery program. This comes in many forms from simply making duplicate copies of data on external media like tapes, CDs to more complex coordinated back-up strategies that leverage a mix of software and hardware to replicate critical information automatically. However, an often recognized danger of data duplication and back-up is the vulnerability of the copied data and media. Surprisingly, most businesses store back-up data in the same physical location as the original data. So, while the back-up data can be helpful in the event of a simple hardware failure, a more devastating event like a fire can easily destroy both the original and duplicate data. A common solution to this is storing duplicate media offsite, either directly or by using a physical media storage services. Again, while this is a better solution, it still leads to more complicated operational procedures.

Hardware Redundancy

The next practice of preparing to deal with disasters is to guard against hardware failure through the use of redundant hardware. In the case of a single information resource such as a server, this is normally done with redundant components inside the server case such as redundant hard drives either using a simple mirroring approach or more complex RAID array using several drives. This way if one piece of hardware fails, the server can usually continue functioning while IT staff repair the issue with new hardware. The next level of redundancy involves actually replicating the servers in pairs or groups, which is commonly referred to as clustering. Again, the servers in these clusters will likely include individual component redundancy such as hard drives. So, with multiple redundant machines running on multiple redundant components, it is highly unlikely that multiple hardware failures would occur simultaneously, making this a very stable approach. However, like the example of data redundancy, most companies that use hardware redundancy to aid in disaster recovery generally keep such hardware resources in the same physical location. Like before, a catastrophic event like a fire or a flood will cripple operations, regardless of the hardware redundancy.

Network Redundancy

In order to ensure continual access to information resources, many organizations deploy redundant networks in addition to the data and hardware redundancies discussed previously. Primarily this is done with the access networks that organizations use to access centralized resources from multiple locations or via the Internet. In practice, redundant circuits are used to connect to the Internet so that should one go down or become congested there is a secondary path for users like customers to access the company's information services. For greatest security against downtime these redundant circuits should also be done through different carriers providing access services. With this approach operations are even resistant to downtime that might result from issues with one particular network carrier. However, like previously stated, these redundant networks still usually end up accessing resources in one single location, leading to a remaining risk of single on-site resources.

The Remaining Threat

It should be clear by now that there is one major risk with most approaches to dealing with potential disasters, and that is that in most cases *all the resources are in the same location*. Localized redundancies are crucial in preparing to deal with disasters, but creating geographical redundancy is next step in moving toward actual disaster avoidance.

Disaster Avoidance

It goes without saying that the best way to re-establish operation after a disaster is to avoid any interruption in operation to begin with, essentially avoiding the disaster. While disasters are actually unavoidable, enterprises and even small businesses can take advantage of a distributed network model born in the Cold War Era to create systems that are truly resilient to service interruption. In the 1960s the US Department of Defense with the aid of several leading universities developed a defense information network designed to withstand a nuclear attack on the United States. As many readers of this document will know, this networking approach eventually evolved into today's Internet, a collection of many interconnected public and private networks all sharing transport, resources, and services. The primary innovation was in its transport mechanism via packets, but the distributed design was also key in ensuring continuity of defense systems during an attack. While the Internet has grown over the past few decades spurred on by the advantages of its packet switching transport approach, only recently has the decentralized design re-emerged with the idea of benefiting corporate applications.

Multi-Locations

To truly be resilient to disasters, localized events must be contained. To do this, information resources need to be replicated across several locations. The simplest approach to this could be for a company that has its information resources located in California to expand to a second mirrored facility in New York. This way if an earthquake on the West Coast were to interrupt operations, users could continue accessing the resources on the East Coast. There are also network advantages to such an approach, but it is also key to remember that such an expansion must employ all of the traditional redundancy methods previously examined to guard against service interruption. Ultimately, locations can be expanded to suit the geographies of business operations and users as well as considering the natural disaster risks of each location.

The limiting factor to a multi-geography approach is ultimately cost. Maintaining multiple data centers spread across the globe is certainly not possible for small business, but is also a struggle for even most of the largest global enterprises.

Cloud Services

The notion of services and applications that are ubiquitously available via the Internet has received much praise and adoption over the past few years under the banner of *software-as-a-service (SaaS)*. Access networks have finally reached a point in both geographical availability and bandwidth where business, organizations, and consumers can access resources on the Internet at speeds that only a few years ago were not even possible on a local network. As such, user expectations in terms of responsiveness and availability are easily met with some of the previous locally-based services that have begun to move into the cloud.

While there continues to be much debate over the suitability of SaaS, the general notion is there are certain functions and applications that can be better served to the user through a location agnostic approach by using the ubiquitous network of the Internet. These functions can also better served by pooling overhead and resources in a shared operation that is solely focused on service delivery and development, rather than using an internal support department whose competencies lie outside of the core of the business. More simply, the SaaS model allows companies to focus on their actual business while achieving similar or improved services at reduced cost from scale economies. It is the goal of such economies from resource pooling and scale to allow even the smallest businesses to enjoy the security of a multi-location approach in an effort to thwart data loss and service interruption from disasters.

Conclusion

In the examination of traditional and emerging approaches to resisting and recovering from disasters, it is clear that any strategy should give careful consideration to the basic redundancies for securing data and ensuring hardware and network operability. However, these critical practices are merely the first steps in working toward true disaster avoidance. Only a strategy including resources at multiple diverse geographies can provide a solution with operational continuation should a catastrophic event occur at one location.

Room for Discussion

At Puremail, we are committed to the advantages of a cloud-based delivery of software services. However, we recognize that there are many misnomers, challenges, and deficiencies in the current SaaS model that need to be addressed. In a forthcoming publication, we will further examine the merits of Saas and cloud-based services as well as a discussion of these challenges, including: the resilience of hosted models, backup vs. mirroring, and the inclusion of legacy data.