the Availability Digest

www.availabilitydigest.com
@availabilitydig

Active/Active Data Centers

December 2013

The life we lead today is powered more and more by technology. In addition to the micro-technologies we find in our microwaves and automobiles, we depend extensively on the macro-technologies of massive data centers. These data centers help us to manage our money, make investments, control our health, shop online, communicate with family and friends, play the lotteries, and complete our education.

IT services have now become so important to us that we expect them to be available around the clock, on a 24x7x365 basis. A decade ago, if a service was down for a few hours, that was often expected. This amount of downtime was important to provide periodic planned windows for upgrades and maintenance. However, the amount of acceptable downtime today has moved from hours to minutes and even to seconds for many applications. For these applications, planned downtime is even unacceptable. Downtime can cost an organization significant amounts of money. Perhaps even worse, in this day of instant communication with the likes of Twitter, a company's reputation can be seriously damaged if it experiences downtime that negatively impacts its customers.

The requirement for extreme availability has led to the deployment by many organizations of active/active data centers, in which no component failure or even the loss of an entire data center will affect the services that are being provided. These new data centers have been exceptionally reliable.

Forrester Research and the Disaster Recovery Journal jointly conducted a survey of top management to determine their views on today's availability requirements.[1] The survey was commissioned by EMC Corporation. The findings of this survey are summarized below.

Disaster Recovery, High Availability, and Continuous Availability

Today's service availability expectations have created a hierarchy of availability solutions:

Disaster Recovery (DR) means that the IT service can be restored eventually following an outage. Recovery is a manual operation. Typically, a periodic copy of the database is kept offsite, and backup systems are provided. To recover from an outage, the database is loaded onto a backup system, the applications are loaded, the system is tested, and it can then be put into service. Recovery can take hours to days. The possibility of a failover fault, in which the backup system cannot be put into service, is significant.

High Availability (HA) means that recovery can be achieved in a time frame ranging from a few minutes to a few hours. An availability of five 9s (less than five hours per year of downtime) is generally considered high availability. Recovery is highly automated. High availability systems typically replicate the database to a standby system at a remote site. Thus, in the event of an outage, the database does not need to be

[1] How Organizations Are Improving Business Resiliency With Continuous Availability, *Forrester Research Report*; 2013.

loaded. The applications must be brought up, they must connect to the database, and the system must be tested before it can be put into service. Failover faults are also a worry for high-availability systems.

Continuous Availability means that recovery is so fast that no one notices that there was an outage, or at least they were not inconvenienced by it. Continuous availability is achieved via active/active systems in which a live copy of the application is running on two or more systems. The system databases are kept synchronized via data replication. A transaction can be sent to any system in the application network and be processed properly. If a system fails, all that needs to be done is to route all transactions to surviving systems. Failover faults are not a problem because the surviving systems are known to be working. Furthermore, planned downtime is eliminated because any one system can be taken offline for maintenance purposes without affecting service availability.

The Forrester Research Study

Every week there is news of another organization suffering a major outage (many of which are reported in our Never Again series). Human error, a botched upgrade, an unsuccessful recovery from a power failure, or a failover fault can bring a system to its knees, often for hours or days.

The Forrester/DRJ study found that 82% of organizations felt that the risk to IT systems was rising. Their top concerns were technology dependency (48%), business complexity (37%), and extreme weather (36%).

Technology Dependency

Technology for the data center is expanding at an ever increasing rate. Multicore servers are getting faster and greener, requiring less power, less cooling, and less floor space. Storage capacities are increasing. Deduplication technology is revolutionizing archive procedures. Networks are becoming faster and more reliable. The sophistication of operating systems and applications are improving.

This means that data centers are frequently updating their hardware and software. In order to meet their SLAs, these upgrades must be accomplished in ever-diminishing maintenance windows. Ultimately, the maintenance windows disappear, and upgrades must be made without degrading service.

As a consequence, any mistakes made during a carefully planned upgrade can lead to significant extended downtime. Since upgrades are occurring frequently, upgrade errors can destroy compliance with SLAs.

Business Complexity

There used to be a one-to-one relationship between a business process and the application that supported it. No more. Today, dozens of applications interoperate to implement a business process. Even worse, these applications are maintained by a variety of organizations, from internal IT to partners and cloud providers. If any of these applications should go down, the business process is impacted or even disabled.

In this environment, the failure of a seemingly non-critical application can disrupt an entire business process. Consequently, more and more applications are being deemed to be mission critical or business critical. Because their SLAs now require recovery in minutes or even seconds, disaster-recovery and high-availability strategies are no longer adequate. Continuous availability is required.

Extreme Weather and Natural Disasters

The organizations surveyed felt that extreme weather is the most common cause of a major business interruption, followed closely by power failures. Extreme weather like Superstorm Sandy has the ability to take out entire data centers. Furthermore, it can lead to cascading events like power failures, fuel shortages, network outages, and floods.

Continuous Availability

As the result of tightening SLAs and the increasing risks for outages, many organizations are moving to continuously available data-center architectures. In these organizations, disaster recovery and high availability look like they did years ago. Backup images are replicated to a remote site where standby systems are available, Recovery times are measured in hours or days, and failover faults are a continuing concern.

The new look for mission-critical and business-critical applications is continuous availability via active/active architectures.[2] In an active/active system, multiple geographically distributed systems actively run the same application. A transaction can be sent to any of the systems and be processed properly. If a system should fail, all that needs to be done is for its transactions to be routed to surviving systems. It is known that these systems are operating properly because they are currently processing transactions. Therefore, there is no concern about failover faults; and there is no management delay while it is decided whether it would be faster to recover the failed system or fail over to its backup.

In order for multiple systems to be processing transactions for a common application, they must all have the same view of the application database. Though this could be accomplished by having the systems connect to a common remote database, network latency makes this infeasible. Rather, each system has its own local copy of the database. The database copies are kept synchronized by data replication. Whenever a change is made to one copy of the database, it is immediately replicated to all other databases.

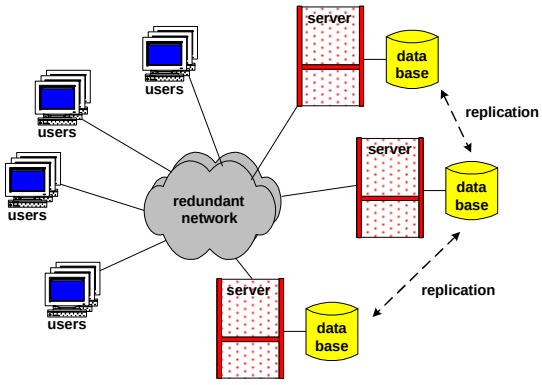

An Active/Active Network

Another advantage of active/active systems is that the systems do not have to be identical. They can even be from different manufacturers. In an active/passive environment, the production and backup systems typically have to be identical. This can be a problem to maintain and leads to many failover faults in these architectures.

Some applications cannot run in a distributed environment. Only one copy of the application can be active at any one time. Even in this case, an active/active architecture can be used to advantage. The application is run in only one system, but the changes that it is making to its database are replicated to the other system. Should the primary system fail, all further transactions are routed to the backup system, which will take over transaction processing immediately. It can be assured that the backup system is operating properly by sending it occasional test messages. This configuration is known as a "sizzling-hot standby."

An alternative to active/active systems is x86 virtualization. In this architecture, multiple virtual machines (VMs) run on a single physical host server. A VM is represented by an executable stored on common disk storage accessible by all physical host servers. To instantiate a VM, all that is required is to boot its executable (including its guest operating system and all applications) onto a physical host server.

VMs can be rapidly restarted should they fail and can be moved to surviving physical host machines should their physical host fail. VMs can even be moved across data centers if their executables have been replicated between data centers.

[2] What is Active/Active?, *Availability Digest*; October 2006.
http://www.availabilitydigest.com/public_articles/0101/what_is_active-active.pdf

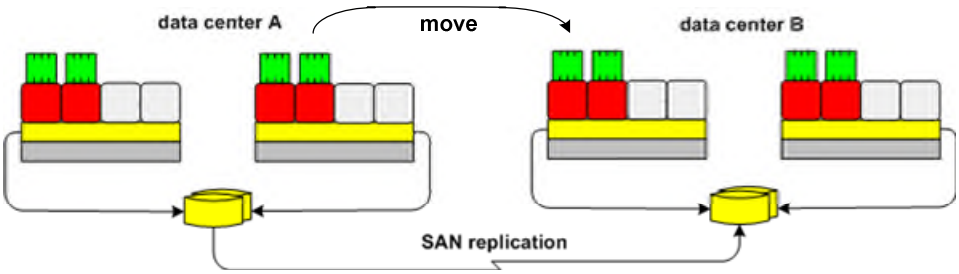

Both of these architectures achieve significantly higher availability than disaster-recovery architectures without the cost of maintaining an idle backup data center or one that is running deferrable workloads. Production workloads can be operating at all data centers.

The big difference between active/active systems and virtualized systems is recovery time. Active/active systems recover in seconds and provide true continuous availability. Virtualized systems recover in minutes, and their recovery times place them more in the high-availability category than in the continuous availability category.

Adoption of Active/Active Data-Center Architectures

The Forrester/DRJ study found that 44% of the responding organizations had adopted an active/active architecture. Twelve percent had gone fully active/active, while 32% had opted for a sizzling-hot standby architecture.

Of those organizations that had gone active/active, 89% agreed that the active/active architecture had helped them combine their DR and HA architectures into a single architecture. They no longer had to invest separately in HA for localized failures and DR for site failures, with the added burden of testing them separately.

Eighty-six percent of organizations that had gone active/active said that the move had reduced their downtime and that they could easily extend the availability benefits to business-critical applications.

Fifty percent of those respondents that have not yet gone active/active said that they are not confident that their DR solution is scalable or that it will achieve their recovery objectives. Fifty-four percent struggle to maintain up-to-date DR environments because of the rate of change in the configurations of their production systems.

Fifty-eight percent of all responding organizations believe that today's off-the-shelf technology for building active/active systems is mature, and 46% of those organizations that have not yet gone active/active are seriously investigating it.

Summary

The organizations that were surveyed lacked confidence in existing DR solutions. They avoid invoking a DR plan unless absolutely necessary. They face the decision of whether it is better to wait out the current outage or to risk a failover fault while trying to bring up the backup system.

Organizations can move from a shaky DR environment to a solid continuously available environment by adopting active/active architectures. These architectures will withstand the loss of critical components or even entire data centers without compromising the availability of IT services. This is certainly a significant journey for most organizations.[3] However, the journey is aided by developing a roadmap that outlines the changes to IT processes, architectures, and technology that must be made, and then proceeding one step at a time.

[3] One Bank's March Towards Active/Active, *The Connection*; September/October 2012.