

the *Availability Digest*

www.availabilitydigest.com
[@availabilitydig](https://twitter.com/availabilitydig)

Dark DR – Avoid Its Costs with Active/Active

August 2017

Traditional disaster recovery (DR) methods utilize an active/passive architecture, one in which there are two redundant servers. One server is actively processing the application, and the other acts as a backup system ready to take over should the production system fail. The two systems typically are located remotely from each other to avoid a dual failure due to some local disaster. From an operations viewpoint, the backup server remains ‘dark’ until it is needed. As such, this architecture is commonly called “Dark DR.”



In contrast, the servers in an active/active system both are processing transactions. Should one server fail, all transactions are simply routed to the surviving server. Of course, both servers must have sufficient capacity to accommodate the full operational load; but in normal operation, the reduced workload on each server results in performance improvement.

Some applications cannot run in a distributed environment such as an active/active configuration. These applications must run in their own server. It is these applications that give rise to Dark DR to provide redundancy to protect against a production-server failure.

Dark DR

The majority of organizations face considerable loss of revenue and/or reputation if their services to customers go down. Companies expect that at some time they will face a disaster that will take down their production systems. Therefore, organizations turn to redundant systems to provide ongoing services to users, even in the event of a system failure. The classic redundant architecture utilizes an active system and a backup system, typically separated by geography.

In this architecture, the active production system processes the application workload and replicates database changes to the backup system to keep it synchronized with the production system. The backup system sits by idly, ready to take over should the production system fail. Because the backup system is not performing any work, it is ‘dark’. That is why active/backup architectures are called “Dark DR.”

Limitations of Dark DR

A major issue with Dark DR is that switching over from a failed production system to the dark backup system can take minutes to hours. During this time, users have no access to their applications. Even worse, the switchover may fail and is known as a ‘failover fault.’

‘Configuration drift’ is yet another concern when the need arises to switch to the backup system. Should the primary server fail, the backup server may lack all of the information it requires to take over. Often this occurs because IT staff overlook the backup server when they apply updates to the production server. These missing pieces must be identified and migrated to the backup server to restore operations, a process that can take hours.

If the backup server is brought online successfully, the primary server's normal workflow must be redirected to the backup site. This may require a significant amount of manual reconfiguration requiring two IT teams – one at each site. The effort seriously can impede the activities of the primary site's IT staff, whose job it is to return the primary site to operation.

Costs of Dark DR

In order for a failover to be successful, it is essential to know that the backup system is fully operational. Thus, running a full-capacity redundant system in a secondary site represents a necessary yet considerable expense. In addition to the hardware and software costs of the otherwise idle system, there are several additional expenses:

- Extra maintenance – IT staff must ensure that the backup site is always operational. Best practices recommend to staff that they periodically switch production to the backup site to make sure it is working and that there will be no failover fault if the backup is needed. During the switchover time, which could be minutes to hours, users will have no access to the application and are down. If there is a failover fault (i.e., the backup system will not come online), users may be down even longer.
- Extra staffing – IT staff is required at both sites to support upgrades, one set of staff to make upgrades to one system and one to keep the other system running.
- Downtime costs – downtime creates a potentially huge cost for organizations, both in terms of immediate business loss and long-term reputation.
- Lost business – slower application performance due to overloaded servers can result in lost revenues as shoppers abandon their shopping carts due to long lag times.

Active/Active

Rather than build DR structures that must fail over to a backup system if the production system fails, organizations need to design for continuous availability. If a system fails, the failure should be transparent to the users of the applications. Active/active systems fulfill this requirement.

What is an Active/Active System?

An active/active system utilizes two or more servers that all are actively processing transactions. Whenever a transaction changes the database, that change is replicated to the other servers in the application network so that all databases remain synchronized. As such, a transaction can be routed to any server in the application network and can be processed in the same way that it would have been had it been routed to another server.

Active/active architectures provide many advantages to an organization:

- Improved asset utilization because each server is providing application processing.
- Seamless scalability because additional servers can be added easily to the active/active network.
- Dramatically higher uptime because if a server fails, all that needs to be done is to route transactions to surviving servers.
- Improved end-user experience because of the absence of user downtime and because of the faster response due to reduced server utilization.

Other Advantages of Active/Active

Technical Advantages

Active/active systems enable a smooth failover. Operations transition from the failed server to the other server with no interruption in service. Server failures are transparent to the users.

Maintenance can be provided on one server while the other continues in operation.

Businesses can cut costs by moving workloads between servers in different locations in response to changing cost factors (energy, rent, etc.).

Applications can handle more traffic by adding servers to the active/active system.

Cutting workload levels in locations creates more capacity for serving traffic growth.

Security improves because IT can patch a vulnerability on demand rather than having to wait for the next maintenance window.

Economic Advantages

By spreading traffic load across multiple systems, organizations put less strain on servers, thereby extending their functional lives.

Reduced site use means lower expenses. An organization does not need as much hardware in each location.

Increased application performance often leads to enhanced revenues. By speeding up performance on e-commerce sites, customers are less likely to abandon their shopping carts.

Maintenance costs are lower because tasks can be performed during work hours rather than in the middle of the night. One server in the active/active system can be taken down for maintenance while the other server continues to provide application services to the users.

With little to no downtime required for maintenance, organizations can increase revenues that otherwise would have been lost during the maintenance downtime.

Where Does HPE Stand?

With all the disadvantages of Dark DR, it remains a necessary architecture to provide redundancy for critical applications that cannot run in an active/active environment. These applications run in a production server that is backed up by a passive server, ready to take over application processing should the production server fail. Since the passive server is undertaking no application processing, it is a dark disaster recovery system. In addition, an active/passive architecture is simpler to implement when compared to the relative complexity of running applications in active/active's distributed environment.

HPE supports Dark DR via several data replication products, including its own RDF replicator and HPE Shadowbase from Gravic, DRNet from NTI, and GoldenGate from Oracle. Some of these products also support active/active architectures, thus avoiding Dark DR if so desired.

Summary

Active/active systems have many advantages over active/passive systems. The primary difference is that there is no 'dark' server sitting idly by, waiting to fill in for a failed primary server. Both (or all) servers are actively processing transactions, leading to greater capacity and faster response times.

Acknowledgement

Some information for this article was taken from the following source:

"The hidden cost of Dark DR: The economic argument for active/active operations," ITProPortal; June 14, 2017.