FEATURE SHEET

CARBONITE an opentext company | WEBROOT an opentext company

Failover and failback

How to eliminate planned and unplanned downtime for critical systems

There are many reasons why businesses would need to reroute incoming traffic away from one server and redirect it to another one. The first reason most people think of is an interruption to a server or application due to a hardware failure, power outage or natural disaster. A more common reason that gets less attention is downtime due to planned upgrades or periodic maintenance. Regardless of the cause, businesses need uninterrupted access to important applications and sources of data. Carbonite® Availability uses a process called "failover" to reroute traffic from a source server to a secondary one to ensure continued access in the event of an outage, whether planned or unplanned. This same process is available in our data migration solution, Carbonite® Migrate where it is referred to as a "cutover."

Health monitoring

Carbonite Availability, with automatic or push-button failover, can reduce or eliminate both planned and unplanned downtime. Once protection jobs are created in the central management console, the software sends all byte-level changes made to the source to the secondary server in real time. When configuring the solution, an administrator sets parameters for monitoring the source server in order to detect a failure.

The console gives administrators the ability to define how frequently to attempt to communicate with the source, how long to continue unsuccessful attempts and how many failed attempts constitute a failure.

Once conditions are met, one of two things will happen:

1. Automatic failover if you've configured it this way
2. Notification that the source has failed so you can initiate a manual failover

Failover and cutover

Whether triggered automatically or manually, once the process is initiated, the secondary server assumes the identity of the source, and user and application requests are automatically rerouted to the replicated environment. For unplanned failovers, users can continue accessing the secondary environment as long as they need to while IT resolves the underlying issue. Once the original source server is repaired or replaced, replication is initiated from the failed over target back to the original source. Once the two systems are synchronized, users can be routed back to the original source again. This process is referred to as "failback," where the target releases the identity it assumed during failover and the source reclaims it. Reversing replication before failback reduces user downtime by allowing users to continue accessing their data on the failed over target during the resynchronization process. Once failback is complete, user and application requests are no longer routed to the target, but back to the original source.

Planned vs. unplanned downtime

Disasters tend to generate attention-grabbing headlines and media coverage. They also have a tendency to expose the flaws in a data continuity plan. But disasters typically represent just a small fraction of the events that cause downtime. Planned downtime creates a real drain, and periodic backup tools and methodologies designed purely for disaster recovery are inadequate to minimize these outages.

Causes of planned downtime:

- Routine maintenance
- OS upgrades
- Application testing
- Hardware upgrades
- Data backups
- Server consolidation
- Infrastructure upgrades

Nondisruptive testing

For both Carbonite Availability and Carbonite Migrate, the software enables you to perform nondisruptive failover testing in an isolated environment with a snapshot of the target system's data so there's no impact to production. This helps administrators ensure functionality of the replicated environment prior to a live failover or cutover.

Customization options

In addition to automated and manual failover activities, administrators can customize failover by running scripts on the source and target as part of the process. Scripts may contain any valid OS command, executable or batch file. Examples of functions specified in scripts include:

- Stopping services on the target before failover that are unnecessary when it becomes the source
- Stopping services on the target that need to be restarted with the source's machine name or IP address
- Starting services or loading applications that are waiting for failover to occur
- Notifying the administrator before and after failover occurs

Recovery objectives

The replication and failover processes behind Carbonite high availability and migration solutions enable IT managers to meet very aggressive recovery objectives for data loss (Recovery Point Objective) and downtime (Recovery Time Objective). Replicating each source server change to a secondary server in real time, and at the byte level, is not only highly efficient over the network, but it allows IT administrators to achieve a sub-second RPO. The ability to quickly failover to that secondary server in the event of a failure and prevent planned downtime allows IT administrators to achieve RTOs of seconds to minutes. With Carbonite high availability and data migration solutions, businesses can ensure greater resiliency for critical systems and stay on top of technology updates without disrupting users.

Causes of unplanned downtime:

- Human error
- Hardware failure
- Software failure
- Unprotected storage or disk failure
- Malicious users
- Local disaster (e.g. fire, storm)
- Regional disaster (e.g. earthquake, flood)

Contact us to learn more – Carbonite US
Phone: 877-542-8637
Email: carb-data_protection_sales@opentext.com

About Carbonite and Webroot

Carbonite and Webroot, OpenText companies, harness the cloud and artificial intelligence to provide comprehensive cyber resilience solutions for businesses, individuals, and managed service providers. Cyber resilience means being able to stay up and running, even in the face of cyberattacks and data loss. That's why we've combined forces to provide endpoint protection, network protection, security awareness training, and data backup and disaster recovery solutions, as well as threat intelligence services used by market leading technology providers worldwide. Leveraging the power of machine learning to protect millions of businesses and individuals, we secure the connected world. Carbonite and Webroot operate globally across North America, Europe, Australia, and Asia. Discover cyber resilience at carbonite.com and webroot.com.

© 2020 Open Text. All rights reserved. OpenText, Carbonite, and Webroot are each trademarks of Open Text or its subsidiaries. All other trademarks are the properties of their respective owners. FS _ 102820