





CARBONITE[®]
an **opentext**™ company

WEBROOT[®]
an **opentext**™ company

Carbonite[®] Availability

Technical overview






Executive summary

The availability imperative

The pace and competitiveness of global commerce require a flexible and resilient IT infrastructure. Any downtime can have serious repercussions for your organization's profitability, reputation and revenue. Achieving the level of high availability required to support mission-critical systems puts increased demands on IT and data center professionals.

Carbonite Availability *Powered by DoubleTake* is a proven, simple and scalable high availability and disaster recovery software solution. Real-time, asynchronous, byte-level replication efficiently and securely replicates entire servers or select data to another server located anywhere—even in the cloud—so your business can quickly and easily recover from an outage. Network independence and bandwidth efficiency allow Carbonite Availability to be run on any network without the need for dedicated resources or special hardware. Easy to manage and designed to scale to complex server environments, it keeps your business running 24×7 at a price you can afford.





True real-time replication

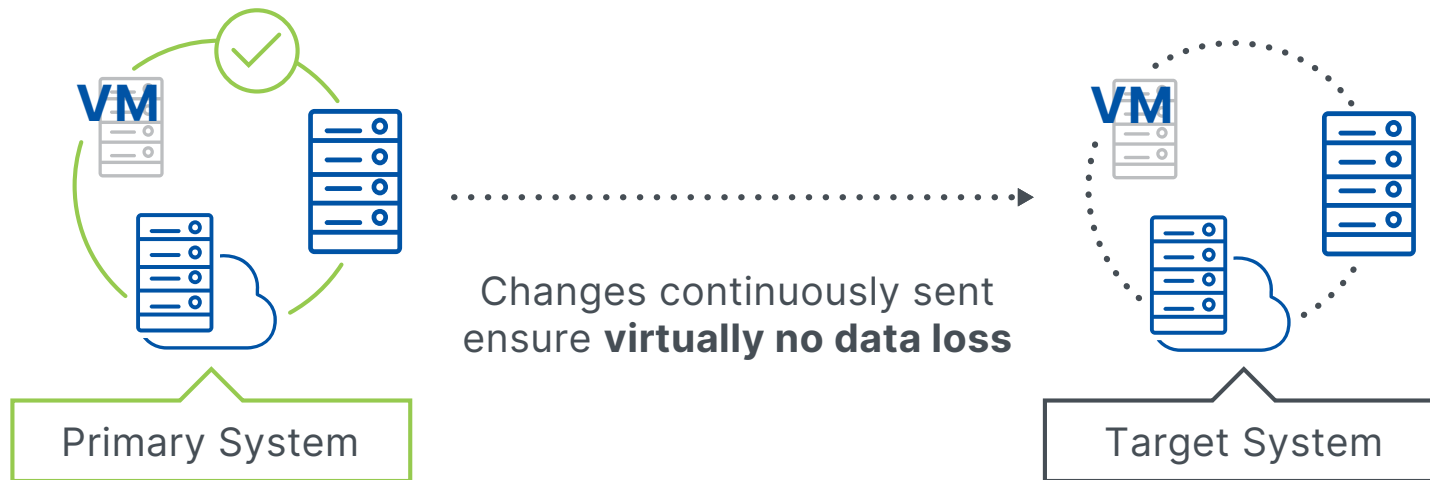
More efficient and better protection


Carbonite Availability features patented, real-time, asynchronous, byte-level replication. It replicates data immediately and continuously to a secondary server as fast as bandwidth and processing allow. Asynchronous, real-time replication is ideal for crossing any geographic distance because it can flexibly queue and send replicated data as bandwidth and latency over distances fluctuate throughout the day, without holding up the production servers.

Our byte-level replication is very efficient over any distance because, unlike some other products, it sends only changed bytes rather than entire blocks. Additionally, Carbonite Availability reduces the amount of data on the network by providing three levels of intelligent data compression during replication.

Bandwidth throttling can further optimize the use of limited network resources at busy times of the day. When bandwidth is limited, Carbonite Availability will intelligently queue any data that cannot be sent, but it will continue to try to send data from the queue as quickly as bandwidth allows, keeping data synchronized between the source and target as fast as possible.

Sequential replication preserves write order to keep important application and file data in a continuously consistent state. This is especially important for database applications, such as Microsoft SQL Server or Microsoft Exchange, which use transactional database technologies. By always maintaining data integrity, Carbonite Availability allows you to fail over to the most current data whenever the need arises. Open-file mirroring and replication capabilities allow you to replicate active files without taking them offline.





Robust protection

Reliably meet all your recovery needs

Uninterrupted operations are imperative for many organizations. Carbonite Availability protects the accessibility of data and applications.

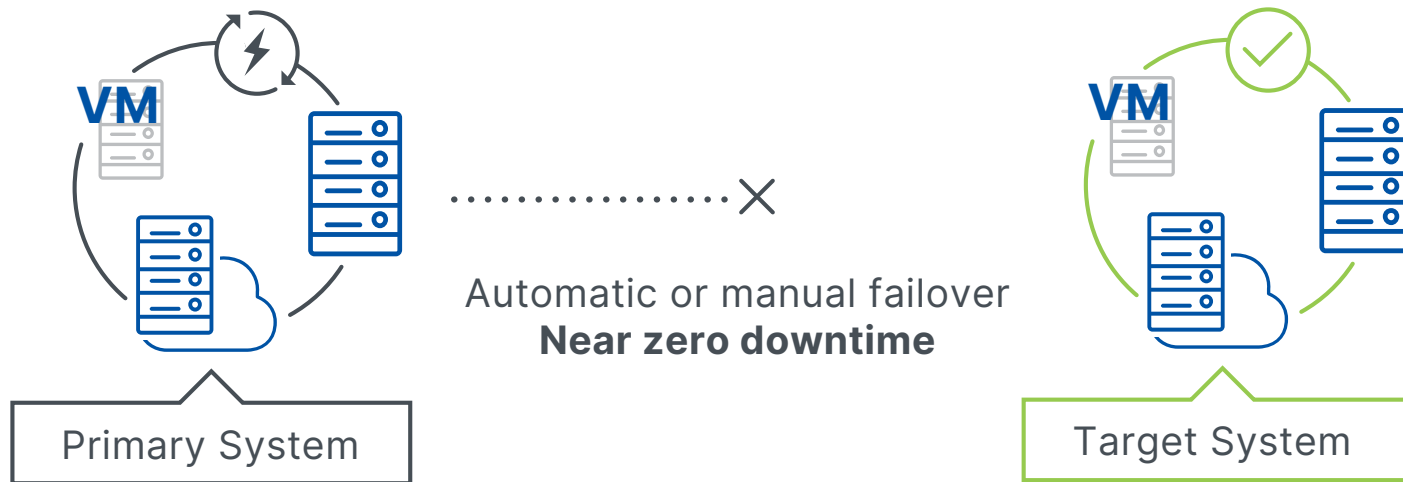
The comprehensive, reliable high availability and disaster recovery capabilities of Carbonite Availability can address your recovery needs. It can be as granular as you need, replicating specific files, entire applications or full servers.


In the event of an outage, failover to the waiting secondary server is easy. A full server failover can be accomplished in seconds or minutes. Carbonite Availability monitors the behavior of the production environment and can automatically take corrective action.

Or you can choose to initiate an automated failover process on demand. When the time is right, you can fail back to the original server, or a replacement server, by performing an automated failback with push-button simplicity.

Regardless of whether the production data, applications or server being protected reside on a physical, virtual or cloud server, it can be recovered to any physical, virtual or cloud server. When recovering to a virtual machine, Carbonite Availability can create and auto-provision the recovery environment.

In the case of data corruption issues, the point of failure is typically not the most appropriate recovery point. Instead, you may need to recover to a point before the corruption occurred. Carbonite Availability has you covered with multiple recovery point options.





Flexibility

Support the full mix of servers in today's data centers

Efficient, real-time background replication prevents data loss and improves recoverability in any Windows or Linux environment.

Modern data centers are intricate for many reasons. For example, their complexity may have evolved due to mergers and acquisitions or simply through the acquisition of new technologies over time to meet new business demands. Unlike Carbonite Availability, many competitor's products support only one part of the modern data center, specializing in a single virtualization platform or providing a one-size-fits-all approach to data and application availability.

Carbonite Availability delivers technologies that replicate data efficiently, while eliminating the barriers between physical, virtual and cloud platforms. It works within and across these diverse computing platforms to provide protection throughout your data center.

Many organizations also use a variety of operating systems, including Windows and Linux. Carbonite Availability delivers comprehensive solutions for each of these operating systems.

Carbonite Availability can protect your entire data center, even if it includes a mix of differing hardware, storage and/or operating system versions. It works on all leading virtualization, private cloud and public cloud platforms and accommodates all business and system applications, including Exchange, SQL Server, SharePoint, Oracle, MySQL and more.

The ideal replication configuration is not the same for all organizations. That's why Carbonite Availability supports a wide range of configurations, such as one-to-one, one-to-many, many-to-one and cascading. Carbonite Availability can also protect cluster workloads to different physical locations without using shared storage.

Furthermore, Carbonite Availability is a highly scalable solution. From the smallest of data centers through to the largest and most complex enterprise information architectures, Availability fulfills your needs.


Supported platforms

Carbonite powered by EVault

- Linux
- Windows
- VMware
- Hyper-V
- AIX/Solaris/HP-UX
- iSeries
- Microsoft SQL, Exchange and SharePoint
- Oracle DB

Carbonite powered by DoubleTake

- Linux
- Windows
- VMware
- Hyper-V
- Microsoft SQL
- Amazon Web Services, Azure and any cloud-based VM



Simplicity

Set up quickly, run continuously, manage easily

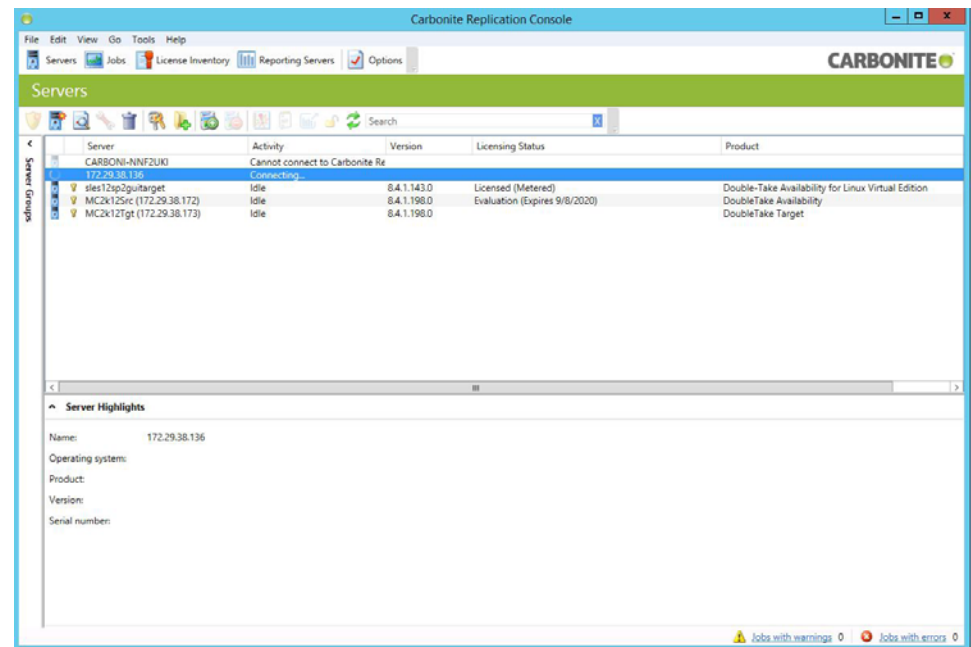
Highly efficient, change-based, real-time replication eliminates the risk of data loss, without straining system resources.

Carbonite Availability automates the setup and configuration of real-time protection and availability management for datasets, business-critical applications and full system states through advanced management features. Using full-server protection, administrators do not need to have an in-depth understanding of the applications or server configurations to enable protection.

Through a streamlined and unified console, you can quickly learn to manage and monitor all the powerful capabilities of Carbonite Availability. An enterprise dashboard provides real-time visibility and management of your physical, virtual and cloud environments, with features that allow you to monitor the health of your primary and failover servers, perform recoveries and failbacks, check event logs and manage your license inventory. Single-screen monitoring allows you to sort, filter and monitor data on the health of servers, check the currency of your target system, and view event logs in one intuitive display.

To assist you in optimizing your configuration, the Carbonite Availability diagnostics job simulates replication traffic and estimates the amount of bandwidth needed for efficient replication between the production and backup servers.

The console further simplifies the management of your Carbonite Availability environment by forwarding events via SNMP to other management tools and providing SMTP email alerts directly to your inbox.




Affordability

Reduce total cost of ownership

Carbonite Availability makes HA/DR administration fast and easy so your IT personnel can focus on what matters most—effective, efficient and profitable operations.

Carbonite Availability minimizes the total cost of owning, implementing and maintaining the highest possible level of data and application availability and disaster recovery readiness. It does so by:

- Taking advantage of existing hardware, storage and network infrastructure to eliminate the need to invest in special equipment
- Replicating only changed bytes and, optionally, compressing data before transmission to minimize the need to purchase extra network capacity
- Providing a simple, intuitive console to minimize the strain on IT staff and reduce the need for specialized training
- Including automatic features that reduce the need for manual administration to, on average, just minutes a day
- Offering flexible licensing models for perpetual ownership or pay-as-you-go DRaaS to provide a solution that fits your budget



Integration and extensibility

Comprehensive SDK

The Carbonite software development kit (SDK) is a collection of application programming interfaces (APIs), documentation and code samples that provides developers with the tools they need to rapidly build solutions that integrate with Carbonite products. It allows them to integrate Carbonite HA/DR capabilities into applications seamlessly. For example, they might use the SDK to integrate the status of Carbonite Availability into their monitoring consoles, whether those consoles are custom-developed or purchased from third parties. An online SDK portal contains the API documentation, sample code, tips and forums. It helps developers accelerate application integration.

The SDK includes SOAP APIs, PowerShell commands and RESTful APIs. You are free to choose the programming language and APIs that you're most comfortable with to automate tasks or to implement your own software using the high availability, disaster recovery and migration services of Carbonite. The SDK is open to all customers and partners that wish to automate tasks or implement their own software solutions using Carbonite services.

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.

CARBONITE an opentext company | Replication Products Documentation Library

Availability and Migrate SDK

The SDK is a collection of APIs, documentation, and code samples that provides the necessary tools for developers to build solutions that integrate with Availability and Migrate. The file types are HTML, CHM, RAR, and ZIP.

Version 8.3.0

- [REST HTML help file](#)—HTML help file for the REST API.
- [REST .chm help file](#)—Compiled help file for the REST API.
- [REST C# examples](#)—Job creation examples using C#.
- [REST Ruby examples](#)—Job creation examples using Ruby.
- [SOAP HTML help](#)—HTML help for the SOAP API.
- [SOAP .chm help file](#)—Compiled help file for the SOAP API.
- [PowerShell](#)—Job creation examples using PowerShell.

Version 8.2.0

- [REST HTML help file](#)—HTML help file for the REST API.
- [REST .chm help file](#)—Compiled help file for the REST API.
- [REST C# examples](#)—Job creation examples using C#.
- [REST Ruby examples](#)—Job creation examples using Ruby.
- [SOAP HTML help](#)—HTML help for the SOAP API.
- [SOAP .chm help file](#)—Compiled help file for the SOAP API.
- [PowerShell](#)—Job creation examples using PowerShell.

Version 8.1.1

- [REST help file](#)—Compiled help file for the REST API.
- [REST C# examples](#)—Job creation examples using C#.