



Technology Overview
New Features
BackupEDGE 3.0

We're constantly asking our clients what tools our products need to serve them better.

Introduction to the New Features in BackupEDGE 3.0

BackupEDGE has a long history of providing reliable data protection for many thousands of users. As operating systems, storage devices and usage needs and tendencies have changed over the years, it has continuously met the challenge of providing inexpensive, stable backup and disaster recovery on a variety of UNIX and Linux platforms.

Clients routinely find new and clever ways to utilize products. Storage devices have taken on new and exciting features, and incredible capacities. Products designed years ago had built-in limits that were thought to be beyond comprehension. Today, these limits are routinely exceeded. The need for data security is even more apparent.

To continue to meet the evolving needs of our clients, we are always asking what features of our products they find most useful, what improvements we can make, and what new requirements they have. We've used this knowledge to map out new product strategies designed to anticipate the needs of the next generation of users, systems and storage products.

This has resulted in the creation of *BackupEDGE 3.0*, with a combination of internal improvements, new features and enhanced infrastructure designed to become the backbone of a new generation of storage software.

Summary of Major Changes and Additions

BackupEDGE 3.0 adds:

- MySQL Hot Backups.
- SharpDrive™ support (03.00.03).
- Blu-ray Disc™ (BD-RE) support.
- New Scheduling System.
- Device Quotas.
- Archive retention times with "lazy" reclamation.
- Disk-to-anything-to-anything support.
- Multiple backups per medium on tape, Blu-ray Disc, DVD+RW, DVD-RAM, and REV.
- Updated Web Interface.

BackupEDGE 2.3 added internet cloud storage and many new operating systems:

- Amazon Simple Storage Service (S3) is now supported as a storage Resource.

Microlite Corporation
2315 Mill Street
Aliquippa PA 15001-2228 USA
Tel: 724-375-6711
Fax: 724-375-6908
email: sales@microlite.com
web: www.microlite.com

BackupEDGE 2.2 added security improvements and a licensing change:

- A new FTP/FTPS client with many security improvements
- Relaxed FTPS rules. An encryption license is no longer required to use FTPS
- FTP/FTPS works on more server types
- SSL libraries have been updated
- A Cancel Jobs facility has been added to EDGEMENU

BackupEDGE 2.1 added many new capabilities over its lifetime:

- Transparent Media™ Technology
- Network Attached Storage (NAS) support with **ftp** backups
- Backups to removable HD or dedicated archive directories
- Command line interface access to Resources
- Multi-Volume Fast File Restore / Instant File Restore
- Access Control List support under Linux, OpenServer 6 and UnixWare 7
- Media-free PXE booting for disaster recovery under Linux, OpenServer 6 and UnixWare 7
- LVM2 support for disaster recovery under Linux 2.6.x kernels
- Support for SCOoffice Server 4.1 and 4.2
- Web Services support for Sun Java 5 (1.5)
- Hot Plug device support
- WORM tape detection and support
- Tape media capacity reporting
- A Cancel Jobs facility has been added to EDGEMENU

BackupEDGE 2.0 improved on almost every facet of its predecessor, as well as providing a variety of new features. Additions and improvements included:

- Choice of Character, Java™ and Web Services Interfaces
- Support for Modern Storage Devices
- Data Encryption (Optional)
- Enhanced Data Compression
- Unlimited Pathnames
- Additional Archive Integrity Checking
- Enhanced Data Format
- Features-based Licensing

Of course, as usual different build levels of the products support more Linux and UNIX releases and supplements, fix bugs and have minor enhancements.

Transparent Media Technology

Transparent Media means that *BackupEDGE* eliminates distinctions between different kinds of backup media. Unless there is a specific hardware or operating system limitation, *BackupEDGE* provides complete functionality for all media and backup types, including:

-
- High performance backups
 - Maximum data integrity through full file checksumming
 - Fastest possible access to files/directories for restore
 - Multiple archive per medium on all re-writable storage types (except CD-RW and DVD-RW).
 - Availability for disaster recovery on platforms supported by *RecoverEDGE*.
 - There are no trade-offs when switching media types, even when using the CLI (command line interface).

No other product provides the same flexibility when it comes to choosing the right storage system for any particular application.

Network Attached Storage

BackupEDGE has an ftp/ftps client which is built directly into the backup engine and treated as a standard storage "Resource". This means that virtually any server or storage appliance, running any operating system, can be used as a backup device, as long as it is running an ftp server and available via a network connection. It is now possible to:

- Use any Network Attached Storage (NAS) device to reduce backup windows and consolidate archives.
- Create your own inexpensive storage server.
- Have two servers backup to each other for redundancy, over your intranet or the Internet.
- Create a star or ring topology, allowing each server to store archives to another server.
- Consolidate all your server backups by sending them to a data center, eliminating the expense of local storage devices.
- Increase the frequency of data protection. Perform multiple differential/incremental backups per day.
- Set quotas per Resource. Expired archive will only be erased when space utilization nears the quota.

Unique Features

- NAS backups are segmented automatically, and are not limited by either "ulimit" or file size limitations on either the local system or the backup server. The high performance compression built into *BackupEDGE* reduces network bandwidth and storage space requirements. All of the other *Transparent Media* features apply.
- NAS devices are treated similarly to autochangers, allowing simple backup rotations and retirement policies. Whenever a restore is requested from a NAS, all current archives are listed and the user may choose the archive to be restored.
- Non-scheduled backups can be given a "Slot Name", or unique storage identification.
- NAS management is included. NAS backups can be viewed and removed at any time. "Volume Size" for a NAS resource is actually a storage quota, or maximum storage allotment for all archives created by the Resource.

- With ftps, the transport layer itself can be encrypted, protecting your data during transit over the intranet or internet.

Amazon Simple Storage Service (S3) Backups

Amazon.com is a web retailer with a large, Internet-based computing infrastructure. *Amazon Web Services™ (AWS)* is a division of *Amazon.com* that provides additional storage and service offering using this infrastructure.

One service AWS offers is the **Amazon Simple Storage Service**, otherwise known as **Amazon S3**, or simply **S3**.

Think of **S3** as a storage server with virtually unlimited storage space and bandwidth, and high availability. To read more about the benefits of the service, browse to <http://aws.amazon.com/s3>.

BackupEDGE 02.03.00 and later treats **S3** as a standard storage Resource, with many similarities to *FTP Backups* (Network Attached Storage). It is a fully functional Resource, capable of:

- Full or partial system backups.
- Checksum and full bit level archive verification
- Instant File Restore™
- Bare metal disaster recovery with *RecoverEDGE*.
- Quota use. Allows placing a maximum limit on storage space used / paid for.

As with any other *BackupEDGE* Resource, servers can be set up to perform complete Master, Differential and Incremental system backups. Individual backup subsets (we call them backup Domains) may be created as a supplement to locally hosted backups.

Storage pricing is very reasonable. rates can be found at: <http://www.microlite.com/s3/>. Billing comes directly from Amazon Payments.

Directory Backups

Backups may be sent to removable HD / flash media (operating system permitting). *BackupEDGE* will manage mounting and unmounting the media as required. Removable media backups may be used for disaster recovery. Flash cards / removable HDs cannot be made bootable, though. You must use CD / DVD / REV (etc.) boot media with them.

Backups may also be sent over NFS or SMB mounts in this fashion. However, disaster recovery does not support these transports, so ftp is the recommended way to perform remote backups.

Backups may additionally be sent to a directory on an existing filesystem. *BackupEDGE* will not attempt to mount or unmount any devices for this case. Again, these backups may be quite useful for targeted snapshots, but are not suitable for disaster recovery, as they're stored on local storage and may be erased by the disaster recovery process!

All of the other *Transparent Media* features apply.

Command Line Interface

We recognize that despite the power and flexibility of the **edgemenue** and Scheduler interfaces, people have custom uses that require archives to be created from the **/bin/edge** command line.

BackupEDGE has the exclusive ability to create archives to Resources in addition to files. For example:

/bin/edge cvf dvd0 /bin

would automatically create an archive on whatever appropriate cd or dvd media was inserted into a DVD drive. The right media strategy, preparation (blank, format, etc.), compression and volume size would be selected automatically by the Resource settings.

/bin/edge cvf rev0 /bin

would do the same for an Iomega REV device.

/bin/edge cvf url0 /bin

would automatically create an ftp backup, negotiating link type, passwords, compression, etc. from Resource settings.

Additional command flags are available for fine control of compression, slot name, etc. No other product has such a capability.

Indexed File Access

BackupEDGE has had quick file access (the ability to quickly find and restore individual files and directories) for many years. We've called this capability Fast File Restore™ for tape backup, and Instant File Restore™ for optical media and Iomega REV backups. We support indexed access capability to archives consisting of multiple volumes and archive with multiple archives per volume, as well as to network, directory and D-2-D backups.

User Interface Flexibility

BackupEDGE has a state-of-the-art user interface that can be rendered three different ways:

- in graphical mode as a *Web Service* from any client system supporting Sun Java 1.4.2 or later, such as a Windows PC.
- in graphical mode on X11 consoles or clients equipped with Sun Java 1.4.2 or later.
- in character mode on system consoles, dumb terminals or xterm clients.

The user interface runs in all three modes using the same compiled program. There are no operational differences between the interfaces, meaning that there is no re-training required, or loss of features when switching between interfaces. Of course, the user has full access to GUI-type features such as buttons and mouse clicks when running graphically.

The Web Services interface provides a secure data link and enables authorized users to maintain a *BackupEDGE* installation from anywhere in the world without compromising system security. The only requirement is a Java equipped web browser.

Microlite is the world leader in enabling our partners the opportunity to save money by deploying modern, inexpensive storage devices and media.

Modern Device Support

No other product supports and controls the full gamut of storage devices enabled by *BackupEDGE*.

- Tape Drives
- Autochangers, Loaders and Libraries
- All Blu-ray Disc, DVD and CD
- All Iomega REV™ devices, including the Autoloaders.
- NAS, removable hard drives, cartridge-based drives and flash devices.

Microlite is particularly proud to have recognized the early potential of optical and REV devices for use as reliable backup devices under UNIX and Linux. We were the first company to support these media types, allowing our partners to build more cost-competitive systems. The addition of NAS/HD/Flash backups provides even more options. Our performance and media handling is still far and away the best-in-class.

Data Encryption

The need for data privacy and security continues increase, and *BackupEDGE* is ready to respond to that need.

BackupEDGE has a license option enabling public key data encryption. Users may choose to encrypt data by specifying filenames, directories or file extension types. Encryption uses the well documented AES encryption algorithm with 256 bit keys, which are further protected by RSA 2048 bit encryption. Encryption is performed at the file level, providing the following benefits...

- only data that needs to be protected is encrypted.
- overall performance stays high as only critical files are subject to CPU intensive encryption.
- full compatibility with our bit-level verify, new file checksum verify, indexing, quick file access and disaster recovery features is maintained.
- Each encrypted file is pre-compressed using the powerful ZLIB libraries to ensure that no space is lost due to the inability of tape hardware compression implementations to compress encrypted data.

Optionally, a user may choose to encrypt an entire archive (except the file headers), although this is not recommended.

The methodology used by *BackupEDGE* for data encryption will be published on our web site to assure users that standards are being followed and that no "back doors" or other security holes are in place. Users of encryption should be aware of the potential consequences of lost or stolen keys or pass phrases before utilizing this new technology. A separate whitepaper describes our encryption capabilities in more detail.

Enhanced Data Compression

The increasing popularity of CD, DVD, REV and NAS backups, as well as the proliferation of new tape devices from vendors like Seagate that do not include hardware compression in the interests of affordability, has placed a whole new emphasis on the importance of software data compression. *BackupEDGE* has the industry standard ZLIB compression algorithm incorporated into its data format.

With our enhanced compression, newer backup Resources becomes far more useful, and legacy tape devices get a whole new lease on life.

Microlite Corporation
2918 Mill Street
Newport News, VA 23601-1178, USA
Tel: 724-375-6711
Fax: 724-375-6908
email: sales@microlite.com
web: www.microlite.com

There is a standard suite for testing compression ratios called the *Calgary Corpus Data Set*. **BackupEDGE** default compression using the *Calgary Corpus Data Set* tests at **2.9:1**. At its lowest setting (there are 9 available compression settings) it is **2.7:1** and at its highest setting it is over **3:1**. By comparison, our legacy products test out at **2.1:1**.

This means that, at the compressibility of the data used by the *Calgary Corpus Data Set*, our legacy products could get **73.5GB** of data onto a 35GB REV, while current **BackupEDGE** would by default be able to store **101.5GB** on the same media. On a 4.7GB DVD, the relative results would be **9.87GB** for the older product, and **13.63GB** for current. Your results will be based on your data mix, and will be different.

We've also incorporated a method for doing streaming compression. This results in faster performance, especially when compressing very large files, and zero temporary space requirements.

Compression and encryption work together seamlessly.

Compression and encryption have a natural synergy. Encrypted files confuse the compression algorithms built into modern tape drives. When encountering files with seemingly random data, the compression algorithms either begin to compress negatively or temporarily disable themselves. Either way, net storage space and performance are impacted. By using our new software compression before encryption, it is possible to maximize available storage space while maintaining reasonable performance.

Of course, archives compressed with legacy versions of **BackupEDGE** can still be restored and uncompressed.

Archive Integrity Checking

A checksum can be added to each file in a **BackupEDGE** archive. This allows an archive integrity check to be made with great confidence months or even years after the archive has been created.

Newer storage technologies require better verification techniques to ensure archive integrity.

The checksum guards against medium errors. For example, if an archive was verified successfully years ago, and one now wishes to restore data from that archive, it is desirable to know that no data has changed while the medium on which the archive resides was in storage. To do this, **BackupEDGE** uses the checksum to make sure the data is identical to what was verified earlier.

The checksum does **not** tell you if the data is an accurate reflection of the filesystem data on which the archive was based; it only checks that the archived data has not changed since it was written. **If** the data was verified successfully against the original filesystem data, **then** the checksum can be used to check if the archive still records that correct data. It cannot tell you if the archived data matches what was on the filesystem initially.

The checksum can be computed and verified without being able to decrypt the archive data; one can verify that the archive is as good as it was when it was written without having to decrypt the data itself.

Note that the checksum does **not** protect against intentional tampering.

Virtually Unlimited Pathnames

Network connected systems now have larger pathname requirements. In addition, people use what can only be described as sentences when creating filenames for new documents.

Removing legacy limitations results in a longer product lifetime as systems and capabilities expand.
Microlite Corporation
Alhambra, CA 91801-7248 USA
Tel: (714) 375-8711
email: sales@microlite.com
web: www.microlite.com

The tar data format as described in IEEE 1003.1-1993 had a natural limit of 99 characters per pathname/filename. *BackupEDGE* 01.00 and 01.01 extended this to 170 characters. *BackupEDGE SS* 01.02 further extended this limit to 400 characters for filenames, while still having a 170 character limit for links and symbolic links.

BackupEDGE 2.x and later have been architected with a **5,000 character pathname limit**, which essentially removes all limitations on pathname or link size. Imagine a pathname that takes up 2½ standard 80x25 data screens and you'll get the idea.

Enhanced Data Format

As we began mapping out all of the new capabilities we listed, it became apparent that we would need to enhance the data format. The old format was too limited to provide an easy mechanism for incorporating these features reliably.

Four primary goals were set for the replacement of this critical part of the product.

- **Standards-Based.** The product must continue to be based upon proven data storage standards, allowing some form of data restoration even without *BackupEDGE*.
- **Backward Compatibility.** That is, the ability to read and restore older *BackupEDGE SS* archives.
- **Command Set Compatibility.** Many of our clients still choose to write their own scripts, and we wanted to maintain maximum compatibility. In fact, we've improved it greatly.
- **Extensibility.** We wanted to build a platform with the ability to add new capabilities quickly and easily. For instance, we've recently added support for ACLs to our Linux, OpenServer 6 and UnixWare 7 products.

The new data format in *BackupEDGE* is not only state-of-the-art today, it is **designed to be easily extended** as system features require.

The New Data Format - PAX from IEEE 1003.1-2001

The data format incorporated into *BackupEDGE 2.x* and later is based on the specifications of the most recent version of the Portable Archive Interchange (**pax**) standard as defined in IEEE 1003.1-2001 by The Open Group and the IEEE. In reality this format is actually the newest generation of the **tar** format. It provides a known, reliable data format definition with the provisions we needed to extend it into a full fledged data protection system.

The format is similar enough to the legacy tar format that existing tar programs can read it, excluding those files which use long filenames or other features that tar does not understand. Free-software and public domain readers for IEEE 1003.1-2001 archives are readily obtainable (**pax**, **star**), so actually requiring tar to read a *BackupEDGE* archive is unlikely. Of course, the IEEE 1003.1-2001 standard document itself is also available.

While *BackupEDGE 2.x* and later provide features not found in these readers, none of these extensions rely on proprietary technology. Most are based on freely-available, published algorithms. The remainder are well-documented on the Microlite web site, to ensure that your data is recoverable even if a copy of *BackupEDGE* is not available.

Microlite Corporation
2315 Mill Street
Aliquippa PA 15001-2228 USA
Tel: 724-375-6711
Fax: 724-375-6908
email: sales@microlite.com
web: www.microlite.com

Features Based Licensing

New extensions to the License Manager allow features to be enabled or disabled easily. The first feature to be licensed via the new License Manager is data encryption. The feature will be licensed at an additional charge. As is our policy, all features, including encryption, are enabled during the 60 day evaluation period that commences when the product is first installed, providing plenty of time to test the full product.

Additional Features and Limitations

Performance and System Overhead

BackupEDGE has been the subject of extensive performance profiling, during both backup and verify.

Standardized Pattern Matching

Specifying pathnames and wildcards is consistent throughout the program and user interfaces.

Virtual File Processing

This is performed as a single pass process.

Limitations

BackupEDGE 3.x can read all archives created by older product versions. However, since it can do so much more than the old products, all archives created by the new release can not be restored by older releases.

Platforms

BackupEDGE is available for the following platforms...

- Linux® - Any version on an AMD64 or EM64T processor with a 2.6.x kernel, including Red Hat Enterprise Linux, Novell SUSE Linux Enterprise Server, Mandriva, CentOS, Debian, Ubuntu and more. Supported versions are listed on the [Linux Support Tables](#) section of the *Microlite Web Site*.
- Linux - Any version on an Intel IA32 processor with a 2.4.x or 2.6.x kernel, including Red Hat Enterprise Linux, Novell SUSE Linux Enterprise Server, Mandriva, CentOS, Debian, Ubuntu and more. Supported versions are listed on the [Linux Support Tables](#) section of the *Microlite Web Site*.
- SCO OpenServer 6.0.0
- SCO OpenServer 5.0.5-5.0.7 and 5.0.7V
- SCO UnixWare 7.1.4