Product Documentation

S3Utils Documentation

S3Utils Release 1.0.0
4 Dec 2024

Table of Contents

Welcome to the documentation for S3Utils™ by NuWave Technologies. You may navigate through the documentation using any of these methods:

- The links in the sidebar to the left
- The quick links shown below
- The search bar in the page header

If you are a new S3Utils user, please start with How To Begin. If you have comments on this documentation or questions about the product that are not answered here, feel free to submit your comment or question at our Support Center.

How To Begin

This section describes how to proceed if you are a first time S3Utils user. S3Utils users typically fall into one of the categories shown below. The sections below describe how each category of user should proceed.

- Administrators - An administrator is anyone who is responsible for installing and managing the S3Utils processes. Systems/Operations managers usually fall into this category.
- Developers - A developer is anyone responsible for building applications that use S3Utils. Developers often need to manage S3Utils on development and test systems so may also be considered Administrators.
- Others - Anyone who is neither an Administrator or Developer but has general interest in the product.

Administrators

- Read Introduction to S3Utils.
- Install S3Utils as described in Installing S3Utils.
- Start S3SERV as described in Starting S3SERV.

Developers

- Read Introduction to S3Utils.
- If also acting as an Administrator, install S3Utils as described in Installing S3Utils.
- If also acting as an Administrator, start S3SERV as described in Starting S3SERV.
- Read Getting Started in the User Guide.

Others

- Read section Introduction to S3Utils.

Introduction to S3Utils

This following sections describe concepts and features of S3Utils. The information is intended for S3Utils users, and those seeking a high level view of how S3Utils. HPE NonStop and application programming experience may be helpful but is not required.

What is S3Utils

Since the introduction of AWS Simple Storage Service (S3), S3 compatible object services have become highly popular, with each of the major cloud services providing S3 compatible storage, and a number of "S3 only" storage services also available.

These services provide low cost, durable storage, and are used for both archiving of file data and transferring data to other parties or services. However, while there are a number of utilities available to integrate with S3 object storage, none are available in HPE NonStop Server Guardian environment. Of those that may be available in the OSS environment, none support online transaction processing (OLTP), the model used by most NonStop Server applications.

S3Utils provides a complete solution for integrating with S3 storage. With a small set of program files, and simple configuration, which is compatible with AWSCLI configuration, S3Utils provides integration with S3 storage in both the batch and OLTP Guardian environments out of the box.

Product Components

S3COM

The S3COM utility provides a command line interface to S3 storage which is compatible with the Guardian COM program model. This allows the program to be driven by batch scripts written in TACL.After some simple configuration, customers can begin reading and writing data directly to S3 storage. The utility also provides the ability to list, delete, or move objects within the object store.

S3SERV

The S3SERV process can be deployed as a standalone server or TS/MP serverclass and provides integration with S3 block storage services via a simple IPM based application programming interface. This allows existing or new NonStop applications to integrate with S3 storage as part of their online transaction processing.

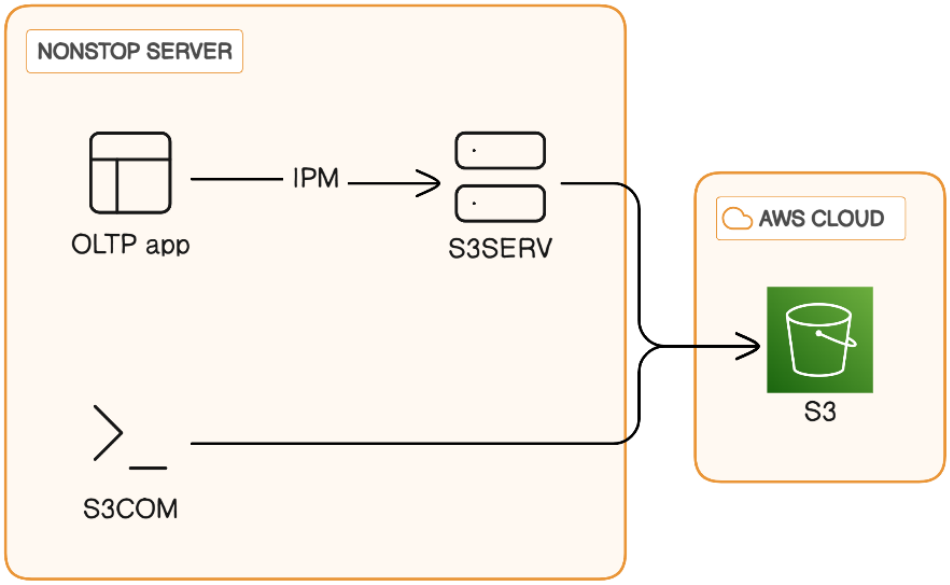

Frequently Asked Questions

Does S3COM support all AWS CLI S3 commands?

No. Bucket management commands, and commands that are problematic due to Enscribe file system naming limitations, such as sync, are not supported. The commands currently supported are:

- ls - list S3 objects
- cp - copy S3 object
- rm - delete S3 object
- mv - move S3 object

How does S3COM authenticate with S3?

S3COM uses the same "config" file storage mechanism as the AWS CLI.

Does S3COM support authentication using AWS IAM Identity Center (previously AWS SSO)?

No. The first product release will not support SSO. It may be available in a future release.

How does S3COM handle the substantial differences between the Enscribe file system and the hierarchical object storage method found in S3 buckets?

Due to the limitations for the Enscribe file system, S3COM limits object file names to 8 characters and recursive operations are limited to one level. S3COM does not attempt to normalize or map long S3 object paths to 8 character Guardian file names.

Does the S3SERV process support multiple concurrent requests (is it "multi-threaded")?

No, S3SERV processes a single S3 request at a time. Multiple concurrent requests can be achieved by configuring the S3SERV pathway serverclass appropriately to handle the required number of concurrent requests.

User Guide

This section contains information on using S3Utils.

Getting Started

This section provides an overview of how to get starting with S3Utils. If you have not done so already, please read the What is S3Utils section before continuing.

S3Utils is designed to facilitate user interaction with the S3 object store, offering two primary modes of operation:

1. **Interactive Mode with S3COM CLI:** Ideal for users who prefer a command-line interface. The S3COM program adheres to the COM program model, allowing for scripting through TACL macros.
2. **OLTP Mode with S3SERV:** Tailored for integrating S3 object store access within OLTP applications.

To get started, we recommend the following steps:

Step 1: Create Your Configuration

Set up the necessary configuration files on the NonStop Server. Since S3Utils is compatible with the AWS CLI configuration format, you can directly copy your existing AWS CLI configuration and credentials files to the NonStop Server.

Step 2: Verify Access to S3 Using S3COM

Test your setup with the S3COM utility. Ensure you can access the S3 object store using your credentials. Basic commands like 'LS' and 'CP' can be used to confirm the operational environment. For detailed instructions, refer to the section Using S3COM.

Step 3: Integrate Your OLTP Application with S3SERV

For OLTP-based access to the S3 object store, integrate your application with S3SERV. Guidance and detailed steps for this integration are available in Using S3SERV.

Configuring S3Utils

In order to use S3Utils, settings for accessing the S3 compatible object store must be configured. S3Utils configuration closely matches that of AWS CLI configuration. Information on AWS CLI configuration can be found here:

Configure the AWS CLI

The configuration information consists of the following:

- **Credentials** - identify which user is access the S3 object store. Credentials are created by the S3 service provider and provided to the S3 service user.
- **Other settings** - Any other settings required to complete the request. Some settings have default values if not provided, other have no effect if not provided.

Configuration and Credentials Preference

Credentials and settings may be supplied by multiple methods. Some methods take precedence over others. The settings take precedence in the following order, with 1 being the highest.

1. **Command line options** - Options supplied with the S3COM command such as REGION, or PROFILE. Or startup options for the S3SERV process.
2. **Environment variables** - Settings supplied via Guardian PARAMs or DEFINEs
3. **Credentials file** - The file AWSCRED. This file is search for in the following location, in this order:
   - The location specified by the AWS^SHARED^CREDENTIALS^FILE param.
   - The subvol specified by the =_DEFAULTS define.
   - The subvol of the running S3Utils program files.
4. **Configuration file** - The file AWSCONF. This file is search for in the following location, in this order:
   - The location specified by the AWS^CONFIG_FILE param.
   - The subvol specified by the =_DEFAULTS define.
   - The subvol of the running S3Utils program files.

Configuration and Credentials File Format

Separating configuration and credentials into distinct files enables multiple users or applications to utilize the same credentials while maintaining individual settings. It is

considered best practice to store credentials and shared settings in the credentials file, and all other settings in the configuration file.

The configuration and credentials files are structured into various sections. Each section, marked by a name, contains a set of related settings. These sections persist until a new section begins. It's possible to have multiple profiles and sections within these files.

The format of these files is straightforward:

- Sections are identified by names within brackets, like [default], [profile user1].
- Within a section, settings are formatted as setting_name=value.
- Lines beginning with a hashtag (#) are comments and are not processed.

In the AWSCONF configuration file, profile sections are named [default] or [profile user1]. In the AWSCRED credentials file, they are named [default] or [user1]. Avoid using 'profile' in the credentials file.

The following example shows related configuration and credentials files.

AWSCRED Sample

```
[default]
aws_access_key_id=AKIASDFRTHTJDFHD2EFF
aws_secret_access_key=l38ufsaHlrasdfe88af3oijad0N5evWAQ26K

[user1]
aws_access_key_id=AKIAXFJDHSRTSDFHSTRHF
aws_secret_access_key=l38ufsaHljasdfao88jadsfjoaweijfiQ26K
```

AWSCONF

```
[default]
region=us-west-2

[profile user1]
region=us-east-1
```

Environment Variables

Supported Environment Variables

AWS^CONFIG^FILE

Specifies the location of the file that the AWS CLI uses to store configuration profiles. The default path is `AWSCONF` in the `_DEFAULT` define location, or the `S3COM` program subvol if not found.

AWS^SHARED^CREDENTIALS^FILE

Specifies the location of the file that the AWS CLI uses to store access keys. The default path is `AWSCRED` in the `_DEFAULT` define location, or the `S3COM` program subvol if not found.

AWS^ACCESS^KEY^ID

Specifies the AWS access key used as part of the credentials to authenticate the command request.

AWS^SECRET^ACCESS^KEY

Specifies the AWS secret key used as part of the credentials to authenticate the command request.

REGION, AWS^REGION

Specifies the AWS Region to send requests to for commands requested using this profile.

- You can specify any of the Region codes available for the chosen service as listed in AWS Regions and Endpoints in the *Amazon Web Services General Reference*.

Configuration and Credentials File Settings

Supported Configuration Options

aws_access_key_id

Specifies the AWS access key used as part of the credentials to authenticate the command request. Although this can be stored in the `AWSCONF` file, we recommend that you store this in the `AWSCRED` file.

Can be overridden by the `AWS^ACCESS^KEY^ID` environment variable. You can't specify the access key ID as a command line option.

```
aws_access_key_id = AKIAIOSFODNN7EXAMPLE
```

aws_secret_access_key

Specifies the AWS secret key used as part of the credentials to authenticate the command request. Although this can be stored in the `AWSCONF` file, we recommend that you store this in the `AWSCRED` file.

Can be overridden by the `AWS^SECRET^ACCESS^KEY` environment variable. You can't specify the secret access key as a command line option.

```
aws_secret_access_key = wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
```

region

Specifies the AWS Region to send requests to for commands requested using this profile.

- You can specify any of the Region codes available for the chosen service as listed in AWS Regions and Endpoints in the *Amazon Web Services General Reference*.

You can override this value by using the `AWS^REGION` environment variable, `AWS^DEFAULT^REGION` environment variable, or the `REGION` command line option.

```
region = us-west-2
```

endpoint_url

Specifies an alternate endpoint for API requests.

Can be overridden by the `ENDPOINT^URL` environment variable.

```
endpoint_url = http://localhost:5000
```

For information on the possible AWS S3 configuration options, see Configure the AWS CLI.

Using S3COM

The S3COM command line interface (CLI) can be used manage S3 objects in a manner similar to the AWS CLI. Commands may be entered interactivity or supplied with a TACL macro or obey file.

S3COM supports the following commands and objects:

```
Commands:

!                ALLOW            CP               ENV
EXIT             FC               HELP             HISTORY
INFO             LS               MV               OBEY
PAGESIZE         RM               VALIDATE         VOLUME

Objects:

LICENSE
```

The commands used to interact with the S3 Object store are:

- CP - Copies a local file or S3 object to another location locally or in S3.
- LS - List S3 objects and common prefixes under a prefix or all S3 buckets.
- MV - Moves a local file or S3 object to another location locally or in S3.
- RM - Deletes a S3 object.

The S3 commands take one or two path arguments. The first path argument always represents the source reference. The second path, if provided, represents the destination. The path argument may be a *LocalPath* or *S3Uri*.

- LocalPath - Represents a Guardian Enscribe file name. This name may be fully or partially qualified. If partially qualified, the full name will be resolved using the value of the _DEFAULTS DEFINE.

- S3Uri - Represents the location of an S3 object, prefix, or bucket.

S3COM is self documenting using the HELP command. For example:

```
TACL> run s3com
S3COM - 1.0.0
Copyright (c) 2023 NuWave Technologies, Inc. All rights reserved.
S3COM 1-> help
Help is available on the following commands and objects:
Commands:

!                ALLOW            CP               ENV
EXIT             FC               HELP             HISTORY
INFO             LS               MV               OBEY
PAGESIZE         RM               VALIDATE         VOLUME

Objects:

LICENSE

Enter HELP <command> | <object> for more information.
S3COM 2-> help cp

CP Command

Copies a local file or S3 object to another location locally or in S3.

  CP <source>
     , <destination> [ ! ]
     [ , PROFILE <profile-name> ]
     [ , REGION <s3-region> ]
     [ , DEBUG ]

<source>

  The source local path or S3Uri.

<destination> [ ! ]

  The destination local path or S3Uri.

PROFILE <profile-name>

  Use a specific profile from your credential file.

REGION <s3-region>
```

```
    The region to use. Overrides config/env settings.

DEBUG

  Enable debug logging.

S3COM 3->
```

CP - Copy

CP Command

Copies a local file or S3 object to another location locally or in S3.

CP <source>
, <destination> [!]
[, PROFILE <profile-name>]
[, REGION <s3-region>]
[, DEBUG]

[, DRYRUN]

[, CA-BUNDLE <local-path>]

[, NO-VERIFY-SSL]

<source>

The source local path or S3Uri.

<destination> [!]

The destination local path or S3Uri.

PROFILE <profile-name>

Use a specific profile from your credential file.

REGION <s3-region>

The region to use. Overrides config/env settings.

DEBUG

Enable debug logging.

DRYRUN

Displays the operations that would be performed using the specified command without actually running them.

CA-BUNDLE <local-path>

The CA certificate bundle to use when verifying SSL certificates. Over-rides config/env settings.

NO-VERIFY-SSL

By default, the AWS CLI uses SSL when communicating with AWS services. For each SSL connection, the AWS CLI will verify SSL certificates. This option overrides the default behavior of verifying SSL certificates.

Examples

```
tacl> run s3com cp myfile, s3://bucket-name/nonstop-files/myfile

tacl> run s3com cp s3://bucket-name/nonstop-files/myfile, myfile !

tacl> run s3com cp s3://bucket-name/nonstop-files/myfile, s3://bucket-name/nonstop-
files/myfile2

tacl> run s3com cp s3://bucket-name/nonstop-files/myfile, s3://bucket-name/nonstop-
files/myfile2, DRYRUN

tacl> run s3com cp s3://bucket-name/nonstop-files/myfile, s3://bucket-name/nonstop-
files/myfile2, CA-BUNDLE $VOLUME.SUBVOL.FILENAME

tacl> run s3com cp s3://bucket-name/nonstop-files/myfile, s3://bucket-name/nonstop-
files/myfile2, NO-VERIFY-SSL
```

LS - List

List S3 objects and common prefixes under a prefix, or all S3 buckets.

LS [<S3Uri>]
[, RECURSIVE]
[, PROFILE <profile-name>]
[, REGION <s3-region>]
[, DEBUG]

[, CA-BUNDLE <local-path>]

[, NO-VERIFY-SSL]

<S3Uri>

An optional S3Uri. If omitted, all user owned buckets will be listed.

RECURSIVE

Command is performed on all files or objects under the specified directory or prefix.

PROFILE <profile-name>

Use a specific profile from your credential file.

REGION <s3-region>

The region to use. Overrides config/env settings.

DEBUG

Enable debug logging.

CA-BUNDLE <local-path>

The CA certificate bundle to use when verifying SSL certificates. Over-rides config/env settings.

NO-VERIFY-SSL

By default, the AWS CLI uses SSL when communicating with AWS services. For each SSL connection, the AWS CLI will verify SSL certificates. This option overrides the default behavior of verifying SSL certificates.

Examples

```
tacl> run s3com ls
2023-01-18 23:13:00 bucket-1
2023-02-25 20:43:00 bucket-2

tacl> run s3com s3://bucket-1/
2023-12-06 22:13:00          15 myfile.txt

tacl> run s3com ls , CA-BUNDLE $VOLUME.SUBVOL.FILENAME

tacl> run s3com ls , NO-VERIFY-SSL
```

MV - Move

Moves a local file or S3 object to another location locally or in S3.

MV <source>
, <destination> [!]
[, PROFILE <profile-name>]
[, REGION <s3-region>]
[, DEBUG]

[, CA-BUNDLE <local-path>]

[, NO-VERIFY-SSL]

<source>

The source local path or S3Uri.

<destination> [!]

The destination local path or S3Uri.

PROFILE <profile-name>

Use a specific profile from your credential file.

REGION <s3-region>

The region to use. Overrides config/env settings.

DEBUG

Enable debug logging.

CA-BUNDLE <local-path>

The CA certificate bundle to use when verifying SSL certificates. Over-rides config/env settings.

NO-VERIFY-SSL

By default, the AWS CLI uses SSL when communicating with AWS services. For each SSL connection, the AWS CLI will verify SSL certificates. This option overrides the default behavior of verifying SSL certificates.

Examples

```
tacl> run s3com mv myfile, s3://bucket-name/nonstop-files/myfile

tacl> run s3com mv s3://bucket-name/nonstop-files/myfile, myfile !

tacl> run s3com mv s3://bucket-name/nonstop-files/myfile, s3://bucket-name/nonstop-
files/myfile2

tacl> run s3com mv s3://bucket-name/nonstop-files/myfile, s3://bucket-name/nonstop-
files/myfile2, CA-BUNDLE $VOLUME.SUBVOL.FILENAME

tacl> run s3com mv s3://bucket-name/nonstop-files/myfile, s3://bucket-name/nonstop-
files/myfile2, NO-VERIFY-SSL
```

RM - Remove

Deletes an S3 object.

RM <S3Uri>
[, PROFILE <profile-name>]
[, REGION <s3-region>]
[, DEBUG]

[, CA-BUNDLE <local-path>]

[, NO-VERIFY-SSL]

<S3Uri>

The object to be deleted.

PROFILE <profile-name>

Use a specific profile from your credential file.

REGION <s3-region>

The region to use. Overrides config/env settings.

DEBUG

Enable debug logging.

CA-BUNDLE <local-path>

The CA certificate bundle to use when verifying SSL certificates. Over-rides config/env settings.

NO-VERIFY-SSL

By default, the AWS CLI uses SSL when communicating with AWS services. For each SSL connection, the AWS CLI will verify SSL certificates. This option overrides the default behavior of verifying SSL certificates.

Examples

```
tacl> run s3com rm s3://bucket-name/nonstop-files/myfile

tacl> run s3com rm s3://bucket-name/nonstop-files/myfile, CA-BUNDLE
$VOLUME.SUBVOL.FILENAME

tacl> run s3com rm s3://bucket-name/nonstop-files/myfile, NO-VERIFY-SSL
```

Using S3CLIENT

The S3CLIENT process can be deployed as a standalone server or TS/MP serverclass and provides integration with S3 block storage services via a simple IPM based application programming interface. This allows existing or new NonStop applications to integrate with

S3 storage as part of their online transaction processing. This section contains information needed to integrate NonStop applications with S3CLIENT.

If you're looking for information on how to install or manage S3Utils refer to the Administrator Guide. Before reading the Developer Guide it is recommended that you read the Introduction to S3Utils.

Invoking S3CLIENT Requests

Once the S3CLIENT process is configured and running, any NonStop application can access S3 by sending the process a predefined IPM request and waiting for the response. Each IPM request includes a header containing a 2-byte request code. The request code indicates which S3CLIENT operation should be invoked by the IPM request. The remainder of the IPM consists of fields that correspond to the parameters of the selected S3CLIENT request.

For example, the `s3-copy-object-rq` DDL for the "copy S3 object" request is used to invoke a copy one S3 object to another:

- Initialize the s3-copy-object-rq structure to 0 values.
- Set the s3-copy-object-rq.rq-hdr.rq-code to S3-RQ-COPY-OBJECT.
- Set the s3-copy-object-rq.source-object to the name of the S3 object to copy from.
- Set the s3-copy-object-rq.target-object to the name of the S3 object or bucket to copy to.
- Send the IPM to the S3CLIENT process using the appropriate Guardian system procedure call.

Once the S3CLIENT process completes the S3 service request, it returns the response in an IPM. This response IPM also contains a standard S3CLIENT header which contains a 2-byte response code. If the response code is zero, the remainder of the IPM consists of the appropriate reply IPM. If the response code is non-zero, then the application should treat the response as an error.

Error Handling

Every reply from the S3Serv process includes the S3Serv reply header, followed by user data if the request was successful or error information if an error occurred, as shown in previous subsections.

The `rp-hdr.rp-code` field should be interpreted as follows:

- If `rp-code` contains S3-SUCCESS, the request was successful and user data follows the S3Serv reply header.
- If `rp-code` contains S3-INFO, the request was successful but S3Object truncation occurred while the response was being processed. In this case the contents of the INFO-CODE field will be set to S3-INFO-TRUNCATION.
- If `rp-code` contains S3-ERROR then an error occurred while processing the request. The error-code field in the S3Serv s3-error-rp reply indicates the cause of the error.

IPM Definitions

S3Serv supports IPM definitions for the following operations:

- Copy S3 objects.
- Delete S3 objects.
- Put a BLOB contained in the request IPM as an S3 Object.
- Get an S3 Object and return it as a BLOB in the reply IPM.

Every request and reply exchanged with S3Serv includes a header. Request and reply IPM layouts are described using NonStop DDL-style notation. Pre-generated formats are available in NonStop C and COBOL (1974).

Request Header

The request header identifies the specific request to S3Serv, along with a timeout value. The request header contains a code to identify the request.

Request Code

```
DEF S3-RQ-CODE-ENUM                    TYPE ENUM BEGIN.
  89  S3-RQ-COPY-OBJECT                VALUE 1.           ! Copy an S3 object.
  89  S3-RQ-GET-BLOB                   VALUE 2.           ! Retrieve an S3 object
as a BLOB.
  89  S3-RQ-PUT-BLOB                   VALUE 3.           ! Store a BLOB as an S3
object.
  89  S3-RQ-REMOVE-OBJECT              VALUE 4.           ! Delete an S3 object.
END.
```

Request Header

```
DEF S3-HEADER-RQ.
  02  RQ-CODE                          TYPE S3-RQ-CODE-ENUM.
```

```
  02  RESERVED                          TYPE CHARACTER 30.  ! reserved, must be
binary 0.
END
```

Reply Header

The S3Serv reply header is the same for every reply. The reply header contains a code to indicate the result of the request.

Reply Code

```
DEF S3-RP-CODE-ENUM                     TYPE ENUM BEGIN.
  89  S3-SUCCESS                        VALUE 0.            ! Success.
  89  S3-INFO                           VALUE 1.            ! The request completed
with warnings.
  89  S3-ERROR                          VALUE 2.            ! An error occurred.
END
```

Info Code

```
DEF S3-INFO-CODE-ENUM                   TYPE ENUM BEGIN.
  89  S3-INFO-TRUNCATION                VALUE 100.          ! Get BLOB truncation
occurred.
END
```

Reply Header

```
DEF S3-HEADER-RP.
  02  RP-CODE                          TYPE S3-RP-CODE-ENUM.
  02  RP-INFO                          TYPE S3-INFO-CODE-ENUM.
  02  RESERVED                         TYPE CHARACTER 28.
END
```

Error Reply

If the reply code value `S3-ERROR`, then an Error Reply is returned instead of a request-specific reply. The list of error codes and the layout of the Error Reply are shown below.

Error Codes

```
DEF S3-ERROR-ENUM                       TYPE ENUM BEGIN.
```

```
  89  S3-ERROR-HEADER                 VALUE 1.
  89  S3-ERROR-REQUEST-CODE           VALUE 2.          ! Unrecognized request
code.
  89  S3-ERROR-INVALID-SOURCE-NAME    VALUE 3.          ! The source name is
invalid.
  89  S3-ERROR-INVALID-TARGET-NAME    VALUE 4.          ! The target name is
invalid.
  89  S3-ERROR-OBJECT-NOT-FOUND       VALUE 5.          ! The referenced S3
object was not found.
  89  S3-ERROR-FILE-NOT-FOUND         VALUE 6.          ! The referenced NonStop
file was not found.
  89  S3-ERROR-ACCESS-VIOLATION       VALUE 7.          ! An S3 access violation
occurred.
  89  S3-ERROR-OBJECT-NOT-EMPTY       VALUE 8.          ! Attempt to delete an
non-empty S3 bucket.
  89  S3-ERROR-REQ-NOT-SUPPORTED      VALUE 9.          ! Attempt to copy a
NonStop file to another Nonstop file.
  89  S3-ERROR-S3-SVC-ERROR           VALUE 10.         ! An error occurred with
the S3 service.
  89  S3-ERROR-S3-SVC-UNAVAILABLE     VALUE 11.         ! The S3 service is not
available.
  89  S3-ERROR-FILE-IO-ERROR          VALUE 12.         ! An error occurred
reading or writing to a NonStop file.
  89  S3-ERROR-INVALID-COUNT          VALUE 13.         ! An invalid item count
or blob length was specified.
  89  S3-ERROR-INTERNAL-ERROR         VALUE 500.        ! An internal error
occurred in S3Serv.
END
```

S3 Object Copy Operations

Object copy operations copy the contents of an S3 objects to or from NonStop Enscribe files. The “copy object” request is used to copy an S3 object to the NonStop, copy an S3 object to another name or bucket in S3, or to copy a file from the NonStop to an S3 object.

Copy Request

```
DEF S3-COPY-OBJECT-RQ.
  02  RQ-HDR                         TYPE S3-HEADER-RQ.
  02  SOURCE-OBJECT                  TYPE CHARACTER 256.
  02  TARGET-OBJECT                  TYPE CHARACTER 256.
END
```

Copy Reply

```
DEF S3-COPY-OBJECT-RP.
  02  RP-HDR                          TYPE S3-HEADER-RP.
END
```

Copy Request Processing

S3Serv evaluates the supplied names to determine the copy operation to perform:

Source Name	Target Name	Operation
S3 URI	S3 URI	Copy an S3 object to the S3 destination.
S3 URI	NonStop file name	Copy the S3 object to the target NonStop file, replacing the file if it already exists.
NonStop file name	S3 URI	Copy the contents of the designated NonStop file to the target S3 object, replacing it if it already exists.
NonStop file name	NonStop file name	Not supported.

BLOB Copy Operations

BLOB (Binary Large Object) copy operations copy the contents of S3 objects to or from the S3SERV request or response IPM.

Get S3 Object

The Get Object request is used to retrieve the content of an S3 object and return its content to the requester in the reply IPM. S3SSERV retrieves the S3 Object contents and stores as much of it as will fit `TARGET-BLOB`. The data length is returned `TARGET-BLOB-LEN`.

Get S3 Object Request

```
DEF S3-GET-BLOB-RQ.
```

```
  02  RQ-HDR                          TYPE S3-HEADER-RQ.
  02  SOURCE-OBJECT                   TYPE CHARACTER 256.
END
```

Get S3 Object Reply

```
DEF S3-GET-BLOB-RP.
  02  RP-HDR                          TYPE S3-HEADER-RP.
  02  TARGET-BLOB-LEN                 TYPE BINARY 32 UNSIGNED.
  02  TARGET-BLOB                     TYPE CHARACTER 32000.
END
```

Put S3 Object

The “put object” request is used to store content received in the IPM to an S3 object. S3SERV attempts to store `SOURCE-BLOB-LEN` bytes from `SOURCE-BLOB` to the S3 Object.

Put S3 Object Request

```
DEF S3-PUT-BLOB-RQ.
  02  RQ-HDR                          TYPE S3-HEADER-RQ.
  02  TARGET-OBJECT                   TYPE CHARACTER 256.
  02  SOURCE-BLOB-LEN                 TYPE BINARY 32 UNSIGNED.
  02  SOURCE-BLOB                     TYPE CHARACTER 32000.
END
```

Put S3 Object Reply

```
DEF S3-PUT-BLOB-RP.
  02  RP-HDR                          TYPE S3-HEADER-RP.
END
```

Removing S3 Objects

Remove operations delete S3 objects from the object store.

Remove S3 Object Request

```
DEF S3-REMOVE-OBJECT-RQ.
  02  RQ-HDR                          TYPE S3-HEADER-RQ.
  02  TARGET-OBJECT                   TYPE CHARACTER 256.
END
```

Remove S3 Object Reply

```
DEF S3-REMOVE-OBJECT-RP.
  02  RP-HDR                          TYPE S3-HEADER-RP.
END
```

Listing S3 Objects

Listing operations list information about S3 Buckets and Objects.

List Objects Request

```
DEF S3-LIST-OBJECT-RQ.
  02  RQ-HDR                          TYPE S3-HEADER-RQ.
  02  OBJECT-NAME                     TYPE CHARACTER 256.
  02  MAX-RP-COUNT                    TYPE BINARY 32 UNSIGNED.
  02  RECURSIVE                       TYPE BINARY 16.
END
```

List Objects Reply

```
DEF S3-CONTENT-ITEM.
  02  CONTENT-NAME                    TYPE CHARACTER 256.
  02  CONTENT-ETAG                    TYPE CHARACTER 256. !Includes dbl-quotes
  02  CONTENT-LAST-MODIFIED-GMT       TYPE BINARY 64.     !TNS julian timestamp
(GMT)
  02  CONTENT-SIZE                    TYPE BINARY 64.     !Unsigned
  02  STORAGE-CLASS                   TYPE CHARACTER 32.
END

DEF S3-PREFIX-ITEM.
  02  PREFIX-NAME                     TYPE CHARACTER 256.
  02  FILLER                          TYPE CHARACTER 304.
END
```

```
DEF S3-BUCKET-ITEM.
  02  BUCKET-NAME                    TYPE CHARACTER 256.
  02  BUCKET-CREATION-GMT            TYPE BINARY 64.     !TNS julian timestamp
(GMT)
  02  FILLER                         TYPE CHARACTER 296.
END

DEF S3-LIST-ITEM.
  02  ITEM-TYPE                      TYPE CHARACTER 1.   !"B": bucket, "C":
content, "P": prefix
  02  FILLER                         TYPE CHARACTER 1.
  02  ITEM                           TYPE CHARACTER 560.
  02  CONTENT                        TYPE S3-CONTENT-ITEM
                                      REDEFINES ITEM.
  02  BUCKET                         TYPE S3-BUCKET-ITEM
                                      REDEFINES ITEM.
  02  PREFIX                         TYPE S3-PREFIX-ITEM
                                      REDEFINES ITEM.
END

DEF S3-LIST-OBJECT-RP.
  02  RP-HDR                         TYPE S3-HEADER-RP.
  02  ITEM-COUNT                     TYPE BINARY 16.
  02  ITEM-LIST                      TYPE S3-LIST-ITEM
                                      OCCURS 0 TO 1000 TIMES
                                      DEPENDING ON ITEM-COUNT.
END
```

Administrator Guide

This section contains information on installing and managing S3Utils.

- Installing S3Utils
- S3SERV Configuration

Installing S3Utils

These instructions apply to new installations. Note that you may install the product more than once (in separate subvolumes) on the same NonStop system if you wish.

> Before proceeding with the installation process refer to the Release Notes for important information about the current release of the software, particularly the Installation Prerequisites.

The software is delivered in a single PAK archive containing the following files:

- S3COM - The S3Utils CLI program.
- S3SERV - The S3Utils server program.
- S3SRVDDL - DDL definitions for use with S3SERV.
- S3SRVCOB - COBOL definitions for use with S3SERV.
- S3SRVH - C language definitions for user with S3SERV.
- ZSTARTPW - Sample CLIENT process Pathway configuration startup macro

The process is performed from a TACL prompt. These are the steps:

- Obtain a Product License
- Download the Release Package
- Install the Release Package
- Install the Product License

Obtain a Product License

A license is required to run S3Utils. If this is the first time you are installing the product on your NonStop system, you may obtain a free 30-day trial license at the Trial License Center. Otherwise, a license will be provided to you upon purchase of the product. Do not request a trial license until you are ready to begin using the product, as the 30-day timer

starts on the day the license is issued. Note that you may install S3Utils without a license, however the product will not function until a valid license is installed.

Download the Release Package

The software is distributed as a single PAK archive file. The software may be downloaded from our Software Download Center. Transfer the PAK file to your NonStop system using binary transfer mode.

Install the Release Package

Unpak the release package PAK archive file into an empty subvolume.

```
> UNPAK <pakfile> ($*.*.*), vol <installation-subvolume>, myid, listall
```

where <pakfile> is the name of the release package PAK archive file you transferred earlier.

Install the Product License

Transfer the license file you received by email to your NonStop system using text or 'ascii' file transfer mode. The license must be stored in an 'edit' file in one of the following locations:

- the file LICENSE in the S3Utils installation subvolume
- the file LICS3UTL in the S3Utils installation subvolume
- the file $SYSTEM.NUWAVE.LICS3UTL

The S3Utils programs validate the license at startup and periodically while running. Any license errors will be written to the event log or the TACL command line. Note that if the license file content is altered in any way, the license will become invalid. For more information see Product Licensing.

S3SERV Configuration

- Using Configuration Files
- Starting the S3SERV Process
- S3SERV Command Line Reference

Using Configuration Files

Some process configuration options support dynamic configuration through the use of configuration files. Instead of providing the configuration options on the process startup command line, the configuration file name is supplied and the configuration options are read from the file. The --monitor option can be used to enable change monitoring for the configuration files. When changes are detected the options are reloaded from the file. Configuration files are Enscribe EDIT files and use the "INI" file format. The following commands may use configuration files:

Command	Description
--log	Logging configuration

The following sections describe the configuration options for these commands:

- Configuration File Format
- Log Configuration

Configuration File Format

Configuration files are Enscribe EDIT files and uses "INI" file format.

Options

Options consist of name, value pairs delimited by an equals sign (=). The value may be enclosed in quotes to preserve whitespace when necessary.

Sections

Parameters are grouped into named sections. A section header consists of the a section name enclosed in square brackets ([ ]). All parameters found after a section header and before the next section header or end of file are associated with that section header. Section headers may include a process name to indicate that the section applies to a specific process. In this case the section name consists of the process name and section name delimited by a colon (:).

Comments

A hash (#) character indicates the start of a comment, which continues to the end of the line. All text between the hash and end of line is ignored.

Blank Lines

Blank lines are ignored.

Example

```
# This section configures generic logging options
[log]
file=zzlog
level=info   # log at info level.
format=text
```

Log Configuration

Dynamic logging configuration is activated using the --log startup option and specifying the log configuration file location. *Note that the file name specification is preceded by a '+' character.*

```
run S3SERV --log +$vol.subvol.logcfg
```

Configuration Reference

The log configuration must be preceded by a *[log]* header

Option	Descriptione
file	The name of the log file. If not fully qualified then the value of the _DEFAULTS define is used to complete the file name.
level	The level value may be "error", "warning", "info", or "debug" and controls the type of information that is output to the log destination. The "error" level produces the least output while the "trace" level produces the most output. The default value is "info".

Option	Descriptione
format	The format value may be "text" indicating that the log events should be output as text strings or "event" indicating that the log events should be output in EMS event format.. If not specified the default value is "text".

Examples

```
# Log to a file
[log]
file=$data1.logs.zzlog
format=text
level=info
```

```
# Log to $0
[log]
file=$0
format=event
level=info
```

Starting the S3SERV Process

The S3SERV process runs as either a Pathway serverclass or a standalone Guardian server process. The advantage of the Pathway option is Pathway's server process management and load-balancing features. In either case, S3SERV supports a number of startup *command-line-options* that control various aspects of its operation. The same options are supported in both Pathway and standalone configurations, specified as server PARAMs or command-line-options, respectively. The object file name for the S3SERV process is “S3SERV” and is located in the installation subvolume.

Starting S3SERV a Standalone Server Process

S3SERV may be started by running the S3SERV program from TACL.

```
tacl> run s3serv / run-options / command-line-options
```

Examples

The example starts the S3SERV process as $MYS3 in CPU 0.

```
tacl> run s3serv / name $mys3, cpu 0 / --log $zhome info --standalone
```

More Information

For complete information about *run-options* and *command-line-options* see S3SERV Command Line Options.

Starting S3SERV as a Pathway Server

A sample S3SERV startup macro, ZSTARTPW, is included with the release. Before using for the first time, copy or rename the file to a new name (e.g. FUP DUP ZSTARTPW, STARTPW) so it won't be overwritten by future releases. The sample macro should be suitable for most installations, but it will be necessary to modify certain parameters for your environment. To use the macro: RUN STARTPW

The S3SERV process may be configured as a Pathway server, which allows you to take advantage of Pathway's process management and load-balancing features. The configuration options are the same as those for standalone process configuration, however in the Pathway environment individual options are supplied as server PARAMs. The following is a sample pathway server configuration. Note that the server settings here (createdelay, deletedelay, maxservers, etc.) are merely examples, not requirements or recommendations.

```
reset server
set server cpus 0:1
set server createdelay 0 secs
set server deletedelay 60 secs
set server highpin on
set server linkdepth 1
set server maxservers 6
set server maxlinks 20
set server numstatic 0
set server program S3SERV
set server tmf on
set server debug off
set server param log "$0 info text"
add server S3-SERVICE
```

You may also provide the options in a command file by setting the SERVER STARTUP value to @<command-file>. The following is a sample pathway configuration using this technique.

```
reset server
set server cpus 0:1
set server createdelay 0 secs
set server deletedelay 60 secs
set server highpin on
set server linkdepth 1
set server maxservers 6
set server maxlinks 20
set server numstatic 0
set server program S3SERV
set server tmf on
set server debug off
set server startup @config
add server S3-SERVICE
```

S3SERV Command Line Reference

S3SERV Program Options

@<command-file>

Reads command line options from <command-file>. Options specified on the command line override any duplicates specified in the file. At most, one '@' option may used. The file itself cannot contain an '@' option (i.e., no nesting).

--license <file-name>

The name of an existing edit file containing the S3Utils product license. If this option is omitted, the license file is located according to Product Licensing rules.

--log [{ <destination> | * } [level [format]] | +<log-config-file>]

Specifies the process log location, the level, and the log event format, or the location of a log configuration file. The destination value may be a process name, a file name, or the asterisk (*) character. If the asterisk is used then the log output is directed to the home term of the process. The level value may be "error", "warning", "info", or "debug" and controls the type of information that is output to the log destination. The "error" level produces the least output while the "debug" level produces the most output. The format value may be "text" indicating that the log events should be output as text strings or "event" indicating that the log events should be output in EMS event format. If omitted, the default is "--log * info text". See Using Configuration Files for information about logging configuration files.

--standalone

The presence of this options causes the process to ignore close messages. In order to function properly in the Pathway environment, the CLIENT process, by default, matches open and close messages and terminates when all clients have closed the process. This option can be used to prevent the process from terminating when it is run as a standalone process.

Remarks

All command-line-option names and values are case-insensitive except where noted. If multiple occurrences of the same command line parameter are encountered, the setting of the last occurrence is used.

When the -log option format is set to event, EMS events will be sent to the output device with the following EMS Subsys ID:

Z-OWNER	"NUWAVE"
Z-NUMBER	6
Z-VERSION	Product major version

Release Notes

The Release Notes contain important information about the software. Review all sections before installing or upgrading the product.

Installation Prerequisites

S3Utils requires the following HP NonStop Server software:

- One of the following NonStop System RVUs:
  - TNS/E: H06.09 or later, J06.03 or later.
  - TNS/X: L15.08 or later.
- HP NonStop TCP/IP

Although not required, S3Utils integrates with these products, if installed:

- HP NonStop TS/MP (Pathway). Large message support is available when supported by TS/MP.
- HP Data Definition Language (DDL or DDL2).
- HP TCP/IP Parallel Library, or TCP/IPv6. S3Utils currently supports IPv4 mode only.

Upgrade Considerations

None.

Cumulative Change Log

1.0.0 - 4 December 2024

- This is the first GA release of S3Utils™

Product Licensing

S3Utils requires the installation of a license. The license is stored in an edit file on the NonStop Server that is read by the product in order to validate the installation. A license is a multi-line text block similar to that shown below:

```
-----BEGIN LICENSE-----
product=s3utils
systemNumber=012345
expiration=6/1/2016
supportExpiration=6/1/2016
restrictedLicense=no
signature=172
T2M23bpqxRzyEgt2cETwRb8j5DMkGdVzyM2WZTRsssoIZXooc6
ROAtJdOevToYeX10/UPCSVZJynUZO/uf7MSxVn2FnA5LH1g87g
3F9wDvajQNdiRxrX6rHThNovxH+dKOXKb+nzkGZx47PXF4izA1
W4GJhmAJi9n7z6x5WZEyE=
-----END LICENSE-----
```

When the license is received from NuWave Technologies it must be copied to an edit file on the NonStop System *exactly* as provided. The product programs search for the license file in the following locations and in the following order:

- The file "LICENSE" in the S3Utils program subvolume
- The file "LICS3UTL" in the S3Utils program subvolume
- The file "$SYSTEM.NUWAVE.LICS3UTL"

License Expiration

The license contains two expiration date fields to be aware of:

expiration

This field contains the date that the license will expire, or "none" if the license never expires. For licenses that expire, once the expiration date has passed the software will continue to function but will log license violation errors to the process log. For licenses that don't expire, the software will run indefinitely.

supportExpiration

This field contains the date that your current software maintenance agreement expires. Once this expiration date has passed, the software will continue to run with full

functionality (subject to the *expiration* date), but you may not upgrade the software to a version that was released after the *supportExpiration* date. Software with a release date after the *supportExpiration* date will shutdown immediately with a license error. The release date for a particular software version can be found in the Upgrade Guide or the product VPROC.

> When you renew your support agreement we will email you a new license. You do not need to install the new license until the next software upgrade.

Third Party Software Licenses

S3Utils incorporates code from several open-source projects. The licenses for these projects are included below:

- Apache Portable Runtime
- Apache Portable Runtime Utility Library
- Jansson JSON Toolkit
- json-schema-generator
- cURL
- libxml2 XML Toolkit
- OpenSSL Toolkit
- Perl Compatible Regular Expressions

Apache Portable Runtime

Copyright (c) 2011 The Apache Software Foundation.

This product includes software developed by The Apache Software Foundation (http://www.apache.org/).

Portions of this software were developed at the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign.

This software contains code derived from the RSA Data Security Inc. MD5 Message-Digest Algorithm.

This software contains code derived from UNIX V7, Copyright(C) Caldera International Inc.

http://apr.apache.org

http://www.apache.org/licenses/LICENSE-2.0

Apache Portable Runtime Utility Library

Copyright (c) 2011 The Apache Software Foundation.

This product includes software developed by The Apache Software Foundation (http://www.apache.org/).

Portions of this software were developed at the National Center

for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign.

This software contains code derived from the RSA Data Security Inc.

MD5 Message-Digest Algorithm, including various modifications by Spyglass Inc., Carnegie Mellon University, and Bell Communications Research, Inc (Bellcore).

http://apr.apache.org

http://www.apache.org/licenses/LICENSE-2.0

Jansson JSON Toolkit

Copyright (c) 2009-2014 Petri Lehtinen <petri@digip.org>

http://www.digip.org/jansson/

https://github.com/akheron/jansson/blob/master/LICENSE

---

json-schema-generator

Copyright (c) 2014 krg7880

https://github.com/krg7880/json-schema-generator/blob/master/LICENSE

---

cURL

Copyright (c) 1996 - 2014, Daniel Stenberg, <daniel@haxx.se>.

http://curl.haxx.se/libcurl/

http://curl.haxx.se/docs/copyright.html

---

libxml2 XML Toolkit

Copyright (C) 1998-2012 Daniel Veillard. All Rights Reserved.

http://www.xmlsoft.org/

http://opensource.org/licenses/mit-license.html

OpenSSL Toolkit

Perl Compatible Regular Expressions

S3Utils Documentation

Release 1.0.0

Document

Classification: Proprietary and Confidential

Type: Customer Documentation

Author: NuWave Technologies, Inc.

Document Version: 1.0.0

Release Date: 4 Dec 2024

Contact: NuWave Technologies

Address 301 Edgewater Pl, STE 100, Wakefield, MA 01880

Email: support@nuwavetech.com

Web: https://support.nuwavetech.com/support/home

support@nuwavetech.com

NuWave