

1 Introduction

ncfc is a high-performance archival and synchronization tool for large file trees. It compresses data on the fly, optionally encrypts it, and stores archives on local disk or cloud object storage. It is aimed at multi-terabyte sensor, camera, and engineering datasets.

2 Features

- Domain-oriented compressors for multi-channel 8/10/12/16-bit and packed raw data
- Automatic per-segment format detection; optional fixed format and row width for known layouts
- Multi-threaded scan, compression, encryption, and I/O pipeline
- Optional XChaCha20-Poly1305 authenticated encryption (AEAD)
- Per-file versioning with append-oriented updates
- Separate binary index: fast listing and metadata queries without reading bulk data
- To reduce cloud object count and request costs, files are automatically combined and archived in large fixed-size chunks (about 1 GB)
- Incremental index sync after each completed part; safe resume after interruption
- Subcommand CLI: add, get, list, daemon
- Smart glob include/exclude patterns for deep directory trees
- Optional per-operation metadata tags; JSON listing for automation
- Stream extraction to stdout (`--pipe`) for pipeline workflows

3 Storage backends

- Local paths and mounted filesystems, including NFS
- AWS S3, Azure Blob and Google Cloud Storage
- S3-compatible endpoints, including MinIO and Ceph
- SFTP

4 System requirements

- Linux, 64-bit
- CPU with AVX2
- About 200 MB RAM per thread

5 Encryption and keys

- Passphrase-protected master key file
- Background daemon caches the key in memory to avoid repeated passphrase entry

6 Compression levels

- Level 0: no compression, maximum throughput
- Level 1: fast compression, when speed is more important than compression ratio
- Level 2: stronger compression, default for long-term retention

7 Typical results

- 30–60% smaller storage footprint on raw sensor-style data compared to uncompressed files (data-dependent)
- 100-200 MB/s compression speed per thread
- Compression ratio and throughput depend on content, CPU, storage, and network; benchmark on representative data is recommended

8 Licensing models

- **single**: one deployment instance
- **site**: any number of instances at one site
- **project**: use within a named project
- **company**: company-wide use

9 Contact

For more information:
info@nanocode.com
+49 6226 7876840