



Elastifile

ECFS v2.7.5 General Availability

Build 2.7.5.4

CloudConnect v2.7 General Availability

Build 2.7.1.0

Release Notes

[Product Overview](#)

[Implementation Architecture](#)

[Enterprise Grade Management and Storage Feature Set](#)

[Key Capabilities Provided in Version 2.7.0](#)

[Supported Deployments & System Requirements](#)

[Key Capabilities Provided in Elastifile CloudConnect Version 2.5.1.](#)

[Elastifile CloudConnect System Requirements](#)

[Fixed Issues in ECFS 2.1.0.10](#)

[Fixed Issues in ECFS 2.5.0.6](#)

[Fixed Issues in ECFS 2.5.0.13](#)

[Fixed Issues in ECFS 2.5.1.18](#)

[Fixed Issues in ECFS 2.5.2.2](#)

[Fixed Issues in ECFS 2.7.0.13](#)

[Known Bugs & Limitations](#)

[The following section lists known limitations in Elastifile Version 2.7.0.](#)

[General Limitations](#)

[High Availability Limitations](#)

[Asynchronous Disaster Recovery](#)

[Protocol Limitations](#)

[Datapath Limitations](#)

[Backend Store Limitation](#)

[Networking Limitation](#)

[Management & Control Limitations](#)

[CloudConnect version 2.5.1.1 known issues](#)

[General Limitations](#)

Product Overview

The Elastifile Cloud File System (ECFS) is an all-flash, primary storage solution for Next-Gen SDDC (Software Defined Data-Centers) and Cloud Deployments. By augmenting commodity hardware or default cloud instances, with intelligent software, Elastifile designed an all-flash elastic scale-out software defined storage solution. The Elastifile Cloud File System supports both transactional and batch workloads with unparalleled efficiency. Combining the flexibility of a software only solution with advanced all flash storage optimizations, the Elastifile Cloud File System delivers flash performance while reducing storage costs. The Elastifile storage solution is delivered as a **software only** (BYOH - Bring Your Own Hardware), all-Flash, **distributed file-system** serving as an enterprise grade scale out primary storage. The Elastifile Cloud File System converges compute and storage resources, scales linearly, and expands elastically providing tens of millions of IOPS at consistently low latency.

Elastifile CloudConnect enables efficient on-going transport of data from on-premises file systems into S3 compatible Cloud. CloudConnect mounts any legacy file system on the source side, and sends ongoing space-efficient compressed and deduplicated data for migration and retention in the cloud as object storage.

CloudConnect identifies changes in the file system and supports versioning of stored file systems, while sending only changed data between versions. Data can be kept in the cloud in object format, viewed or read in-place as a file system, restored back to the source file system, or can also be checked out on demand in the cloud into ECFS (Elastifile Cloud File System) for use as a high performance distributed file system that retains the original file structures.

Implementation Architecture

The Elastifile Cloud File System is deployed by placing an Elastifile controller on each physical server or cloud instance in an Elastifile cluster. Each Elastifile controller performs multiple functions, such as exposing NFS to the guests or external clients, and providing metadata, control services, and/or managing physical flash resources on the server where Elastifile controller resides.

The controllers in the Elastifile cluster replicate and distribute the data. The controllers expose a unified distributed file system accessible through any of the cluster nodes. The controller is managed through a central management pane.

Enterprise Grade Management and Storage Feature Set

The Elastifile Cloud File System provides enterprise grade, highly available management and storage features optimized for an all flash deployment. Files are grouped into data containers. Each data container's attributes can be managed by policy. Advanced storage features can be assigned to data containers using policies, including data reduction functionality such as deduplication and compression, as well as system resiliency through data replication. Data containers can be exposed as NFS shares through multiple exports providing the enterprise with full flexibility to deploy Elastifile NFS shares either directly to guest VMs, containers or any bare-metal client. Logical Data protection using Read-Only snapshots and Disaster Recovery protection using Asynchronous Replication can be defined at the Data Container level.

Key Capabilities Provided in Version 2.7.5

Elastifile Version 2.7.5 is intended for enterprise deployment and in cloud deployment providing the following key capabilities:

1. Dedicated storage with External NFS deployment model
 - a. High availability and load balancing of clients connections
 - b. Support for L2 and L3 client network
2. GCP/AWS deployment
3. Disaster Recovery using Asynchronous Replication between two or more Elastifile systems
 - a. Multi Replication services support for HA and performance
 - b. DC to DC level replication
 - c. Each site can host active DC and replicated DC
 - d. Replication direction switch
 - e. Create a test image from a snapshot
 - f. Multiple recovery points
4. Highly Available Management and monitoring system
5. Data reduction (compression and deduplication)
6. Scalability from 3 to 100 nodes
7. Read Only Snapshot
8. Non Disruptive Upgrade
 - Upgrading from version 2.0.2.3, 2.1.0.10, 2.5.0.6, 2.5.0.15, 2.5.1.10, and 2.5.1.18, 2.5.2.2, 2.5.2.3, 2.7.0.13, 2.7.5.1 is supported (DSM to DMS, GCP to GCP, and AWS to AWS)
 - 2.0.2.3, 2.1.0.10, 2.5.0.6, 2.5.0.15 - please contact support before upgrade
9. Hot add/remove nodes
10. 2 or 3 way internal data replication
11. Soft/hard quotas for data containers
12. Suite of enterprise grade high availability features:
 - a. Recovery from loss of nodes
 - b. Recovery from loss of a device
 - c. NIC failures (in multipath mode)
 - d. Data recovery from complete cluster power loss
13. eMRI log collection mechanism

Supported Deployments & System Requirements

The following deployment models are supported for testing with Version 2.7.0 GA:

- **Dedicated Storage Mode**

A designated set of dedicated storage nodes provide storage resources while exposing external NFS interface. The NFS Virtual IP can be accessed by any external client regardless of the client deployment mode (Physical, Virtual, Container). A highly available self load balanced mechanism ensure clients connectivity to the front end process is maintained.

- **AWS/GCP Deployment**

Utilizing cloud resources Elastifile combines predefined compute and storage devices into a scalable, highly available file system offering the full capabilities as on premessis Elastifile File System

The following table summarizes the key capabilities and system limitations provided in this version:

Feature	Version 2.7.5.4
Environment	
Interface	NFSv3
Deployment Scale and Capacity	
Number of nodes supported in Dedicated Storage Mode	3 to 96 ¹ (4 storage nodes recommended)
Number of devices (per node)	Dedicated Storage node: 12 devices ²
Max device capacity	4 TB
Max capacity per node	Dedicated Storage node 48 TB
Reliability & Availability	
Replication	Global setting at setup 2 way or 3 way ³
Multipathing	Global setting at setup: full multipath
Networking failure	Supported in full multipath mode
Storage device failure	Supported ⁴
Node failure	Any node can fail without loss of data
Cluster failure/power outage resilience	Supported
Call home	Supported
eMRI	Supported
Data Reduction	
Compression	Optimized
Deduplication	Optimized

¹ For deployments over 20 nodes please contact Elastifile support

² 2,4,8,12 devices were tested

³ 3 Way replication requires a minimum of 5 nodes in cluster.

⁴ Full rebuild will initiate including the remaining node storage devices.

Compactization (data reduction capacity reclamation)	Supported
Number of Cores Used for Virtual Controller⁵	
Dedicated storage node	Up to 20 ⁶
Minimum Memory	
Dedicated storage node	128GB
Local SSD for DSM Controller	
Local SSD for DSM Controller	256GB
EMS Requirements	
Number of cores	4
RAM	8GB
SSD	64GB
Replication Service Requirements (minimal)	
Number of cores	4
RAM	16GB
SSD	32GB
1 bidirectional replication per core, each core requires 4GB RAM. Number of cores can be increased.	
Management Model	
Managed objects	System, data containers, exports, hosts, devices
Dynamic add/remove node	Supported ⁷
Soft/Hard quotas	Supported
Basic analytics and statistics	System, node, data container, client VM
Max Parameters	
Max number of files/folders ⁸	2 [^] 64
Max files/subdir per folder ⁹	500,000

⁵ Numbers indicated supported production core counts. Additional core deployment options available in lab mode for testing only.

⁶ Installation will use all cores available in the physical server, only dual socket 6 cores, dual socket 8 cores and dual socket 10 cores are approved. Please contact support for any other CPU configuration.

⁷ Multiple node additions can be requested in one operation, otherwise processed management operations (such as add node) must complete before a new operation can be started.

⁸ Design limit. Actual number of files may be limited by cluster capacity.

⁹ File count dependant on length of filenames. Limit may be lower if all file names are long (e.g. 18,000 files of 255 bytes length filenames, and 500,000 files if all file names are of 20 bytes length).

Max active files per node ¹⁰	64,000
Max number of connections per node	2000
Max number of GUI sessions	3
Max number of data containers	999
Max number of exports per node	1000
Max length of filename/directory name	255
Max file size	2PB

Key Capabilities Provided in Elastifile CloudConnect Version 2.7.1.

Version 2.7.1 of Elastifile CloudConnect allows the installation of CloudConnect as a separate VM providing the following capabilities

- **Any source file system** - Simply mount a snapshot of your legacy source file system using NFS to the CloudConnect writer for migration of its data to the cloud
- **Very efficient delivery and storage** - Elastifile CloudConnect writer compresses and deduplicates all data before shipping and storing. Version diffs are identified and only new data is shipped for storage in the cloud. Dedup is global per bucket regardless of source file system reducing cloud storage costs.
- **Direct connection to S3 compatible Object storage** - There is no need to run any VM in the cloud to receive the data. CloudConnect writer directly writes to S3 compatible cloud storage providing TCO efficiency and flexibility.
- **Multiple object storage supported** - The following object storage are supported
 - IBM Object Storage (Cleversafe)
 - HGST ActiveScale
 - DellEMC ECS
 - Cloudian Hyperstore
 - Swiftstack
- **Restore single files or complete file system** - Stored data can be viewed and read as a file system while data is in place as objects. Full versioning is supported enabling to restore either entire file systems or specific files from any version.
- **Differential Checkout** - Per user selection, only new or changed files are checkout, user can select what to do with deleted files (delete or archive)
- **Fuse Client** - Mount the stored data without copying any file, view, list, run diff commands and copy individual files.
- **Graphical UI** - easy to use GUI to manage your data, view archived versions and

¹⁰ 200 per front end core

Elastifile CloudConnect System Requirements

Feature	Version 2.7.1.0
VM Requirements	
Virtual Cores	2(Higher is better for performance)
Memory	4GB (higher is better for performance)
Disk Size	64GB (larger for large data sets)

Note: Please make sure to start and configure NTP before connecting to the cloud provider.

Fixed Issues in ECFS 2.1.0.10

[EL-4721] Delete or truncate a large file might take a long time and impact service

Deleting or truncating a very large or sparse file can take a long time, sending multiple delete request might overflow the queue and hang the system

[EL-5755] Deleting a DC with snapshots gives a general error

When deleting a DC that contains snapshots the system reports back “Failed to delete the data container”

[EL-6700] Listing locked files does not work

The CLI command to show and release locked files does not work.

[EL-6588] GUI does not enable SR-IOV

GUI marked SR-IOV as enabled but does not change the ESX configuration

[EL-6489] Ownership recovery is stuck in OID gathering loop

In some rare cases if a node was fenced before lockdown recovery the ownership recovery process might get stuck in a loop of OID gathering.

[EL-6572] GCP deployment might fail with a message on vHead Authentication

Due to a race condition, virtual controller is created but the system fails to set the authentication credential. Please remove the nodes and run the deployment process again.

Fixed Issues in ECFS 2.5.0.6

[EL-7911] Network IP assigned in the wrong order

Client network IP were assigned in the wrong order preventing access over routed networks.

[EL-6419] Add SNMP configuration to a deployed system

SNMP configuration can be added to a system after deployment and does not require re-deploy

[EL-6388] Call Home failed to resolve service name

The call home feature was unable to resolve the service name as failed to send the daily reports

[EL-7987] Add node operation failed due to invalid device path

Multiple request to add the same node caused the original add node to fail. Added protection to prevent multiple requests to the same task

[EL-5131] There is no option to add IP address when adding a node in DSM mode

If add node command fails, please add additional physical IP address in the System Configuration screen

[EL-5763] Services Performance tab show the wrong number of services

The system counts storage nodes and EMS as total number of working services, when moving to the performance tab only the storage nodes are shown but the total count still show the number including EMS

[EL-7429] EMS failover during lockdown might not present accurate system status

DC delete operations done in close proximity to system lockdown followed by EMS failover might not be reflected in the secondary EMS

Fixed Issues in ECFS 2.5.0.13

[EL-8787] EMS fails to communicate with EMS that is set up with HTTPS

Replication from a version 2.5 system to a system running version 2.5.1 (new or upgraded) will fail since newer system is using HTTPS. This fix is a must for customers who wishes to continue replication while upgrading the systems.

[EL-8656] Replication fails if a hard-link file is modified

If a file that has hard-link is modified after replication, next replication will fail.

[EL-8633] In Some cases replication service does not remove deleted files from target

Files that exists in both directory and subdirectory and are deleted from source might remain in the target system.

[EL-8627] In some cases replication does not work with symlinks

If a symlink exists in the source it might cause the replication to fail.

[EL-8555] Add node might fail in system after NDU from 2.1

Adding a node in system that was upgraded from 2.1 might fail.

[EL-8271] Wrong affinity settings can cause system lockdown

In some rare cases, wrong affinity setting can cause system lockdown while performing system maintenance functions.

[EL-8135] NDU got stuck after EMS upgrade

After EMS upgrade the NDU process didn't proceed to upgrade the system nodes.

[EL-8478] High system utilization and multiple snapshots can cause uneven data distribution

In systems with high utilization and high number of snapshots there is a risk of uneven data distribution. Chances are reduced as the number of devices in the system increases.

[EL-7917] Default GW is lost after system reset

System nodes lose the default GW configuration after system force reset

Fixed Issues in ECFS 2.5.1.18

[EL-2875] Statistics collection may be suspended after failures in creation of new data containers

After a failure in the creation of a new data container, statistics collection may become suspended and performance indicators will appear as zeros in both the management console and the CLI.

[EL-3995] False Status after system re-activation

After successful system re-activation, all devices are marked as down. The state changes once I/O are restarted.

[EL-5673] EMRI collection from CLI might fail

EMRI collection from CLI might fail with session timeout message if collection takes over 1.5 hours. The EMRI creation is still going on and logs can be retrieved. GUI EMRI collection is working

[EL-5781] elfs-cli system help contains encoded characters

The elfs-cli system help contains ' instead of single quote

[EL-5886] Multiple tasks can slow down the GUI

Issuing multiple tasks such as snapshot creation might slow down the GUI and cause empty screens in the dashboard.

[EL-5861] Events of data container and export creation are re-sent after EMS failover

In some cases events of data container and export creation are re-sent as if they were created after the EMS failover.

[EL-7370] UI shows no I/O when interconnect network is down

Statistics collection is done on one of the interconnect networks (randomly selected). If this network is down due to switch failure, the UI will not show I/O for 40 seconds until the statistics collection service failover to the second network

[EL-7341] Controller physical IP is not updated to secondary EMS

When adding a new physical IP via CLI the secondary EMS is not updated. This does not have impact on system operation but if the node is removed the specific IP is not added back to the pool and will require to be added again if another node is added.

[EL-7270] Device count in UI is wrong after EMS failover

In System View the device count after EMS failover might show some devices as failed.

[EL-7446] Wrong DC information in CLI

When a DC is created with dedup and compression the CLI status shows it was created with the wrong policy name. UI shows the correct DC characteristics

[EL-7651] Graceful Shutdown waits for all operations to complete

System graceful shutdown waits for all system operations such as ownership recovery to complete before shutting down the system

[EL-8115] After cold upgrade EMS failover and failback might fail

After a cold upgrade from version 2.1.0.10 and Data Containers creation and deletion, an EMS failover and Failback might fail, please contact Elastifile support to recover the EMS

[EL-8087] Replication to multiple site does not work

Configuring replication for multiple sites, overwrite the site SSH token in the replication service and might interrupt the previous establish replication.

[EL-8020] Name sorting is not supported in client view

Clients are shown as IP address and sorting by client name is not supported.

[EL-7792] ACL on passive Data Container remain active

When making a DC passive, for example after a replication collision, the existing ACL remains active and client can access and write data. Please ensure you prevent clients from writing in parallel to the replication service copies.

[EL-8122] Capacity report is zeroed if both EMS are shutdown and restarted

Shutting down both EMS and starting them up, first the previous secondary and next the previous primary will result in zero capacity report.

[EL-7874] System health does not reflect system inability to sustain another failure

If a system lost a controller and it is now below the minimal number of controllers per the FTT level, the system health will be OK (when rebuild is done) even though the system cannot sustain another failure

[EL-7985] Rebooting the replication service will fail the replication

A manual restart of the replication is required.

[EL-8010] DC name in dashboard is capitalized

DC names are shown in dashboard with first capital letter even if created with all lower-case

[EL-7205] Deleting a replication connection leave the remote connected

When deleting the replication connection from the source side the remote side will shown as connected

[EL-8727] DC replication status changed without user intervention

When switching replication roles between sites in some rare cases after EMS failover the DC in the site reports their previous roles (active/passive)

[EL-5846] Ownership recovery doesn't have a progress bar

Ownership recovery progress bar is not displayed as part of the recovery stages.

[EL-7945] Add node failed with invalid device path message

In some rare cases the client UI sends the add node command multiple times causing the add node operation to fail.

[EL-8089] Snapshot is not added to snapshot list if created via RestAPI

GUI does not update the snapshot list if the snapshot is added in another interface such as the RestAPI

[EL-8335] Remote replication site status is not reported correctly

When remote replication site service is turned off the source site still reported it as connected.

[EL-8499] ctime is not update when creating or deleting a link

St_ctime was not updated when creating or deleting a link.

[EL-8603] Leftover control tasks prevented the replication from finishing

In some rare cases due to another issue, leftover control tasks prevented the replication from finishing and required manual intervention.

[EL-6309] Call home does not retry on failure

If call home function does not succeed it does not retry to send the logs

[EL-8684] EMRI report list is not ordered according to creation date

EMRI report list was not ordered correctly.

[EL-8723] DC list screen did not refresh after DC deletion

The list still showed the deleted DC

[EL-8748] EMRI menu name changed

EMRI menu option is called "Log Collection"

[EL-8821] DC creation takes long time to complete

Other background operations prevented DC creation until they were completed

[EL-4609] EMS failover took long time

EMS failover took close to 2 min.

[EL-8512] System lockdown when deleting a DC

In some rare cases the system tries to validate if a DC is empty while it is being deleted which cause a panic.

Fixed Issues in ECFS 2.5.2.2

[EL-10507] DC data is not synced after EMS failover

In some rare cases changes to a Data Container are not synced and disappear after EMS failover.

[EL-10427] Time sync issue on storage nodes

Storage nodes' time sync service (chronyd) was mis-configured and nodes didn't sync with EMS time.

[EL-10415] Too many clients connection caused a storage node panic

Storage node didn't reject new connections over the limit of connections per node.

[EL-9914] File and Dir lock CLI command caused EMS to stop responding

A bug in "get_dir_lock" and "get_file_lock" caused the EMS to stop responding.

[EL-10206] AWS networking fragmentation prevented cluster recovery

Internal node traffic got fragmented by AWS networking.

[EL-10192] Wrong server name was returned in response to SM_MON

The system returned elfs instead of the configured service name or IP address in the response to SM_MON notification

[EL-9781] File locking issue

Multiple file locking with repeated lock request caused the locking queue to be filled, left files locked without client awareness and caused storage node panic

Fixed Issues in ECFS 2.7.0.13

[EL-4737] Reformat system might not work in all cases

In some rare cases if the system needs to be reformatted (deleting all data) but is missing some critical configuration information, the re-format will not work and the system needs to be re-installed.

[EL-5067] Dedicated Storage server installation failed with wrong BIOS settings

When installing a node as dedicated storage server the BIOS/UEFI settings must be set to legacy. UEFI setting is not supported

[EL-5535] Nodes are reported as fenced during NDU

During the NDU process each node is being upgraded and is reported as fenced during this stage, when the upgrade is done the node is automatically added to the system

[EL-5649] EMS installer script for DSM allow the selection of vCenter deployment

When installing EMS in DSM mode (even as a VM) user should not select vCenter deployment mode.

[EL-5779] System statistics might show negative numbers after a large data set deletion

Deleting a large data set might show negative performance numbers for the data container

[EL-8194] It is possible to define a broadcast IP in the host physical IP screen

Please ensure to use only host IP addresses in the Client Network configuration screen

[EL-10104] Username and password are not optional in register node cli command

The command help shows the username and password as option but they are mandatory for the command execution.

[EL-10532] System lockdown due to bucket overflow

A very high dedup rate and a high rate of rewrites can cause, in a rare case, a bucket overflow that lead to system lockdown.

Fixed Issues in ECFS 2.7.5.4

[EL-12000] DSM deployment EMS password is mandatory

To avoid user lockout or default password the user must provide the EMS password in DSM deployment.

[EL-11957] DSM deployment failed

In some cases the installer tried to execute the same command twice which failed the installation.

[EL-11919] Exports are not shown in UI after EMS failover during NDU

During NDU some exports that exists in the system are not shown in the UI after the EMS failover phase.

[EL-11918] DSM deployment failed when EMS IP is not provided

The DSM deployment failed when during the first EMS installation the EMS IP was not provided.

[EL-11813] EMS installation as Hyper-V VM failed

Installer timeouts were not correct and failed the EMS installation.

[EL-11788] GCP deployment, mount example

The mount example at the end of the DC creation was improved with additional steps and instructions.

[EL-11787] Set default quota for DC

DC creation has default quotas of 1000GB soft and 1500GB hard quotas. The quotas can be changed by the user.

[EL-11786] Public DC permissions

Public DC is now created with default “No User Mapping” permissions allowing root users to access the namespace without any additional change.

Known Bugs & Limitations

Elastifile deploys a version based a roadmap for the addition of functions and features to the Elastifile Cloud File System. Certain functions and features of the Elastifile Cloud File System may be scheduled for upcoming releases and not included in Elastifile Version 2.7.0. The Elastifile roadmap is subject to change without notice.

The following section lists known limitations in Elastifile Version 2.7.5.

General Limitations

- Hyper Converged mode is not tested in this version and it is not approved for production or upgrade
- Version 2.7.5 assumes homogeneous configurations. In dedicated storage mode, all storage nodes must be identical.
- Elastifile clusters in Version 2.7.5 are limited to 96 nodes (testing limitation).
- It is not possible to install two ECFS system in the same GCP VPC. To install two systems in the same project please refer to the following KB:
[how-to-install-2-or-more-elastifile-cloud-file-system-in-the-same-project-on-gcp](#)

High Availability Limitations

- Failure domains are not supported.
- In version 2.7.5, for systems that are configured with Fault To Tolerate = 2, a failure of a second node during the rebuild process will stop the rebuild process and restart a new one.
- If a system that is configured with Fault To Tolerate = 2 is down to 4 controllers (due to previous failures or intentional removal) the FTT will be changed to 1
- Pulling a device from a live controller might result in a controller failure and not a device failure due to the underlying OS limitations.
- When working in GCP using Round Robin DNS, new clients will not be able to resolve the NFS service name if the EMS is not available. Please do not change the EMS VM auto-restart settings

Asynchronous Disaster Recovery

- RPO depends on the replicated data set, amount of changes and the throughput of the connection between the two sites.
- Bi-directional replication is supported but may impact the minimal possible RPO.

- Concurrent replications is limited to one bidirectional replication per replication service core. I.e. the minimal server configuration can support up to 4 concurrent replications. There is no enforcement of this limitation. User are advised to add more replication services or increase the core count and RAM in case more concurrent replication are needed.

Protocol Limitations

- NFS v3 is the only protocol supported for Version 2.7.5 (SMB and Object are not supported).
- LDAP integration is not implemented.

Datapath Limitations

- Replication (protection level) can be configured system wide on setup only.
- Only replication levels 2 and 3 are supported.
- Dedup has performance impact in Version 2.7.5

Backend Store Limitation

- Hot add or remove of devices (SSDs) are not supported.

Networking Limitation

- NIC failure support and multipathing is supported as a global setting only either every node is configured for multipath or none are configured.
- No RDMA support.

Management & Control Limitations

- Deploying two EMS instances for High Availability is required in DSM (Cloud deployment requires only one). Control and management console VM must be running to initiate a rebuild process during failure and must remain running during the rebuild process.
- Correction of errors during setup requires manual intervention.
- Graceful shutdown does not provide a progress indication.
- Management supports English only (internationalization/unicode is not supported).
- Data containers and export names are limited to 80 characters.
- Only Chrome, Firefox and EDGE browsers are supported for the web management interface.

In addition to the above, Elastifile Version 2.7.5 includes the following known bugs:

[EL-2592] The system status is reported incorrectly during a graceful shutdown

After the user requests a graceful shutdown, the management console indicates the system as “shut down” before the actual shutdown process is completed. Review the system logs to confirm the graceful shutdown has completed before physically powering off any node.

[EL-3081] Data displayed for certain drives in the management console may be incomplete

In certain environments, not all drive parameters may appear in the drive display on the management console. In particular, NVMe drive capacity may not appear in the display. The missing information can be found using the vCenter management console.

[EL-2909] Deployment will fail to complete if network disconnections occur during the deployment process

If network connectivity fails during the deployment process, the deployment will not correctly complete. Full network availability is required during the deployment process.

[EL-3138] Deployment may fail when using the drive TRIM feature

If a new deployment is attempted on drives that have previously used the TRIM feature, the deployment may fail. Contact Elastifile support before deployment if you intend to use the drive TRIM feature.

[EL-3060] System statistic is incorrect in case of time sync issues

If the server and browsers times are not in sync, the system statistics might be incorrect or missing.

[EL-3869] Switching from lab to production deployment require re-deploy

Switching from lab configuration to production configuration is not supported and requires re-deployment of the Elastifile System.

[EL-3993] E-mail notifications are not sent over secure SMTP (port 465)

E-mail notification are not working when configured to be sent over secure SMTP (port 465). Port 25 and Port 587 are working correctly.

[EL-3083] Deploying a system while a host is still being installed might fail

Running the deployment while one of the hosts is still being installed might fail., Wait for all hosts installations to finish before deploying the system.

[EL-4685] It is not possible to remove the / export if another export was created

When creating a new export a directory is created under the root export, removing the export does not remove the directory thus leaving the root with data. Please remove the directory via NFS interface before deleting the root export.

[EL-4993] Statistics and Events gap during EMS failover

In case of EMS failover there will be a short gap in the statistics and events due to the asynchronous sync mechanism in current version

[EL-4900] Dedicated Storage Hosts info is missing

Host status and NIC Bandwidth is missing from the relevant screens and CLI commands.

[EL-4591] The system does not prevent installing second EMS on same node

The system does not prevent or alert if the user is trying to install the second EMS on the same ESXi as the first EMS. Please ensure to install EMS on different physical hosts

[EL-5521] EFI boot devices are not supported

Do not select an EFI boot device or system partition as the node target boot device.

[EL-5471] ECDB device failure will prevent future use of the node as ECDB target

If an ECDB device fails the node will not be used as ECDB target. If the system size is close to the minimal nodes number required per installation, future failures might interrupt service. Please remove the node, fix the device and add the node back to the system

[EL-5219] When using DSM deployment please ensure to format any other boot device

Node might boot from another boot device if exists and not rejoin the system automatically

[EL-5654] EMRI collection in initial deployment stages does not work on nodes logs

If the deployment fails before nodes are initiated there is no option to collect logs from nodes. Only EMS logs are needed for support at this stage.

[EL-5339] NDU event is missing

There is no event when DNU starts and finish

[EL-5800] Hard quota can be modified to lower value than actual data

It is possible to change the DC's hard quota to a lower level than the actual data written to the DC. Future I/O to the DC will not be possible.

[EL-5802] Data rebuild progress bar is cleared before rebuild is done

The data rebuild progress bar in the UI is cleared but the system is still at risk for few more minutes. The data rebuild ended events is sent with the right timestamp

[EL-5526] Slow provisioning operations during Non Disruptive Upgrade

In some cases users might experience slow reaction time of provisioning operation during Non Disruptive Upgrade

[EL-5828] Event might appear in wrong order in the GUI

Events are sorted by ID and not timestamp and might appear in the wrong order in the GUI

[EL-7683] Reported compression is less than expected

When testing compression rate with VDBANCH, the reported compression is slightly less than expected due to internal metadata overhead

[EL-8040] Wrong Controller indication in the UI

In some cases if a Dedicated Storage Node fails, it will be shown in the UI with ELFS-DR role.

[EL-10050] Replication service needs HTTP port open

While the replication service is working over HTTPS and SSH the initial connection is done over HTTP and redirected to HTTPS.

[EL-9892] Adding a replication service give a false error

When adding a replication service via the UI there is a false indication that the operation failed but the replication service is added

[EL-9760] Wrong reporting on replication pair

If the source site is down the target site will report that the RPO is not met but will not give an indication of the replication site status.

[EL-8587] When switching replication direction status is not updated

When switching the replication direction the status will be shown as down until the next RPO cycle

[EL-9760] Wrong reporting on replication pair

If the source site is down the target site will report that the RPO is not met but will not give an indication of the replication site status.

[EL-8281] Async Replication might copy a redundant moved file

A file move is not identified at the source and the replication service will copy it to the remote site and delete it afterward.

[EL-8280] Async Replication might copy a redundant linked file

A hard link is not identified at the source and the replication service will copy it to the remote site and delete it afterward.

[EL-8279] Async Replication can only replicate up to 1M hardlink per DC

Over 1M hard links in a DC will fail the replication

[EL-8278] The max directory tree depth of the Async Replication is 1024

The replication will fail if the directory tree depths is over 1024

[EL-8266] Trying to add controller and removing one in the same operation might fail

In some rare cases trying to add a controller and removing one in the same operation might result in the removed operation staying in a “pending removal” status.

[EL-7898] EMS failure hides the progress of “Remove Node” operation

An EMS failure during remove node operation will hide the progress of the operation.

[EL-7683] Compression values of ECFS and VDbench are not aligned

VDBench numbers do not take into consideration the differences in the compression algorithms

[EL-7043] SMTP fails with Office 365 and Google

The workaround for Google was to enable "lesssecureapps"

[EL-6643] In AWS a node with local SSD cannot be restarted

Due to the ephemeral nature of the local SSD, you need to remove the node and add a new one

[EL-6136] During NDU the progress bar is misleading

The progress bar indicate 100% done after upgrading the EMS server but once the nodes upgrade starts the bar goes back to 0% and start reporting the progress of the nodes upgrade

[EL-10015] Peering success notice does not contain the peer side info

When peering sites for Async Replication the success notice does not contain the peer side information and it is appeared empty.

[EL-11569] GCP subnet size is not validated

When installing in GCP please ensure that the selected subnet has enough free IP address to support the cluster size (EMS and minimal 3 Storage Nodes) and allows for future growth and added services (Replication Service nodes)

[EL-11551] Snapshot deletion looks like it is stuck

If there is a big file deletion in the background, the deletion of a snapshot containing this file might look like it stuck for a long time. The snapshot deletion will progress once the file deletion background processes finish

[EL-11297] EMS does not show system as in lockdown if all controllers are turned off

If all controllers are turned off the EMS will not show the system as in lockdown state.

[EL-11022] Async replication might fail if number of files per directory is close to the max

During replication the Async Replication first add new files and then delete the removed ones. If the number of files is close the max number of files per directory the operation might exceed the max number and fail.

[EL-10771] EMS will appear as “Unreachable” during deployment

When deploying a cluster via the GUI, the EMS will appear as “Unreachable” in the background but the deployment process will continue. The EMS will be active when the deployment is done.

[EL-11019] Can't peer a site after removal

In some rare cases when removing a peer site it is not possible to add it again. Rebooting the EMS and removing the peer site again helps solves the issue.

[EL-10890] GCP NDU double the cluster resources

During NDU the cluster resources will double, please ensure your project has enough quota before starting the NDU.

[EL-10543] Cold-Upgrade is not supported via the GUI

Only CLI command is supported

[EL-11633] GCP/AWS DNS Round Robin might not load balance traffic

DNS Round Robin does not force the client to use a specific IP address and as result might not load balance the clients evenly across all storage nodes

[EL-11669] Can't delete a DC if snapshot deletion is not done

Snapshot deletion will happen immediately to prevent client access to the snapshot, but there are background processes that may take some time to finish the snapshot deletion. In this case a DC might look empty but DC deletion will not be possible.

[EL-11453] Can't deploy a DSM cluster with multiple IP

During the deployment of DMS cluster all nodes should have only one IP address. Additional IP address can be added after deployment.

[EL-11560] Wrong IP at the end of DSM deployment

At the end of the DSM deployment the user is instructed to use the EMS physical IP address and not the VIP. User should always use the VIP.

[EL-11290] EMS deployment in VMware requires paravirtual devices

EMS in version 2.7.x requires all the devices to be set to Paravirtual mode. Please do not upgrade your system before changing this settings.

CloudConnect version 2.7.1.0 known issues

The following section lists known limitations in Elastifile CloudConnect Version 2.71.0.

General Limitations

- CloudConnect assume a highly available file system as the source. If the FS is not available during a commit the commit will not be resumed and a new one should be initiated when the FS returns
- CloudConnect v2.7.1.0 was tested with 100M files and directories
- Deleting a version is allowed only from a single cloudconnect instance while there are no other check in operation to the same bucket

[ELCC-173] GUI stuck if connection fail

Refresh the browser if GUI seems to be stuck when connection fails.

[ELCC-297] Credentials are needed when creating a second connection to Cleversafe bucket

When creating a second connection to a cleversafe bucket the user is requested to re-enter the credentials.

[ELCC-297] Wrong indication for abort operation

Aborting a checkout operation, using the GUI, returns “operation not aborted” even though operation is aborted

[ELCC-279] Temporary files are not deleted when operation fails

Files in /elastifile/ccweb/elcc_tmp/ are not deleted when operation fails. User is advised to check disk space and clear files.

[ELCC-278] Mounted directories are not deleted when operation fails

Directories under /elastifile/conf/mnt are not deleted when operation fails.

[ELCC-264] CCstore delete does not remove bucket in Cleversafe

Bucket and references are not removed from Cleversafe when issuing a CCstore delete command

[ELCC-341] Archived file location is not shown correctly in the UI

When selecting to archive deleted files the files location is shown with /.elcc/ where it should be /_elcc_/