Elastifile

ECFS v3.2.1 GA

Build 3.2.1.39

Release Notes

Product Overview

Implementation Architecture

Enterprise-Grade Management and Storage Feature Set

New Key Capabilities Provided in Version 3.2.1

Supported Deployments & System Requirements

- GCP Configurations

Fixed Issues in ECFS 2.x

Fixed Issues in ECFS 3.0.0.15

Fixed Issues in ECFS 3.0.1.4

Fixed Issues in ECFS 3.1.0.13

Fixed Issues in ECFS 3.1.0.22

Fixed Issues in ECFS 3.1.0.27

Fixed Issues in ECFS 3.1.0.31

Fixed Issues in ECFS 3.2.1.8

Fixed Issues in ECFS 3.2.1.19

Fixed Issues in ECFS 3.2.1.20

Fixed Issues in ECFS 3.2.1.30

Security Fixed Issues in ECFS 3.2.1.30

Fixed Issues in ECFS 3.2.1.37

Security Fixed Issues in ECFS 3.2.1.37

Fixed Issues in ECFS 3.2.1.39

Known Bugs & Limitations

The following section lists known limitations in Elastifile Version 3.2.1.

General Limitations

High Availability Limitations

Asynchronous Disaster Recovery

Protocol Limitations

Datapath Limitations

Backend Store Limitation

Networking Limitation

ClearTier Limitation

SnapShot Limitation

Management & Control Limitations

Known limitations

Product Overview

The Elastifile Cloud File System (ECFS) is an all-flash, primary storage solution for Cloud Deployments and Next-Gen SDDC (Software Defined Data-Centers). By augmenting default cloud instances or commodity hardware, 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 scale linearly and expand 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.
In version 3.2.1 CloudConnect is an integral part of the solution and is used for Object Tiering (alpha preview).

Implementation Architecture

The Elastifile Cloud File System is deployed by placing an Elastifile controller on each cloud instance or physical server 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.

Integrating file storage and object storage (alpha - preview)

Most enterprise applications are designed to interface with a file system. These applications expect to see file and directory structures and they expect standard file system protocols (e.g. NFS). At the same time, however, object storage delivers the lowest cost-per GB-per-month and has established a well-deserved reputation as the "cheap and deep" cloud storage layer.

Copy data management

Most enterprises expect to retain periodic snapshots of data state for recordkeeping and to facilitate data restoration/recovery. In the cloud, however, the cost of the underlying storage infrastructure must be carefully considered as snapshots can rapidly accumulate. To enable

retention of the desired data state, while still complying with IT budget constraints, enterprises need modern mechanisms to cost-effectively store snapshots in the cloud.

New Key Capabilities Provided in Version 3.2.1

Elastifile Version 3.2.1 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/On-Prem deployment
3. Disaster Recovery using Asynchronous Replication between two Elastifile systems
    a. Multi Replication services support for HA and performance
    b. DC to DC level replication
    c. Replication direction switch
    d. Create a test image from a snapshot
4. Highly Available Management and monitoring system
5. Data reduction (compression and deduplication)
6. Scalability from 3 to 32 nodes
7. Read-Only Snapshot
8. Non-Disruptive Upgrade
    - Upgrading from version 3.0.1.4/31.0.8 is supported (DSM to DSM and GCP to GCP)
    - Earlier versions - please contact support to get the correct upgrade path
9. Hot add/remove nodes
10. 2-way internal data replication
11. Soft/hard quotas for data containers
12. A 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
14. ClearTier Object Tiering (Alpha-preview)
    a. Policy-based tiering of cold data to an object store of your choice.
    b. “Infinite” snapshot retention in the object store
15. Multi Availability Zone support
16. Single Zone HA utilizing persistent disks protection increasing storage utilization and reducing costs.

Supported Deployments & System Requirements

The following deployment models are supported for testing with Version 3.2.1 GA:

- **Dedicated Storage Mode**
  A designated set of dedicated storage nodes provide storage resources while exposing an 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.
- **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-premises Elastifile File System

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

Feature	Version 3.1.0
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] GCP: see GCP configuration table
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] GCP: 1-way replication on persistent disks 2-ways on single availability zone 2- ways on multiple availability zone
Networking failure	Supported in DSM full multipath mode
Storage device failure	Supported[4]
Node failure	Any node can fail without loss of data

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

Cluster failure/power outage resilience	Supported
Call home	Supported
eMRI	Supported
Data Reduction	
Compression	Optimized
Deduplication	Not Performance Optimized
Compactization (data reduction capacity reclamation)	Supported
Number of Cores Used for Virtual Controller[5]	
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. The number of cores can be increased.	
Management Model	
Managed objects	System, data containers, exports, hosts, devices
Dynamic add/remove node	Supported[7]

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

[6] The 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.

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	2^64
Max active files per node	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
Max number of snapshots on primary tier	255
Max number of snapshots on object tier	2559

The following are the default cloud configurations supported in version 3.2.1.39

GCP Configurations

Configuration	Instance Type	# CPUs	Total RAM (GB)	# Devices	Capacity per Node (TB)	# Min Nodes
EMS	n1-standard-4	4	15	1	0.100	1
Replication Agent	Custom	4	16	1	0.100	1
Small (Persistent-SSD)	Custom	4	32	4	0.7	3
Medium (Persistent-SSD)	Custom	4	42	4	4	3
Medium Plus	Custom	8	64	4	4	3

[8] Design limit. The actual number of files may be limited by cluster capacity.

(Persistent-SSD)						
Large (Persistent-SSD)	N1-highmem 16	16	104	5	20	3
Local SSD	Custom	16	96	8	3	3
Small Local SSD	Custom	4	42	3	1.125	3
Standard PDs	Custom	4	64	4	4	6
Small Standard PDs	Custom	4	32	2	1.3	3

Note: Starting from version 3.2.1 the minimal node configuration is enforced in the GUI. To install a system with standard PD with less than 6 nodes (in single zone high-availability) please contact Elastifile support.

Fixed Issues in ECFS 2.x

For fixed issues in versions before version 3.0.0 please refer to older copies of the release notes

Fixed Issues in ECFS 3.0.0.15

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

It is possible to change 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-5526] Slow provisioning operations during Non-Disruptive Upgrade

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

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

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

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

During the deployment of the DSM cluster, all nodes should have only one IP address. An 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. The user should always use the VIP.

[EL-10050] Replication service for Async DR or ClearTier 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-11297] EMS does not show the 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-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.

Fixed Issues in ECFS 3.0.1.4

[EL-13465] Commit to object failed on unexpected value from GCS

In some rare cases, an unexpected return value from GCS caused the commit to fail.

[EL-13449] listing a .object might fail with I/O error

Listing the content of a directory in .object sometimes returns an I/O error

[EL-13407] Missing port 443 in GCP FW configuration

Deployment of a replication agent in GCP failed due to missing port 443 in the FW configuration

[EL-13389] Changes in file’s metadata might fail a commit to the object

Changes in the file's metadata such as changes to a number of hardlinks can cause a commit to object to fail.

[EL-13350] Async DR failed to delete remote snapshot

Async DR failed to delete remote snapshot thus reached the max snapshot per DC and failed the synchronization process.

[EL-13339] Failed to add a node in case there are stopped nodes in the cluster

If there are stopped nodes that are part of the cluster, adding a new node will fail due to communication errors with the stopped nodes.

[EL-13297] Data commit to the object store failed on missing port

Port 12121 is missing from the FW configuration and causing the commit action to fail.

[EL-13231] Object tier creation failed if the system name had spaces in it

Auto configuration of the object tier is using the system name as part of the path but spaces are not allowed in object tier name.

[EL-13226] Recovering a system while marking one of the nodes as failed didn’t work

Marking a node as failed-node to allow the system to recover without it didn’t work as expected.

[EL-13220] Can't delete a version from the object if the file size is an exact multiplication of 2MB

Trying to delete a version that contains a file whose size is exactly a multiplication of 2MB, will fail with an error.

[EL-13219] Empty version deletion failed

Trying to remove an empty version (empty snapshot) from the object store failed.

[EL-13218] Cooling data didn't behave according to the cooling policy

The data calculation after the scan returned wrong values due to a wrong variable type, which caused some misses when cooling data.

[EL-13183] Cooled files where marked as dedup even if they were changed

In some rare cases, a changed file was marked as a duplication of the older file version in the object store even if the file changed.

[EL-13178] Upgrading the replication agent cause the existing connection to fail

The upgrade of the replication agent didn't restart the service and caused the existing connection to fail.

[EL-13145] Deployment of a system with more than 50 devices failed

System deployment didn't support a cluster with more than 50 devices

[EL-13123] Incremental rebuild might risk data integrity

In some rare cases when operations that reuse object keys, such as snapshot deletion or file truncate, occur while one node is down, the incremental rebuild might assume that two objects with the same key are indeed the same and will not rebuild them.

[EL-13117] It is not possible to remove a deleted replication agent from the system

If all of the system replication agents are turned off and deleted, they'll appear in the system as unreachable without the option to remove them and add new ones.

[EL-13065] Object GW is serializing requests

All the requests from a single front end core to the object store are serialized by the Object GW impacting performance

Fixed Issues in ECFS 3.1.0.13

[EL-13102] Can install new EMS or Replication Agent after NDU

The NDU process moves the RPM repository to a different location. Please move the file to the directory with the version name.

[EL-13363] symlink are not cooled

Symlinks are not cooled by design. Trying to cool a list of files that includes symlinks will not show a failure if the symlinks are not cooled

[EL-13934] Small Std PD configuration does not fit into projects with a default quota

In GCP the small std PD configuration does not fit in the minimal capacity of 4096GB of a default project quota.

[EL-13656] When ClearTier policy is set to 100% primary, no scan should run

When the ClearTier policy is set to 100% primary there are no files to cool and there is no need to run the scan process.

[EL-13597] Adding replication agents does not add them to the active RA list

When adding replication agents they are not added to the active list and the storage nodes are not using them as GW to the object store.

[EL-13596] I/O error when listing an empty object snapshot

Running “ls -l” on an empty snapshot in the object store resulted in I/O error

[EL-13569] EMS required project.owner permissions

Permission requirement was changed to project.editor.

[EL-13557] It is not possible to delete the upgrade tar file from the UI

It is not possible to delete the upgrade tar file after a failed NDU attempt.

[EL-13536] Events less than critical where not shown in the dashboard

The dashboard event icon was only showing critical event masking warning level events.

[EL-13536] Potential deadlock on a truncate operation

In some rare cases, a truncate operation might get into a deadlock due to wrong order of locking and truncate resource requests operations.

[EL-13507] Privileged clients flag is not clear

An explanation was added to make it clear that Privileged Clients are no_root_squash clients.

[EL-13503] Trying to recover a system with one failed node didn’t work

Recovering a system with one failed node (delete a node in GCP) even after marking it as failed didn’t work

[EL-13492] portmap requests where answered by UDP that was blocked by customer FW

The system sent the UDP answers from the nodes private IP and not the VIP and was blocked by the FW rules.

[EL-13423] Removing a GCP node from the cluster didn’t delete the actual instance

Instances where left in the project after being removed from the cluster and required manual intervention.

[EL-13374] Node encounter panic if both replication agents are down and the client accessed a file in the object store

In some rare cases, this could lead to a system lockdown.

[EL-13325] Old versions tar files were not deleted from the upload directory

Systems run out of space after multiple upgrades due to old versions of the tar files that were not removed automatically.

[EL-13292] Potential NFS frontend starvation when running multiple operations on a single file

In some rare cases when running multiple attributing setting and write operations on a single file the front end NFS core run out of resources.

[EL-13198] Cluster recovery failed if one replication agent was unreachable

System activation failed if one or more of the replication agent were unreachable at that time

[EL-13177] UI blocked all ClearTier operations if only one replication agent was active

ClearTier requires high availability of the replication agents for normal operation but the UI should not block ClearTier configuration changes if one is down.

[EL-12766] DSM deployment does not preserve the management interface config

Rebooting will cause the management interface config to be lost.

[EL-12764] There is no way to remove a failed EMS from the UI

The remove EMS option is greyed out in the UI and there is no option to remove a failed EMS instance.

[EL-13875] Replication Agent deleted during system reactivation

As a result of a status update race condition, the replication agent was deleted from the system during system reactivation.

[EL-14303] VIP was set with the wrong netmask

The VIP settings were/24 instead of /32 making it impossible to set the VIP on the same network as the storage nodes IP.

[EL-13884] gsutil does not work from EMS

gsutil does not work from the EMS and does not allow to upload the logs to an external bucket.

[EL-14125] Core dump logs filled the replication agent disk

Too many core-dumps logs filled the replication agent disk causing it to fail

[EL-14123] Replication agent was created with the wrong memory config

Replication agent was created with 16GB RAM instead of 26GB RAM

[EL-14122] Cooling operation was marked as failed to fast

The EMS marked a cooling operation as failed while it was still in progress

[EL-14120] Running multiple snapshot cooling operations on the same DC failed

Running multiple snapshot cooling operations on the same DC is not supported. EMS should restrict the operation to one snapshot at a time

[EL-14100] ECFS node panic when adding or removing replication agent

In some rare cases, the ECFS node working with the object GW in the replication agent got panic when the replication agent was removed.

[EL-14089] ECFS node panic during multiple rSync operations

In some rare cases while uploading a high amount of files using rSync or similar to the same directory from multiple clients cause an ECFS node panic

[EL-14060] Node removal didn’t delete the actual instance in a multi-AZ deployment

Removing a stopped node in a multi-AZ deployment did not delete the actual node instance

[EL-14241] ClearTier was limited to 1M files in a directory

The number was increased to support 10M files in a single directory

[EL-14060] Node removal didn’t delete the actual instance in a multi-AZ deployment

Removing a stopped node in a multi-AZ deployment did not delete the actual node instance

[EL-14247] Add an option for an unmanaged load balancer

Add an option to support unmanaged routing rules or ILB including the option to upgrade such system

[EL-12796] Load balancer option is not working in shared VPC

Customers should use DNS round robin in shared VPC mode. The Load Balancer option is available in the UI but is not recommended.

Fixed Issues in ECFS 3.1.0.22

[EL-15042] Cold Upgrade didn’t restore Replication-Agent mounts

Running cold upgrade didn’t restore existing RA mounts and required manual intervention.

[EL-15037] Cold Upgrade failed because of Replication-Agent umount failure

Cold Upgrade failed to execute umount command on the RA and could not continue the operation

[EL-14972] Async-DR source and target site reports different RPO status

In Async-DR the source site reported meeting RPO while the target/remote site reported RPO exceeded.

[EL-14971] ClearTier delete operation failed with error cloud key doesn't exist

In some cases, when an object key was changed in a previous data traversal the delete operation of cold data fails and return an error.

[EL-14967] Async-DR failed to resume after system recovery from lockdown

In some cases, when an async-dr task is in an init state and the system went to a lockdown state, the task was not resumed after the system recovery.

[EL-14964] Cold upgrade EMS server is stuck in the startup state

After a cold upgrade, the EMS server was stuck in the startup state for a long time due to a stuck DB query.

[EL-14923] RA are not sharing index file and reading data from the object store

Running cooling operation using different RA was inefficient and forced each one to create a new index file by reading data from the object store

[EL-14866] Deleting snapshot from the object store failed with an error

In some cases, trying to delete a snapshot from the object store, failed with an error "missing in action".

[EL-14851] Replication agent had too many open connections to the object store

The replication agent kept too many open connections to the object store causing unnecessary load on the server.

[EL-14844] EMS sent too many operations to the RA

EMS didn't limit the number of operations is allocating to each RA causing load and imbalance

[EL-14841] RA opened a new TCP connection for each RPC command

The RA had too many open connections holding up to 1000 at a time.

[EL-14840] EMS recovery took 30 min

EMS restart took 30 min in a system with a high number of objects.

[EL-14827] ClearTier - internal export creation flooded the log

Internal export creation and deletion flooded the log hiding real system events

[EL-14819] ClearTier - version commit reported wrong dedup numbers

The version commit report did not report on multiple objects that were deduped

[EL-14767] ClearTier - object tier performance shows no activity when reading from the object

Reading from the .object was not shown in the system dashboard object tier performance tab

[EL-14762] ClearTier - DC snapshot view all like took too long to load

Loading the snapshot view all screen took too long in case of a high number of snapshots.

[EL-14760] ClearTier - the max number of RA was limited to the number of storage nodes in the system

It was not possible to add more RA than the number of storage nodes in the system. Current limit changed to 10 RA

[EL-14756] ClearTier - event log got full with failed to remove snapshot messages

If all the RA's in the system are busy the EMS sent snapshot delete tasks that failed and loaded the event log.

[EL-14754] ClearTier - readdir from .object didn't return the cold snapshots

In some rare cases, readdir from .object didn't return all the available cold snapshots.

[EL-14724] ClearTier - RA can't find a bucket since the version was deleted

A file stub was pointing to a version without a specific bucket when the version was deleted a new RA was unable to find the file via the version number and failed with an exception.

[EL-14712] ClearTier - EMS kept sending tasks to the same RA even when others are free

EMS kept sending tasks the first RA that looked free even if it was unable to perform them and didn't utilize other RA

[EL-14686] RA configuration - c-state was enabled

C-state was enabled in the RA configuration causing it to wait on CPU cycles from the hypervisor

[EL-14616] Create snapshot task was not restored after ECS failure

In case of ECS restart some snapshot creation failed and never created.

[EL-14574] ClearTier - NDU didn't restart services running on the RA

As a result, some services kept running with the old version

[EL-14543] ECS failed due to snapshot delete and update

ECS got concurrent operations to remove and update a snapshot and failed.

[EL-14446] NDU - NDU failed to remove old storage nodes

When working with manually created load-balancer, EMS was unable to update the routing rules and failed to remove the old storage nodes

[EL-14398] RA load was too high

The RA load was too high when running 4 parallel cooling tasks. The new value is set to 2.

[EL-14386] Log was flooded with cooling errors after NDU

Event log was flooded with "Cool_DC_Failed=5061" error messages after NDU

[EL-14368] ClearTier - the current log should be compressed

RA was compressing logs when they are rotated. To save disk space the current log of the RA should be compressed as well.

[EL-14364] Events - internal clients mount/umount should not be in the UI

Internal clients mount and umount were flooding the logs and hiding real events in the UI

[EL-14221] Log was flooded with cooling errors after NDU

Event log was flooded with “CREATE_VERSION_FAILED = 5046” error messages.

[EL-14763] Async-DR, continuous failures of incremental syncs

Incremental syncs kept failing due to a timeout issue.

[EL-14475] [EL-14388] ClearTier log was running in debug mode

The ClearTier log was running in debug mode consuming too much disk space

[EL-14387] ClearTier core dump consumed disk space

The ClearTier core dump consume too much disk space and cause the RA to hang

[EL-15104] Snapshot scheduler: Setting the scheduler delete or cool to 0 doesn’t work

O value was ignored by the backend and didn’t stop the cooling or the deletion

[EL-15087] ECS allowed 2 tasks with the same ID

ECS allows 2 tasks with the same ID to enter failing to identify they are the same task, while one is running, the other completes with an error, as it is completed it was released by EMS leaving the first one without an owner.

[EL-15073] ClearTier: multiple version creation failures

Multiple retries to move a snapshot to the cold tier failed.

[EL-15137] Async DR: Target site does not show an event on missing RPO

The target site shows missing RPO in the UI but does not have an event for it.

Fixed Issues in ECFS 3.1.0.27

[EL-15595] Async DR: file creation and directory deletion race condition

In some rare cases, a race condition between file creation and directory deletion caused the async process to panic.

[EL-15590] Async DR: EMS restart can cause replication to fail and never resume

In some rare cases, if EMS is restarted between marking the replication end time stamp and marking the task as done, a wrong last replicated snapshot might be created on each site causing replication to fail.

[EL-15225] Async DR: Target site does not report RPO as exceeding

Target system does not create an event in case RPO is exceeded.

[EL-15118] Async DR: Failback was not working according to spec

Failback didn't copy the last synced snapshot to the DC before starting the reverse direction replication.

[EL-15216] Async DR: REST API allowed creating a DC pair without RPO or role

It was possible to create a DC pair without setting a role for any DC or set an RPO.

[EL-15158] Async DR: UI didn't give an indication of DC becoming active

During DC activation there was no indication in the UI allowing multiple attempts.

[EL-15099] Async DR: EMS ignore ECS errors when trying to create an export

EMS was trying to create an export for the Async DR process but ignored the errors coming from the ECS and tried to continue with the replication.

[EL-14931] Async DR: Last replicated timestamp disappeared from UI

At the beginning of every sync cycle, the last replicated timestamp disappeared from the UI.

[EL-14898] Async DR: SQL error when trying to remove an export

In some cases trying to remove an export resulted with a SQL error " Mysql2::Error: Lock wait timeout exceeded" retry of the operation after 5 minutes succeeded.

[EL-15141] ClearTier: Readdir returns corrupted file name on a virtual directory

Readdir, but not readdir+, might return corrupted file name on a virtual directory such as .object or .snapshot in case the packet contains over 40 entries.

[EL-15029] Async-DR: UI does not report on replication error

The UI does not report on replication error until the time for the next sync

Fixed Issues in ECFS 3.1.0.31

[EL-15676] ECS handling of parallel operations

In some rare cases, when multiple concurrent tasks of the same nature were sent to the ECS, it did not clean resources fast enough it might get to a state of no resources.

[EL-15655] Async DR: missing traces in log

Async DR traces are missing from collected log after Async DR operation failure

[EL-15645] Async DR: EMS ignores export creation errors

When preparing a snapshot for replication the EMS ignores export creation errors and continues with the operation which might fail if the permissions are not correct.

[EL-15685] Async DR: replication failed with no LRCP

In some rare cases of frequent and multiple ECS and RA failures the last replicated snapshot is lost and the replication fails.

Fixed Issues in ECFS 3.2.1.8

[EL-13159] It is possible to change capacity while the system is in lockdown

EMS allowed the user to add/remove capacity while the system is in lockdown which can cause problems during recovery attempts.

[EL-13211] Adding RA does not add it to the UI

If all of the system's RA were deleted from the Google Console, but not from the system DB. Adding more RAs does not add them to the UI.

[EL-13158] Performance graphs blink out of the UI

In some cases, the performance graphs in the dashboard disappear from the UI.

[EL-13943] Dual failure of EMS and storage node might need manual reactivation

When working in a single zone HA mode failure of the EMS and one of the storage nodes will cause the system to enter a lockdown mode and require a manual reactivation

[EL-12907] System with single zone HA does not recover from double failure

If a failure occurs during the recovery cycle the system didn't retry the recovery.

[EL-13115] Replication agent looks like it was removed while it is still on

The system prevents removing the last RA in case ClearTier is operational, the GUI has a false indication that the RA is removed.

[EL-13077] Failure in a persistent disk caused the system to be in a loop without indication

Failure in a data persistent disk, while running in single zone HA, caused the system to be in a recovery loop without any indication to the user.

[EL-14183] ClearTier: Failed to cool a snapshot with space in the name

Snapshot cooling failed if the snapshot had space in the name.

[EL-13976] Multi-AZ: system with 6 nodes was installed in 4 zones

Multi-AZ system with 6 nodes was installed in 4 zones instead of 3.

[EL-13930] logs are not rotated correctly

Old logs were not rotated correctly and consumed disk space on the EMS.

[EL-13284] Add support to NDU with manually managed Load-Balancer

NDU failed when the cluster was installed with manually managed load-balancer.

[EL-13626] New system is created with ECS panic in the log

On initial system start-up ECS tried to load before the boot.conf was ready and panicked, resulting in an unnecessary log entry.

[EL-14317] NDU didn't finish due to Java boundaries error

NDU didn't finish and old nodes were not removed due to Java boundaries error.

[EL-14315] Nodes were not removed at end of NDU

The no-name-server configuration got reset causing node removal stage of the NDU to fail.

[EL-14313] Remove the log-collection "Auto/Manual" column in the UI

The value is never set and the column is always empty.

[EL-14179] Cluster reactivation failure after node removal

In a rare case of cluster reactivation after node failure during another node removal (test procedure), the cluster failed to restart.

[EL-13048] GUI: Window flash during EMS upgrade

During the EMS upgrade stage of the NDU the screen keeps on flashing and not showing any progress.

[EL-14455] NLM lock break CLI does not support IP addresses

Lock break CLI command ("elfs-cli system break-lock") supports only MAC address and not IP addresses.

[EL-14236] NLM sub range release left orphan locks

When a client unlocks a subrange of previously lock request, npal_locks does not copy the entire metadata of the lock to the new lock, which leaves an "orphan" lock, which cannot be released.

[EL-14029] ECS crashes when a node is turned off

ECS crashed when trying to get the NIC status of a node that was turned off.

[EL-13864] Terraform deployment failed when trying to set "lb_vip" to null

Terraform validation failed on available VIP verification when trying to set Elastifile load balancer.

[EL-14046] Changes to the NLM grant request mechanism when working with a VIP

Cluster outbound grant message was sent from the wrong IP addresses when working with a VIP and multiple clients.

[EL-13614] Snapshots: Multiple deletions of snapshots failed

The system prevents multiple snapshots operations on the same DC.

[EL-13060] ClearTier: Actual statistics of data tiering policy appear as N/A

If there are no files in the system or when all the files from the primary storage are cooled the statistics will be shown as N/A

[EL-14307] ClearTier: Object GW crashed on multiple accesses to a zero blub

A race condition might occur if two operations are done in parallel on the same zero blub causing the object GW to crash.

[EL-14213] EMRI: logs do not contain the EMS logs

Automatic EMRI logs created automatically when a node is fenced didn't contain the EMS logs.

[EL-13872] Stopped node deletion failed

Trying to delete a node that was already stopped failed.

[EL-13090] ClearTier: System might show snapshot size while there are no snapshots

During file creation, the system might report snapshot size while there are no snapshots in the system. The statistics are reported correctly after some time.

[EL-13068] ClearTier: bucket deletion didn't work

Marking the bucket for deletion when deleting a DC didn't delete the bucket.

[EL-13067] ClearTier: GUI data size didn't show the object size

Data size column in the DC list UI didn't include the object data size.

[EL-12989] Single Copy recovery might fail due to device error

In rare case, auto system recovery failed due to high latency when adding a device.

[EL-14768] ClearTier: Snapshot export can be deleted by the user

The internal snapshot used by the cooling process can be deleted by the user

[EL-14258] ClearTier: Snapshot deletion from the object store failed

Snapshot deletion error was not handled correctly leaving deletion operation in progress.

[EL-14074] Multiple events are sent to the UI

Multiple "Calculated new data distribution map" are sent to the UI.

[EL-14003] Performance graph misses 10 min of data after NDU

Performance graph does not show data for about 10 min after NDU.

[EL-13894] Async DR: failed to add RA if all other RA failed

If all the RA in the system are in failed status, it was impossible to add a new RA.

[EL-15641] Async DR: GUI error message hides the OK button

In some cases when the error message is too long it hides the make active OK button.

[EL-15505] Local disk, the system removed the wrong node

The system failed to remove a node with failed disks and removed another node instead.

[EL-15406] EMRI process can fill the system disk

EMRI process was using the /tmp dir and can fill the system disk.

[EL-15280] Async DR: all sync exceeded RPO but showed as synced

Sometimes all DC are showed as synced even though they've all exceeded RPO.

[EL-15213] ClearTier: readdir on .Object should be answered locally

Readdir on .object was sent to the GW instead of being answered from the local file system.

[EL-15168] Async DR: UI dialog stays open after clicking stop

Stopping replication keeps the dialog open without disabling the button making it unclear if the button was click or not.

[EL-15027] Async-DR: setting changes are not shown in the UI until a refresh

Replication setting changes are not shown in the UI until it is refreshed.

[EL-14875] System deployment failed when customer defined FW rules manually

The system didn't check if the FW rules were already defined.

[EL-14874] ls on .object takes a long time

Running metadata operations on .object takes a long time.

[EL-15762] ClearTier: running cooling operation after lockdown failed

Running the same cooling operation after the system recovered from a lockdown failed.

[EL-15972] FW rules should support elfstop ports

Elfstop ports should be enabled by default to support debuging.

[EL-15840] ecs events are queued and cause system lockup

The event suppression mechanism starts to slow when the event log is big (over 2.5M entries) and as result events are queued and can cause a node panic or even a system lockdown.

[EL-14684] long-running CLI commands might end with a timeout

Running cli commands that take too long might end with a timeout and "401 Unauthorized" error message.

[EL-15688] Async DR: Cold upgrade fails when Async DR operations are in the background

The cold upgrade stops the cluster causing async DR operations to fail which cause the RA to be unresponsive.

[EL-15703] ClearTier: failed node's tasks are overloading another node

When a node fails all the tasks were moved to another node instead of being redistributed.

[EL-13124] Cloud instances are deployed without delete protection flag

Users are advised not to delete any Elastifile instance that is still in use by the cluster.

[EL-12941] New nodes are shown as faulty

When adding new nodes to the system, they'll appear as faulty until the rebalance process is done, since they do not have a full copy of the data.

[EL-13870] Add capacity in AWS with LB configuration might not work

In some cases adding capacity to a system in AWS when working with network load balancer might not add the node to the cluster.

[EL-15490] Async DR: It is possible to create export on passive snapshot but not remove it

It is possible to create an export on a snapshot on the passive data container but it is not possible to remove it as long as the DC is passive.

[EL-15727] ClearTier: DC list does not show DC size if object store is not defined

In case the object store is not defined in a system the DC list does not show the total DC size.

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

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

[EL-6749] Used capacity is wrong after EMS stop and start

Stopping the EMS server can cause a used capacity drift.

[EL-13599] During NDU all the new nodes looks as faulty

The system adds new nodes as part of the NDU process but marks them as faulty until the NDU is done.

[EL-13886] File attribute change failure

In some rare cases changing multiple files' attributes failed.

[EL-14147] ECS panic during initial deployment

Due to the deployment order, ECS started and panic since the system was not ready yet, causing an unnecessary log creation.

[EL-14293] NDU process does not add the cluster hash to the RA

When upgrading from a system that was not supporting the cluster hash routing rules, the process didn't update the system RA with the new hash.

[EL-14298] Showmount didn't return the full list of exports

Showmount was limited to 64KB and didn't return the full list. The new limit will be 600 exports

[EL-14309] GUI: rebuild look stack at 25%

When trying to remove a node from a system with 90% utilization, the remove process fails as expected but the GUI shows the rebuild as stuck at 25%.

[EL-14331] Adding 4 nodes caused a system lockdown

Trying to add 4 nodes to a system caused a system lockdown due to map distribution limitation.

[EL-14442] Async DR failed due to FW rule

Even though Async DR is using TLS and port 443, it is looking for port 80 due to backward compatibility. Port 80 was missing from the automated FW rules causing Async DR to fail.

[EL-14653] EMS DB reached 10GB

EMS DB size grows due to empty snapshot statistics records.

[EL-14711] ClearTier:Object store synchronous operation might fail

In some cases creating a version or linking a DC to a bucket failed

[EL-14718] Snapshot Scheduler: changing the cooling time to 0 didn't stop cooling

Zero value was not treated as never but as an immediate cooling. Added the option to specify "Never" in the UI

[EL-14757] ClearTier: RA should handle more snapshot delete operations

As snapshot delete operations are light in resources, the RA should be able to handle more operations in parallel.

[EL-14759] Internal export events should be hidden from the UI

Create and delete export for ClearTier and Async DR should not overload the GUI.

[EL-14805] Using an invalid hash during the deployment issue multiple errors

Using an invalid hash during the installation resulted in service account scope errors.

[EL-14808] Insufficient service account permissions were reported as warning

Insufficient service account permissions, that prevented deployment, were reported as warning only.

[EL-14810] Quota validation test

Quota validation test was removed from the first stage and is done before the cluster is deployed.

[EL-14868] ClearTier: Allow multiple concurrent operations on a single DC

User shall be able to delete multiple cold snapshots on a DC or delete a snapshot while another is being cooled.

[EL-14873] NLM: Grant UDP messages didn't reach the client

If client listening UDP port changed the NLM grant message never reached it.

[EL-15001] ClearTier: GUI percentage has too many digits

GUI shows too many digits after the decimal point.

[EL-15034] Snapshot scheduler not responding

On DSM system after EMS failover it is possible that the DC status won't be set, causing snapshot scheduler creation to hang.

[EL-15080] EMRI: log collection should include eSync traces

Collecting logs didn't include the eSync traces.

[EL-15190] Disabling export does not disable cache

Disabling an export does not invalidate the caches allowing some I/O to continue being served.

[EL-15192] Async DR: Full replication after lockdown recovery

The async process started full sync after lockdown recovery instead of an incremental one.

[EL-15314] Async DR: System didn't prevent writing during role change

The system allowed to write to the active cluster while it was being demoted to passive.

[EL-15631] MySQL deadlock errors in the logs

The system had multiple MySQL deadlock error in the log.

[EL-15638] Sort by replication role was not working

In the GUI, trying to sort by replication role didn't work

[EL-15668] Cold upgrade failed if EMS had NFS mounts

NFS mounts on EMS prevented the cold upgrade from running. Elastifile mounts are removed and recreated, user mounts are removed. User should not manually mount NFS from the EMS.

[EL-15758] Large write failure may cause node fencing

In some rare case a large write failed and when retried caused a node to fence.

[EL-15787] NDU caused the name server setting to reset

No name server settings were reset after NDU.

[EL-13657] ClearTier: default buckets are changed to multi-regional

When setting ClearTier using the default one-click settings the object store bucket is created as multi-regional.
To create a regional bucket, please use the manual mode and specify the region.

[EL-14597] ClearTier: reduce the number of saved rotated logs

Replication agent disk might get filled due to the number of logs being retained

[EL-14459] ClearTier: adjust the level of logs

Replication agent disk might get filled due to logs using debug level settings

[EL-14517] ClearTier: avoid reading the stderr file if it too big

CCWEB process in Replication Agent failed due to a large log file

[EL-14478] ClearTier: Change waiting period for ELCC status

CCWEB process waiting status should be increased to 120 seconds to avoid timeout if ELCC is busy.

[EL-13623] ClearTier: Trying to cool/warm a large directory can cause OOM

Directory scan was held in memory and failed if number of files was very large (around 2M files)

Fixed Issues in ECFS 3.2.1.19

[EL-16601] ClearTier: Warm DC left files on the cold tier

Warming a full DC using the warm DC command left some file on the cold tier due to limitation in number of open files per Replication Agent
Note: If during the warm DC command, the user performs move on existing data it may cause files to be left on the cold tier.

[EL-16591] ClearTier: unaligned big writes over cold data caused a system lockdown

Large unaligned writes, over 1.5MB, over cold data caused a system lockdown.

[EL-16584] FW rules missing local subnet

Automatic FW rules during system deployment didn't open the local subnet for all traffic, forcing the user to add NFS client tag to all clients.
Note: If this FW rule is too open, you can remove it and use tags instead.

[EL-16561] Create Snapshot command during node failure should be blocked

In single zone HA mode, any operation should be blocked if a node failed, until the system recovers

[EL-16543] NDU failed if password had special characters

Temp password created for the service included special characters that failed the EMS upgrade.

[EL-16601] ClearTier: Warm DC task failed

In some rare cases the task to warm DC failed when trying to warm a DC with high number of files

[EL-16467] ClearTier: Colling files faild

Cooling files that ended with backslash "\" failed.

[EL-16466] ClearTier: large DC scanning was very slow

Scanning a large DC with over 2M files was very slow.

[EL-16458] ClearTier: Change Default Policy

Default cooling policy is set to be based on files' access and modify time and not on system utilization percentage.

[EL-16628] Async DR: failed to replicate a DC when cluster size is over 20 nodes

When the cluster size is larger than 20 nodes trying to replicate a DC got an error in the replication agent.

[EL-16828] ClearTier: failed to delete a cold snapshot

When a file appeared in several directories in a cold snapshot, deleting this snapshot may fail.

[EL-16817] NDU: NDU failed if replication agent was patched manusaly

In some cases if the RA is patched manually the NDU process may fail.

[EL-16812] ClearTier: All logs were saved in a single folder

All the replication agent logs were saved in the same folder.

[EL-16808] eMRI: process didn't collect RA logs

RA logs were not collected during the eMRI log collection process.

[EL-16675] EMS recovery process didn't work

EMS recovery process didn't work and could result in node deletion

Fixed Issues in ECFS 3.2.1.20

[EL-16974] Async DR: ECS restart can load tasks that can't run

In case of multiple background tasks, an ECS restart can reload tasks that has other dependencies and cannot be completed, thus preventing other tasks from running.

[EL-16941] Upgrade logs are not published to stackdriver

Upgrade logs are not pushed to stackdriver even if configured.

[EL-16935] Failed to add RA after upgrade

The RA host and availability zone were not stored in the internal file system and become empty after upgrade, thus preventing adding additional RAs.

[EL-16934] RA is not deleted from internal filesystem

Deletion of RA only removed it from the EMS and not the internal file system and can cause an upgrade process to fail.

[EL-16923] ClearTier: Snapshot cooling failed on unsupported filename encoding

If a snapshot had files with non UTF-8 characters the snapshot cooling failed with not clear error. Non UTF-8 files will be added to the exception list.

[EL-16883] Unmanaged LB VIP was not persistent

If a storage node in a system with an unmanaged load balancer was rebooted the VIP configuration was removed the node got stuck in a fence loop.

Fixed Issues in ECFS 3.2.1.30

[144984164] ClearTier: Snapshot cooling failed during directory scan

During a scan of a large directory the cooling processed timed out and failed to cool the snapshot

[144678963] DSM: Unclear error during deployment

If the deployer failed to install the EMS YUM repository the error was not informative enough

[146419901] UPgrade - UI does not show indication of running upgrade

[144134099 EL-15758] Large write failure may cause node fencing

In some rare case a large write failed and when retried caused a node to fence.

[145965357] EMS froze due to 500k log entries

Log entries that were collected for the datadog service (now deprecated) caused the EMS to freeze

Security Fixed Issues in ECFS 3.2.1.30

The following are security fixes included in this release. While there are no known exploits, customers are advised to upgrade.

[144613366] Remove obsolete private keys from the image

[144128707 EL-16334] Remove internal users from image

[145898615] UI does not work after MacOS upgrade

After upgrading macOS to Catalina, the UI stopped working due to unsupported configuration of the self signed certificate

[145383065] Potential buffer overflow in XDR

[145740602] Potential denial of service in XDR due to wrong state of a packet

[145388559] Potential out of boundaries in NLM server

[146534086] Sun RPC crashed as result of fuzzing test

[146533207] Potential overflow condition

[144206263, 144481370, 144206484] Outdated jquery module in web UI

Fixed Issues in ECFS 3.2.1.37

[148434422] Checkpoint version_id advanced too fast

An improvement implemented to prevent checkpoint version_id collision in case of node fencing cause the uniqe_id service to advance the version too fast and cross the 4 billion threshold. In some rare cases this can cause issues in the system during and after upgrade.

[148342150] Snapshot: added validation to the snapshot external ID

Overlap between a read only snapshot ID and the read/write snapshot ID caused a problem when trying to delete a data container

[148340057] Delete Data Container (DC) failed but the DC was marked as valid

A DC deletion failed due the bug 148342150 but the role-back operation marked the DC as valid.

[148323840] Snapshot: rollback process didn't clean up the created objects

When a snapshot creation failed the rollback process didn't clean up all the objects.

[147185896] Internal load balancing was set to core and not to node

Default internal load balancing was set to core affinity and not to node affinity, causing a core overload and performance impact.

[149199363] DSM: Error in parsing the CPU interrupt file

DSM deployment may fail since there was an error parsing the CPU interrupt file

[149459186] Snapshot: create a directory inside a read/only snapshot caused a SW panic

System panicked when a copy process tried to create a directory inside the ./snapshot directory

[149359301] A failed autorecovery prevented the system from running a second recovery

System auto-recovery failed after the first attempts but a second one didn't run

[151030251] Disable incremental rebuild for versions that do not support version_id leapfrog

[151030468] Don't use the leapfrog version_id if version is not updated to supported version

In a rare case, after upgrading the version didn't flip to the new version but the version_id service was using leapfrog and progressed too fast.

Security Fixed Issues in ECFS 3.2.1.37

The following are security fixes included in this release. While there are no known exploits, customers are advised to upgrade.

[146622155] Potential SQL injection in data_containers API via order argument

Fixed Issues in ECFS 3.2.1.39

[153237966] NLM operation impact cluster performance

Excessive use of NLM operation from over 500 clients caused an increase in create_file operations and resulted in high latency and impired cluster performance

[155580201] File locks might leak and caused a DB lockout

As a result of a race on the nlm aux-file the lock operations on the aux-file are not performed in the same order as sent to the sunrpc nlm server from one client machine and might lead to lock leak.

[155686622] Non-Disruptive Upgrade didn't flip the version when upgrade is done

Due to a timeout on the remove node process, the Non-Disruptive Upgrade (NDU) didn't flip the version at the end of the NDU

[155539308] Race condition when adding a node can cause data integrity issue

A rare race condition in the initial phase of adding a node can cause a data integrity issue.

[148993626] Inconsistent node list after Snapshot Officer failure

If the node hosting the Snapshot Officer process is fenced the SO node list becomes inconsistent since it contains the fenced node.

Known Bugs & Limitations

Elastifile deploys a version based on 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 3.2.1. The Elastifile roadmap is subject to change without notice.

The following section lists known limitations in Elastifile Version 3.2.1.

General Limitations

- Hyper-Converged mode is not tested in this version and it is not approved for production or upgrades
- Version 3.2.1 assumes homogeneous configurations. In the dedicated storage mode, all storage nodes must be identical.
- Elastifile clusters in Version 3.2.1 are limited to 96 nodes (testing limitation) for in cloud deployment of over 32 nodes please contact Elastifile support.
- Force-Reset should not be used in production, if you need to do force-reset in your testing, please ensure to umount all the client and re-mount after the reset is done.

High Availability Limitations

- Failure domains are not supported in the Dedicated Storage Mode and AWS deployments, in GCP multi availability zones are supported.
- Failure to tolerate = 2 is not supported in GCP deployments
- If a system that is configured with Failure 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
- When working in GCP using Single Zone HA the service recovery time will change according to the system utilization and change rate, customers are encouraged to test this feature before moving to production

Asynchronous Disaster Recovery

- RPO depends on the replicated data set, amount of changes and the throughput of the connection between the two sites. Minimal support RPO is 30min
- Bi-directional replication is supported but may impact the minimal possible RPO.
- Concurrent replications are 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. Users are advised to add more replication services or increase the core count and RAM in case more concurrent replications are needed.
- Async DR creates two snapshots per RPO, if the snapshot shipping time is longer than half the RPO time, momentary RPO miss will be reported.
- If snapshots deletion takes a long time and the async replication process has a backlog of 10 snapshots to delete, per DC, it will pause the replication. This can happen in a very high change rate such as loading data to the cluster.
- Timestamps in the UI are translated to the local time of the user's browser. Snapshots names are using UTC timestamps

Protocol Limitations

- NFS v3 is the only protocol supported for Version 3.0.1 (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 a performance impact in Version 3.1.0

Backend Store Limitation

- Hot addition or removal of devices (SSDs) are not supported.

Networking Limitation

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

ClearTier Limitation

- Only one object store is supported
- There is no option to disconnect the object store if there is a DC connected to it.

- Upon enabling the service the system will reconnect the DC to the object store automatically.
- The cooling policy is global but it is possible to disable it per DC
- ClearTier require two replication agents to ensure data accessibility
- In DSM installing the replication agents for ClearTier require management network.
- Deleting the associated object tier buckets, via the cloud console, without moving the data back to the primary storage will result in data loss.
- Concurrent deletion or creation of snapshot in the object store is not supported and will result in an error.
- If there are not enough replication agent resources a manual cooling operation might fail

SnapShot Limitation

- Snapshots per DC are limited to 255 on the primary storage
- Snapshots per DC are limited to 2559 on the object store
- If snapshot limit is reached newer snapshot will not be created unless older snapshot are deleted either manually or by the scheduler
- There is no automatic snapshot rotation

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.
- Progress indication of long tasks is progressing according to the performed tasks and not according to the tasks length.
- During In Cloud NDU the cluster resources will double, please ensure your project has enough quota before starting the NDU.

Known limitations

In addition to the above, Elastifile Version 3.2.1 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 notifications 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 the case of EMS failover, there will be a short gap in the statistics and events due to the asynchronous sync mechanism in the current version

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

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

[EL-5339] NDU event is missing

There is no event when DNU starts and finish

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

The progress bar indicates 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-11551/EL-11573] 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-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-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 a 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-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 these settings.

[EL-13121] Storage node appear as faulty during recovery

When a storage node rebooted, it will appear as faulty even though it is in a recovery stage.

[EL-13114] System will not show version mismatch between EMS and Storage Nodes

If the upgrade process was interrupted before the storage nodes are upgraded the system will not show a mismatch indication but may fail to reactivate the system if needed. Please run the NDU again to complete the process.

[EL-13081] ClearTier cooling process jumps from 66% to 100%

The progress indication is counting steps in the process and not the actual data being cooled.

[EL-13080] Can't create object tier without API access

If the Replication Agent is installed without public IP the VPC it is installed on must be configured for private IP API access.

[EL-13002] No indication when replication agent rebooted

The system does not have an indication in case the RA was rebooted

[EL-12972] Object tier performance does not show cooling operations

The object tier performance view shows only the operation by clients to data that is on the object tier. Cooling operations are done by the replication agent against the primary tier and will be shown as read operations on it.

[EL-12746] the Main dashboard shows only primary storage statistics

When data is cooled the effective utilization might go down. Raw utilization will show the actual used capacity on the disks which might be higher due to metadata and background delete operations.

[EL-12744] Object tier alert level does not raise an event

Setting the object tier alert level does not raise an event in case it was crossed

[EL-12739] Snapshot scheduler only impact future snapshots

The snapshot scheduler retention policy only impacts snapshots created after the change

[EL-12646] There is no indication of the bucket name in the UI

Each Data Container is associated with a different bucket in the object store but there is no indication of the bucket name in the UI

[EL-11843] Region is not validated

The object store region name is not validated, inputting the wrong region will result in failure to activate object tiering.

[EL-11574] readdir+ does not fill the advertised payload

The return response of a readdir+ command does not return the full buffer that is advertised.

[EL-11232] Adding and removing the same storage node might fail

In DSM if a node is removed, reinstalled and added back the system with the same IP address the operation will fail due to SSH security. In order to solve this clean the known hosts' list in the EMS SSH

[EL-13354] Long operations are stopped in case of a node failure in Single Zone HA mode

In case of a node failure when working in a single zone HA mode, all long operations such as snapshot cooling, and deletion will be stopped.

[EL-13242] Adding capacity while a full zone is down will not add all the requested capacity

When working in Cross-Zone HA mode, during a full zone failure, adding capacity that should add nodes in the failed zone will only add capacity in the remaining zones.

[EL-13204] New DC wizard is using the last parameters from the last command

The new DC wizard is using the parameters from the last command even if it was executed via Rest-API, Compression, Dedup, soft quota, and hard quota are impacted.

[EL-13535] ClearTier cooling process might miss files age by 24 hours

The cooling operation is running once a day, files that didn't meet the cooling criteria in the last scan will be cooled in the next cycle, this might result in 24 hours delay in the cooling.

[EL-13970] NDU is not supported with OSlogin enabled

In GCP enabling or disabling OSlogin, after cluster installation will fail NDU and eMRI collection, please contact Elastifile support for a workaround.

[EL-15078] ClearTier: multiple GW connection errors when RA is restarting

When all the system replication-agents are restarted there are multiple events on GW client connection is down and ILM client connection is down.

[EL-15041] ClearTier: RA without an internal mount will still get cooling tasks

If an RA has a mount error and can't use the shared index file it will still get cooling tasks that will take a longer time.

[EL-15033] Async-DR: restarting target EMS results in a critical false error

The EMS of the target system falsely reports after a restart that all replication services are down

[EL-14959] Async-DR: User can delete internal shares used for replication

It is possible to delete internal shares that are used for replication which will result in Async-DR failure.

[EL-14615] ClearTier: no progress indication for cooling operation

Cooling operation have no progress indication

[EL-14525] ClearTier: Failed cooling task may block other tasks

In a rare case of a failed cooling tasks, it might block other tasks from being executed.

[EL-14008] System is created with errors in the event log

Events regarding nodes disconnection appear during system deployment and NDU

[EL-15157] Async DR: Stopping replication will result in error

Stopping replication without letting the last replication to finish will result in error reported on the last replication.

[EL-15193] Async DR: Starting Async replication while loading data can cause RPO miss

Loading data to a cluster is considered a very high change rate and can cause async replication to miss RPO or even to hang

[EL-15133] Async DR: Target site does not report on site connection status

The target site does not report a connection issue with the source site

[EL-15095] Async DR: Target site status is unclear

If the target site reports RPO is OK the status is clear but if the RPO is missed, the actual status is unclear and needs to be calculated from the last snapshot timestamp if exists

[EL-12020] Async DR: Can't remove site when first sync failed

If the first sync failed there is no option to delete the target DC since it is not active and there is no snapshot to make it active from. Use the CLI to force promote the DC and delete it.

[EL-15604] Async DR: Failover with switch role failure allow access to inconsistent data

If a failover with switch roles operation fails, there is a chance that the target data container may become active with inconsistent data. Please run the command again and verify the success before writing data and replicating it to the other site.

[EL-15597] Async DR: Create an export on a snapshot used for async replication

It is not possible to create an export on a snapshot created by the Asynchronous replication, on the active data container, as long as the internal export exists.

[EL-15546] Async DR: Node restart can hurt RPO

If a node is restarted multiple times the asynchronous replication load becomes unbalanced and eventually hurts the ability to meet RPO.

[EL-15531] Async DR: Wrong error while trying to make a DC passive

Trying to make a data container passive, while the system is down, returns an unclear error "Failed to save the record".

[EL-15526] Async DR: Difference between UI and replication log

In some cases where the sync operation takes one second longer than the RPO the log on the source system will report the RPO as exceeded while the log on the target site will report the status as sync.

[EL-15518] Async DR: Running failover command on multiple DC may fail on some

It is not possible to run the switch role command from the CLI on more than 4 data containers in parallel, using REST API and doing them in serialized fashion will work.

[EL-15508] Async DR: UI is not responsive during failover operation

During multiple operations of switch role, the UI (including GUI, REST API, and CLI) might be slow to respond.

[EL-15458] Async DR: Failover operation finish but there are background operations

Switch role UI and CLI operations finish but there are background processes that are still running until both sites get to the desired state

[EL-15446] Async DR: Targe DC is marked as connected when the source site is down

The passive (target) site is shown as connected even if the active (source) sites EMS is disconnected.

[EL-15436] Async DR: After system reactivation, the system reports meeting RPO

If one site was reactivated (from a failure or shutdown) the reactivated site will show all async status as OK while the other site will show the Meeting RPO as failed. After a short while, the statuses are synced again.

[EL-15433] Async DR: Misleading status in the replication log

The log entry in the site that is being made passive is reported as “promoting” instead of “demoting”.

[EL-15432] Async DR: Replication log and GUI are not synced

In some cases where the sync operation takes one second longer than the RPO the log on the source system will report the RPO as exceeded while the log on the target site will report the status as sync.

[EL-15425] Async DR: RPO looks OK when sites are disconnected

In some cases when the sync operation is suspended or disconnected, the source site shows the connection as disconnected (in red) and the RPO as OK, the more severe status is the correct one.

[EL-15414] Cold Upgrade can’t start if node removal is in progress

It is not possible to run a cold upgrade if there are other maintenance tasks running in the background

[EL-15557] Async DR: snapshot are not deleted if failover fails

The snapshots created fro the failover process are not deleted if the process fails. Once replication resumed the snapshots will be cleared.

[EL-15717] Async DR and ClearTier: It is possible to move the PR snapshot to the object tier

The PR snapshot used for Async DR role change can be selected to be moved to the object tier. The operation eventually fails.

[EL-15716] Snapshot Scheduler: GUI does not reflect the correct status of multiple scheduler configurations

It is possible to configure multiple snapshot schedulers per a single DC via the CLI but the GUI does not reflect the status correctly, including if the schedulers are active or not.

[EL-12899] Application type DC wizard supports only IP address

When creating an application type DC the wizard only allows IP addresses and not subnets. Later when editing the DC subnets are supported.

[EL-13155] ClearTier: need to run the auto cooling command twice

The first run of the auto cooling command only scans the cluster data, the second run will start cooling data.