CLEONDRIS

Cleondris Technical Implementation Note | TIN-2090

# SnapGuard Technical Deep Dive

**(Working Copy, 2023-10-01)**

Cleondris GmbH, Switzerland

Ocotber 2023

## IMPORTANT

The information given in this technical implementation note represents current internal planning for Cleondris and can be subject to further changes without further notice. As such, this document is subject to change and may be changed by Cleondris at any time without notice. The information is not intended to be binding upon Cleondris to any particular course of business, product strategy and/or development.

## Table of Contents

# 1 Introduction

SnapGuard is a Cleondris software product which helps customers to protect their valuable CIFS data on shared NetApp Data ONTAP systems. The product features a unique FPolicy based CIFS firewall that is dynamically hooked into ONTAP and stops attacks from malicious clients.

SnapGuard can be installed as a virtual appliance (OVA) within an hour and can be attached to NetApp Data ONTAP systems easily. It does not require architectural changes to existing NetApp ONTAP setups.

This quick start document describes the architecture of SnapGuard, the requirements for software hardware and network, as well as the involved communication flows. There is also a section dedicated to the involved workflow when performing an initial SnapGuard setup.

## 1.1 What is SnapGuard from Cleondris?

SnapGuard is the first product worldwide which can actively protect NetApp NAS data from access by malicious clients.

No agents are needed on the clients, as SnapGuard directly connects to the storage system and can monitor all accesses. The NAS protocols CIFS and NFS (version 3 and 4) are supported.

SnapGuard can - depending on customer requirements - trigger various actions if it detects that a client is showing undesirable behavior:

- Stop all further accesses of the client
- Trigger an emergency (out-of-order) snapshot of the affected NetApp volume
- Alarming of external monitoring systems via e-mail, SNMP or Syslog (direct splunk integration)

Monitoring possibilities

- SnapGuard supports the NetApp native FPolicy and EVTX audit log mechanisms (see below)
- SnapGuard can check changed Office files for encryption/damage

## 1.2 What is NetApp FPolicy?

NetApp FPolicy is a mechanism within NetApp Data ONTAP that allows you to include an external "FPolicy" server that is consulted by ONTAP for selected data accesses. The connection to the FPolicy server can be either asynchronous (ONTAP informs about a data access, but does not expect a response) or synchronous (ONTAP informs about a data access, but does not release it until the FPolicy Server responds).

It is important to understand that NetApp itself does not offer FPolicy server products. NetApp only provides the FPolicy communication mechanism, and selected implementation partners (like Cleondris) can then implement an FPolicy server. Originally FPolicy was developed for

quota management, HSM and continuous file indexing. In 2016 Cleondris was the first partner worldwide to develop an FPolicy server with ransomware protection.

The communication between an ONTAP data SVM and an FPolicy server is handled by a NetApp proprietary protocol, the required TCP connections are established by the data LIFs (interfaces) from the SVM to the FPolicy server.

In larger installations, several FPolicy servers are usually used, on the one hand to distribute the load, on the other hand to ensure that the responsible FPolicy server for a SVM is “close” (latency-wise).

For SnapGuard, we use the term "FPE" (FPolicy Engine) for the FPolicy server mechanism, this is also in line with the ONTAP Console, where the external servers are called "engines".

The Cleondris FPE component (which implements the FPolicy server) can also be used in different ways:

- For small and/or test installations, the FPE integrated in SnapGuard can be used - the SnapGuard Server serves as the FPolicy server for all SVMs
- For larger installations, dedicated servers (Windows or Linux) can be used on which the Cleondris DMT Software runs. The Cleondris DMT Software (Data Manager Tools) is a scaling solution, which also contains a built-in FPE.

In principle, an FPE can be assigned to several SVMs, but in the case of large data volumes, it is advisable to use one dedicated FPE per SVM.

SnapGuard has a highly scalable architecture in which hundreds of SVMs and FPEs can be managed simultaneously. As the system grows, more DMT servers can be added at any time, thus increasing the number of available FPEs.

### How is FPolicy configured?

Customers using SnapGuard do not need to configure anything on ONTAP. SnapGuard has full ONTAP integration and automatically creates the required FPolicy configuration (e.g. which volumes need to be monitored and how, what are the IP addresses of the FPEs, etc.) based on the protection settings in the SnapGuard web interface.

## 1.3 What is EVTX Audit Logging?

NetApp Data ONTAP has a built-in auditing mechanism that stores audit data in the so-called "EVTX" format. These EVTX files can be viewed with the Windows Event Viewer.

The audit logging parameters (which accesses are monitored) can be configured by the administrator either in the filestem (using SACLs) or (invisible to clients) at the ONTAP volume level ("SLAG" - Storage Level Access Guard). The EVTX data is then collected per SVM in a dedicated volume.

The advantage of EVTX Logging over FPolicy is the lower latency effect and direct integration with ONTAP. However, using only native EVTX has the disadvantage that no rules can be defined on ONTAP how to handle unwanted access (e.g. deleting many files).

SnapGuard offers the following features for EVTX files handling:

- There is an integrated EVTX viewer which greatly simplifies the viewing of audit logs. Unlike the Windows Event Viewer, it can be filtered directly by path, username or volume. Thanks to the integration with ONTAP, it is also possible to view EVTX logs located on protected volumes where CIFS access is not possible.
- SnapGuard can monitor the EVTX log tail and trigger actions similar to the integration with FPolicy.

## 1.4 What is Snapshot Scanning?

Snapshot scanning allows for a firewall to periodically check the indexes of a volume for patterns in a pattern pool. This can be used in air-gap backup scenarios or to check if new patterns are present in existing data. To use snapshot scanning for a volume, several things need to be configured:

- A pattern pool must be used to manage the patterns.
- The snapshot scanning option in the ruleset which protects the volume must be enabled.
- An index for this volume must be configured.

If all these conditions are met, after a new index of a snapshot is created, it is checked for the patterns in the pool. The matches found are displayed in the pattern pool view in the “Files” and “Trend” column. The “Files” column indicates the number of files in the last index that match this pattern. The “Trend” column indicates how many files matching this pattern were created in the last three days.

## 1.5 What is a SnapGuard Firewall?

A SnapGuard Firewall defines how accesses to the volumes of an SVM should be closely monitored. From a technical point of view, the SnapGuard Firewall is a sub-component of an FPE.

The following settings are made in a firewall:

- Protection settings for SVM volumes. A so-called ruleset can be assigned to each volume (see below).
- SID resolution for the SVM. FPolicy only communicates SIDs on accesses, the FPE must resolve the SIDs to readable usernames using a domain controller connectio
- File verification settings. If the firewall should monitor changed files, a CIFS user is needed to allow the firewall to access the SVM.

- Local firewall exceptions. In large installations often a small number of rulesets are used on many SVMs. By defining exceptions on SVM level (e.g. "allow arbitrary access from server "10.1.2.44" on this SVM) it is avoided that specific rulesets have to be created for each SVM.

The protection settings for an SVM are always set as a firewall configuration. The FPE required for the protection can be changed at any time, SnapGuard then informs the new FPE about the SVM to be protected and reconfigures ONTAP so that the new FPE is used as FPolicy Server.

## 1.6 What is SID resolution?

The events generated by the NetApp system (file create, directory delete, etc. etc.) also contain information about the user who initiated the action. However, the user information (in case of SMB) is only available in the form of a so-called SID, which is short for “security identifier”. In the Microsoft security architecture, every account or group has a unique SID. To be able to show the respective human readable username or group name, a SID must be looked up in the active directory.

If SID resolution is not configured, SnapGuard can only show the numerical SID for each access.

To configure SID resolution, SnapGuard must be configured with the hostname of a domain controller and with the username and password of a domain user. SnapGuard is then establishing a SMB2 connection to the domain controller and authenticates the connection with the given user and password. This user does not need to have any special permissions, every member of a domain is allowed to contact a domain controller and ask for the translation between human readable usernames and SIDs (the so called “LSARPC” protocol, which is available on the hidden IPC$ share).

By default, SnapGuard is using NLMP/NTLMv2 to authenticate the SMB2 connection to the domain controller. If NTLMv2 is not desired, Kerberos is available, too. When using Kerberos, SnapGuard is first requesting a Kerberos ticket from the domain controller (via TCP port 88) which can then be used for SMB2 access to the same domain controller. If the active directory account of the user is configured with pre-authentication, then an intermediary step is involved, where SnapGuard encrypts a challenge of the Kerberos server to prove that it knows the user’s password.

The Kerberos ticket is encrypted by the domain controller using the user’s password, so only SnapGuard can decrypt it. This Kerberos ticket is then being used during the authentication of the SMB2 connection.

Please note: since the target service in a Kerberos ticket is always specified as a FQDN (fully qualified domain name), the correct hostname of the domain controller must be known, otherwise SnapGuard cannot request a matching ticket.

## 1.7 What is a Pattern Pool?

A Pattern Pool is a mechanism to manage lists of patterns and define which patterns should be blocked and which patterns are in legitimate use. Each pattern in a pool passes through different states. From newly entered to blocked, allowed or ignored. The different states are the following:

- **New**
  This is the state all newly imported patterns have. New patterns can eighter be entered manually or they can be fed from an automatically updated pattern list. When a pattern is imported, the source list and the date are recorded along with the pattern.
- **Checking Index (requires index license)**
  When a pattern enters the checking index state, a process is started that searches all available indexes of the volumes that this pool protects to find existing files that match the pattern. The hits this generates are stored and can be explored afterward.
- **Monitoring**
  A pattern in the monitoring state is evaluated on every client operation on volumes protected by this pool. If an operation matches the pattern no action is taken but this hit is stored in a database. The duration of this phase can be configured for each pool.
- **Decision Required**
  After the index check and the monitoring phase a pattern enters the decision required state. In this state an operator can decide, based on the statistics obtained during the previous phases, if a pattern should be included in the blocking list or if it should be ignored or explicitly allowed. SnapGuard can be configured to skip this state if the patterns had no hits in the previous phases.
- **Blocked, Allowed or Ignored**
  At the end of the pipeline, all samples end up in one of these three categories. File operations, on a volume that is protected by a respective ruleset, that match patterns in the blocked category are blocked. Patterns in the allowed list are explicitly allowed and patterns in the ignored list are not considered when evaluating a file operation.

### 1.7.1 Details about "Blocked", "Allowed" and "Ignored"

- Patterns in the "Blocked" state are unwanted patterns. This is to prohibit the use of certain file names.
- In certain cases it must be possible to override patterns in the"Blocked" category (e.g. to allow the file name "Test.locky" although "*.locky" is actually in the "Blocked" state). For this purpose the "Blocked" state can be overridden. When checking requests, if a filename matches a pattern on the "Allowed" category, no more checks for a "Blocked" pattern will be done.
- Patterns in the "Ignored" state are patterns that should not be processed by the firewall engine. The reason for the "Ignored" state can be an external pattern source that cannot be influenced. For example, if an external service provider provides a pattern list that contains the pattern "abc.pdf", but we do not want to use it, we could simply delete the

pattern from the pattern pool. However, during the next import process, the pattern "abc.pdf" would be re-imported, and we would have to delete it again. By assigning the state "Ignored", the pattern is no longer imported (it is already in the pool), but also does not lead to evaluation effort during pattern matching (since only "Monitored", "Allowed" and "Blocked" need to be checked, see below).

**How to protect a volume using a pattern pool**

To protect a volume using a pattern pool some configuration needs to be done. First the pool must be filled with patterns, this can be done manually in the pool directly or using a predefined pattern list which also can autoupdate from a remote list. After that, a firewall ruleset must be created. This ruleset must include at least one rule which has the "Match patterns from pool" checkbox activated. This ruleset is then assigned to a volume in a firewall. If there are multiple pools, it is also possible to select which of them should be used in this firewall.

## 1.8 What is a Firewall Ruleset?

A firewall ruleset defines the protection rules that should be applied to a volume. A ruleset can be assigned to any number of volumes, but each volume has a maximum of one ruleset (a volume without an assigned ruleset is not protected).

A firewall ruleset consists of the following components:

- A set of On-Access FPolicy rules. The On-Access FPolicy rules determine which file operations should be monitored by FPolicy and how to react to them.
- A set of EVTX rules (optional). The EVTX rules determine whether to monitor the tail of the EVTX log and how to respond to it.
- Snapshot Scanning settings. Enabling snapshot scanning settings configures the firewall to check volumes with a configured index against patterns in a pattern pool on a regular basis.
- Post-Access File Verification settings (optional). Enabling file verification will check Office data for corruption on volumes where the ruleset is assigned.

The main difference between EVTX and FPolicy is the integration in the data path: FPolicy is inside the data path, while EVTX is fully asynchronous and may observe events up to a minute later, due to the fact that the EVTX log is flushed by ONTAP on a regular basis and detection of changes in the log is not instantaneous. Furthermore, FPolicy can be fully controlled via SnapGuard, while EVTX needs to be configured at the ACL level of the filesystem.

**Structure of FPolicy rules in a ruleset**

When an FPolicy message is received from ONTAP, the firewall within the FPE determines the ruleset associated with the volume and scans the rules one by one.

Each rule has two areas:

- A definition of the type of accesses the rule should be triggered for (CIFS/NFS Operation Filter, Hostname Filter, User Filter, Path/Pattern Filter, Arrival rate of similar accesses by the user)
- A definition of which action(s) should be triggered when the rule "triggers" (block access, create snapshots, trigger external alarms)

**Emergency Snapshots**

FPolicy rules can contain the "Emergency Snapshot" action, which will be executed in parallel to other actions (e.g., blocking client access or firing an external event).

Emergency snapshots are created on the supervised volumes and can act as an additional restore point (typical standard snapshot policies do not create more than one snapshot per hour).

Emergency snapshots have the following properties:

- The name of an emergency snapshot is "**SGEE_ES**.yyyy-MM-dd-HHmmss", where "yyyy" is the current year ("2021"), "MM" the current month (01 – 12), "dd" the current day in the month (01-31), and "HHmmss" the current hour, minute and seconds in 24-hour format (e.g., "052356" or "231102").
- Emergency snapshot creation is rate limited. No more than one snapshot per minute per volume is created.
- The number of emergency snapshots on a volume is limited. If there are already 10 emergency snapshots, the oldest will automatically be deleted before a new one is created.
- These snapshots will then remain unless someone deletes them manually or new emergency snapshots occur (due to rotation as described above).
- Emergency snapshot creation does not block incoming SMB client requests: these snapshots are made asynchronously, i.e., possibly a few seconds after an attack starts.
- The default rate-limitation and max. snapshot count can be modified with advanced options, please contact Cleondris support if you would like to adjust these settings.

**How Rules are applied**

For each incoming client request the firewall applies the ruleset that has been assigned to the volume that hosts the data being accessed.

1. Rules of the ruleset are evaluated from top to bottom.
2. If several rules match a request, then the actions of the last matching rule (i.e., bottom most) are applied.
3. There is one exception: in case a rule matches and has the "*In case of match, ignore other rules*" checkbox checked then the following rules are not evaluated and the actions (if any) of the matching rule are executed immediately. Please note that this

property is called "*last rule*" in the ruleset overview screen as well as in the ruleset editor of earlier releases.

4. Rules with a "*rate filter*" additionally track the arrival rate of a user's requests that are covered by the filters defined in the rule (file operation(s), filename include/exclude patterns etc.). The rule "matches" only while the arrival rate exceeds the defined threshold.

The above is true for both "FPolicy" and "EVTX" rules.

Please consider the following example ruleset:

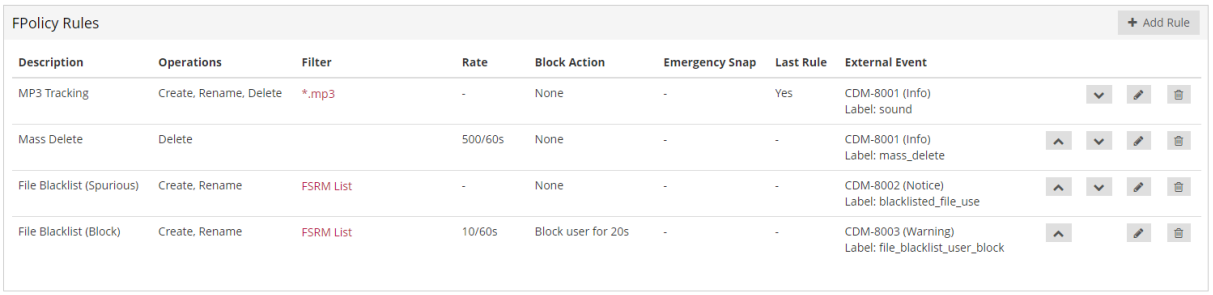


- The first rule tracks the creation/movement or deletion of MP3 files. An external event is generated (Cleondris event 8001, with custom label "sound"). This allows to collect such data, e.g., in a Splunk database. In case of a match, the evaluation stops with this rule, no further rules are evaluated ("last rule" property).
- The second rule monitors users that delete many files in a short timeframe. An external event is generated (8001, with custom label "mass_delete"). Note: since the first rule that tracks .mp3 usage is set to be the last rule, mass deletes of .mp3 files won't be tracked by this rule.
- The third rule monitors users that create or rename files with a filename that is on the "FSRM List" pattern. If such an access is happening, an external event is generated (8002, with label "blacklisted_file_use").
- The last rule monitors the same accesses as the third rule but applies a rate filter. If a user creates or renames files using a filename on the "FSRM" pattern list only occasionally, the third rule applies. However, as soon as the request arrival rate goes beyond 10 requests per 60 seconds, the fourth rule applies, and its actions override the third rule ("last matching rule wins"). Then, the user is blocked for 20 seconds and an external event of type 8003 with a custom label "file_blacklist_user_block" is generated.
- Note: no matter whether ".mp3" is on the "FSRM" pattern list or not, the evaluation of ".mp3" accesses stop with the first rule, due to its "last rule" property.

## 1.9 Rules that are linked with the Pattern Pool

The patterns in the pattern pool are processed in the firewall engine as follows:

- Basically, the patterns are only used in rules that have the "Match Patterns from Pool" checkbox set.

- The rules that use the Match Patterns from Pool " option can currently only be operated in blocking mode (the "None" blocking mode option is not available).
- They are therefore blocking rules, which block unwanted patterns based on the patterns in the pattern pool. However, you can set a threshold and a blocking mode to customize the behavior.

Each pattern in the pattern pool has an assigned state (New, Ignored, Index Check, Monitoring, Decision Required, Allowed, Blocked).

When evaluating firewall rules, only the patterns in the "Monitoring", "Allowed" and "Blocked" states are evaluated:

- If the rule matches one or more patterns in the "Monitoring" state, then this information is noted for the corresponding patterns ("Hit" in the Pattern Pool screen). No external event will be sent, no matter what is noted in the rule.
- If the rule applies to a pattern in the "Allowed" state, then the rule is skipped, no external event is sent, no matter what is noted for the rule. However, the statistics will be updated ("Hit" in the Pattern Pool screen).
- If the rule applies to a pattern in the "Blocked" state, and an (optional) rate level is reached, then the rule will result in a block. If the external event is switched on for the rule (CDM-8001- CDM-8005), an external event is triggered. Additionally, the statistics will be updated ("Hit" in the Pattern Pool screen).

The category "Allowed" is needed to override patterns in "Blocked" in exceptional cases.

Example: "*.locky" is on "Blocked" list, but "SpecialT*.locky" is on "Allowed" list. Then "*.locky" files are always blocked, unless the filename is e.g. "SpecialTmp.locky".

## 1.10 What is the Self-Service UI (SSUI)?

If a user is blocked due to triggering a firewall rule which a block action and a minimum block time has specified in the rule, the user cannot further proceed until either the block time has expired, or an administrator removes the block entry in the SnapGuard user interface.

In case of a setup with many protected SVMs that are part of different domains (e.g., in a provider setup), the administrators of these domains typically do not have access to the central SnapGuard appliance and therefore cannot unblock users in their SVM, unless they contact an admin of the SnapGuard appliance and ask the admin to unblock the user in their SVM.

For these scenarios, it is possible to enable a minimal UI (the ”SSUI”, self service UI) directly on the FPE (which runs anyway inside the network of the SVM). This UI allows selected admins of the domain to login, check the current block list, and unblock users.

The SSUI can be enabled on a per-SVM level and it is possible to specify selected users and/or AD groups that can login. Also, one can restrict the IP range that has access. The SSUI typically

runs on port 8080 on the FPE and uses an auto-generated self-signed certificate (i.e., one must access the SSUI using the URL "https://hostname-of-fpe:8080").

# 2 Best Practices for Firewalls

Once the SnapGuard software has been installed and there is access to an ONTAP system, the implementation of firewalls and rulesets can now begin.

To do this, you need to think about the following:

- Which data (SVMs, volumes) should be protected and what is the purpose of the protection (detect/block ransomware attacks, detect users accessing honeypot data, prevent working with certain file extensions, ...)
- Which actions should the system react with if unwanted behavior is detected?
- How should be alarmed?

## 2.1 Deciding on a Protection Strategy

When deploying SnapGuard, the security officer needs to decide on various aspects of the protection strategy.

### 2.1.1 Active vs Passive

Active protection blocks unwanted access automatically and as quickly as possible. Passive protection monitors the system in the same way but does not actively intervene. It merely informs the administrator and leaves further steps to him. With Cleondris not only both variants are possible, but also mixed forms. On the one hand, Cleondris allows a separate strategy for each protected ONTAP volume, on the other hand it is also possible to use an intermediate form of active and passive (i.e., triggering of emergency snapshots which have no influence on the users, but allow a very good restore point).

### 2.1.2 Undesirable User Actions

Typically, customers use SnapGuard to protect themselves against the following 3 scenarios:

- Detection of ransomware behavior which refers to the known ransomware filename patterns.
- Detection of new ransomware that does not attract attention by using a known filename pattern, or even does not change the filenames at all and only destroys data.
- Detection of Mass-Delete, to detect early the (even accidental) deletion of large amounts of data before the data runs out of retention.
- Detection of user access to honeypot files.

***Roadmap info: is working on further scenarios which can easily be recognized in the future:***

- Detection of random data movements in the file system (next-gen ransomware without encryption)

- Detect users who read large amounts of data, even over long intervals (data theft).
- Detection of incorrect movement of directories (user errors due to mouse glitches, a very common problem of Windows Explorer when opening a folder and moving the mouse while a mouse button is pressed).
- Simplified and automatic deployment of honeypots.

## 2.2 FPolicy Engines

### 2.2.1 Availability

If - for whatever reason – NetApp ONTAP loses connectivity between an SVM and its assigned FPE (FPolicy engine), then the originating SMB traffic won't be blocked, but rather FPolicy processing is suspended (i.e., no protection) until the connection has been re-established or until Cleondris has assigned a new FPE to the SVM. Low network latency and high network stability between the FPE and the data interface(s) of assigned SVMs is crucial.

### 2.2.2 Scalability

In larger installations with many requests, or when there is network latency between the data SVMs and the SnapGuard built-in FPE (FPolicy engine), additional FPEs should be used.

Typically, as a rough estimate, a single FPolicy engine should be planned per 1000 concurrently monitored operations per second, no matter whether these originating SMB sessions are targeting one SVM or are the sum of different SVMs.

In the case of multiple FPEs, it may be a tedious task to spread the load of the SVMs by configuring for each SVM the corresponding FPE. As a solution to this, it is possible to build so called FPE pools (inside Cleondris, ONTAP is not aware of this) and then assigning the pool to the SVMs (instead of a single FPE). Cleondris will then automatically spread the different SVMs over the FPEs in the assigned pool.

Each FPE can be a member of multiple pools. Pool members can be based on direct assignment or by using tags (the tags can be edited in the FPE as well as in the FPE pool detail screen).

Tags can also be used to influence the selection of an FPE inside a pool. If you use metro cluster, and have two physical sites each having some FPEs, then it may make sense to use a shared pool of FPEs for both sides, however, it may be needed to "tip" the SVMs residing on one cluster (location) to prefer the "local" FPEs. If a SVM changes its location (e.g., during a metro cluster failover or after SVM-DR), Cleondris will then automatically reconfigure the destination ONTAP cluster by choosing an FPE for the SVM which is closer to the new cluster. This mechanism can be controlled by using the same tags in the NetApp ONTAP cluster setup in the Cleondris UI.

Please note, even when using FPE pools, at a given point in time a given SVM is always connected to one FPE. Cleondris may extend the product in the future by allowing FPE load-balancing (one SVM being connected to more than one FPE, FPolicy events will then be sent

by ONTAP using a round-robin approach to the FPEs). Allowing multiple concurrent FPEs for a single SVM is not a trivial task, since the Cleondris firewall allows to track the incoming rate of requests on a per-user basis: if the sequence of requests of a user is sent to multiple FPEs in a round-robin manner, then these servers need also to share the rate counters with very low latency.

## 2.3 Integration with External Event Monitoring Systems

No matter whether one uses an active or passive strategy, and regardless of the actions being monitored, it is imperative that SnapGuard be integrated with the alerting system of the IT or security infrastructure.

SnapGuard offers the following options:

- E-mail (only for testing purposes or very small installations)
- SNMP (v1/v2 traps/notifications)
- Syslog (SnapGuard has native Splunk and CEF integration)

Most customers choose the Splunk integration, which allows to send events from SnapGuard via syslog to a Splunk server. Since SnapGuard natively supports Splunk's CIM (Common Information Model of the Splunk Enterprise Security product), the information from SnapGuard direct can be correlated with other data sources in Splunk.

## 2.4 Sending Regular Reports

Regular reports on the status of a firewall can be sent by e-mail. The recipients of the e-mails are managed centrally in the setup menu. A different recipient can be specified for each firewall. For the intervals at which a report is sent, either predefined periods can be selected or a user-defined cron schedule can be used. The report contains an overview of all protected volumes of this firewall and a list of files that match the patterns defined in the pattern pool.

## 2.5 Example Ruleset

### 2.5.1 Ruleset Structure

- File verification should be enabled to protect against unknown ransomware. An external log event shall be enabled (“event level”) in case a user is detected that produces defect files.
- A FPolicy rule should be created that monitors file “create” and “rename” and matches filename patterns from the FSRM ransomware list. The patterns should be imported over a global pattern list in the setup section and referenced in the rule. To reduce spurious false positives, a rate filter of at least 10 per 60 seconds should be used. An external log event shall be enabled in case a user is detected that works with filenames on the known ransomware list.
- A FPolicy rule may be created that monitors file “delete” operation and uses a rate filter of at least 200 per 60 seconds. A ruleset containing such a rule should not be applied

to volumes containing user home directories. An external log event shall be enabled in case a user is detected that deletes many files in shared directories. The external log event can be set to "info" or "notice", as this is typically not very urgent because deleted data can be easily restored from a snapshot if required.

- A FPolicy rule may be created that monitors file "open" and is limited to special honeypot directories using a filename include pattern that gives the name of such a directory that has been created and populated by the security team with "interesting" looking data. An external log event shall be enabled in case a user is detected that tries to read files in the honeypot directory.

In case FPolicy is not desired or needed, e.g., because the system is already configured to audit all accesses to an EVTX file, then the above 3 rules can also be implemented using EVTX rules. Please note that EVTX rules with enabled event generation trigger events 8041-8045.

### 2.5.2 Active vs Passive

A ruleset that shall actively protect the volumes it is being assigned to, must use a block action in the corresponding rules. Typically, "Block User" with a timeout of 60 seconds is used.

If user blocking is not desired, but a good restore point in case of an anomaly, the "Create Emergency Snapshot" option should be used.

In case some volumes need to be protected actively, and others in passive mode, at least two different rulesets must be created accordingly.

## 2.6 Example Event Configuration

SnapGuard is based on the general Cleondris architecture, which contains a large set of pre-defined events. For the purposes of ransomware detection, the events with numbers 8001 – 8045 were implemented. These external events will be generated by FPolicy based rulesets that have detected an anomaly. Most customers trigger the external events 8001, 8002, 8003, 8005 from their rulesets. These 4 events are basically identical, they only differ by the event ID and the default syslog priority that has been assigned. This due to external alarming systems that may not be properly able to distinct between the different received syslog levels. To circumvent this, Cleondris has created 4 events (with different event numbers) that can be triggered by FPolicy rules. By using distinct event numbers, one can more easily distinguish between different event priorities in the external monitoring system. Furthermore, the customer may assign a custom "label" inside each rule. The label will then also be sent as part of the external event.

The Syslog content of these custom events contains the following fields (in case of Splunk format):

- cdm_id: "8001" (or 8002, 8003, 8005)
- action: Splunk CIM action, either "allowed" or "blocked"

- category: "fpolicy"
- count: The ingress rate
- dest: The NetApp ONTAP SVM
- dest_cluster: The NetApp ONTAP cluster
- dest_volume: The NetApp ONTAP volume
- dm: "malware"
- file_name: The filename
- file_name_new: The new filename (in case of a rename)
- file_path: The path to the file
- file_path_new: The new path to the file (in case of a rename)
- method: The SMB method
- rule: The firewall rule (contains the label, otherwise the numeric index in the ruleset)
- signature: Detailed reason
- src: The host of user that is initiating the action
- user: The user initiating the action
- vendor: "Cleondris"
- product: "CDM"
- product_version: Cleondris product version

For a detailed event description, please consult the separate documentation of the Cleondris event system.

In a typical scenario, the customer…

- Enables the global triggering of events 8001-8005 in the SnapGuard event setup
- Adds a Syslog receiver in the SnapGuard event setup and sets it to the "Splunk" format

Inside Splunk, the security team may use the "cdm_id" and/or the label in the "rule" field to differentiate less important events (e.g., a user us deleting data, but there is no hurry, since there is still enough time to restore from a recent snapshot) to more important events (e.g., a user is creating/rename massive amounts of files using filenames on the ransomware list) and act accordingly. When forwarding event data or when opening a ticket, the fields "dest", "dest_cluster", "dest_volume", "user", "rule" and "signature" are very helpful and should be included.

# 3 Architecture

In smaller environments, SnapGuard can be easily deployed as a standalone appliance. However, this kind of deployment is not suitable for large-scale setups with many NetApp Data ONTAP filers or clusters (possibly distributed over different sites), where ten thousand of concurrent CIFS sessions need to be tracked. Therefore, SnapGuard has a built-in tiered architecture that allows to tackle the following problems:

- Adaptation to the workload: scaling with large amounts of concurrent CIFS/NFS sessions

- Low-Latency: no impact on client experience
- High-Availability: system must continuously available, even if parts are failing or being upgraded
- Ease-of-use: centralized manageability of all involved components

## 3.1 Tiered Architecture of SnapGuard

The tiered architecture consists of the following components:

**Tier-1**: A central SnapGuard installation, based on the Cleondris appliance, that aggregates the views of all attached NetApp Data ONTAP systems. In smaller setups, this is the only Cleondris software component needed, as the built-in FPolicy engine ("FPE") can be used.

**Tier-2**: Local FPolicy engines ("FPE") that are connected to nearby NetApp filers. Customers need to install the Cleondris DMT software ("DMT", Data Manager Tools) on dedicated servers, which can then host the FPolicy engines. This is the only part of the architecture that relies on low-latency connections.

**Note: in small and/or testing deployments, the built-in FPolicy engine (FPE) in the Cleondris Tier-1 appliance can be used, an no Tier-2 servers need to be deployed.**

Thanks to the clear separation between the central SnapGuard server and the assigned FPolicy engines (no low-latency connection required between Tier-1 and Tier-2), the architecture can easily be extended for setups where a customer has many locations, e.g., different branch offices in a country or region.

Many customers start with a standalone SnapGuard deployment for testing purposes or small installations. As the number of attached NetApp ONTAP systems grows, the system can easily be scaled by installing the DMT software on dedicated servers. In case of unexpected growth of the CIFS traffic on a NetApp Data ONTAP cluster, additional Tier-2 FPolicy engines, based on hosts running the DMT agent, can be added at any time.

**Detailed Description of the involed Tiers**

**Tier-1** serves as a single point of management of presence and management for the different locations of the customer. A Tier-1 SnapGuard installation connects to one or more assigned NetApp Data ONTAP clusters and the assigned FPolicy modules. Each SnapGuard installation contains the actual configuration for a complete customer location ("which volumes on which NetApp cluster shall be monitored by which FPolicy modules"). The SnapGuard instance pushes the actual FPolicy configuration to the Tier-2 FPolicy modules, followed by the NetApp clusters which are informed about the volumes to monitor and the assigned FPolicy modules. The central SnapGuard installation regularly checks the status of the FPolicy modules (suspicious activity, dynamic blocking of clients, etc.). It is the responsibility of the SnapGuard installation to send monitoring events (Syslog, SNMP, E-Mail) to the respective monitoring system. There is no HA needed on this level: Tier-2 (DMT) based FPolicy engines cache all escalation events locally until a "parent" SnapGuard installation is fetching them. In case the central SnapGuard

installation is being updated (i.e., unavailable for a couple of minutes), polling of Tier-2 events is suspended and continues once the updated central SnapGuard installation starts again. Hence, even in the case of an unavailable Tier-1 SnapGuard instance, events will never be dropped, only the reporting via Syslog/SNMP/E-Mail may be delayed for a short time.

**Tier-2** is the actual workhorse: one or more FPolicy engines (running as part of a DMT installation) are connected to the nodes of a NetApp cluster and monitor all CIFS activity. Preferably, there is more than one FPolicy engine per NetApp Data ONTAP cluster. This allows to spread the load of the SVMs over multiple FPolicy engines. If needed, a set of FPolicy engines can also be shared between different clusters (in case the load per cluster is not very high or in case of MCC (metro cluster) or SVM-DR, where SVMs may be "migrated" between clusters).

## 3.2 Required Software Components

All SnapGuard components run on virtual or physical hardware. SnapGuard (Tier-1) is typically installed as a virtual appliance on VMware ESX, however, it can be installed on pre-existing CentOS or RHEL 7 server as well.

In larger environments, the FPolicy engine built into the Tier-1 is typically not used, instead FPolicy engines are dynamically deployed on Tier-2 servers. For this to work, the target servers (physical or virtual servers) need to have an installation of the Cleondris DMT software ("Data Manager Tools").

The DMT software is available for Windows Server (64-bit) and Linux (64-bit).

- In case of NetApp cDot clusters ("Clustered Data ONTAP"), DMT agent installations on both Windows and Linux are supported (the nodes of the NetApp cluster send messages using a proprietary message protocol via multiple TCP connections to the FPolicy engine running on the DMT agent).
- In case a legacy NetApp 7-Mode system needs to be monitored, the respective FPolicy module must be deployed on a DMT agent running on a Windows machine (the monitored NetApp controller is using a CIFS connection to report FPolicy events to the FPolicy engine).

Please refer to the following table for deployment options:

| Tier | Software | Deployment Options |
|---|---|---|
| Tier-1 | SnapGuard | • Cleondris Virtual Appliance on ESX<br>• RHEL/CentOS 7.x/8.x 64-bit (virtual or physical) |
| Tier-2 (optional) | DMT | • Windows Server 2008+ (64 bit)<br>• RHEL/CentOS 7.x/8.x 64-bit (virtual or physical) |

The Cleondris virtual appliance is distributed by Cleondris as an .OVA file (less than 600MB). The image is based on Linux, however it includes a console based configurator (similar to the yellow ESXi console interface), therefore no Linux knowledge is needed to install or maintain

the software. The complete installation and initial configuration can be done in less than 15 minutes.

The Cleondris DMT software for Windows comes as an MSI based installer with Windows Installer intdgration. The size of the software is about 50 MB. The DMT software can be installed and configured within minutes.

The Cleondris DMT software for Linux comes as a tar.gz archive, which includes an automated setup script. The DMT software can be installed and configured within minutes.

Please note that SnapGuard is internally based on the well-established Cleondris Data Manager software (backup & restore for NetApp Data ONTAP, available as a virtual appliance since 2010) and the dynamically launched, external FPolicy engines (FPE) rely on the well-established Cleondris Data Manager Tools software (DMT, available since 2014).

## 3.3 Hardware Requirements

Depending on how the software components are installed, there are slight differences in the hardware / base OS requirements:

| Software | Hardware Requirements |
|---|---|
| **SnapGuard**<br>Cleondris<br>Virtual Appliance | • Min. 1 virtual 64-bit CPU<br>• Min. 8 GB memory<br>• 32 GB disk space (single VMDK)<br>• Min. 1 virtual network interface (min. 1 GigE) |
| **SnapGuard**<br>RHEL 7/8 64-bit<br>CentOS 7/8 64-bit | • Min. 1 virtual or physical 64-bit CPU<br>• Min. 8 GB memory<br>• 32 GB disk space<br>• Min. 1 network interface (min. 1 GigE)<br>• PostgreSQL 9.6+ installed |
| **Cleondris DMT**<br>Windows Server<br>2008-2019 | • Min. 1 virtual or physical 64-bit CPU<br>• Min. 4 GB available memory for the DMT software<br>• Min. 1 GB disk space for the DMT software<br>• Min. 1 network interface (min. 1 GigE) |
| **Cleondris DMT**<br>RHEL 7/8 64-bit<br>CentOS 7/8 64-bit | • Min. 1 virtual or physical 64-bit CPU<br>• Min. 4 GB available memory for the DMT software<br>• Min. 1 GB disk space for the DMT software<br>• Min. 1 network interface (min. 1 GigE) |

## 3.4 Networking Requirements

There are different message flows in the system:

- Administrators need to be able to manage SnapGuard Tier-1 installations
- SnapGuard Tier-1 need to communicate with ONTAP cluster management interfaces

- SnapGuard Tier-1 needs to communicate with Syslog, SNMP or E-Mail servers
- SnapGuard Tier-1 needs to communicate with Tier-2 DMT instances (running the FPolicy modules)
- ONTAP cluster nodes and 7-mode controllers need to report FPolicy events to FPolicy Engines running either on SnapGuard or on dedicated, external servers (DMT)
- FPolicy Engines running on either SnapGuard or DMT need to communicate with AD servers (SID resolution) and optionally with Clustered Data ONTAP SVM data interfaces (in case the file verification feature is used)

Please refer to the following communication matrix:

| Source | Target | Protocol | Use |
|---|---|---|---|
| Administrator | SnapGuard | HTTPS over TCP Port 443 | Webinterface |
| SnapGuard | ONTAP Cluster Management | HTTPS over TCP Port 443 | Management via NetApp ZAPI protocol |
| SnapGuard | Syslog/SNMP/E-Mail Servers | UDP Port 514, TCP Ports 161, 25/465 | Event Notification |
| SnapGuard | AD (Active Directory) Servers | MSRPC/SMB2 via TCP port 445 | SID/username lookups (if built-in FPE is used) |
| SnapGuard | ONTAP SVM Data Interface | SMB2 via TCP port 445 | Optional, if built-in FPE shall use file verification |
| SnapGuard | DMT | HTTPS over TCP Port 7683 | Management/Deployment via Cleondris proprietary protocol |
| DMT | AD (Active Directory) Servers | MSRPC/SMB2 via TCP port 445 | SID/username lookups |
| DMT | ONTAP SVM Data Interface | SMB2 via TCP port 445 | Optional, if built-in FPE shall use file verification |
| **Clustered Data ONTAP** | | | |
| cDot ONTAP SVM Data-LIF | FPolicy Engine running on DMT or SnapGuard | NetApp FPolicy via TCP port 10000-10099 | FPolicy Events via NetApp proprietary protocol |
| **7-Mode Data ONTAP** | | | |
| 7-Mode vFiler Interface | FPolicy Engine running on DMT | SMB1 or SMB2 via TCP port 445 | FPolicy Events via NetApp proprietary protocol over CIFS/DECRPC |
| FPolicy Engine running on DMT | 7-Mode vFiler Interface | SMB2 via TCP port 445 | Registration for 7-Mode FPolicy via NetApp proprietary protocol over CIFS/DECRPC |

Cleondris GmbH
Buckhauserstrasse 17
CH-8048 Zürich
Switzerland

https://www.cleondris.com