



Dell DR Series Appliance Cleaner Best Practices

Dell Engineering
June 2016

Revisions

| Date | Description |
|-----------|-----------------|
| July 2016 | Initial release |

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2016 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.



Table of contents

| | |
|---|----|
| Revisions..... | 2 |
| Intended audience..... | 4 |
| 1 Deduplication and cleaning..... | 5 |
| 2 How to determine if the cleaner is keeping up..... | 7 |
| 2.1.1 Using the stats --cleaner command..... | 7 |
| 2.1.2 Understanding cleaner statistics..... | 8 |
| 3 Cleaner adjustments & best practices..... | 9 |
| 3.1 Adjusting the cleaner runtime and schedule..... | 9 |
| 3.1.1 Cleaner considerations..... | 11 |
| 3.1.2 Scenario 1: Cleaner analysis and adjustments..... | 11 |
| A Additional cleaner commands..... | 13 |
| A.1 Create a cleaner schedule..... | 13 |
| A.2 Examine the cleaner..... | 13 |
| A.3 Modify the cleaner schedule..... | 13 |
| A.4 Running the cleaner manually..... | 14 |
| B How to run the cleaner..... | 16 |
| C Estimating cleaning work currently in progress..... | 17 |
| D If the cleaner runs behind..... | 19 |
| E Cleaner monitoring and adjustment flow chart..... | 24 |



Intended audience

The DR appliance cleaner is process that is configured to run efficiently and effectively out of the box with no tuning or adjusting required. Only in extreme cases will the DR cleaner possibly may need calibration. This document is intended only for DR appliances which are exposed to the following extreme cases:

- Ingesting or deleting 100TB's or more per week
- Ingest occurs 24x7
- When poor savings is experienced or reported
- Full cleaner pass not finishing once per week

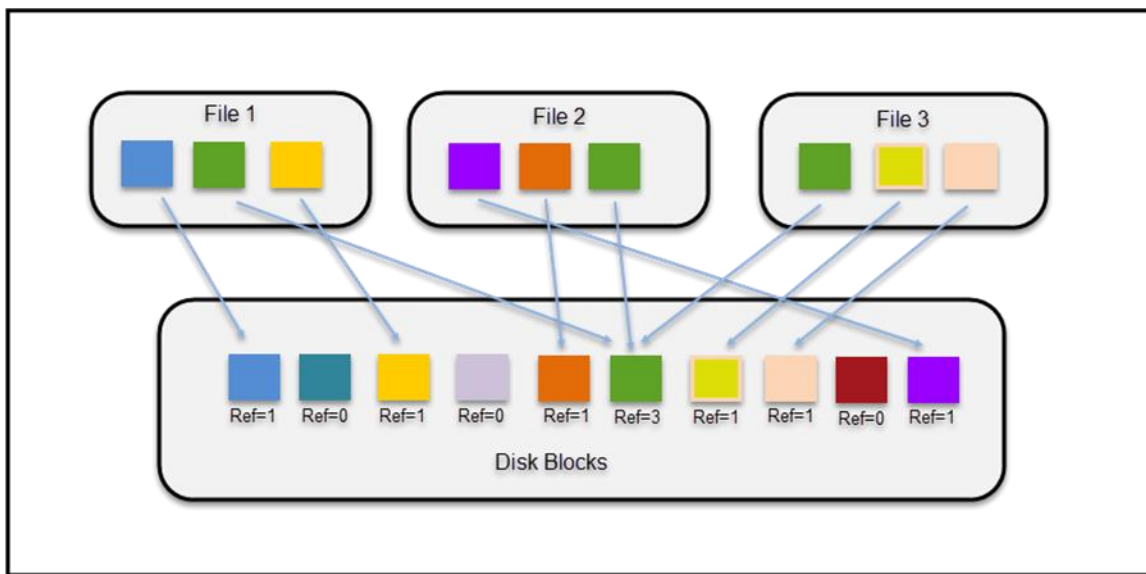


1 Deduplication and cleaning

Data deduplication refers to a technique for eliminating redundant data across all files in a data set. Large amounts of data can be significantly reduced saving costs and resources. The DR implements a variable block sliding window deduplication engine to produce industry leading results.

Each file stored in a DR appliance is a blockmap consisting of pointers to its chunks of data saved on the filesystem. When a file is stored in the DR appliance, a chunking process is applied such that the file is segmented into variable sized chunks. Each chunk is fingerprinted and examined in the DR's deduplication dictionary to see if it has been discovered before. If the DR system has encountered the chunk already, the files block map is updated to point to the already existing chunk in the system and the chunks reference count is increased. If a chunk is unique to the DR system, the chunk's fingerprint is inserted into the dictionary, the chunk is written to the filesystem, the files blockmap is updated to point to it and the chunks reference count is set to 1. Thus, files stored in the DR might point to unique blocks, non-unique blocks that are shared with other files in the system, or a combination of unique and non-unique blocks.

When a file is deleted, the files chunk reference counts are decremented by 1.



The DR cleaner plays a critical role to a DRs capacity because it updates chunk reference counts and reclaims space when chunk references are equal zero.

DR appliances are shipped with the cleaner setting to run automatically during idle time. In most cases, this setting allows more than enough time to update chunk references and reclaim space. In extreme cases where the DR cleaner is not permitted complete a full pass at least once per week, the suggestions in the documentation should be applied to the cleaner settings.



This document explains how to determine if the cleaner is keeping up, how much is required for the cleaner to finish, how to adjust the cleaner schedule and runtimes, how to determine what the cleaner is doing and best practices on what to set.

Note: To use this document, ensure the DR appliance is running 3.2.0.2 or 3.2.6.1. Build 3.2.0418.1a and future releases can also be used with this document.



2 How to determine if the cleaner is keeping up

Dell recommends that a full cleaner pass complete once every 7 days or less. To determine if the cleaner is keeping up, a quick inspection of the cleaner statistics 'Last Completion time(s)' variable is described below.

2.1.1 Using the stats --cleaner command

The **stats --cleaner** command displays the current running cleaner progress and the amount of time taken to complete its latest full pass. The 'Last Completion time(s)' value indicates that the cleaner is keeping up when it is less than or equal to 7 days.

```
Example - 1
administrator@ptcdr6000 > stats --cleaner

Last Run:
Last Files Processed      : 17083
Last Bytes Processed     : 123775927756481
Last Bytes Reclaimed     : 9165802475224
Last Start Time          : 03/07/16 00:32:20
Last End Time            : 03/10/16 00:27:17
Time To Completion(s)    : 198394.00

Current Run:
Start Time               : 03/10/16 00:28:20
Files Processed          : 4446
Bytes Processed          : 25376233292316
Bytes Reclaimed          : 43427629102602
Phase 1 Start Time      : 03/10/16 17:03:58
Phase 1 Records Processed : 0
Phase 1 End Time        : 03/10/16 17:03:58
Phase 2 Start Time      : 03/10/16 17:03:58
Phase 2 Records Processed : 755129948
Phase 2 End Time        : 03/10/16 17:03:58
Phase 3 Start Time      : 03/10/16 17:03:58
Phase 3 Records Processed : 641439
Phase 3 End Time        : 03/10/16 17:05:34
Phase 4 Start Time      : 03/13/16 23:46:38
Phase 4 Records Processed : 21514332
Phase 4 End Time        : 03/13/16 23:46:38
Phase 5 Start Time      : Unknown
Phase 5 Records Processed : 0
Phase 5 End Time        : Unknown

Last Completed Run:
Last Completed Start Time : 03/06/16 19:21:42
Last Completed End Time   : 03/09/16 19:12:23
Last Completion time(s)   : 258897.00

Cleaner Pending Work:
Estimated Logical Bytes left : 66390640
```

The 'Last Run' values describe the last time a cleaner phase was run. Because individual cleaner phase run times do not assist us to determine full pass cleaner runtimes, these values will be ignored in the following

The DR cleaner process includes 5 phases to sort and process logs and free up disk space. Space is freed up only during phases 3 and 4. Space is not freed up during phase 1, 2 and 5. Because of this, the 'Estimated Logical Bytes left' value may become idle for prolonged periods of time. Since the phases do not describe cumulative runtimes, the 'Current Run' section do not help determining cleaner full runtimes and will be ignored in future.

The Last Completed Run values provide last cleaner full run details and are helpful to determine if and how much cleaner adjustments are needed.

Note: If the cleaner is running and it is not reporting to reclaim space, that does not mean there is an issue or problem with the cleaner. When cleaner is running, somecleaner phases do not reclaim data space. Only when the cleaner is executing in phases 3 and 4 is space reclaimed and the 'Estimated Logical Bytes left' is updated.



2.1.2 Understanding cleaner statistics

The 'Last Completion time(s)' variable provides the latest full pass cleaner runtime. In Example-1 above, the cleaner completed its latest full pass consuming 258897.00 seconds or 2.99 days. To convert from seconds to days, take the total seconds and divide by 86,400 ($60 * 60 * 24 = 86,400$). Thus, $258,897 \text{ seconds} / 86,400 = 2.99 \text{ days}$. Because the cleaner completed its latest job within a week's time frame (7 days), the cleaner is not falling behind.



3 Cleaner adjustments & best practices

Under rare and extreme cases, the DR cleaner may run longer than 7 days. If so, apply the following cleaner best practices in order when adjusting cleaner runtimes.

- 1. If the cleaner is set to run automatically (default setting) schedule the cleaner to run for 40 or more hours per week.
- 2. Schedule the cleaner to run in long durations instead of smaller short durations.
- 3. If the cleaner does not finish in 40 hours, schedule cleaning to run in parallel with replication.
- 4. If the cleaner does not finish in 40 hours, schedule cleaning to run during low periods where there is not a full backup load.

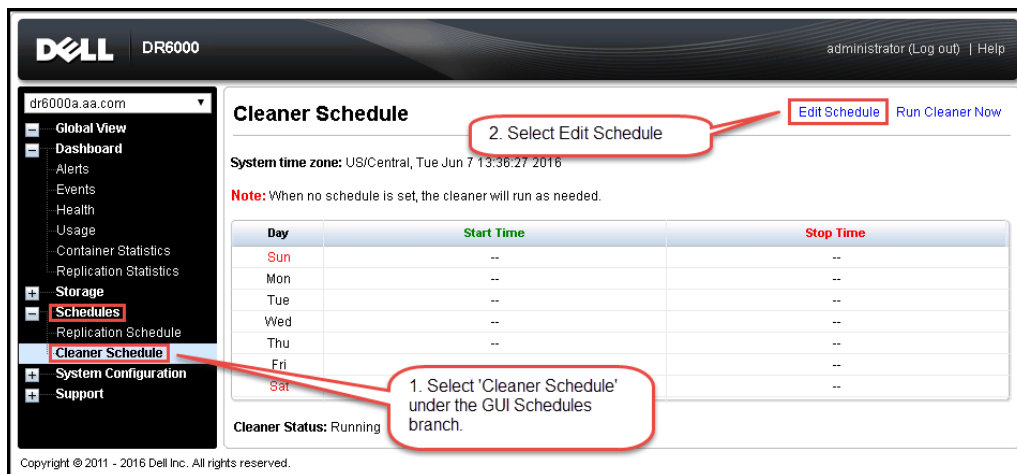
See Appendix-C & D for additional details and examples.

3.1 Adjusting the cleaner runtime and schedule

If the value of 'Last Completion time(s)' exceeds 7 days, the cleaner needs adjusting with additional runtime. Begin by scheduling the cleaner to run 40 hours / week during non-peak times. Scheduling can be applied in both the GUI and CLI.

The following example schedules the cleaner to run 40 hours per week:

1. On the **Cleaner Schedule** page, click **Edit Schedule**:



- For each day, apply a cleaner runtime schedule, then click **Set Schedule**.

| Day | Start Time | Stop Time |
|-----|------------|-----------|
| Sun | 00 : 00 | 24 : 00 |
| Mon | 20 : 00 | 04 : 00 |
| Tue | 01 : 00 | 04 : 00 |
| Wed | 02 : 00 | 04 : 00 |
| Thu | 03 : 00 | 04 : 00 |
| Fri | 04 : 00 | 04 : 00 |
| Sat | 05 : 00 | 04 : 00 |

Buttons: Cancel, Reset, **Set Schedule**

- Inspect the newly configured cleaner schedule.

DELL DR6000 administrator (Log out) | Help

dr6000a.aa.com

Cleaner Schedule

[Edit Schedule](#) | [Run Cleaner Now](#)

System time zone: US/Central, Tue Jun 7 13:34:52 2016

Note: When no schedule is set, the cleaner will run as needed.

| Day | Start Time | Stop Time |
|-----|------------|-----------|
| Sun | 00:00 | 24:00 |
| Mon | 20:00 | 04:00 |
| Tue | 20:00 | 04:00 |
| Wed | 20:00 | 04:00 |
| Thu | 20:00 | 04:00 |
| Fri | 20:00 | 04:00 |
| Sat | -- | -- |

Cleaner Status: Done

Copyright © 2011 - 2016 Dell Inc. All rights reserved.

- After a week (7 days), run the stats `--cleaner` command to examine cleaner's 'Last Completion Time(s)' value. If the cleaner completed a full pass within 7 or less days, no more cleaner adjustments are necessary. If the cleaner has not yet completed a full pass or has taken beyond 7 days to complete, increase the cleaner schedule runtimes and repeat this step.



3.1.1 Cleaner considerations

1. Additional cleaner runtime is required as more data is ingested or begins to age out.
2. When the cleaner is triggered to run by manual, scheduled or forced methods, the cleaner is run at the same priority as ingest and replication.

3.1.2 Scenario 1: Cleaner analysis and adjustments

Raymond has been successfully running his DR appliance at his computing center for the past 3 years. Because his DR has recently been reporting poor savings and performance is slow, he feels his cleaner is not completing within the Dell recommendation of a 7 day timeframe.

Inspect Raymond’s cleaner results are shown below:

```

Scenario - 1
administrator@Raymond> stats --cleaner

Last Run:
Last Files Processed      : 478410
Last Bytes Processed      : 49584323092
Last Bytes Reclaimed      : 29304996443
Last Completed Run:
Last Completed Start Time : 05/09/16 06:13:52
Last Completed End Time   : 05/17/16 00:23:43
Last Completion time(s)   : 666060.00

Cleaner Pending Work:
Estimated Logical Bytes left : 39987409

administrator@Raymond>
  
```

Since Raymond’s cleaner is not completing within 7 days, $(666,060 / 86,400 = 7.7 \text{ days})$ he will schedule his cleaner to run for 40 hours during the week.

Raymond’s Current DR Schedule of Events:

| Weekdays | 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 |
|-------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| Full | | | | | | | | | | | | | | | | | | | | | | | | | |
| Incremental | X | X | X | X | X | X | | | | | | | | | | | | | | | | | X | X | X |
| Replication | | | | | | | | | | | | | | | | | | | | | | | | | |
| Cleaner | | | | | | | | | | | | | X | X | X | X | X | X | X | X | | | | | |
| Weekends | 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 |
| Full | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X |
| Incremental | | | | | | | | | | | | | | | | | | | | | | | | | |
| Replication | | | | | | | | | | | | | | | | | | | | | | | | | |
| Cleaner | | | | | | | | | | | | | | | | | | | | | | | | | |

Raymond would like to determine if his adjusted cleaning schedule is adequate for his environment, so he waits one week before he analyses his cleaner runtimes. Raymond then runs the stats --cleaner command.



```
Scenario - 1
administrator@Raymond> stats --cleaner

Last Run:
Last Files Processed      : 6590326
Last Bytes Processed     : 3980203365
Last Bytes Reclaimed     : 2286403410
Last Run Time             : 05/16/16 16:18:18

Last Completed Run:
Last Completed Start Time : 05/17/16 07:01:00
Last Completed End Time   : 05/23/16 13:46:43
Last Completion time(s)   : 585960.38

Cleaner Pending Work:
Estimated Logical Bytes left : 8778223

administrator@Raymond>
```

Raymond adjusted cleaning schedule successfully completes in 6.78 days which meets Dells recommended full cleaner pass of 7 or less days.



A Additional cleaner commands

A.1 Create a cleaner schedule

Using the CLI:

```
schedule --add -day <Day of the Week> --start_time <HH:MM> --stop_time <HH:MM> -  
-cleaner
```

Example

```
administrator@dr6000a > schedule --add --day Sunday --start_time 06:00 --  
stop_time 22:00 --cleaner  
Successfully updated Cleaner schedule.  
administrator@dr6000a >
```

A.2 Examine the cleaner

Using the CLI:

```
schedule --show --cleaner
```

Example

```
administrator@dr6000a > schedule --show --cleaner  
Cleaner schedule:  
  
          Start    Stop  
Monday    06:00    17:00  
Tuesday   06:00    17:00  
Wednesday 06:00    17:00  
Thursday  06:00    17:00  
Friday    06:00    17:00  
Saturday  --         ---  
Sunday    --         --  
administrator@dr6000a >
```

A.3 Modify the cleaner schedule

Using the CLI:

```
schedule --add --day Sunday --start_time 06:00 --stop_time 22:00 --cleaner
```



Example

```
administrator@dr6000a > schedule --add --day Sunday --start_time 06:00 --stop_time 22:00 --cleaner
Successfully updated Cleaner schedule.
administrator@dr6000a >
```

```
schedule --delete --day <Day of the Week> --cleaner
```

Example

```
administrator@dr6000a > schedule --delete --day Sunday --cleaner
Successfully updated Cleaner schedule.
administrator@dr6000a >
```

Other cleaner schedule help is defined by running the following command:

```
Schedule --help
```

A.4 Running the cleaner manually

The cleaner can be manually run and is a onetime run and will complete a full cleaner pass very quickly as it does not pause for any DR activities until the cleaning process has completed.

Using the CLI:

```
maintenance --filesystem --reclaim_space
```

Example

```
administrator@dr6300a > maintenance --filesystem --reclaim_space
Successfully started cleaner.
administrator@dr6300a >
```

Using the GUI:

DELL DR6000 administrator (Log out) | Help

dr6000a.aa.com

- Global View
- Dashboard
- Alerts
- Events
- Health
- Usage
- Container Statistics
- Replication Statistics
- Storage
 - Containers
 - Replication
 - Encryption
 - Clients
- Schedules
 - Replication Schedule
 - Cleaner Schedule
- System Configuration
 - Networking
 - Active Directory

Cleaner Schedule

[Edit Schedule](#) | [Run Cleaner Now](#)

System time zone: US/Central, Fri May 27 12:04:30 2016

Note: When no schedule is set, the cleaner will run as needed.

| Day | Start Time | Stop Time |
|-----|------------|-----------|
| Sun | -- | -- |
| Mon | 15:00 | 18:00 |
| Tue | 22:00 | 23:00 |
| Wed | -- | -- |
| Thu | -- | -- |
| Fri | -- | -- |
| Sat | -- | -- |

Cleaner Status: Done

1. Select Cleaner Schedule

2. Select Run Cleaner Now



B How to run the cleaner

There are multiple ways in which the cleaner process is run:

- **Scheduled** – The cleaner is triggered to begin processing and continues until the cleaning work has been completed or the scheduled window ends. If cleaning has completed within the scheduled window, it will not start again until the next scheduled window. For large cleaning jobs, the cleaner will span multiple schedule windows picking up where it left off until the cleaning job has completed. Note: The cleaner job does not pause except in-between cleaning schedules.
- **Manual** – When the cleaner process is manually started, it will continuously run until the cleaning process has finished. Note: the cleaner will not pause until the cleaning job has completed.
- **Automatically (default)** – The cleaner process automatically starts if it detects that there are:
 - No running data ingest
 - No running replication processes
 - Two minutes of system idle time since the latest data ingest

Once replication or data ingest begins, the cleaner pauses until the three conditions for auto start become true again. It does not matter if a cleaner schedule has been defined, the DR cleaner will run during idle times when there is work to be done. In other words, 50 hours of idle time equals 50 hours of cleaner time.

- **Forced Run** – When the DR is running lower than 1TB of free space or when internal cleaner logs (PL) are queued to 1 day of cleaner work the cleaner is triggered by the DR to begin. Here the cleaner will not pause until a full cleaner pass has been completed.

At a minimum, the DR appliance will need to run through a full cleaner pass at least once per week. Dell recommends the cleaner process to be run 6-8 hours per day or 40-50 hours per week. Running the cleaner for long contiguous periods of time is more efficient than many short periods of time. For example, the cleaner will free up disk space more quickly if it is run for 16 hours each day for three days, instead of 7 hours each day for the week.

When ingestion and/or replication complete, the cleaner will automatically start. When replication and/or ingestion stops, the cleaner will automatically resume.

If deleting data to free up disk space as quickly as possible is desired, manually running the cleaner is recommended.



C Estimating cleaning work currently in progress

The 'Estimated Logical Bytes left' value provides current active cleaner progress. It provides an estimate of how much work is left for the current active cleaner to complete a full pass. When this value reaches 0, the cleaner has completed a full pass.

```
Scenario - 1
administrator@Raymond> stats --cleaner

Last Run:
Last Files Processed      : 6590326
Last Bytes Processed     : 3980203365
Last Bytes Reclaimed     : 2286403410
Last Run Time             : 05/16/16 16:10:16

Last Completed Run:
Last Completed Start Time : 05/17/16 07:01:00
Last Completed End Time   : 05/23/16 13:46:43
Last Completion time(s)   : 585960.38

Cleaner Pending Work:
Estimated Logical Bytes left : 8778223

administrator@Raymond>
```

To determine the current state of the cleaner, run the following command:

Stats --system

```
Example - 2
administrator@ptcdr6000 > stats --system
Capacity Used          : 16585.6 GiB
Capacity Used in GB   : 17808.692
Capacity Free         : 7787.3 GiB
Capacity Free in GB  : 8361.520
Read Throughput      : 0.00 MiB/s
Write Throughput     : 63.96 MiB/s
Current Files        : 33
Current Bytes        : 39149050480731
Post Dedupe Bytes    : 17686624192813
Post Compression Bytes : 17686624192813
Post Encryption Bytes : 0
Post Encryption Bytes in GiB : 0.0 GiB
Compression Status   : Done
Cleaner Status       : Pending
Encryption Status    : Disabled
Total Inodes         : 48
Bytes decrypted      : 0
Dedupe Savings       : 54.82 %
Compression Savings  : 0.00 %
Total Savings        : 54.82 %
administrator@ptcdr6000 >
```



The DR cleaner operates in three states:

- **Done** – When the cleaner is done and there is nothing more to process.
- **Running** – The cleaner is currently running.
- **Pending** – The cleaner has work to do, but is currently waiting to run.

In Example – 2 above, the cleaner is placed in “Pending” mode because there is cleaning yet to be done, and the DR cleaner is outside of its scheduled window of operation and the DR is not idle, but is busy with ingest or replication.



D If the cleaner runs behind

When a full cleaner pass runs beyond the 7 days recommended by Dell, the cleaner is considered to be in cleaning debt. For example, if a full cleaner pass completes in 9 days, the cleaner has 2 days of cleaning debt and cleaner runtimes need to be immediately increased. Cleaning debt occurs when the DR ingests more data than can be cleaned. Depending on the rate of cleaner debt accumulation, the DR can quickly fill up eventually halting backups. When cleaner runtimes are increased, the cleaning debt can be immediately addressed and eliminated over time.

Follow the cleaner best practices outlined in Section-3 of this document.

Example

Cindy manages backups at a pharmaceutical company which includes multiple remote DR appliances replicating to a production site. Cindy notices that her production site DR appliance is rapidly and unexpectedly running out of free space. Since she has not paid attention to the cleaner process in quite some time, she feels her DR cleaner has fallen behind and is not freeing up disk capacity fast enough.

At her production site, Cindy currently holds a retention policy of 12 weekend full backups and 15 daily incrementals and her backup schedules are described in the following table.

| Weekdays (Mon – Fri) | 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 |
|-------------------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| Full | | | | | | | | | | | | | | | | | | | | | | | | | |
| Incremental | X | X | X | X | X | X | X | | | | | | | | | | | | | | | X | X | X | X |
| Replication | | | | | | | | | | | | | | | | | | X | X | X | X | | | | |
| Cleaner | | | | | | | | X | X | X | X | X | X | X | X | X | X | | | | | | | | |
| Weekend (Sat & Sun) | 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 |
| Full | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X |
| Incremental | | | | | | | | | | | | | | | | | | | | | | | | | |
| Replication | | | | | | | | | | | | | | | | | | | | | | | | | |
| Cleaner | | | | | | | | | | | | | | | | | | | | | | | | | |

In this scenario, Cindy will calibrate the production site cleaner schedule such that full cleaner passes meet Dell’s recommended cleaner completion time of 7 days.

Note: Each of Cindy’s remote sites will require a cleaner schedule inspection as well, but in this scenario, only the production site cleaner will be examined.

Cindy examines the current state of her production site cleaner by running the following command:

```
stats --cleaner
```



```
Scenario - 2
administrator@CindyDR63000 > stats --cleaner

Last Run:
Last Files Processed      : 674372
Last Bytes Processed      : 12648375619043
Last Bytes Claimed        : 957231

Last Completed Run:
Last Completed Start Time : 03/25/16 02:45:00
Last Completed End Time   : 04/06/16 12:06:00
Last Completion time(s)   : 1069260.00

Cleaner Pending Work:
Estimated Logical Bytes left : 66390640

administrator@CindyDR63000 >
```

Cindy notes that the latest full pass cleaning process required 1,069,260 seconds or 12.38 days (1,069,260 / 86,400) thus, Cindy's cleaner has fallen behind of 5.3 days (12.38 – 7 = 5.3) in cleaner debt. Cindy must increase her cleaning runtimes such that cleaning debt begins to diminish and eventually achieves a full cleaner pass to complete within 7 days.

How much should Cindy increase her cleaner schedule? If Cindy's DR appliance free space is critically low, she must aggressively schedule the cleaner to recover disk space as quickly as possible by running the cleaner manually such as the GUI's 'Run Cleaner Now' command. In this example, Cindy is not in a critical free space situation, but will need to increase her cleaner runtime ASAP. Without increasing cleaner runtime schedules, Cindy's production DR can quickly consume the remaining amount eventually entering into a critical state.

Cindy's cleaner debt is substantial because it is ~75% Dell's recommended 7-day full cleaner pass (5.3 / 7 = .75). Thus, she decides to initially increase her cleaning time by ~ 75% in hopes of being to reduce the 5.3 days of cleaning debt and to eventually complete a full cleaner pass within a week. A 75% increase of 50 hours is ~ 88 hours (50 * 1.75 = 87.5), so Cindy decides to increase the cleaning schedule up to 90 hours per week, then observe the impacts to cleaner adjustments. Cindy realizes that the 90 hours of cleaner runtime is just an estimate, so she expects multiple adjustments to calibrate the cleaner schedule correctly.

To achieve 90 hours of cleaner run time, Cindy makes her adjustments shown below in the table below:



| Weekdays (Mon – Fri) | 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 |
|-------------------------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| Full | | | | | | | | | | | | | | | | | | | | | | | | | |
| Incremental | X | X | X | X | X | X | X | | | | | | | | | | | | | | | X | X | X | X |
| Replication | | | | | | | | | | | | | | | | | | X | X | X | X | | | | |
| Cleaner | | | | | | | | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X |
| Weekend (Sat & Sun) | 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 |
| Full | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X |
| Incremental | | | | | | | | | | | | | | | | | | | | | | | | | |
| Replication | | | | | | | | | | | | | | | | | | | | | | | | | |
| Cleaner | | | | | | | | | | | | | | | | | | | | | | | | | |

Cindy takes note of the value 66,390,640 in the Estimated Logical Bytes left variable and logs it with the date.

After making the cleaner adjustments above, Cindy waits 7 days to determine if additional cleaner adjustments are required. She hopes to reduce the accumulated cleaner debt, and the weeks' worth of aged out data. Cindy now runs the `stats --cleaner` command after a week when the cleaner adjustments were made:

```
Scenario - 2
administrator@CindyDR63000 > stats --cleaner

Last Run:
Last Files Processed      : 498830
Last Bytes Processed      : 6648375619043
Last Completion time(s)   : 1069260.00

Last Completed Run:
Last Completed Start Time : 03/25/16 02:45:00
Last Completed End Time   : 04/05/16 12:06:00
Last Completion time(s)   : 1069260.00

Cleaner Pending Work:
Estimated Logical Bytes left : 6690640

administrator@CindyDR63000 >
```

Cindy takes note of the value 6,690,640 in the Estimated Logical Bytes left variable and logs it with the date. Cindy notices that the cleaner has made significant cleaning progress but has not completed a full cleaning pass during the past week.

| | | |
|--|------------------|------------|
| | 04/06/16 | 66,390,640 |
| | 04/13/16 | 6,690,640 |
| | Cleaner Progress | 59,848,708 |

Cindy does not adjust the 90 hour per week cleaning schedule and expects a full cleaning pass to complete within the next couple days. After waiting a week (for the next backup to occur), she runs the following command:

```
stats --cleaner
```



```
Scenario - 2
administrator@CindyDR63000 > stats --cleaner

Last Run:
Last Files Processed      : 436667
Last Bytes Processed     : 2340583744
Last Completed Run:
Last Completed Start Time : 04/06/16 13:01:00
Last Completed End Time   : 04/14/16 04:41:00
Last Completion time(s)   : 678927.00

Cleaner Pending Work:
Estimated Logical Bytes left : 66539348

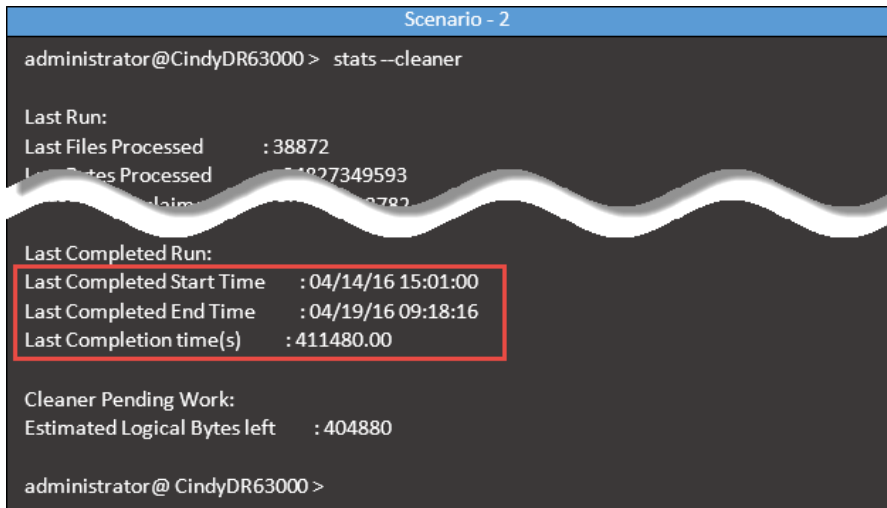
administrator@CindyDR63000 >
```

Notice that a full cleaner pass completed in ~ 7.85 days (678,927 / 86,400) which means Cindy's adjusted 90-day cleaner schedule has nearly eliminated the cleaner total run time from 12.38 to 7.85 days. With such big cleaner debt progress and such little cleaner debt remaining, Cindy believes that the current cleaner settings will accomplish Dells recommended full cleaner pass within the next 7 days during next week's backup cycle. Thus, Cindy does not make any cleaner adjustments and waits 7 days to inspect the cleaner results.



Cindy then runs the following command:

```
stats --cleaner
```



```
Scenario - 2
administrator@CindyDR63000 > stats --cleaner

Last Run:
Last Files Processed      : 38872
Last Bytes Processed     : 127349593
Last Completion time(s)  : 2782

Last Completed Run:
Last Completed Start Time : 04/14/16 15:01:00
Last Completed End Time   : 04/19/16 09:18:16
Last Completion time(s)   : 411480.00

Cleaner Pending Work:
Estimated Logical Bytes left : 404880

administrator@ CindyDR63000 >
```

Cindy has achieved in setting a successful cleaner schedule as her full cleaner pass has completes in 7 days or less. Her full cleaner pass actually completes in 4.76 days (411,480 / 86,400).

Cindy will wait another 3-6 months and will reexamine if her cleaning runtimes.



E Cleaner monitoring and adjustment flow chart

