


ELIMINATE RISK THROUGH DATA PROTECTION MONITORING



Mike van der Steen

Principal Systems Engineer

Dell Technologies

Mike.vandersteen@dell.com



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

[Learn more at \[www.dell.com/certification\]\(http://www.dell.com/certification\)](http://www.dell.com/certification)

Table of Contents

1	Introduction	4
2	Centralized Reporting with DPA	4
2.1	Why it is needed	4
2.2	The stakeholders.....	5
2.3	DPA Infrastructure Components.....	6
2.4	Monitoring of Data Protection Elements/Objects	6
2.5	Smart Groups	8
2.6	Data Collection Retention	9
2.7	The Elements of Running a Report	9
3	Reporting and Dashboards	11
3.1	Reports for the Backup Administrator and Operations Manager	12
3.2	Reports for the Application Owner	17
3.3	Reports for the Risk/Security Teams.....	18
3.4	Summary Dashboards for All	21
4	Flexibility with Report Customization	23
4.1	Modifying a Report Template: Backup All Jobs – No Restarts.....	23
4.2	Adjusting the Merge Condition to Detect Restarts.....	24
5	Working with External Data Sources	25
5.1	Custom Template – The State.....	26
5.2	External Data Source.....	26
5.3	Adding the Data Source	26
5.4	Define the Operator.....	27
5.5	Order and Report	28
6	DPA Analysis Engine	29
6.1	Analysis Policy	29
6.2	Detecting Higher Rate of Change Backups – Example Analysis Policy.....	29
6.3	Protection and Chargeback Policies.....	31
7	Summary	32
8	References	32

Disclaimer: The views, processes or methodologies published in this article are those of the author. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

1 Introduction

At the core of any business or organization, there is an enormous reliance on data to serve customers, drive insights, develop products and much more. Throw into the mix cyber threats from external sources and the unfortunate increase of insider bad actors, it is more important than ever to ensure that the organization's data is protected and recoverable. What makes it even more challenging is the rapid growth of data, the distributed nature of data locations, and the load it places on infrastructure resources and administrators. All this needs to be achieved with a finite budget and resources.

In addition to the management and protection of the data, application owners and the business will assess their data from different risk or exposure points of view. That is, each owner has a level of tolerance for data unavailability through system outages, recovery operations or in a disaster recovery situation. Focusing on the data protection, recovery and disaster recovery readiness aspects of data, providing the relevant information to each party may not be easily achieved by the data protection application itself.

Leveraging a single centralized monitoring and reporting application, independent of the various backup applications by a single vendor or across different vendor applications is required. Dell EMC Data Protection Advisor¹ (DPA) is a centralized monitoring and reporting application that is extremely powerful when leveraging its analysis engine, customization and dashboards.

This Knowledge Sharing article provides an overview of DPA, its reporting capability and insights which can be gleaned from them by different stakeholders. Knowing where to start can be daunting for any reporting-based application, especially when the application has hundreds of built in reports available. To assist administrators in getting started with DPA, this article contains details of reports which can be run and what valuable insights they provide. Also Included is a section on how reports can be customized to meet specific requirements of application owners.

The article concludes with the analysis engine, a powerful capability of DPA which inspects data collected from various sources, runs predefined rules as defined by application owners and reports the findings. The analysis engine is what enables compliance-based reporting, a vital capability that helps organizations determine their risk exposure when it comes to data protection.

2 Centralized Reporting with DPA

DPA is a centralized monitoring application which captures data from a wide range of applications, storing that information in a central database for analysis and enabling a comprehensive single view of the data protection environment. Enabling unified access to data from various components, DPA runs analysis rules constantly, while providing alerting to components and clients which require attention. This automated, unified approach to managing the data protection environment simplifies the process of monitoring and tuning whilst maximizing utilization.

2.1 Why it is needed

DPA is designed to not only monitor Dell EMC data protection applications like Avamar², NetWorker³ and PowerProtect Data Manager⁴, it supports multiple 3rd party vendor applications including CommVault⁶, Veritas NetBackup⁷ and many others. Additionally, the infrastructure used to store and transmit data within the environment can also be monitored, including tape libraries, fibre channel switches, physical or virtual servers, and deduplication appliances like Dell EMC PowerProtect DD Series appliances⁵ (formally known as Dell EMC Data Domain and will be referred to as Data Domain in this article).

Regardless of the size of the organization or the volume of data protected, DPA provides a layer of abstraction to one or more data protection titles used within the organization. This enables multiple stakeholders to obtain information relevant to them, not just from the previous day, but going back months or years to meet specific requirements.

It provides insight not only for data protection administrators and operations managers, but also for application owners, risk/security teams and management. Collectively, the various stakeholders are provided automated reporting and alerting of the data protection environment, enabling them to ensure gaps in data protection are addressed and within compliance/risk-based service level agreements (SLAs).

Knowing what is protected, how often backups are performed, the retention of data and more importantly, what is not protected is extremely important to an organization.

2.2 The stakeholders

Defining all stakeholders and their connection to data protection will vary with each organization. This article will discuss four main stakeholders, each requiring different type of reporting information and insights based on their role.

Backup Administrator

First and foremost, the backup administrator is the primary stakeholder and one that will rely heavily on the reporting information from DPA. Beyond the basic backup success/failure reports, the backup administrator needs to ensure at a minimum that:

- data is protected within specific time periods
- backup and replication jobs are scheduled accordingly to meet recovery point objectives
- the data protection infrastructure is performing optimally
- data protection requirements forecasts meet data growth

By leveraging the reporting and analytical capability of DPA, the backup administrator can operate in a proactive manner and address gaps in data protection before issues arise.

IT Operations Manager

The manager of IT operation teams needs awareness of the various IT environments under their control and the health of those environments. With DPA, the IT Operations Manager can obtain information about the state of the backup environment and the applications protection levels. Reports can be run ad-hoc by the Operations Manager or review the predefined scheduled reports. From a data protection viewpoint, the manager can report to senior management on health of the backup environment or exposure due to protection issues, without needing to rely on the backup administrator.

Application Owner

Most organizations prefer applications and data be protected using centralized scheduling and retention policies capabilities of a backup application. From an application owners' point of view, they are responsible for ensuring applications and data they are responsible for are available to serve the requirements of the organization or their customers.

That they have valid recover points for their application needs to be known and the automated reporting capability of DPA can provide valuable information. Knowing when backups occurred, the duration and if a second copy is available at an alternative location is vital. This will provide them the confidence to recover data within their service level agreements (SLAs) or take action to address any data protection exposures.

Risk/Security Teams

Risk/security teams require insights and visibility over the entire IT infrastructure, of which data protection is one element. With DPA, these teams can obtain information on data protection exceptions which may include clients not protected within a specified service level agreement, delays in have offsite backup copies created, etc. Important to these teams is also knowing what changes have been made within a backup application and by whom. This information is often required to meet compliance-based standards internally or to external organizations.

2.3 DPA Infrastructure Components

The DPA server is comprised of two components, the application and datastore. It is recommended that these components be installed on different servers, as installing DPA application and datastore on a single server is not supported. This article will not discuss the deployment architecture in detail and recommend reviewing the DPA Installation and Administration Guide⁹ or the DPA 6.3 and later Deployment Architecture Guide⁸.

Once installed, the administrator or user of DPA access the application via a HTML-based UI hosted by that server, while the database server stores all data collected from the data protection environment via agents. A DPA agent is installed by default on the DPA application server and for smaller environment, using the DPA server as the only agent may be acceptable; however, larger environments do require multiple agents be installed.

To ensure that the DPA server is not heavily loaded with agent data collection operations, it is recommended to install an agent on each of the backup applications to be monitored. Additionally, there may be a dedicated virtual server used as a remote agent to collect information from other elements to be monitored like Data Domain, tape library and associated switches, data servers, etc. Using multiple agents will ensure data is collected in a timely manner, as by default some data collection requests have a 5 minute data collection frequency.

2.4 Monitoring of Data Protection Elements/Objects

With agents installed, the configuration of data protection elements or objects can commence via the discovery wizard in the inventory section of DPA. A detailed process of how to discover elements of the data protection environment will not be provided in this article. The intention is to highlight the flexibility and customization of DPA that will be of benefit as reporting is provided to the various stakeholders.

There are three main categories where discovered data protection elements are linked to; servers, storage and switches, as shown in Figure 1.

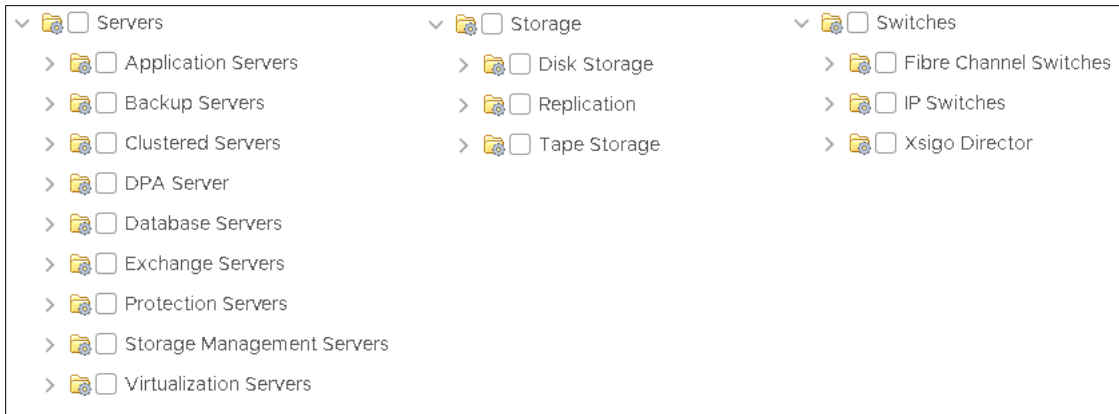


Figure 1- Categorization of Data Protection Elements

Each category is then further divided into smaller categories for granularity and flexibility when reports are run. Figure 2 shows the discovered NetWorker servers and which have been linked to the NetWorker subcategory/group.

A discovered data protection element can be assigned to multiple groups and this becomes important when providing reporting to stakeholders. For example, a stakeholder may require reporting of backup applications based on the environment, backup application type or both. If a backup server was assigned to a group as shown in Figure 2, the required granularity of report per environment is not possible. This is easily achieved by creating two new groups (e.g. Production and Dev | Test) as shown in Figure 3. Reporting can now be undertaken by selecting each environment or by selecting the parent NetWorker group.

Alternatively, assign discovered data protection elements to multiple groups. There are numerous ways to create a group structure that addresses the reporting requirements of the stakeholders and the organization. Set aside time to consider the reporting requirements and customize the group structure to provide maximum flexibility when it comes to creating reports for the stakeholders.

Assigning protection elements or applications to groups is a static assignment and for most environments would work very well. There is an ability within DPA to create groups dynamically and this is achieved via Smart Groups.

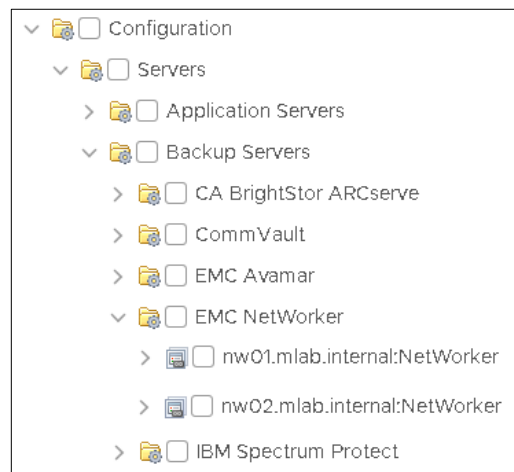


Figure 2 - Discovered NetWorker Servers Linked to EMC NetWorker category

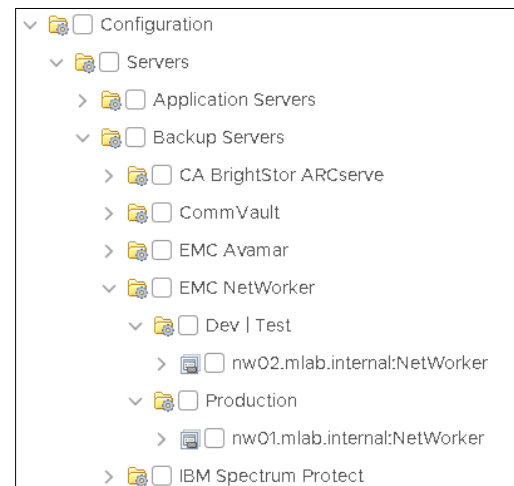


Figure 3 - Customize the Groups to Provide Granular Reporting Capability

2.5 Smart Groups

A Smart Group can be created by a DPA user with administrative privileges and is used to obtain a list of objects by running a report where the results are filtered based on a criterion. When creating a Smart Group, the returned objects list can be based on a single level or multi levels.

In Figure 4, a Smart Group called SQL is created with the following parameters. A single-level Smart Group is selected and based on the Backup Client Configuration report, where the selected scope is the Backup Server, which in this example is NetWorker. Clients are filtered based on a Client Name, where the client name contains “sql” over a time period of the Last Day once a day at 12am.

The final options to select relate to the fields that are to be compared and returned. If the results from the Smart Group do not align with the expected results, it may be due to the fields selected. For this example, the *Children Type* is the Backup Client and the *Server Name* Field is

Server (think of this as the backup server) and the final option is the *Backup Client* Field, which for this example as the Smart Group needs to return the Client Names, the Client Name is selected.

For ease of management, it is recommended to create Smart Groups in a dedicated group/s, rather than saving them within the Servers, Storage or Switches groups located under Configuration. In Figure 5 a new group called Smart Groups was created under Configuration and the newly created SQL Smart Group is saved within this group. This Smart Group is run daily at 12pm and the results are displayed on the right-side pane in Figure 5.

Create Smart Group

Single-level
 Multilevel

CONFIGURE SMART GROUP LEVEL

Objects in Smart Group

Report Name: Backup Client Configuration
Scope: Configuration>Servers>Backup Servers>EMC NetWorker
Time Period: Last Day
Filter: Client Contains "sql"
Children Type: Backup Client
Server Name Field: Server
Backup Client Field: Client

Generation Frequency: 12:00 am

SELECT FREQUENCY

Figure 4 - Single Level Smart Group Configuration

Name	Path	Type	Date Discovered
sql02.mlab.internal	nw01.mlab.internal:NetWorker	Backup Client	2/6/20, 3:29 PM

Figure 5 – Smart Group called SQL

With the Smart Group created and working based on the filtering criteria, reports can now be run where the scope is the smart group, rather than the static groups of Servers, Storage or Switches listed under Configuration.

Smart Groups work well when a consistent naming standard for applications and servers has been implemented within the organization. Using Smart Groups avoids creation of static groups as shown in Figure 3 where ‘Prod’ and ‘Dev | Test’ groups were created. Within environments that are very dynamic,

Smart Groups greatly reduce the maintenance of DPA where clients need to be manually assigned to groups. And, they are extremely useful when reporting of data protection activities is required for application owners.

2.6 Data Collection Retention

DPA enables different retention periods for data collected by agents for any discovered data protection element. The default values assigned to each data metric being collected can be adjusted under the *Data Collection* menu option; select *Defaults* to view all data collection defaults.

In Figure 6 the default settings for the NetWorker client status is shown where data gather will be retained for 13 weeks. Depending on stakeholder requirements, this default value can be adjusted to match the desired reporting requirements. Any data collected older than the specified retention period will be aged out of the DPA datastore.

For NetWorker configuration data, the risk/security team may require a longer period of retention for configuration changes to satisfy their compliance requirements.

It is worth ensuring data being collected is retained for the desired period. Adjusting this in the data collection defaults section will ensure consistency of data captured by DPA.

With the default values configured, DPA provides the flexibility to adjust these parameters for every data element being protected during the discovery wizard process.

The screenshot shows the 'Edit Default Request' interface. It includes sections for 'Request Name' (NetWorker Client Status), 'Retention period' (13 Weeks, Forever checkbox), 'Data Collection Agent' (Local selected, Remote with 'SELECT DATA COLLECTION AGENT' button), 'Credential' (with 'SELECT CREDENTIAL' button), and 'Frequency' (Period 1 Hours selected, Schedule with 'SELECT SCHEDULE' button, Manual). A footer note reads: 'The request will not collect data unless you run it manually'.

Figure 6 - Data Collection Settings

2.7 The Elements of Running a Report

There are three elements required to run a report; the scope, a report template and time period. Expanding the scope reveals the group structure created by DPA during the installation and any customization of those groups.

Group Selection

Select a group, an item or multiple combinations thereof to against which the report will be run. For example, in section 0 the grouping of backup elements to be monitored was discussed. Continuing with that example, Figure 7 shows a series of images (from left to right) demonstrating the flexibility of choice with groups to report on the Production, Dev | Test environments or finally, all NetWorker servers.

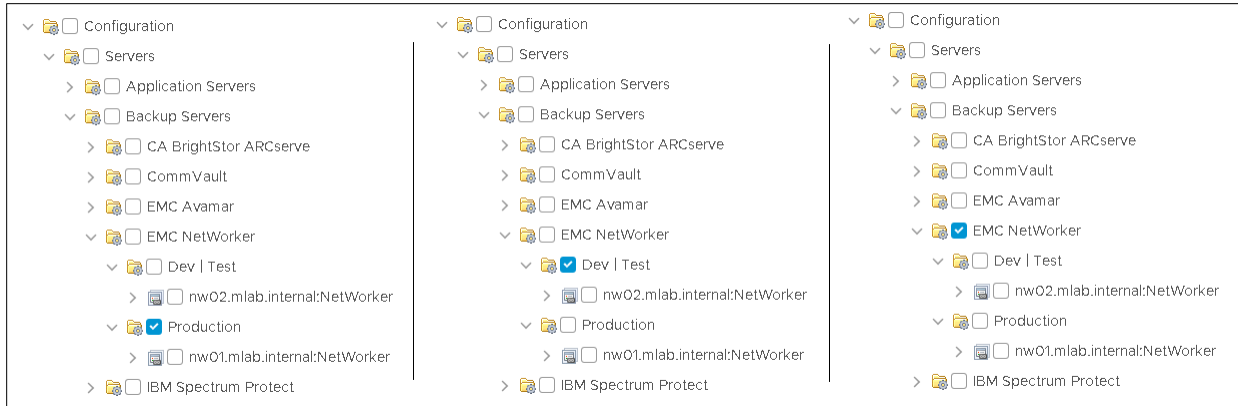


Figure 7 - Flexibility in Selecting Backup Elements with Group Selection

All discovered backup elements that are contained within the selected parent group or sub-groups are used when a report is generated.

If Smart Groups have been created as described in section 0, these are also available to be selected when creating a report.

Report Templates

With the group/s (scope) selected, DPA will present a list of available reports which can be run. DPA filters the list of reports based on the scope of item/s selected, for example, reports only relevant to NetWorker will be displayed if that was selected in the scope. Keep this in mind when generating reports, if the report cannot be found as expected, check the scope selection and adjust as needed.

Looking at the list of available reports, it can be daunting to know which one to select or where to start. In section 3 a sample of reports is provided covering the four stakeholders mentioned in this article. While the sample reports are not extensive, it should provide a solid foundation to extract value from DPA in a short period.

Time Period

The last element that is needed is to define the time period and typically a time period of the last day is selected by default. Under the *System Settings* of DPA, *Time Periods* can be created, modified or deleted as required.

Running a report using the *Last Day* time period may result in the report providing results based on overlapping backup windows, skewing the results which may not be desirable. The time period of Last Day is run from the time the report is run and covers the previous 24 hours.

Edit Time Period > Manage Times > Time Properties

Name	8pm Last Night
Description	
Day of Week	
Day of Month	1 days ago
Month	
Year	
Hour	8 PM
Minute	

Figure 8 - Creating a Starting Time for Last Night's Backup

If the prescribed backup window which starts last night at 8pm and completes by 6am the following day, create a custom time period to allow reports to be created for last night's backups. This time period may not be suitable for every report; however, it does allow the report to focus on a prescribed period. Now, regardless of when the report is run, the time period will match as defined. Creating the report is straight forward and an example is provided for reference

Navigate to the *System Settings* of DPA, then *Time Periods* and select the *Create Time Period* option. Review the start and end time options available in the list and if the desired times are not listed, create a new time by selecting *Edit Times* option. From here, create two new times periods for starting *8pm last night* and for *6am this morning* time. Select *Create* and for last night, set the *Day of Month* to *1 days ago* and then specify the hour or 8pm as shown in Figure 8. For the finishing time, select *Create* and set the hour to 6am as shown in Figure 9.

With the start and end times defined, the time period can be created and leveraged in creating reports. Simply create a new time period, select the newly created start and end times, provide a description as shown in Figure 10 and the time period is ready for use.

Figure 9 - Create a Finish Time Last Night's Backup

Figure 10 - Creation of the Time Period

3 Reporting and Dashboards

With some of the basics covered and a few topics to consider when configuring DPA, lets dive into running a selection of reports and the value it provides to stakeholders. DPA contains hundreds of templated reports and detailed information about these reports is provided in the Report Reference Guide¹⁰.

Reports can be run in an ad-hoc manner or scheduled to run automatically, at any time to stakeholders via email, to a file repository or saved to a SharePoint site. The report formats available from DPA include image based, PDF, HTML or comma separate value (CSV). The image based, PDF or HTML formats are ideal for stakeholders to quickly review the information and action it accordingly as required. Report data sent to a file repository in a CSV format may be beneficial if additional analysis is required to be undertaken.

The sample of reports provided in this section are categorized by stakeholder to highlight the different insights or value that would be gleaned by the individuals. Reports which are relevant to every stakeholder will be included in the backup administrator section and commentary will be provided when the report is relevant to other stakeholders. The backup application selected for these reports is NetWorker protecting a small number of servers in a lab environment with Data Domain as the protection storage.

3.1 Reports for the Backup Administrator and Operations Manager

For the most part, reports run for the Backup Administrator are also valuable to the Operations Manager, except for a few reports. In general, the Operations Manager will require reports that detail the health and volume of data protected, but more importantly the failed backups which may lead to data recovery exposures or risks.

Job Summary

A simple and quick report that provides a summary of the number of successful, failed and active jobs. A sample of the Failed Jobs report (Report template: *Data Protection | Jobs | Job Summary*) is shown in Figure 11 and is ideal for the Operations Manager.

<input type="checkbox"/>	Completed	Succeeded	Failed	Active	Size (GB)	Success Rate (%)
<input type="checkbox"/>	179	179	0	0	54	100

Figure 11 - Job Summary

Admittedly, this report does not provide a great level of detail, however it provides the backup administrator a quick glance of the data protected for the last day in the above example. Referencing this high-level report over time will allow the administrator to quickly determine if there are any unusual anomalies in the number of jobs completed or volume of data processed. If an anomaly is noticed, the administrator can then analyze data from other reports to identify what has occurred.

Failed Jobs

Without a doubt, the most important report that any Backup Administrator will wish to view is the failed backup jobs. This report is also valuable to share with the appropriate applications owner/s as knowing that a backup has failed allows them to investigate the health of the affect server/s. A sample of the Failed Jobs report (Report template: *Data Protection | Jobs | Failed Jobs*) is shown in Figure 12.

<input type="checkbox"/>	Server	Media Server	Group	Client	Policy	Workflow	Action	Job	Status	Error Code	Level
<input type="checkbox"/>	nw01.mlab.internal	nw01	RL	test01.mlab.internal	ztest-RL	backup	backup	C:	failed	255	Full
<input type="checkbox"/>	nw01.mlab.internal	nw01	SQL02-test	sql02.mlab.internal	SQL02-test	Workflow1	backup	MSSQL\$LABDB:AdventureWorks2012	failed	4294967295	
<input type="checkbox"/>	nw01.mlab.internal	nw01	SQL02-test	sql02.mlab.internal	SQL02-test	Workflow1	backup	MSSQL\$LABDB:AdventureWorks2012	failed	1	Full
<input type="checkbox"/>	nw01.mlab.internal	nw01	SQL02-test	sql02.mlab.internal	SQL02-test	Workflow1	backup	MSSQL\$LABDB:StackOverflow2013	failed	4294967295	
<input type="checkbox"/>	nw01.mlab.internal	nw01	SQL02-test	sql02.mlab.internal	SQL02-test	Workflow1	backup	MSSQL\$LABDB:StackOverflow2013	failed	1	Full

Figure 12 - Failed Jobs

For a backup to be successful, the entire backup path and infrastructure must be functional from the storage hosting the application, through the server itself, network, backup server and agent, and destination protection storage. Any one of these elements could impact the success of a backup and configuring DPA to report on failed backups to the application owners will allow them to investigate backup failures from their perspective. Resolving failed backups is best achieved when both the backup and application owner work together.

Uncloned Backups

While it is easy to focus purely on client backup success and failures, don't forget about replication or cloning of backups when multiple sites are protected. The Uncloned Backups report should be run daily and is one that should be provided to all stakeholders. A sample of the Uncloned Backups report (Report template: *Data Protection | Clones | Uncloned Backups*) is shown in Figure 13.

Server	Group	Client	Policy	Workflow	Job
nw01.mlab.internal	Active Director - Basic	ad01.mlab.internal	Active Directory	Basic - 1 Local Copy	DC=MLab,DC=Internal
nw01.mlab.internal	VM Image - Basic	ad02	VM Image	Basic - 1 Local Copy	vm:5035b1ea-6184-fd35-a784-c5adcf0b5f02:vcenter.mlab.internal
nw01.mlab.internal	VM Image - Basic	ddmc	VM Image	Basic - 1 Local Copy	vm:500b8915-c657-da66-2d6c-184d0404fd4:vcenter.mlab.internal
nw01.mlab.internal	VM Image - Basic	dpc	VM Image	Basic - 1 Local Copy	vm:500b104c-c17b-57fb-be21-7d3e5359a46b:vcenter.mlab.internal
nw01.mlab.internal	Virtual Synthetic	jb01.mlab.internal	Virtual Synthetic	Workflow1	WINDOWS ROLES AND FEATURES:
nw01.mlab.internal	Virtual Synthetic	jb01.mlab.internal	Virtual Synthetic	Workflow1	E:
nw01.mlab.internal	Virtual Synthetic	jb01.mlab.internal	Virtual Synthetic	Workflow1	\\?\VOLUME{1FEFB185-9DD6-11E6-93E8-806E6F6E6963}
nw01.mlab.internal	Virtual Synthetic	jb01.mlab.internal	Virtual Synthetic	Workflow1	C:
nw01.mlab.internal	Virtual Synthetic	jb01.mlab.internal	Virtual Synthetic	Workflow1	WINDOWS ROLES AND FEATURES:
nw01.mlab.internal	Virtual Synthetic	jb01.mlab.internal	Virtual Synthetic	Workflow1	DISASTER_RECOVERY:

Figure 13 – Uncloned Backups

Application backups without a clone or second copy places the organization at risk from a disaster recovery perspective. The application owner needs to understand the available recovery point objectives (RPO) and the risk team needs to know from an exposure/compliance viewpoint. Unless disclosed in the data protection policy, the backup administrator needs to address backups without a clone or copy.

The Report Card

Reviewing reports in a tabular format containing lots of information may not enable the backup administrator to detect re-occurring backup related issues. For example, a client may fail in full or in part on a regular basis and for large environments, keeping track of failed clients may prove difficult. Depending on the report, DPA provides the ability to visualize data, enabling an administrator or application owner to identify issues more easily. One of these reports that allows failed backups to be visualized is the Report Card report. A sample of the Report Card report (Report template: *Data Protection | Overviews | Report Card*) is shown in Figure 14.

Node	8 Dec since 11:32	9 Dec	10 Dec	11 Dec	12 Dec	13 Dec	14 Dec	15 Dec	16 Dec	17 Dec	18 Dec	19 Dec	20 Dec	21 Dec
ad01.mlab.internal	█	█	█	█	█	█	█	█	█	█	█	█	█	█
ad02							█	█	█	█	█	█	█	█
ex01.mlab.internal	█	█	█	█	█	█	█	█	█	█	█	█	█	█
jb01.mlab.internal	█	█	█	█	█	█	█	█	█	█	█	█	█	█
jb02.mlab.internal							█	█	█	█	█	█	█	█
linux10.mlab.internal	█	█	█	█	█	█								
test01.mlab.internal	█	█	█	█	█	█	█	█	█	█	█	█	█	█
vcenter.mlab.internal	█	█	█	█	█	█	█	█	█	█	█	█	█	█

Figure 14 – The Report Card

When run over a time period of the last two weeks in the above example, it is easy to identify clients with re-occurring issues. One client has some failed backups every 4th day and this may not be evident if the backup administrator or application owner only review the failed backup report. This report also allows for each identification of clients no longer being protected or those which have been added to a backup policy. In the above example, one client is no longer being protected since the 15th, while two new clients were added to a protection policy on the 14th of the month.

This report may not need to be run every week, but certainly is valuable to run this if there is instability in the backup environment or there have been a lot of configuration changes implemented.

Client Schedules

Continuing the visualization aspect, balancing backup schedules can be challenging for larger environments and the Client Schedule report provides the backup administrator the ability to visualize the backup schedule. A sample of the Client Schedule report (Report template: *Scheduling | Backups | Client Schedule*) is shown in Figure 15.




Figure 15 – Client Schedule

Like the Report Card report, visualizing the client schedules would be performed in an ad-hoc manner or prior to adding new clients to the backup environment. Scheduling client backups during quieter periods in the backup window will ensure that the backup infrastructure is utilized efficiently. This report helps to visualize any scheduling hot spots that result in backup jobs queuing and likely to place a strain on the backup infrastructure.

Longest running clients

Ensuring that all backups are completed within the designated backup window is often undone by a handful of long running clients – large database or file systems come to mind as culprits. Generally, the backup administrator is aware of these troublesome clients, however with large environments identification of these clients may be more difficult. A sample of the Longest Running Clients report (Report template: *Data Protection | Clients | Top 10 Longest Running Clients*) is shown in Figure 16.

Client	Longest Run (hour)
jb02.mlab.internal	2 hours 52 minutes
jb01.mlab.internal	1 hour 28 minutes
nw01.mlab.internal	36 minutes 36 seconds
vcenter.mlab.internal	19 minutes 18 seconds
test01.mlab.internal	17 minutes 7 seconds
ex01.mlab.internal	14 minutes 4 seconds
win03	9 minutes 33 seconds
sql02.mlab.internal	7 minutes 29 seconds
win05.mlab.internal	4 minutes 48 seconds
dpc	2 minutes 21 seconds

Figure 16 – Top 10 Longest Running Clients

Ideally the time taken to complete long running backups should be within the RPOs as defined by the organization or application owner. Every effort should be made to ensure clients are protected as quickly as possible; however, the backup administrator is bound by the throughput and performance of the entire data protection path – client to protection storage.

If a client backup time exceeds the specified RPO, the risk team need to be informed of this situation and action it accordingly. One option is to increase the time period between RPOs which would address this violation, but greatly increase the exposure risks for data recovery. Alternatively, the risk team/organization invests in a solution to address these long backups to achieve backups within the prescribed RPO.

Backup Job Change Ratios

Knowing the daily rate of change for a client, a group of applications or the entire backup environment may be difficult to calculate. This information is often required when calculating the storage requirements of the protection storage platform. All too often, an estimated guess is made which may be acceptable. However, with DPA, this guesswork is removed. A sample of the Aggregate Backup Job Change Ratios report (Report template: *Data Protection | Change Ratios | Aggregate Data Change Ratio*) is shown in Figure 17.

Protected (TB)	Last Incremental Size (MB)	Last Incremental Rate (%)	Average Incremental Size (MB)	Average Incremental Rate (%)
1.729	485	0.027	622.19	0.034

Figure 17 – Aggregate Backup Job Change Ratios

The average incremental size will vary and is dependent on the time period selected. In the above example a time period of the last week was selected and this average incremental backup size can then be compared to the last incremental size.

Breaking down the rate of change per client or group of applications is achieved running the Data Change Ratio by Client report (Report template: *Data Protection | Change Ratios | Data Change Ratio by Client*) is shown in Figure 18.

Server	Client	Protected (GB)	Last Incremental Size (MB)	Last Incremental Rate (%)	Average Incremental Size (MB)	Average Incremental Rate (%)	Total Size (GB)	Size Rate (%)
nw01.mlab.internal	jb01.mlab.internal	221.123	149	0.066	257	0.114	5.832	2.637
nw01.mlab.internal	jb02.mlab.internal	52.032	96	0.18	79.571	0.149	1.41	2.71
nw01.mlab.internal	nw01.mlab.internal	2.495	90	3.523	129.786	5.08	1.136	45.519
nw01.mlab.internal	test01.mlab.internal	122.971	8	0.006	8.5	0.007	0.596	0.484
nw01.mlab.internal	vcenter.mlab.internal	3.205	142	4.327	147.333	4.489	1.441	44.973

Figure 18 – Data Change Ratio by Client

Like the job summary in section 0, reviewing this change ratio report on a regular basis will enable the backup administrator to detect if there are any anomalies compared to the average change ratio. The data change ratio report is also an ideal to share with application owners, enabling them to understand the changed data protected daily.

Protected Backup Capacity

Knowing how much data is protected for a subset of application, by a single backup application or across multiple environments is easily obtained by running the Estimated Protected Backup Capacity Details report. This report looks for the largest backups completed for each client over the time period specified. Running this report with a minimum time period of at least 1 month is recommended. A sample of the Estimated Protected Backup Capacity Details report (Report template: *Status | Backup | Estimated Protected Backup Capacity Details*) is shown in Figure 19.

Server	Client	Capacity Protected (Not MSSQL via NW/Avamar) (GB)	Capacity Protected (Virtual Host) (GB)	Capacity Protected (GB)
nw01.mlab.internal	ad01.mlab.internal	16.208		16.208
nw01.mlab.internal	ex01.mlab.internal	31.045		31.045
nw01.mlab.internal	jb01.mlab.internal	220.916		220.916
nw01.mlab.internal	jb02.mlab.internal	51.825		51.825
nw01.mlab.internal	nw01.mlab.internal	30.499		30.499
nw01.mlab.internal	ad02		52.099	52.099
nw01.mlab.internal	ddmc		240.087	240.087
nw01.mlab.internal	dpc		15.383	15.383
nw01.mlab.internal	mg01		58.097	58.097
nw01.mlab.internal	srs		64.001	64.001
nw01.mlab.internal	win05.mlab.internal		53.09	53.09
nw01.mlab.internal	sql02.mlab.internal			94.717

Figure 19 – Estimated Protected Backup Capacity Details

Resource utilization

When a DPA agent is installed on a backup server, not only does it collect information about the backup and recovery operations, it also can collect performance information about the host. The time taken to complete a backup is dependent on the entire IT infrastructure and knowing if the backup server itself is under strain needs to be known. With DPA, the backup administrator can report on the host. A sample of the Resource Utilization Overview report (Report template: *Performance | Overviews | System Performance*) is shown in Figure 20.




Figure 20 – Resource Utilization Overview – CPU, Memory and Network graphs shown

This report is in fact a dashboard and the CPU, memory and network graphs are shown in Figure 20 and can be generated individually under *Performance* and *Resource Utilization* sections of the report templates.

Resource and performance data can be collected from other data protection elements including Data Domain, tape libraries, switches, storage arrays and applications. Being able to obtain performance metrics from DPA will assist the backup administrator troubleshooting performance related issues.

3.2 Reports for the Application Owner

Some of the reports discussed in the backup administrator section are also relevant to the application owners. In addition to those, other reports are available from DPA which may be of interest to the application and a selection is provided in this section.

Virtual Machines without Protection

Protection of Virtual Machines (VMs) can be undertaken by performing a VM image level backup, using a backup agent or a combination of both. There are several reports available to the VM administrator that should be run periodically to ensure VMs hosted on their platform are protected. Knowing which VMs have a current image level backup ensure that the VMware vCenter is being monitored by DPA and run the VM Protection report. A sample of the VM Protection report (Report template: *Data Protection / Exposure / VM Protection*) is shown in Figure 21.

ESX Server	Vhost	Client	Protected By	Last Successful Backup
esx01.mlab.internal	ad02	ad02	networker	1/12/21 8:00 PM
esx01.mlab.internal	ddmc	ddmc	networker	1/12/21 8:01 PM
esx01.mlab.internal	dpc	dpc	networker	1/11/21 8:03 PM
esx01.mlab.internal	linux01			
esx01.mlab.internal	nw01			
esx01.mlab.internal	nwvproxy01			
esx01.mlab.internal	ppdm03			
esx01.mlab.internal	sql01			
esx01.mlab.internal	sql10			
esx01.mlab.internal	srs	srs	networker	1/12/21 8:01 PM

Figure 21 – Estimated Protected Backup Capacity Details

Knowing that some VMs may be protected by using a backup agent, a subsequent exposure report will list only VMs without any form of backup. This report is called the Virtual Hosts Not Backed Up. This is a very simple report and a sample of the Virtual Hosts Not Backed Up report (Report template: *Data Protection / Exposure / Virtual Hosts Not Backed Up*) is shown in Figure 22.

Host Name
ddve02
nwvproxy01
pp01
pp02
ppdm01
win 2012 r2
win02
win04

Figure 22 – Virtual Hosts Not Backed Up

Application Groups Specific

If consistent naming standards have been used in the environment, consider using Smart Groups as discussed in section 0 when creating reports for Application Owners. An example was provided to filter for client names which contain “sql” and applying the same methodology for any other application type can be created. With Smart Groups, it is easy to provide consistent reporting experience for the application owner without needing to manually add and remove clients as the environment evolves.

Continuing on the SQL theme, all backup jobs for the SQL servers identified from the SQL Smart Group over a time period of the Last Day was run and is shown in Figure 23.

Client	Workflow	Job	Status	Level	Size (MB)	Files	Started	Finished	Duration (minute)
sql02.mlab.internal	Basic - 1 Local Copy	MSSQL\$LABDB:	success	Full	0	2	1/8/21 8:45 AM	1/8/21 8:45 AM	5 seconds
sql02.mlab.internal	Basic - 1 Local Copy	MSSQL\$LABDB:AdventureWorks2012	success	Full	1	3	1/8/21 8:40 AM	1/8/21 8:40 AM	9 seconds
sql02.mlab.internal	Basic - 1 Local Copy	MSSQL\$LABDB:LoggingDB	success	Full	0	7	1/8/21 8:40 AM	1/8/21 8:40 AM	3 seconds
sql02.mlab.internal	Basic - 1 Local Copy	MSSQL\$LABDB:StackOverflow2013	success	Full	161	3	1/8/21 8:40 AM	1/8/21 8:45 AM	4 minutes 56 seconds
sql02.mlab.internal	Basic - 1 Local Copy	MSSQL\$LABDB:master	success	Full	0	3	1/8/21 8:45 AM	1/8/21 8:45 AM	4 seconds
sql02.mlab.internal	Basic - 1 Local Copy	MSSQL\$LABDB:model	success	Full	0	7	1/8/21 8:45 AM	1/8/21 8:45 AM	3 seconds
sql02.mlab.internal	Basic - 1 Local Copy	MSSQL\$LABDB:msdb	success	Full	1	3	1/8/21 8:45 AM	1/8/21 8:45 AM	2 seconds
sql02.mlab.internal	Basic - Logs	MSSQL\$LABDB:	success	txnlog	0	2	1/7/21 5:04 PM	1/7/21 5:05 PM	4 seconds
sql02.mlab.internal	Basic - Logs	MSSQL\$LABDB:AdventureWorks2012	success	Full	1	3	1/7/21 5:00 PM	1/7/21 5:00 PM	10 seconds
sql02.mlab.internal	Basic - Logs	MSSQL\$LABDB:LoggingDB	success	Incr	0	3	1/7/21 5:00 PM	1/7/21 5:00 PM	4 seconds
sql02.mlab.internal	Basic - Logs	MSSQL\$LABDB:StackOverflow2013	success	Full	161	3	1/7/21 5:00 PM	1/7/21 5:04 PM	4 minutes 38 seconds
sql02.mlab.internal	Basic - Logs	MSSQL\$LABDB:master	success	Full	0	3	1/7/21 5:04 PM	1/7/21 5:04 PM	2 seconds
sql02.mlab.internal	Basic - Logs	MSSQL\$LABDB:model	success	Incr	0	3	1/7/21 5:04 PM	1/7/21 5:04 PM	1 second
sql02.mlab.internal	Basic - Logs	MSSQL\$LABDB:msdb	success	Full	1	3	1/7/21 5:04 PM	1/7/21 5:05 PM	5 seconds

Figure 23 - All Jobs for SQL Client

Several reports examples provided in section 0 are also relevant to the application owners, including the Backup Job Change Ratios, Protected Backup Capacity and without question, the Failed Jobs. Providing application owners with visibility to data protection activity and metrics provides them deeper insights into their application.

3.3 Reports for the Risk/Security Teams

Risk and Security teams are more interested in reports that enable them to verify if data protection is within prescribed SLAs or RPOs. For auditing purposes, Risk and Security teams also want to know what changes have occurred within the backup environment

Frequent Recoveries

Data protection is implemented to provide recovery of data, enabling applications to continue serving customers or end users. From a security point of view, being able to report on the recovery operations performed, how often and from which applications is of high value to them. With DPA, the Restore Details report run across the entire environment or a specific group provides details of both failed and successful recoveries. A sample of the Restore Details report (Report template: *Data Protection | Restores | Restore Details*) is shown in Figure 24.

Server	Client	Job	Status	Err Code	Size (GB)	Num Files	Backup Time	Queued	Started	Finished
nw01.mlab.internal	jb01.mlab.internal	E:\	failed	1	0	1	1/11/21 12:00 PM	1/12/21 4:31 PM	1/12/21 4:31 PM	1/12/21 4:31 PM
nw01.mlab.internal	jb01.mlab.internal	E:\	success		0.474	415	1/11/21 12:00 PM	1/12/21 4:32 PM	1/12/21 4:32 PM	1/12/21 4:33 PM
nw01.mlab.internal	jb01.mlab.internal	E:\	success		3.273	9	1/11/21 12:00 PM	1/12/21 5:00 PM	1/12/21 5:00 PM	1/12/21 5:01 PM
nw01.mlab.internal	sql02.mlab.internal	MSSQL\$LABDB:StackOverflow2013	success		47.510		1/12/21 12:00 PM	1/12/21 4:36 PM	1/12/21 4:36 PM	1/12/21 4:46 PM

Figure 24 – Restore Details

While most recovery operations are for valid business reasons, there may be occasions when unusual recovery operations are undertaken and if detected, further investigations can be performed.

Backup SLA summary

IT operations teams are there to serve the organization and with that, a level of service is expected to be maintained. Typically, a data protection policy will include an SLA for backup operations per client, for a group of clients or the entire environment. Within the Service Level Management (SLM) set of report templates, several SLA-based reports can be run for the backup administrator, operations manager and risk/security teams. A sample of the Daily Success Rate report (Report template: *Service Level Management | Backups | Daily Success Rate*) is shown in Figure 25 in a tabular format view.

Object	Success Rate (%)	Noted
EMC NetWorker	99.359	12/11/20 4:23 PM
EMC NetWorker	99.359	12/12/20 4:23 PM
EMC NetWorker	99.535	12/13/20 4:23 PM
EMC NetWorker	100	12/14/20 4:23 PM
EMC NetWorker	100	12/15/20 4:23 PM
EMC NetWorker	100	12/16/20 4:23 PM
EMC NetWorker	100	12/17/20 4:23 PM

Figure 25 – Daily Success Rate

Breaking down the overall success rate per client will allow problematic clients to be identified. An example of the SLA Summary by Client report (Report template: *Service Level Management | Backup | SLA Summary by Client*) run over a longer time period is provided in Figure 26. From this report, only one client has experienced backup issues and affected the overall data protection SLA.

<input type="checkbox"/>	Object	Jobs	Successful	Within SLA	% Successful (%)	% Within SLA (%)
<input type="checkbox"/>	nw01.mlab.internal:NetWorker:ad01.mlab.internal	132	132	132	100.0	100.0
<input type="checkbox"/>	nw01.mlab.internal:NetWorker:ad02	27	27	27	100.0	100.0
<input type="checkbox"/>	nw01.mlab.internal:NetWorker:ex01.mlab.internal	240	240	240	100.0	100.0
<input type="checkbox"/>	nw01.mlab.internal:NetWorker:jb01.mlab.internal	265	265	265	100.0	100.0
<input type="checkbox"/>	nw01.mlab.internal:NetWorker:jb02.mlab.internal	250	250	250	100.0	100.0
<input type="checkbox"/>	nw01.mlab.internal:NetWorker:linux10.mlab.internal	6	6	6	100.0	100.0
<input type="checkbox"/>	nw01.mlab.internal:NetWorker:test01.mlab.internal	169	163	163	96.4	96.4
<input type="checkbox"/>	nw01.mlab.internal:NetWorker:vcenter.mlab.internal	159	159	159	100.0	100.0

Figure 26 – SLA Summary by Client over a Month Time Period

The SLA Summary by Client report uses conditional formatting which can be customized to align with the SLA percentages as defined in the organizations' data protection policy. The report in Figure 26 was customized to highlight clients with an SLA less than 98% in orange and any clients below 95% success rate in red. To adjust the conditional formatting, select the Report Format option and click on Table Format as shown in Figure 27.

<input type="checkbox"/>	Client	Jobs	Successful	Within SLA	% Successful (%)	% Within SLA (%)
<input type="checkbox"/>	ad01.mlab.internal	128	128	128	100	
<input type="checkbox"/>	win03	41	41	41	100	
<input type="checkbox"/>	nw01.mlab.internal	713	713	713	100	

Figure 27 - Table Format

Then, as shown in Figure 28, navigate to Series, Cell Styles and by clicking on the 3 vertical dots for each condition, edit the values to align with the SLA conditions in the backup policy. All reports, regardless of the Table type can be customized.

Report	Table Styles	Series	Advanced
Cell Styles			
⋮	All column(s) in all rows where % Within SLA Less than "95"		
⋮	All column(s) in all rows where % Within SLA Less than "98"		
⋮	All column(s) in all rows		

Figure 28 - Customizing the Series

Backup and restore KPIs

Being able to determine activity within the backup environment over a given time period, whether it is the last day, week, month or year, is key information provided by the Backup and Restore KPIs report. This report is not only relevant to the Backup Administrator and Operations Manager, but also offers historical insights for the risk/security team. A sample of the Backup and Restore KPIs report (Report template: Compliance and Risk Mitigation | Backup and Restore KPIs) run over a period of one month is shown in Figure 29.

Total Backups	5687
Total Restores	4
Backup Success Rate (%)	99.93
Restore Success Rate (%)	75
Backup Data (GB)	575
Restore Data (GB)	4
% of Backups Restored (%)	0.053
% of Data Restored (%)	0.696

Figure 29 - Backup and Restore KPIs

Configuration Changes

Configuration changes in the backup environment take place almost daily and keeping track of these changes will be known by the Backup Administrator, while the Operations Manager and Risk/Security Teams may have no idea what is occurring. Through monitoring the backup servers, DPA can keep track of changes and provide reporting of these changes. Organizations may be obligated to keep track of changes for auditing or security compliance reasons.

Starting with a simple report, Client Configuration Changes provides a high-level overview of clients which have been added, modified or removed. A sample of this report (Report template: Compliance and Risk Mitigation | FDA | Client Configuration Changes) is shown in Figure 30.

Server	Client	Change
nw01.mlab.internal	ex01.mlab.internal	Modified
nw01.mlab.internal	jb01.mlab.internal	Modified
nw01.mlab.internal	jb02.mlab.internal	Added
nw01.mlab.internal	nw01.mlab.internal	Modified
nw01.mlab.internal	search.mlab.internal	Removed
nw01.mlab.internal	sql02.mlab.internal	Modified

Figure 30 - Client Changes

Configuration changes that DPA reports on for NetWorker backup server as used within this article include client, group, policy, jobs, schedule and media, with each reporting on changes relevant to that aspect of the backup server. A sample of Job Configuration Changes report (Report template: Change Management | Backup | Job Configuration Changes) is shown in Figure 31.

Server	Client	Group	Policy	Workflow	Job	Change	Noted	Noted
nw01.mlab.internal	ex01.mlab.internal	Filesystem - Standard	Filesystem	Standard - 2 Local Copies	C:\	Deleted	12/14/20 9:34 AM	12/14/20 9:34 AM
nw01.mlab.internal	ex01.mlab.internal	Filesystem - Standard	Filesystem	Standard - 2 Local Copies	DISASTER_RECOVERY\	Deleted	12/14/20 9:34 AM	12/14/20 9:34 AM
nw01.mlab.internal	isilon01.mlab.internal	Isilon			/ifs/labdata/Text_files	Added	12/14/20 9:34 AM	12/14/20 9:34 AM
nw01.mlab.internal	isilon01.mlab.internal	Isilon	NAS	Standard - 2 Local Copies	/ifs/labdata/Text_files	Deleted	12/14/20 9:34 AM	12/14/20 9:34 AM
nw01.mlab.internal	isilon01.mlab.internal	NAS_Test - Remote Index			/ifs/labdata/Text_files	Deleted	12/14/20 9:34 AM	12/14/20 9:34 AM
nw01.mlab.internal	jb01.mlab.internal	Virtual Synthetic	Virtual Synthetic	Workflow1	All	Added	12/14/20 10:43 PM	12/14/20 10:43 PM
nw01.mlab.internal	jb01.mlab.internal	Virtual Synthetic	Virtual Synthetic	Workflow1	E:\	Deleted	12/14/20 10:43 PM	12/14/20 10:43 PM
nw01.mlab.internal	jb02.mlab.internal	Virtual Synthetic	Virtual Synthetic	Workflow1	All	Added	12/14/20 10:43 PM	12/14/20 10:43 PM

Figure 31 - Job Configuration Changes

The final example relates to backup servers, knowing what changes have occurred needs to be tracked for auditing purposes. In Figure 32 changes to all backup servers being monitored by DPA are listed. A sample of Job Configuration Changes report (Report template: Change Management | Backup | Server Configuration Changes) is shown in Figure 32.

Server	Change	Noted	Difference
nw01.mlab.internal	Added	2/6/20 3:29 PM	
nw01.mlab.internal	Modified	7/6/20 4:33 PM	Version value was NetWorker 19.2.0.1.Build.117 Enterprise Edition, now NetWorker 19.3.0.0.Build.21 Enterprise Edition
nw01.mlab.internal	Modified	12/14/20 10:43 PM	Version value was NetWorker 19.3.0.0.Build.21 Enterprise Edition, now NetWorker 19.4.0.0.Build.25 Enterprise Edition
pp01.mlab.internal	Added	2/14/20 12:26 PM	
pp01.mlab.internal	Modified	4/17/20 3:07 PM	Application value was PowerProtect Data Manager 19.3.0-7, now PowerProtect Data Manager 19.4.0-10. Version value was 19.3.0-7, now 19.4.0-10
pp01.mlab.internal	Modified	11/5/20 8:56 AM	Application value was PowerProtect Data Manager 19.4.0-10, now PowerProtect Data Manager 19.6.0-3. Version value was 19.4.0-10, now 19.6.0-3

Figure 32 – Server Configuration Changes

3.4 Summary Dashboards for All

There are several pre-defined dashboards designed to provide an overview of various elements of the data protection environment in a single view. These dashboards are built by combining individual report templates into a single page. These summary dashboards are of value to the Backup Administrator and Operations Manager when looking holistically across the entire environment. These dashboards can also provide relevant information to Application Owners when setting the scope on their applications only.

Some of the dashboards contain large amounts of information which is not always easy to display within this article and therefore, only selective report elements will be shown.

Backup Client Summary

Starting with the Backup Client Summary report, this report displays several elements of interest, including a summary of the data protection activities, changes in client configurations and the volume of data stored. Shown in Figure 33 is the Backup and Restore Summary element from the Backup Client Summary report (Report template: Overviews | Backup | Backup Client Overview) to provide current and historical data protection activities.

Period	Last Day	Last Week	Last Month
Total Backups	164	1148	5282
Backup Success Rate (%)	100	100	100
Backup Data (GB)	3	29	157
Total Restores	0	0	4
Restore Success Rate (%)	0	0	75
Restore Data (GB)	0	0	4

Figure 33 - Backup and Restore Element from the Backup Client Overview Report

Compliance Overview

Most overviews or dashboards containing lots of information are best viewed via the DPA web UI. The reason for this relates to DPA's report drill down feature. There are reports which allow the user the ability to drill down to gather additional detailed information and the Backup Compliance Overview report (Report template: Overviews | Backup | Backup Compliance Overview) is a good example of this. In Figure 34 several reporting elements are shown, and the Strike Summary and Client Changes reports have the drill down feature. By clicking on the numbers within these reports, additional information is loaded in a separate report tab within DPA.

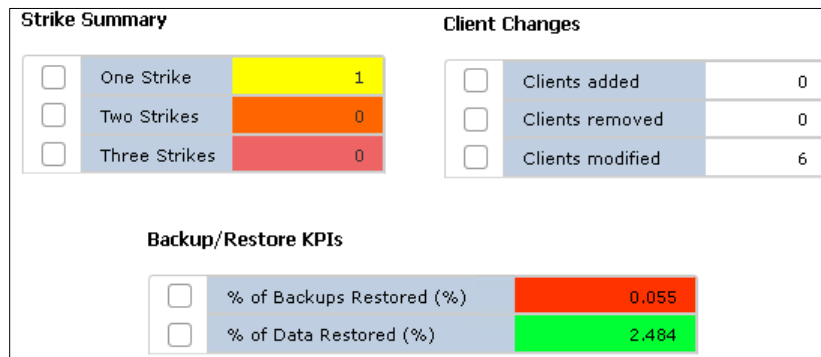


Figure 34 - Selection of Elements from the Backup Compliance Overview

By clicking on the numeric value in the One Strike Summary report element, a One Strike Failed Clients report is automatically loaded as shown in Figure 35 and lists the clients with failed backups in this example. The same can be achieved for the Client Changes report element and Figure 36 shows a list of modified clients over the time period used to for the Backup Compliance Overview report.

<input type="checkbox"/>	Client
<input type="checkbox"/>	sql02.mlab.internal

Figure 35 - Drill Down of One Strike Reporting Element

Server	Client	Differences
nw01.mlab.internal	vcenter.mlab.internal	Value of Backup Type has changed from vProxy to Filesystem
pp01.mlab.internal	dpc	Value of Active has changed from false to true, Value of Client Identifier has changed from c805e9e5-8d54-5739-a5a8-2ad5bb266b49 to 558da17f-5c87-544d-b01e-ad96c2d135a8
pp01.mlab.internal	fw01	Value of Active has changed from true to false, Value of Client Identifier has changed from ae3bb9b3-0558-5f0f-aacf-94b6e6bd4c3c to 9df1662f-8ea0-504d-9e05-af77d3098afc
pp01.mlab.internal	nwvproxy01	Value of Active has changed from false to true, Value of Client Identifier has changed from a1bd4f90-c146-5017-a052-ff7a6cad6c2f to d0cfa989-3bc1-543a-9ac1-025e944a8693
pp01.mlab.internal	vproxy10.mlab.internal	Value of Client Identifier has changed from 8e1ca5c5-686c-5114-8f00-cf6c230dd254 to 52d3d730-6690-5104-add6-9e9d659d64d2
pp01.mlab.internal	vproxy11.mlab.internal	Value of Client Identifier has changed from d68a6eee-7a1b-5de3-8bdf-c045bd247a1c to e382460e-c01e-5ed0-84fb-e4a0a79e7f0f

Figure 36 - Drill Down of Modified Clients

Further detail about reporting, including best practices are found in the DPA Product Guide¹².

4 Flexibility with Report Customization

DPA provides flexibility to customize existing system report templates or create brand new reports to cater to specific reporting requirements. Before diving into DPA to create your own reports, it is recommended to review the extensive list of existing report templates, saving them as your own and then make modifications as required.

4.1 Modifying a Report Template: Backup All Jobs – No Restarts

Where to start in customizing a report will depend on what specific information the Backup Administrator, Operations Manager, Application Owner or the Risk/Security teams wish to know. The reports that the stakeholders need to review typically relate to issues or anomalies within the backup environment. There is a report titled All Jobs – No Restarts (Report template: Data Protection | Jobs | All Jobs – No Restarts) which lists all jobs where a backup job restart did not occur. From a data protection point of view, it is more interesting to know which clients have had restarts in the last day, week or month. Knowing which clients had restarted backup operations may be a sign of underlying issues, even if the restarted job is successful.

With that in mind, a very simple report customization will be shown in this section so that the report lists only clients where backup jobs have restarted. With DPA web UI open, navigate to the *Report Templates* under *Reports*, select the *System Report Templates*. Either scroll through the list of available reports or click on the *filter* icon and enter the name of an existing report. In Figure 37 filtering was used to locate the Backup All Jobs – No Restarts report.

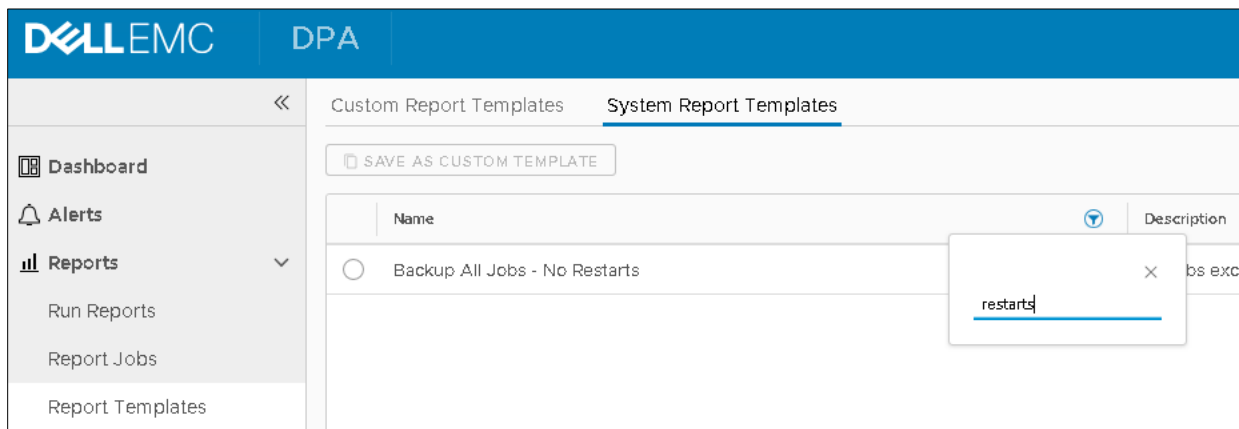


Figure 37 - Finding a System Report Template

Select the report and the *Save as Custom Template*, specify a name for the report and for this example the report name will be changed to *Backup All Jobs – Restart Occurred* as shown in Figure 38. This report

is created by taking several data sources that exist in the DPA datastore and performing operations between the data sources which result in a report being presented.

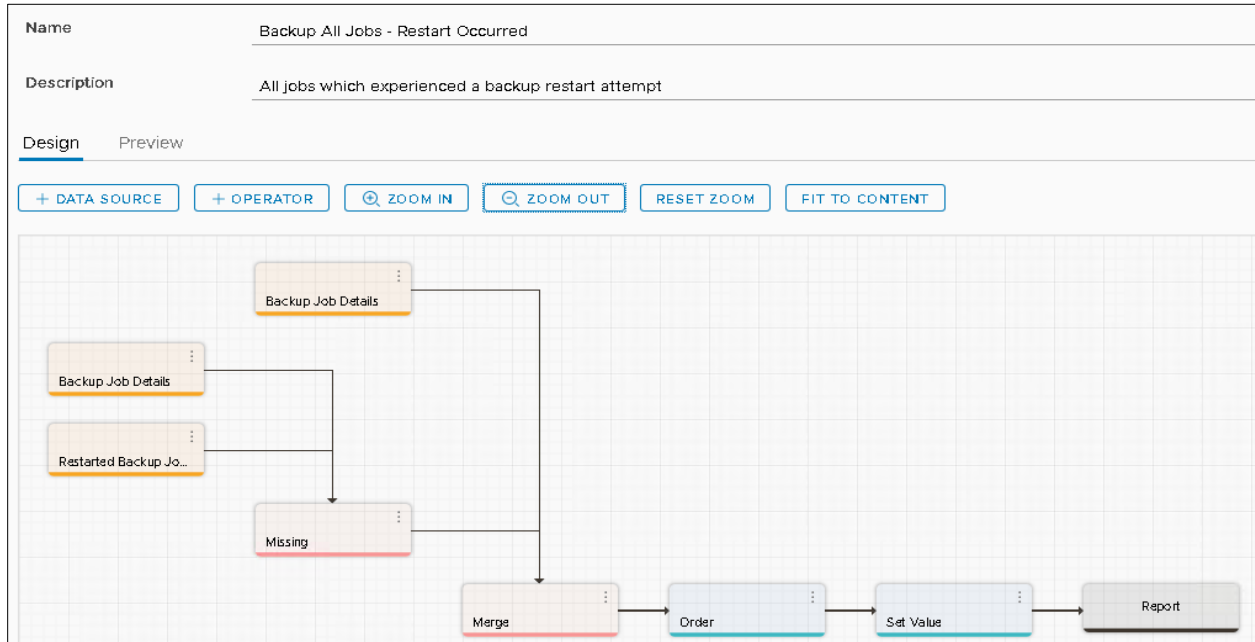


Figure 38 - Report Customization View

4.2 Adjusting the Merge Condition to Detect Restarts

By selecting the Merge operator, the properties of this operator are shown and in expanding the condition, there is nothing set as shown in Figure 39. A condition needs to be set to determine if a job restart did indeed occur. Before this can be done, the field value type (cast) needs to be known. Expanding the Fields section of the Merge operator properties, the Restarted Job field is located, and the field type or value is in the form of text.

If a job has restarted, this field will contain text-based information generated by the backup application and if no job restart occurred, the value in this field will be empty. To have this custom report generate output on for restarted jobs, a *Is Present* condition needs to be set. With the Condition properties expanded as shown in Figure 39, select Edit Condition and add a condition where the *Restarted Job* field *Is Present* as per Figure 40 and select *OK* to save the new condition.

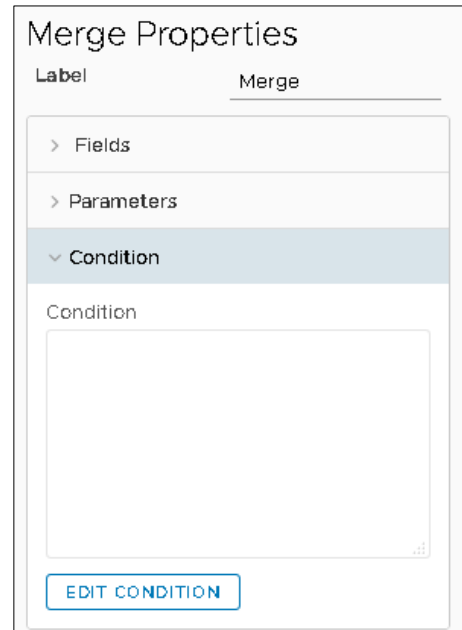


Figure 39 - Merge Operator Properties




Figure 40 - Creating a New Condition

Upon returning to the custom report template view, select *Save* and select *Preview* which is located just below the description of the report. It is now time to test if the modifications that were made to the report work. Select the scope, a time period and click on *Preview*. All going well and if job restarts have occurred, DPA will generate the results and display them accordingly. For the test environment being monitored, there was a single backup job restart as shown in Figure 41.

Server	Action	Job	Status	Err Code	Level	Size (MB)	Start Time	End Time	Job ID	Restarted Job	Group Job Id
nw01.mlab.internal	backup	MSSQL\$LABDB:	failed	1	Full	0	1/20/21 3:36 PM	1/20/21 3:36 PM	1841871	1841865	1841863

Figure 41 - Failed Backup Job that was Restarted

When the output of the report is working as expected, *Save & Close* the customized report. This by no means covers all aspects of creating customized reports. Further details can be found in the DPA Custom Report Guide¹¹.

5 Working with External Data Sources

Most data collected by DPA is achieved using agents and defining the data collection parameters when the discover wizard is used to monitor data protection elements or applications. Data can also be collected from a range of external data sources. A list of available data source types is shown in Figure 42. For example, extracting assets from a Configuration Management Database (CMDB) and cross matching them against discovered protected clients is extremely useful.

In this section an example is provided where external data in the form of a comma separated value (CSV) file, to represent CMDB data, is imported into a custom report and compared to protected clients. While using CSV files is not sustainable or dynamic, using the *Database Query* data source type allows data to be extracted directly from a CMDB.

The primary purpose of this custom report is to determine assets listed in the CMDB which are not protected, enabling the Backup Administrator to address the exposure of these unprotected clients. It can also be used to list clients protected, but not present in the CMDB, allowing the CMDB Administrator to remove old assets and ensuring the accuracy of the CMDB.

Category	Data Source
<input type="radio"/> External	Active Directory Host List
<input type="radio"/> External	Conditional Report
<input type="radio"/> External	Database Query
<input type="radio"/> External	Read XML File
<input type="radio"/> External	ReadCSVFile
<input type="radio"/> External	System Variable
<input type="radio"/> External	Trend Line
<input type="radio"/> External	User Input

Figure 42 - List of External Data Sources

5.1 Custom Template – The State

To make it easier to follow the process of building the custom template for the CMDB reporting, an overview of the end state is shown in Figure 43.

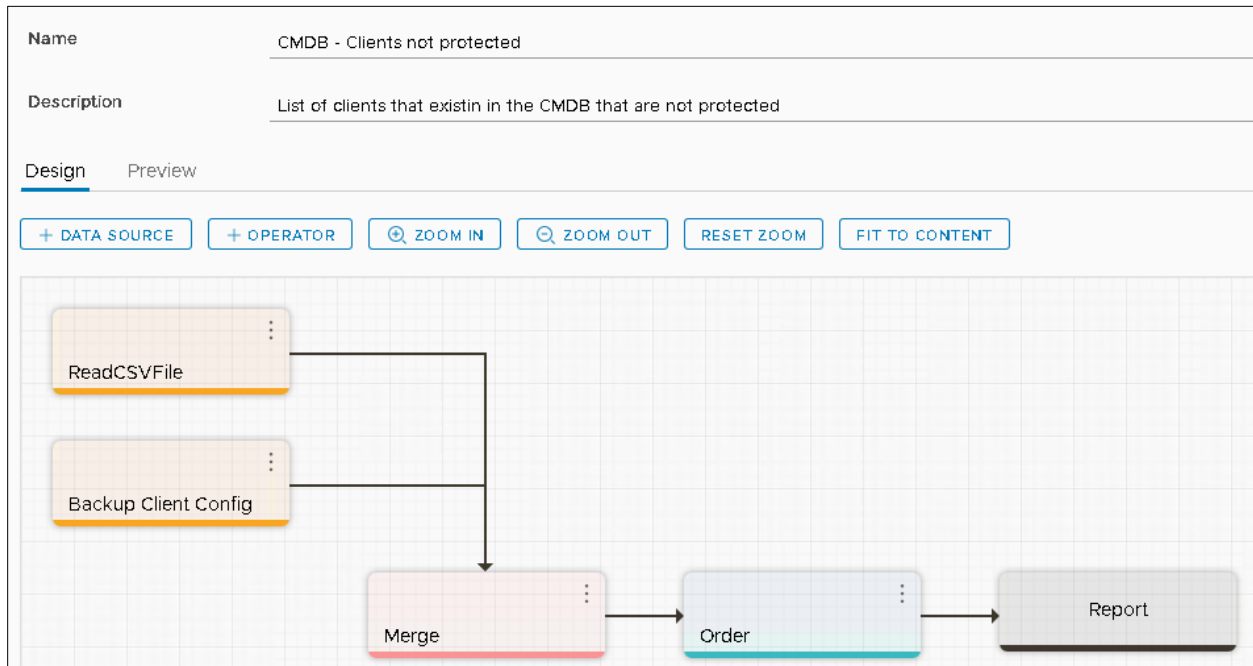


Figure 43 - CMDB - Clients Not Protected Custom Report Template

5.2 External Data Source

By no means is this an extensive or detailed external data source, however, it is designed to show how a custom report is built from scratch and how a data source is used. In this example, a simple CSV file called `hosts.csv` is located at `E:\CMDB\` on the DPA server and contains a list of servers as shown in Figure 44.

```
ad01.mlab.internal
ad02.mlab.internal
test01.mlab.internal
test02.mlab.internal
sql01.mlab.internal
sql02.mlab.internal
jb01.mlab.internal
jb02.mlab.internal
jb03.mlab.internal
```

Figure 44 - List of Clients Representing CMDB Assets

Before we can add an external data source, a custom report needs to be created. Expand *Reports* menu located on the left side of the DPA web UI, select *Report Templates* and select *Create Custom Template*. Provide a name and a description for the custom report. For this example, the custom template will be named *CMDB – Clients not protected*.

5.3 Adding the Data Source

There are two data sources that need to be added, the external CSV file and client backup data collected by DPA. To import data from an external source, select *+ Data Source*, filter the category for *External* and select the *ReadCSVFile*, which will be used to read the `hosts.csv` file and click on *Select*. Next, the backup client data source from DPA needs to be imported. To do this, select *+ Data Source*, filter the category for *Backup* and select the *Backup Client Config*, and click *Select*.

No changes need to be made to the Backup Client Config data source. For the ReadCSVFile data source, the parameters need to be defined as shown in Figure 45.

Fields	
<input checked="" type="checkbox"/>	Include All
<input checked="" type="checkbox"/>	Client
Label	Client
Cast	Text
Unit	
Null Value	
Key Field	<input checked="" type="checkbox"/>

Parameters	
Cast String	string
Field String	Client
Filename	E:\CMDB\hosts.csv
Key String	true
Separator	,

Figure 45 - ReadCSVFile Parameters

The Fields properties will be empty initially and will be populated automatically when the Parameters properties have been completed. For this example, the following details were entered;

- **Cast String** – the field format type; for client names this is set to *string*
- **Field String** – the name of the field for the data being imported. As this will be compared to clients in the DPA data in the Backup Client Config data source, the name of field needs to match and therefore *Client* is entered for this parameter.
- **Filename** – the location of the CSV file to import. As it is stored on the DPA server, a local drive path is provided as E:\CMDB\hosts.csv
- **Key String** – this is a true or false value for each comma separated value in the file. As there is only one value being imported a single *true* parameter is added.
- **Separator** – specifies the separator used in the CSV file. While there is only a single value, a parameter value of the comma is provided to complete all parameter values.

With the parameters entered, the fields properties will be populated. As there is a single field, it is also automatically set as the *Key Field*. Review the Backup Client Config data source properties and fields which have been marked as a Key Field will appear in a lighter colored blue. To confirm that the Client field in the Backup Client Config data source is also a Key Field, expand it as shown in Figure 45 to confirm that the Key Field is marked as Enabled.

5.4 Define the Operator

Now that there are data sources in place, they need to be merged to return only clients not protected as the result. Add a Merge operator by selecting **+ Operator**, filter the *Operator* list for merge, select *Merge* and click on *Select*. With the Merge operator shown on the custom template screen, it needs to be linked to both data sources and this is achieved by dragging data sources to the operator. Using the mouse, select a data source, then drag this over the operator and then release. This will link the elements together and an arrow should now appear

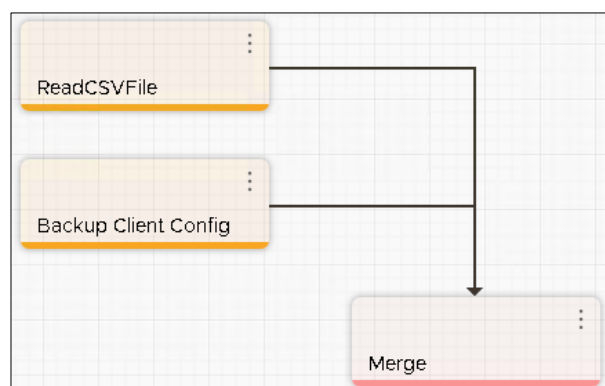


Figure 46 - Linked Data Sources to the Merge Operator

from the data source to the operator. Do the same for the other data source and the result of this linking should be the same as shown in Figure 46.

With the operator now linked to the data sources, a condition or a set of conditions need to be configured. The two data sources are merged based on the *Client* field and a condition needs to be defined where a client's corresponding *Server* field is empty. If the *Server* field is empty, this means that a backup copy of that client does not exist.

Select the Merge operator to display the properties, expand *Fields* properties to view the list of all available fields and ensure that at a minimum both *Client* and *Server* fields are selected. Expand the condition properties to set a condition. Select *Edit Condition* and select the *Server* field with the condition of *Is Missing* as shown in Figure 47. If the server is missing, this means that the client exists in the CMDB and not DPA.




Figure 47 - Merge Operator Properties for Fields and Condition

5.5 Order and Report

Before we can preview the report or results of the Merge condition, it needs to be linked to a report, achieved by dragging the Merge operator over the Report element for the connection to be established. But before this is done, is there a desire to have a field sorted in ascending or descending order by default? If so, an Order operator is needed and this is done by selecting *+ Operator*, filter the *Operator* list for order, select *Order* and click on *Select*. With the *Order* added to the custom template, it needs to be linked so that the Merge operator is linked to Order and then finally to the report element as shown in Figure 48.

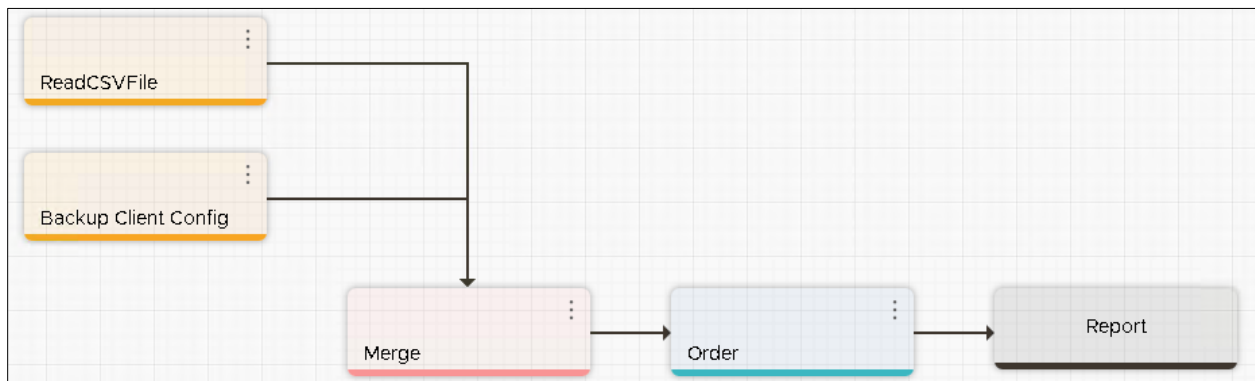


Figure 48 - Linking Merge to Order to Report

Selecting the *Order* operator, its properties are displayed and defining the order of a field is set in the parameter property. Select the field to which order will be applied and whether it is ascending or descending. Once done, select *Save* located in the lower right of the DPA web UI. When the data sources

and operators are linked to the report element, a preview of the results can be viewed. Below the description of the customer template, there is an option to *Preview* the results of the custom template. With the Preview option selected, define the *scope* and *time period* for the report and select the *Preview* button. For this example, the scope of *Backup Servers* under *Configuration | Servers* is selected, and a time period of Last Day is used. A sample report shown in Figure 49 lists clients which exist in the CMDB and where the server is missing from a data protection perspective.

<input type="checkbox"/>	Client
<input type="checkbox"/>	ad02.mlab.internal
<input type="checkbox"/>	test02.mlab.internal
<input type="checkbox"/>	jb03.mlab.internal

Figure 49 - Clients in CMDB not Protected Report Results

6 DPA Analysis Engine

The Analysis Engine used by DPA inspects captured data and compares this data against a set of rules defined by the administrator. This feature of DPA is ideal for capturing issues or anomalies, then automatically providing alerts using several available methods. For example, the analysis engine can be configured to capture failed backups and provide alerting via a Simple Network Management Protocol (SNMP) trap to a ticketing system for action. Or, alert the backup administrator via email if a volume of data processed for a client is x% higher than the previous number of backups, enabling capture of unusual growth or high changes in data volumes on clients.

6.1 Analysis Policy

An Analysis Policy is a collection of one or more rules that is assigned to an object or group. Rules contain the logic and triggers an alert if the condition is met. DPA compares collected data in its database to the conditions in the rule per event or on a schedule. By default, rules are compared based on event, while a schedule-based rule is run periodically to check whether a condition has been met. When a condition is met, actions are performed that can include sending an email, running a script, sending an SNMP trap or writing to the Windows event log.

6.2 Detecting Higher Rate of Change Backups – Example Analysis Policy

A common Analysis Policy that can be created relates to alerting a ticket application for any failed backups. This ticketing application can alert the Application Owner and/or Backup Administrator to address the cause of the failed backup before it is re-run. The example used in this section relates to the volume of changed data protected for a client, based on the average of the last X number of backups and the percentage of change. All stakeholders can benefit from this type of analysis alerting for clients with unusually high-volume rate of change. Typically, high volumes of changed data relate to migration projects, but it could also be the result of unusual behavior as a result of internal or external bad actors.

Creating an Analysis Policy

In this example, an Analysis Policy will be created, with a single Analysis Rule to detect backups larger than average and have it applied to all Servers. Expand *Policies* in the menu located on the left side of the DPA web UI and select *Analysis Policies*. Under *Analysis Policy Library*, select *+ Create Policy* and provide the policy a name as shown in Figure 50.

Create Analysis Policy

Name * Higher rate of change than average Enable Policy

Description Detect clients with 10% rate of change compared to the last 10 backups

Policy Based Actions

Email	No	EDIT POLICY-BASED ACTIONS
Script	No	
SNMP Trap	No	
Event Log	No	

Analysis Rules

[ADD/REMOVE RULES](#) [EDIT RULE PARAMETERS](#) [EDIT SCHEDULE](#) [EDIT ACTIONS](#)

Figure 50 - Creating a new Analysis Policy

Next, actions need to be set, which will be triggered if one or more Analysis Rules are met. At minimum, if no Policy Based Actions are enabled, an event will appear within the DPA Alerts section of the web UI.

A rule needs to be added to this Analysis Policy. Select *Add/Remove Rules* and a list of scheduled and event-based rules are presented. Use the *Filter* for the *Rule Name* to help narrow down a list of rules. For this example, the *Backup Larger than average for events number* is added as shown in Figure 51.

<input type="checkbox"/>	Rule Name	Type
<input type="checkbox"/>	Full Backup Larger than average for time window	Event Based
<input type="checkbox"/>	Backup Larger than average for time window	Event Based
<input checked="" type="checkbox"/>	Backup Larger than average for events number	Event Based
<input type="checkbox"/>	Full Backup Larger than average for events numb...	Event Based

Figure 51 - Selecting Rule for Analysis Policy

With the Rule added, the default parameters for the rule are displayed. For this example, the percentage of the deviation parameter of the rule is set to 10% from the default 50% as shown in Figure 52.

Rule Template Name	Parameter	Type	Value	Default Value
Backup Larger than average for events number	Number of backups	NUMERIC	10	10
Backup Larger than average for events number	Deviation	PERCENTAGE	10	50

Figure 52 - Setting the Rule Parameters

Depending on the application type and size, the deviation percentage could be much smaller. For example, a large file system containing in excess of 100TB may not ever reach 10% deviation and a percentage as small as 2% may be required. In that case, create a separate Analysis Policy to be applied to large file systems.

Applying the Analysis Policy

With the Analysis Policy created, it needs to be applied to a group or an individual object. If a Policy is applied directly to an object, it will take precedence over any policy applied inherited by the group. Selecting the *Applied Analysis Policies* option, all applied policies can be reviewed, modified and new policies applied. For this example, the newly created policy will be applied at the Servers level. This is achieved by expanding the groups and selecting *Servers* as shown in Figure 53.

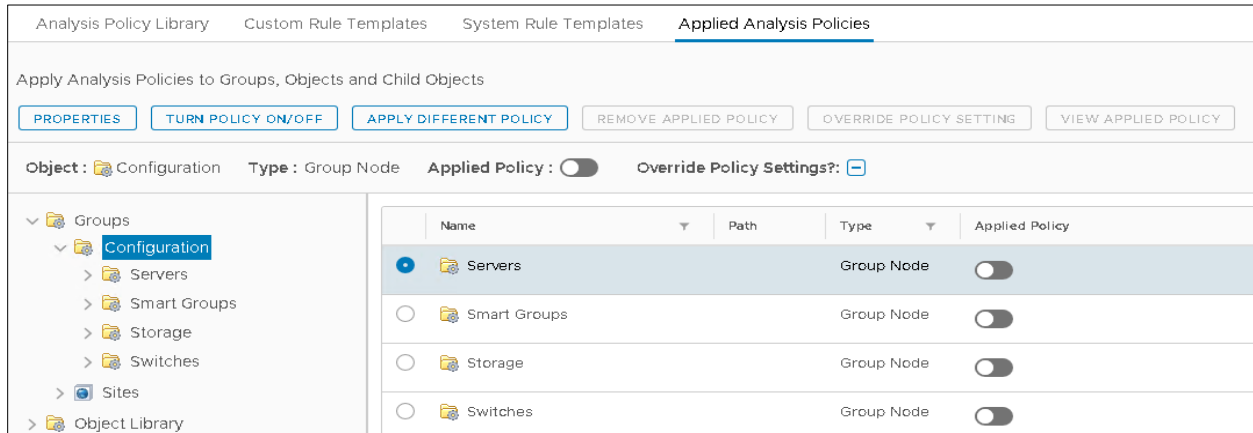


Figure 53 - Applying Analysis Policy to a Group

The Analysis Policy can now be applied to the group by selecting *Turn Policy On/Off*, scrolling through the list of available Analysis Policies, selecting one to be applied and save the changes. In the example shown in Figure 54, the Analysis Policy of *Higher rate of change than average* has been applied to the *Servers* group.

Name	Path	Type	Applied Policy	Source (Where Policy is Applied)	Override Policy Settings?
<input checked="" type="checkbox"/> Servers		Group Node	<input checked="" type="checkbox"/> Higher rate of change than average	Configuration > Servers	<input type="checkbox"/>

Figure 54 - Applied Analysis Policy

To test that the Analysis Policy was applied successfully, a large volume of data was saved to the test01 client, a backup conducted and an alert was seen in DPA as shown in Figure 55.

Severity	State	Message	Object	Child Object	Component	Policy	Count	Last Updated
New		Backup Larger than average for events number	nw01.miab.internal	test01.miab.internal	Backupjob C:	Higher rate of change than average	1	01/22/2021 1:50 PM

Figure 55 - DPA Alert for Analysis Policy Test

This example illustrates only a single Analysis rule that can be applied. To decide which Analysis rule should be implemented, determine what conditions would result in an elevated risk or undesirable condition. Looking for data protection gaps, unusual activity or behavior by using the Analysis Policy will help improve the overall health of the data protection environment.

6.3 Protection and Chargeback Policies

In addition to Analysis Policies, DPA provides two additional policies. Protection Policies relate to expectation of backup and replication operations within the environment and are ideal if reporting against agreed SLAs have been defined. Meanwhile, Chargeback Policy enables the organization to determine

costs in providing backup and recovery operations. Examples of implementing each of these is not provided within this article. Reviewing the DPA administration guide⁹ is recommended.

7 Summary

Understanding the state of the data protection environment, application recovery points, and unprotected data is key in determining the risk that the organization is exposed to. With Data Protection Advisor provides the organization, its administrators, managers and risk/security teams a centralized view of the data protection environment. With automated reporting, alerting, analysis and customization, identifying gaps and unusual behaviors is easily achieved.

With Data Protection Advisor deployed, insights are gleaned, gaps in data protection are revealed and identified risks can then be eliminated.

8 References

1. <https://www.delltechnologies.com/en-au/data-protection/data-protection-suite/data-protection-advisor.htm>
2. <https://www.delltechnologies.com/en-au/data-protection/data-protection-suite/avamar-data-protection-software.htm>
3. <https://www.delltechnologies.com/en-au/data-protection/data-protection-suite/networker-data-protection-software.htm>
4. <https://www.delltechnologies.com/en-au/data-protection/powerprotect-data-manager.htm>
5. <https://www.delltechnologies.com/en-au/data-protection/powerprotect-backup-appliances.htm>
6. <https://www.commvault.com/backup-solutions>
7. <https://www.veritas.com/en/au/protection/netbackup>
8. [https://dl.dell.com/content/docu52326_Data_Protection_Advisor_\(DPA\)_6.3_and_later_Deployment_Architecture_Guide.pdf?language=en_US](https://dl.dell.com/content/docu52326_Data_Protection_Advisor_(DPA)_6.3_and_later_Deployment_Architecture_Guide.pdf?language=en_US)
9. https://dl.dell.com/content/docu101060_Data_Protection_Advisor_19.4_Installation_and_Administration_Guide.pdf?language=en_US
10. https://dl.dell.com/content/docu101058_Data_Protection_Advisor_19.4_Report_Reference_Guide.pdf?language=en_US
11. https://dl.dell.com/content/docu101054_Data_Protection_Advisor_19.4_Custom_Report_Guide.pdf?language=en_US
12. https://dl.dell.com/content/docu101056_Data_Protection_Advisor_19.4_Product_Guide.pdf?language=en_US

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.