

Using EMC Data Protection Advisor with EMC Data Domain Deduplication Storage Systems

Applied Technology

Abstract

EMC® Data Protection Advisor (DPA) provides a comprehensive set of features to analyze data protection operations to ensure that your data is protected and recoverable. Analyzing the backup applications, supporting infrastructure, and target storage, DPA can capture issues so that they can be addressed before a failure. This white paper outlines how Data Protection Advisor operates in conjunction with EMC Data Domain® deduplication storage systems.

June 2010

Copyright © 2010 EMC Corporation. All rights reserved.

EMC 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.” EMC CORPORATION 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 EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number h7115.1

Table of Contents

Executive summary	4
Introduction	4
Audience	4
EMC Data Domain	4
What can DPA provide Data Domain customers?	5
Managing complexity and access	5
Data Domain, VMware, and DPA.....	8
Sample reports.....	8
Resource utilization	8
Deduplication ratio.....	9
Network interface utilization	11
Processor utilization and status	12
RTO and RPO.....	13
Cleaning	13
Status reports.....	14
Data Domain status control panel	14
Disk status.....	14
Fan status.....	15
Thermometer status	16
PSU status.....	16
Battery status	16
Network interface status.....	17
Filesystem status.....	17
VTL status	17
Configuration reports	18
Host and disk configuration	18
Battery configuration	19
Network interface configuration.....	19
VTL configuration	20
File system configuration.....	20
Performance reports	20
Virtual drive performance	20
Disk performance	21
Fileserver performance	21
Network interface status.....	22
Change management reports	22
Multidomain analyses and alerting	24
Conclusion	25
References	25

Executive summary

EMC® Data Protection Advisor (DPA) collects, monitors, analyzes, and reports on information from our customer's entire data protection infrastructure, providing a unified data protection management window across their entire backup and EMC replication investment, accelerating access to information, saving time and money, allowing faster decisions, and improving data protection.

Through support for heterogeneous backup infrastructures including EMC Backup Solutions and support for EMC Replication Solutions, DPA reduces the cost and risk to manage a data protection environment, enabling our customers to get more from their existing investments and increase efficiency.

The DPA graphical user interface (GUI) can present this information in a familiar manner by arranging assets in common sense groups or views such as business unit or geography, greatly enhancing readability and giving the user the ability to perform advanced reporting, troubleshooting, performance management, and capacity planning operations.

EMC Data Domain® support in DPA gathers information about the configuration, status, and performance of Data Domain components. Data Protection Advisor uses SNMP to gather this data from the Data Domain management information base (MIB).

Introduction

DPA monitors a very wide range of components/assets and provides a comprehensive range of reports for each. This white paper seeks to describe DPA and its interaction with Data Domain deduplication storage systems. Because DPA supports a wide range of products and technology, it often uses general terminology. Therefore in this document and within the DPA console, Data Domain Global Compression technology is referred to as "deduplication."

Audience

This white paper is intended for use by backup administrators and operations managers to understand the benefits of using EMC Data Protection Advisor in conjunction with EMC Data Domain deduplication storage systems. As a target storage device used by backup systems to store backup copies of critical data, it is helpful to understand aspects such as system performance, capacity, and availability. DPA maintains a historical record of the systems operation, providing consistent monitoring, alerting, and reporting. Data Domain support in conjunction with data from the backup environment enables a comprehensive view of the operations and health of the backup systems.

EMC Data Domain

Data Domain high-speed, inline deduplication technology has proven to be the key enabler that has helped customers reduce or eliminate the use of tape for backup and recovery; experience longer onsite retention for backup, archive and reference data; and leverage network-efficient replication capabilities for disaster recovery (DR). EMC Data Domain systems are integrated appliances that can support small remote office appliances up to large enterprise data centers.

This white paper focuses on the DD140, DD600 series, and DD880 appliances running Data Domain Operating System (DDOS) versions 4.4.2 to 4.8.

EMC Data Domain Enterprise Manager

All EMC Data Domain deduplication storage systems ship with EMC Data Domain Enterprise Manager, a simple Web-based rich internet application for managing Data Domain systems. With DD Enterprise Manager, you can monitor multiple systems from a central workstation and manage all critical aspects of Data Domain systems, such as the filesystem, replication, EMC Data Domain Boost, NFS, CIFS, and virtual tape libraries, in a simple manner using the GUI. DD Enterprise Manager dashboards provide a

high-level overview of health status and then let you drill down into areas of interest. The operational simplicity of the dashboards helps to reduce administrative costs.

What can DPA provide Data Domain customers?

EMC DPA is more than a reporting tool. It provides the analytical capabilities that provide insight into the entire backup environment from familiar viewpoints. DPA gathers data from backup applications, backup servers, backup clients, network switches, SAN switches, tape libraries, and target storage. This paper focuses on Data Domain environments, specifically the following areas:

- DPA gathers information about the Data Domain system configuration, status, and performance using SNMP to gather data from the Data Domain MIB. The flexibility of the DPA data collection allows EMC to constantly improve Data Domain reporting capabilities in DPA as the Data Domain MIB expands.
- The DPA GUI presents Data Domain system information in a familiar manner by arranging assets in common sense groups or views such as business unit or geography, greatly enhancing readability and giving the user the ability to perform advanced reporting, troubleshooting, performance management, and capacity planning operations.
- Data Domain deduplication reduces the amount of disk storage needed to retain and protect data by ratios of 10 to 30 times, making disk a cost-effective alternative to tape. Data Domain systems are qualified with all leading enterprise backup software and archiving applications and easily integrate into the existing storage infrastructure without change for either data center or distributed office data protection. DPA can monitor both the backup application and the Data Domain storage, providing visibility into both worlds.
- DD Enterprise Manager can monitor and manage up to 20 Data Domain systems on a per-system basis. However, DPA can add value to DD Enterprise Manager in the following ways:
 - Provides the ability to aggregate monitoring information from more than 20 Data Domain systems in a single interface.
 - Provides the ability to tie in the resource utilization from the perspective of multiple backup applications and multiple Data Domain systems from a single interface.
 - Provides visibility into the status and performance of Data Domain replication
- Data Domain systems support replication fan-in from up to 270 remote sites with Global Deduplication across all remote sites. DPA allows the user to gain a global view of all the remote and local Data Domain systems on one screen.
- DPA can be useful for documenting system settings and configurations in multisite deployments, for both Data Domain systems as well as the backup application, OS versions, network infrastructure, and any other backup-related components.

DPA can track the system settings and through its change management function help maintain standardization of multiple systems and reduce time to troubleshoot problems.
- DPA aids in ongoing capacity planning, through regular data collection, and ongoing measurement. DPA can forecast capacity needs, and even alert on a system running out of capacity at some point in the future by leveraging historical data.

Managing complexity and access

Data Domain systems support all leading backup applications and can be used for backup and recovery, protection of enterprise applications (Oracle, Microsoft Exchange, VMware, and others), archiving, and online reference storage. Since these systems can be deployed in many configurations and environmental combinations, it can be difficult to monitor the relationships of all components without knowledge of these components and suitable permissions to all of the different interfaces.

DPA allows the user to view Data Domain system functions and operations in the context of the wider environment, that is, backup application and host, and FC/IP switches and ports.

The multiple layers of a Data Domain user's environment and the resource utilization within it can be difficult to visualize.

The backup application, operating system, virtual devices, physical devices, and underlying SAN / SCSI can have different architectures, and the utilization of these various layers and their inter-relationships, can be difficult to reconcile.

DPA allows the user to create overviews/control panels that are essentially reports with multiple constituent report windows embedded within them. These have a very useful application in Data Domain environments as one could, for instance, graphically show on one screen the utilization of:

- The backup application (for instance, IBM Tivoli Storage Manager)
- TSM management processes and their relationships with storage pools
- TSM storage pools from the TSM perspective
- The underlying disk volumes from the OS perspective
- Data Domain system utilization

The ability to view Data Domain system activity in the same window as TSM server-side processes is key.

The configuration tree in Figure 1 illustrates the wealth of assets that DPA supports.

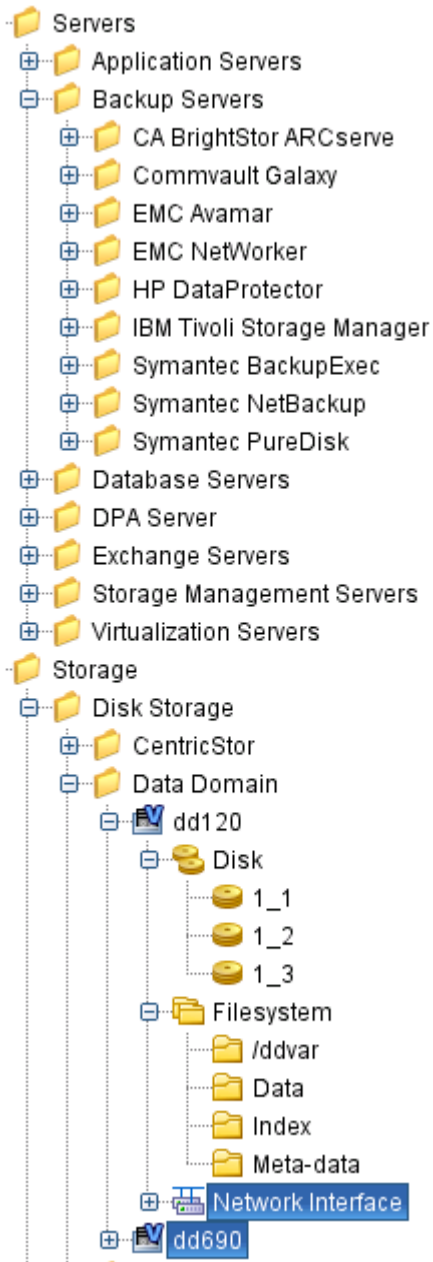


Figure 1. DPA supports many assets

Benefit

DPA provides the user with a common interface for the wide range of Data Domain system configurations in numerous client environments.

Data Domain, VMware, and DPA

The benefits of DPA in complex environments are even more evident in VMware environments. Due to the relative ease with which VMware environments can be expanded, “virtual machine (VM) sprawl” creates increased pressure on the backup environment. By increasing the number of systems sharing a set of physical resources, backups can be constrained by the backup of other VMs on the same host. Due to the redundant nature of VMware backups, these environments are ideal for deduplication technology, usually achieving deduplication ratios above those of conventionally hosted systems.

DPA provides excellent insight into VMware and Data Domain environments, allowing the user to identify bottlenecks, load balance, and view the environment from differing perspectives in one screen.

By viewing Backup Application Schedules along with Resource Utilization reports of multiple VMs and the ESX host, a user can establish if the backup workload is having an undue effect on the physical host. If the server is overloaded from the backup workload, backup schedules can be adjusted accordingly to balance the environment.

If the user is sending disk image backup to a Data Domain system, DPA can be used to ease the process of scheduling the required VM shutdowns. DPA can also monitor the synchronization effects of the VM shutdown scripts and backup activity, keeping downtime or performance impact to a minimum.

Sample reports

The following sections outline some typical reports with a description of the data presented and possible interpretations.

Resource utilization

DPA provides reports that show the utilization of CPU, memory, and network activity. The Resource Utilization control panel presents this information in a single report.

From the screenshot below, it can be seen that there could be some cause for concern as the memory is so high (>80%) and processor utilization is so low this might be worth investigating due to the CPU-centric architecture of Data Domain systems.

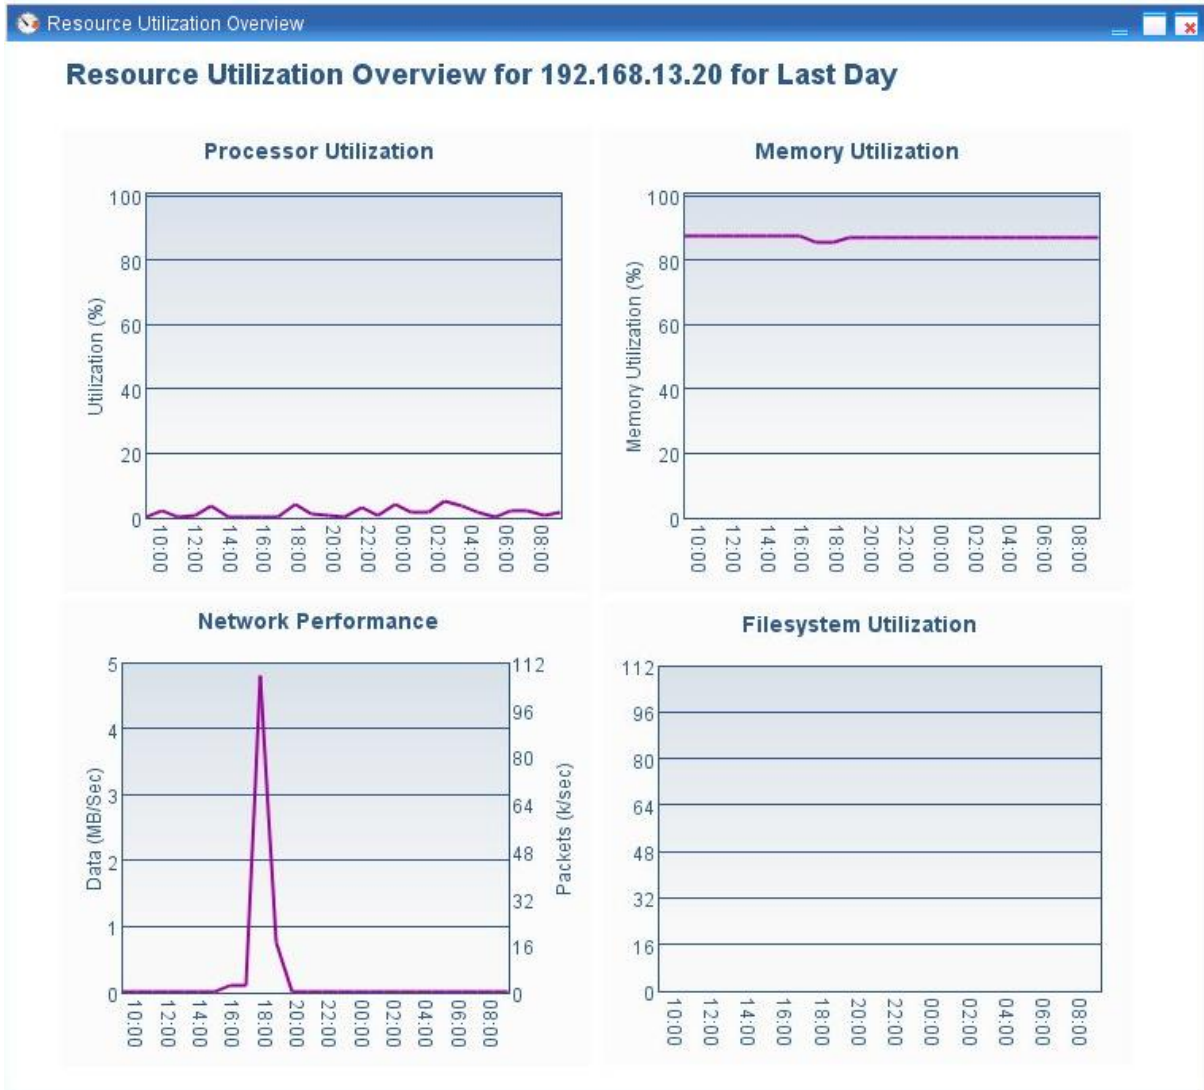


Figure 2. Resource Utilization Overview report

Note: Filesystem utilization is not seen in this report but can be viewed from the perspective of the attached host.

Benefit

DPA allows the user to view the vital signs of the Data Domain system in one window, allowing insight into any imbalance, such as the imbalance between the CPU, network, and memory utilization shown in Figure 2.

Deduplication ratio

The Deduplication Ratio report plots the dedupe ratio that the Data Domain system is achieving over time. This shows the storage savings gained from using Data Domain deduplication storage systems versus traditional backup storage like tape or standard disk.

In a normal environment, one would expect the dedupe ratio to improve over time and eventually plateau. If the deduplication ratio has declined, there may be reason to investigate any changes in the backup data types over the time period and revise provisioning and capacity planning estimates.

In addition, the deduplication ratio can be viewed and compared across the enterprise to identify lower-performing configurations, systems, or locations.




Figure 3. Deduplication Ratio report

Benefit

DPA allows users to track deduplication ratios of a Data Domain system over time.

Network interface utilization

The Network Interface Utilization Summary control panel provides insight into the performance of network interfaces on the Data Domain system. This control panel shows the maximum and average utilization over a reporting period, along with the utilization on a per-interface basis.

A well-balanced system/environment will likely have a similar utilization across all components. Figure 4 shows that two of the six available network ports on this Data Domain system are bearing the majority of the load, which might be worth investigating. This imbalance could mean a variety of things. For example backups may be occurring from different systems on different networks at different times. Other potential problems include hardware failures, unbalanced workload, a backup host is offline, other network activity impacting backups, backups configured for only two networks, and other issues

Although we could spend lots of time analyzing an unbalanced system, it may not be a serious problem. The intent is that DPA provides a starting point to identify possible problems. The “Processor utilization and status” section next provides more information.



Network Interface Utilization Summary for dd120 f...




Figure 4. Network Interface Utilization Summary report

Benefit

DPA allows the user to identify potential load-balancing and functional issues across network interfaces.

Processor utilization and status

The Processor Utilization report shows CPU utilization over time and allows the user to identify possible resource constraints on the Data Domain system. The Processor Status report shows how many CPUs are online and/or offline.




Figure 5. Processor Utilization and Processor Status reports

Benefit

DPA provides CPU utilization detail that allows processor bottlenecks to be identified and addressed.

RTO and RPO

Data Domain systems allow great gains in recovery times by eliminating the wait times required for backup media availability, particularly the time required to mount physical tape. DPA can accurately calculate and report on the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) based on historical data, estimated mount times, and data transfer rates. By setting the Off-Line Data Overhead field (also known as tape recall) to zero we can account just for the time to restore. These accurate predictions are important in ensuring SLAs can be met and that the benefits of deploying a Data Domain system in the environment are evident to the end customer.

Exposure

Restore Time Objective: 10 Minutes

Recovery Point Objective: 1 Hours

Use Historical Restore Performance if Available

Off-Line Data Overhead: 0 Seconds

Assert restore will take: 100 % of time taken to backup the data

Figure 6. DPA reports on RTO and RPO

Cleaning

Data Domain systems require cleaning or “garbage collection” operations to be run regularly, allowing space reclamation to help maintain optimum performance. DPA can be used to ensure that the cleaning processes run within periods of low backup activity.




Figure 7. TSM Server Processes report

Status reports

Status reports show the state of the Data Domain system components at a glance, allowing adverse conditions to be readily identified and rapidly addressed. Analyses can be configured and assigned to alert the user to degradations in the environment that may have adverse effects.

Data Domain status control panel

The condition of vital Data Domain system components – including the disk, network, PSUs, fans, and thermometer – can be viewed in an intuitive manner, providing insight into the health of individual systems and the environment as a whole.

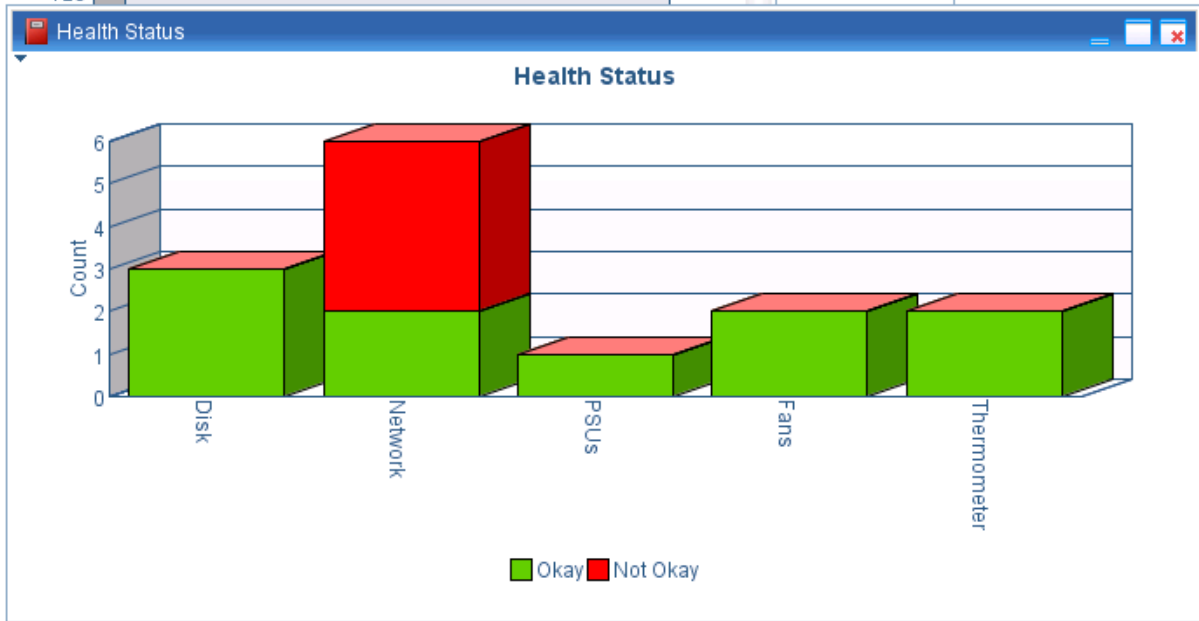


Figure 8. Health Status report

The Network status bar in the figure (second from the left) shows two network interfaces are online and working (in green), while four of the six available are not. This relates back to Figure 4, where two network interfaces are taking all of the workload.

Disk status

This Disk Status report provides detailed information about the underlying disk, allowing the user to readily identify those disks with high operating temperatures, error counts, and other issues, and helping to anticipate failures and inform when they result in deterioration of availability.

The Disk Status report and other reports are shown in the following pages.

Disk Status for 10.64.205.200 for Last Hour											
Hostname	Device	Size (GB)	State	Use	Errors	Name	Temperature (Celsius)	Group Name	Status	Recon Percentage (%)	Resynch Percentage (%)
10.64.205.200	1	372.0	Online		0		35.0	dg0	Ok	0	0
10.64.205.200	10	372.0	Online		0		33.0		Ok	0	0
10.64.205.200	11	372.0	Online		0		32.0		Ok	0	0
10.64.205.200	12	372.0	Online		2		32.0		Ok	0	0
10.64.205.200	13	372.0	Online		0		31.0		Ok	0	0
10.64.205.200	14	372.0	Online		0		33.0		Ok	0	0
10.64.205.200	15	372.0	Online		1		31.0		Ok	0	0
10.64.205.200	2	372.0	Online		0		38.0	dg0	Ok	0	0
10.64.205.200	3	372.0	Online		2		36.0	dg0	Ok	0	0
10.64.205.200	4	372.0	Online		0		35.0	dg0	Ok	0	0
10.64.205.200	5	372.0	Online		0		31.0	dg0	Ok	0	0
10.64.205.200	6	372.0	Online		0		36.0	dg0	Ok	0	0
10.64.205.200	7	372.0	Online		0		34.0	dg0	Ok	0	0
10.64.205.200	8	372.0	Online		0		32.0		Ok	0	0
10.64.205.200	9	372.0	Online		0		30.0		Ok	0	0

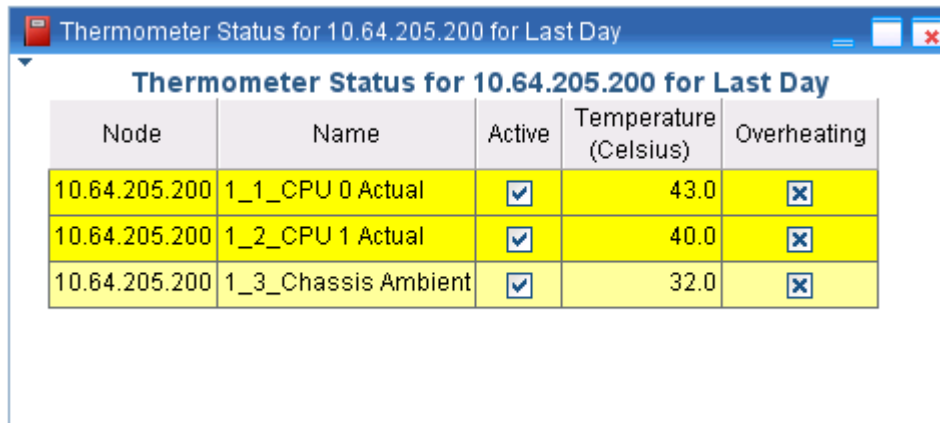
Figure 9. Disk Status report

Fan status

Fan Status for 10.64.205.200 for Last Day					
Node	Name	Active	Level	Status	
10.64.205.200	1_1_Crossbar fan #1	<input checked="" type="checkbox"/>	normal		
10.64.205.200	1_2_Crossbar fan #2	<input checked="" type="checkbox"/>	normal		
10.64.205.200	1_3_Crossbar fan #3	<input checked="" type="checkbox"/>	normal		
10.64.205.200	1_4_Crossbar fan #4	<input checked="" type="checkbox"/>	normal		
10.64.205.200	1_5_Rear fan #1	<input checked="" type="checkbox"/>	normal		
10.64.205.200	1_6_Rear fan #2	<input checked="" type="checkbox"/>	normal		

Figure 10. Fan Status report

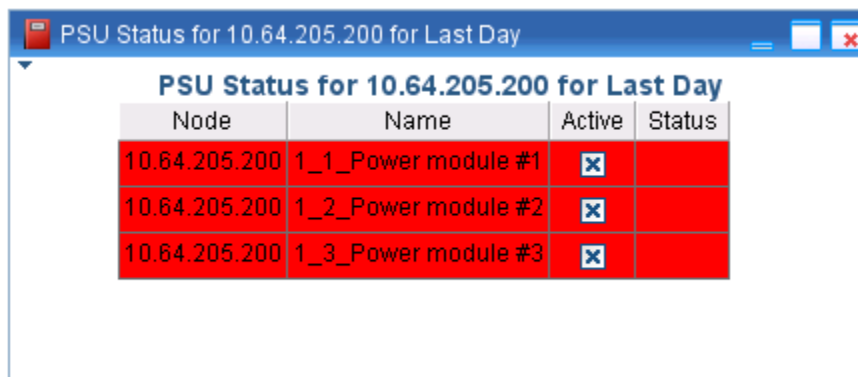
Thermometer status



Node	Name	Active	Temperature (Celsius)	Overheating
10.64.205.200	1_1_CPU 0 Actual	<input checked="" type="checkbox"/>	43.0	<input checked="" type="checkbox"/>
10.64.205.200	1_2_CPU 1 Actual	<input checked="" type="checkbox"/>	40.0	<input checked="" type="checkbox"/>
10.64.205.200	1_3_Chassis Ambient	<input checked="" type="checkbox"/>	32.0	<input checked="" type="checkbox"/>

Figure 11. Thermometer Status report

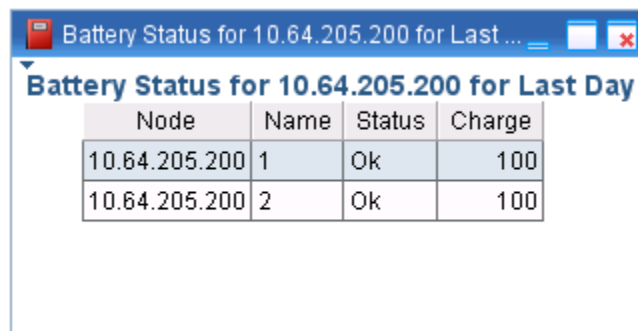
PSU status



Node	Name	Active	Status
10.64.205.200	1_1_Power module #1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
10.64.205.200	1_2_Power module #2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
10.64.205.200	1_3_Power module #3	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Figure 12. PSU Status report

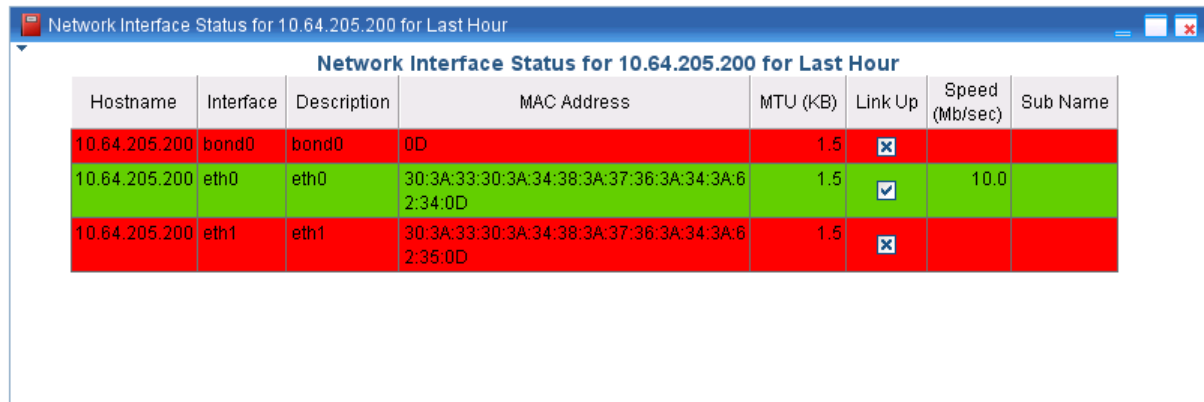
Battery status



Node	Name	Status	Charge
10.64.205.200	1	Ok	100
10.64.205.200	2	Ok	100

Figure 13. Battery Status report

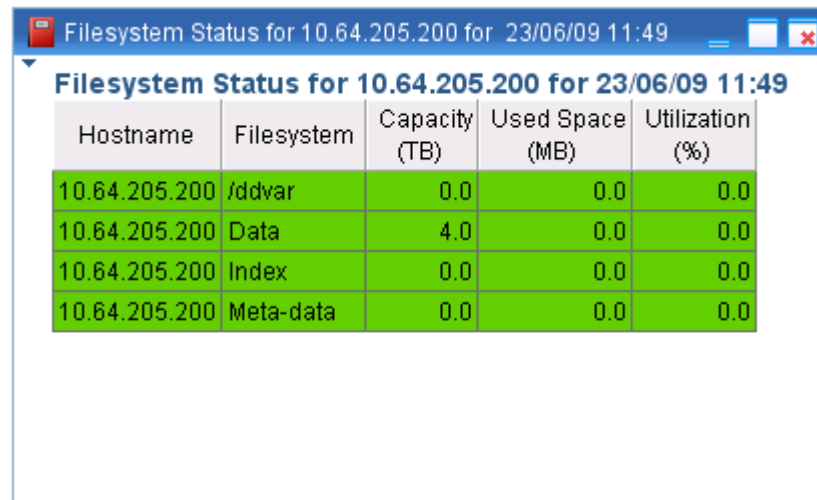
Network interface status



Hostname	Interface	Description	MAC Address	MTU (KB)	Link Up	Speed (Mb/sec)	Sub Name
10.64.205.200	bond0	bond0	0D	1.5	X		
10.64.205.200	eth0	eth0	30:3A:33:30:3A:34:38:3A:37:36:3A:34:3A:6:2:34:0D	1.5	✓	10.0	
10.64.205.200	eth1	eth1	30:3A:33:30:3A:34:38:3A:37:36:3A:34:3A:6:2:35:0D	1.5	X		

Figure 14. Network Interface Status report

Filesystem status



Hostname	Filesystem	Capacity (TB)	Used Space (MB)	Utilization (%)
10.64.205.200	/ddvar	0.0	0.0	0.0
10.64.205.200	Data	4.0	0.0	0.0
10.64.205.200	Index	0.0	0.0	0.0
10.64.205.200	Meta-data	0.0	0.0	0.0

Figure 15. Filesystem Status report

VTL status

This VTL Status report provides detailed information listing of the VTL system status, similar to the Filesystem status report shown above. The VTL Status report is supported in DPA version 5.5 SP1 if you are running Data Domain 4.7.

Benefit

With all of the various status reports available, DPA provides useful insight into the status of the environment and can present this information in an automatically refreshed format that is ideal for bridges/operations rooms.

Configuration reports

The configuration reports display configuration detail on a wide range of components of the Data Domain system. These reports might be useful to identify differences in configurations when investigating differences in function across the enterprise. Any configuration changes can be automatically detected and alerted to the user (see the “Change management reports” section). Likewise, the configuration report can be used to easily identify upgrade candidates based on details like firmware revision.

Host and disk configuration

These reports provide details of the Data Domain system itself, including make, model, version, and host ID as seen in the following screenshots.

Host Configuration for 10.64.205.200 for Last Day

Hostname	Sub Name	Vendor	Product	OS Class	Version	iSCSI Name	Host Id
10.64.205.200		DataDomain	DataDomain	Appliance	3.2.2.4-24897		acsdd460.acs-inc.com

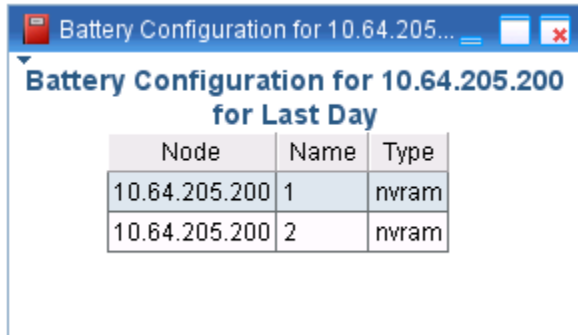
Disk Configuration for 10.64.205.200 for Last Hour

Hostname	Device	Manufacturer	Model	Serial Number	Firmware Revision	Size (GB)	Enclosure
10.64.205.200	1		HDS724040KLSA80	KRFS11RAH2D30C	KFAOAC6A	372.0	
10.64.205.200	10		HDS724040KLSA80	KRFS11RAH2GH1C	KFAOAC6A	372.0	
10.64.205.200	11		HDS724040KLSA80	KRFS11RAH2ZLYC	KFAOAC6A	372.0	
10.64.205.200	12		HDS724040KLSA80	KRFS11RAH2P4RC	KFAOAC6A	372.0	
10.64.205.200	13		HDS724040KLSA80	KRFS11RAH2002C	KFAOAC6A	372.0	
10.64.205.200	14		HDS724040KLSA80	KRFS11RAH2YK9C	KFAOAC6A	372.0	
10.64.205.200	15		HDS724040KLSA80	KRFS11RAH228AC	KFAOAC6A	372.0	
10.64.205.200	2		HDS724040KLSA80	KRFS11RAH2KG6C	KFAOAC6A	372.0	
10.64.205.200	3		HDS724040KLSA80	KRFS11RAH2JMXC	KFAOAC6A	372.0	
10.64.205.200	4		HDS724040KLSA80	KRFS11RAH2ZM8C	KFAOAC6A	372.0	
10.64.205.200	5		HDS724040KLSA80	KRFS11RAH300AC	KFAOAC6A	372.0	
10.64.205.200	6		HDS724040KLSA80	KRFS11RAH2KT8C	KFAOAC6A	372.0	
10.64.205.200	7		HDS724040KLSA80	KRFS11RAH2RW4C	KFAOAC6A	372.0	
10.64.205.200	8		HDS724040KLSA80	KRFS11RAH76L4T	KFAOAC6A	372.0	
10.64.205.200	9		HDS724040KLSA80	KRFS11RAH2W9RC	KFAOAC6A	372.0	

Figure 16. Host Configuration and Disk Configuration reports

In addition to these reports, configuration reports provide details on the VTLs, disks, CPUs, memory, network interfaces, and Fibre Channel interfaces on the Data Domain systems. Other reports are shown next.

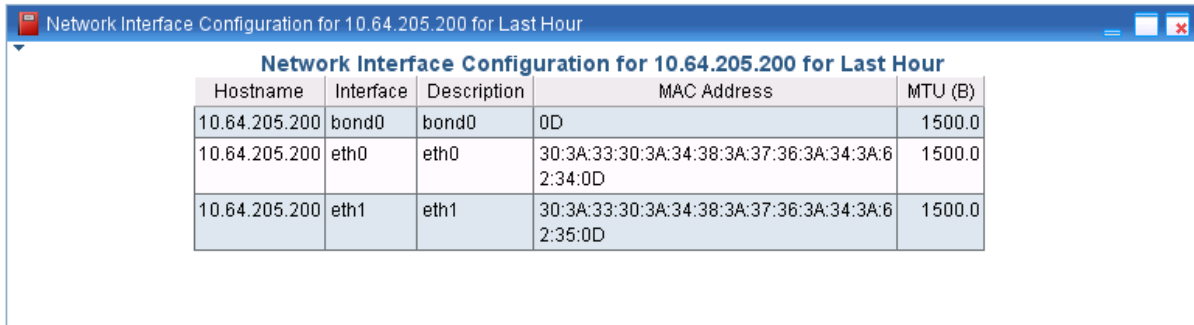
Battery configuration



Node	Name	Type
10.64.205.200	1	nvram
10.64.205.200	2	nvram

Figure 17. Battery Configuration report

Network interface configuration



Hostname	Interface	Description	MAC Address	MTU (B)
10.64.205.200	bond0	bond0	0D	1500.0
10.64.205.200	eth0	eth0	30:3A:33:30:3A:34:38:3A:37:36:3A:34:3A:62:34:0D	1500.0
10.64.205.200	eth1	eth1	30:3A:33:30:3A:34:38:3A:37:36:3A:34:3A:62:35:0D	1500.0

Figure 18. Network Interface Configuration report

VTL configuration

The VTL Configuration report is supported with DPA version 5.5 SP1 and Data Domain Operating System version 4.7.

File system configuration

Filesystem Configuration for 10.64.205.200 for Last Day														
Hostname	Filesystem	Device	Type	Guarantee	Capacity (TB)	Save Volume	Convert UCode	Create UCode	SVO Enable	SVO Allow RMAN	SVO Checksum	SVO Reject Errors	Extent	FS Size Fixed
10.64.205.200	/ddvar				0.0		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
10.64.205.200	Data				4.0		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
10.64.205.200	Index				0.0		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
10.64.205.200	Meta-data				0.0		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Figure 19. Filesystem Configuration report

Benefit

DPA provides one-stop insight into the configuration settings of the entire Data Domain environment.

Performance reports

DPA provides a wealth of performance reports for Data Domain systems and the relationships with the boarder environment.

Virtual drive performance

With dedicated mapping the user could monitor the FC port of the backup host (storage node/media server) in order to establish the performance of the virtual tape drive to which it is attached.

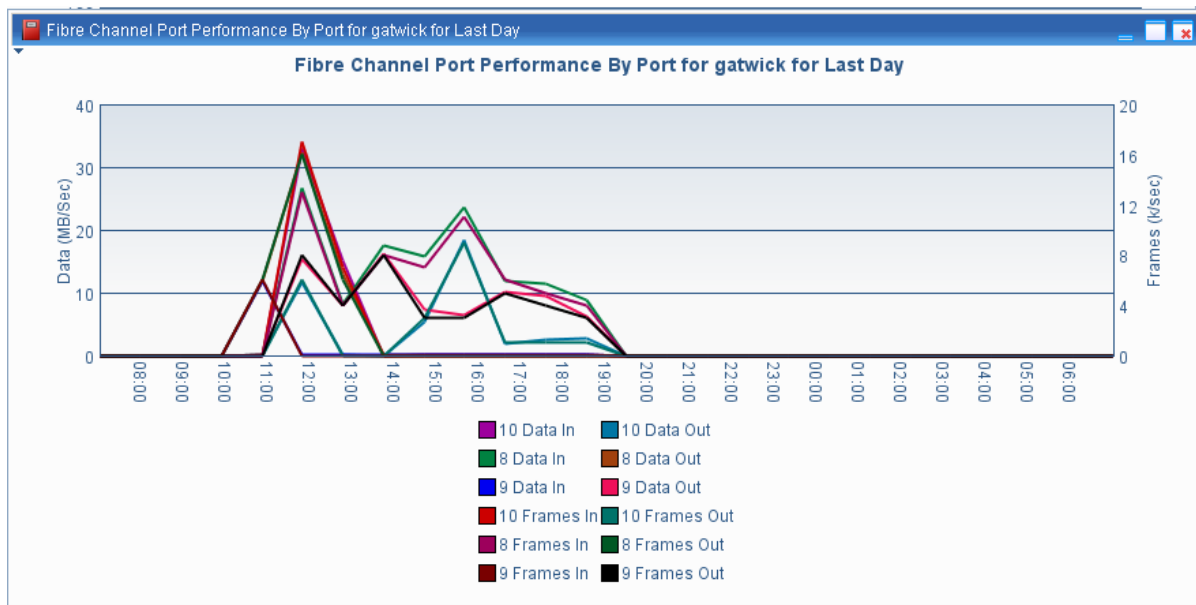


Figure 20. Fibre Channel Port Performance report

Benefit

DPA can use its reporting of FC switches and backup host (storage node) FC HBAs to establish virtual drive performance.

Disk performance




Figure 21. Disk Performance report

Fileserver performance




Figure 22. Fileserver Performance report

Network interface status

This report shows the status of ports within a Data Domain system. Note the fact that a number of the ports are not active, corresponding with the Health Status report earlier. This could mean that these network cards were turned off or that there was a failure within the system.

Hostname	Interface	Description	MAC Address	MTU (KB)	Link Up	Speed (Mb/sec)
dd120	eth0	eth0	00:15:17:1F:EF:EC	1.5	<input checked="" type="checkbox"/>	1000.0
dd120	eth1	eth1	00:15:17:1F:EF:ED	1.5	<input checked="" type="checkbox"/>	1000.0
dd120	veth0	veth0		1.5	<input type="checkbox"/>	
dd120	veth1	veth1		1.5	<input type="checkbox"/>	
dd120	veth2	veth2		1.5	<input type="checkbox"/>	
dd120	veth3	veth3		1.5	<input type="checkbox"/>	

Figure 23. Network Interface Status report

Benefit

DPA gives the user visibility of port status and performance from the perspective of the Data Domain system and also potentially the switch and HBA ports in the wider SAN.

Change management reports

Change management reports reflect any changes in the configuration of Data Domain appliances' VTLs, tape drives, and shares on the Data Domain system. (DD VTL is supported in DPA 5.5 SP1 if running Data Domain Operating System version 4.7.)

These allow the user to identify when any of these assets has been added, modified, or deleted. A real-time and historical record of such changes is important to maintaining system configuration standards, tracking system issues back to changes, and ensuring compliance with regulatory and local strictures.

When any aspect of the data protection environment is misaligned, it may affect daily backup job completion or performance of the system.

The complex layered configuration architecture of environments can make standardization and change management difficult. DPA can alert the user to any changes in configuration or policy across all Data Domain systems in a worldwide enterprise, allowing gold standards to be enforced — reducing diversity, complexity, and cost.

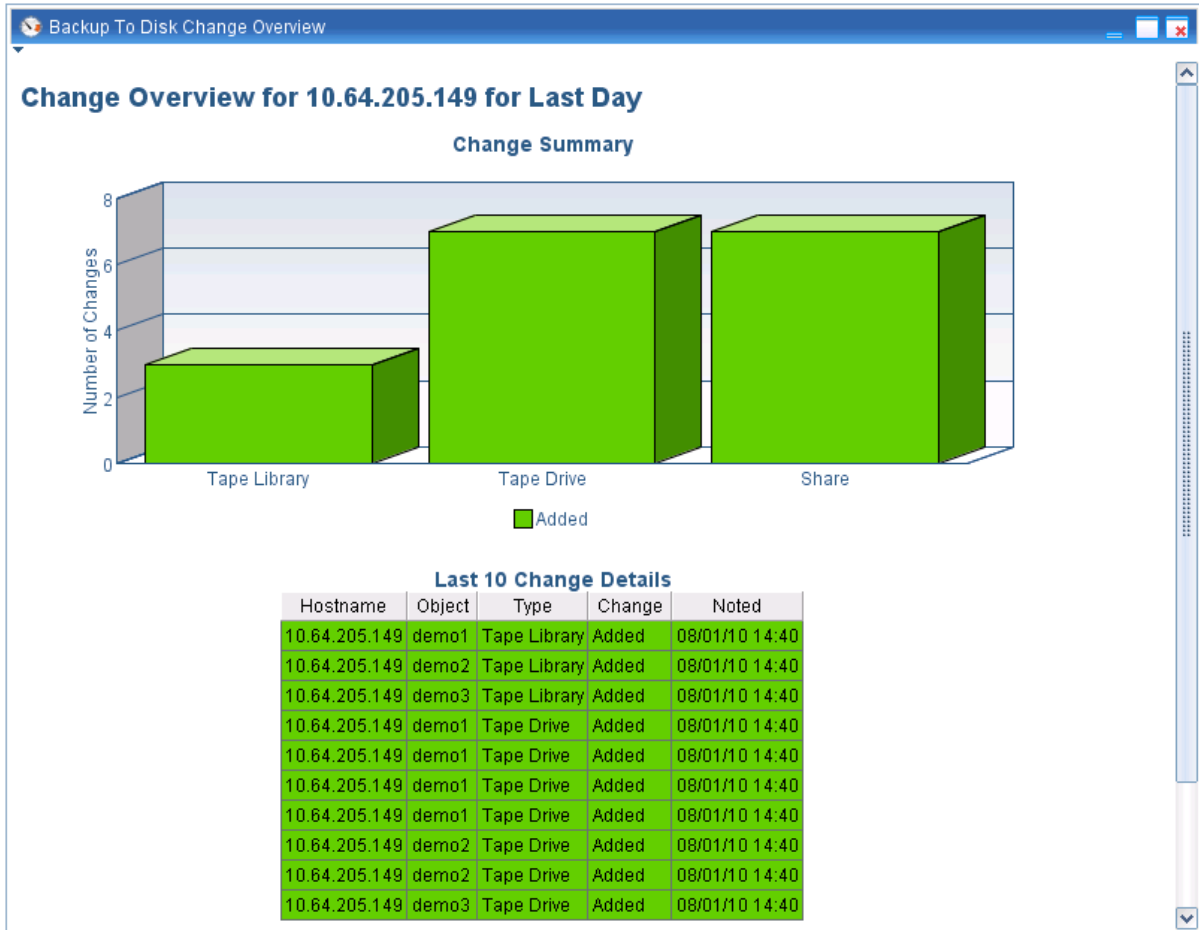


Figure 24. Change Overview

Hostname	Object	Type	Change	Noted
10.64.205.149	demo1	Tape Library	Added	08/01/10 14:40
10.64.205.149	demo2	Tape Library	Added	08/01/10 14:40
10.64.205.149	demo3	Tape Library	Added	08/01/10 14:40
10.64.205.149	demo1	Tape Drive	Added	08/01/10 14:40
10.64.205.149	demo1	Tape Drive	Added	08/01/10 14:40
10.64.205.149	demo1	Tape Drive	Added	08/01/10 14:40
10.64.205.149	demo1	Tape Drive	Added	08/01/10 14:40
10.64.205.149	demo2	Tape Drive	Added	08/01/10 14:40
10.64.205.149	demo2	Tape Drive	Added	08/01/10 14:40
10.64.205.149	demo3	Tape Drive	Added	08/01/10 14:40
10.64.205.149	/backup	Share	Added	08/01/10 14:40
10.64.205.149	/backup/export1	Share	Added	08/01/10 14:40
10.64.205.149	/backup/export2	Share	Added	08/01/10 14:40
10.64.205.149	/backup/ost	Share	Added	08/01/10 14:40
10.64.205.149	backup	Share	Added	08/01/10 14:40
10.64.205.149	share1	Share	Added	08/01/10 14:40
10.64.205.149	share2	Share	Added	08/01/10 14:40

Figure 25. Change Details report

Benefit

In today’s dynamic environments it is essential to keep track of changes. DPA allows the user to readily establish the change history of the Data Domain system assets both in real time and historically.

Multidomain analyses and alerting

A domain refers to each piece of hardware and software in the chain from the source of the client’s data to the destination of the data on a Data Domain system. These multiple layers in the environment can make it difficult to manage and troubleshoot when problems arise. The first reaction is usually that the most recent change caused the problem, when many times the problem was simply uncovered by this latest change. DPA enables a user to look at each domain to locate problems, getting the source of the problem quicker.

DPA is much more than a reporting tool. It provides the analytical capabilities of a skilled operator scanning the backup infrastructure and processes 24x7 to pick up problems and potential issues and alert operations staff before the issues become breaches of data protection policies.

DPA has a comprehensive set of analyses that can run continually, look for exceptions, and alert the user when it finds them. These real-time alerts can identify issues before they escalate and provide a truly proactive capability.

Although reporting is an important area for customers, they need something more proactive to help them meet their data protection requirements. This is where DPA’s Proactive Analysis Engine (PAE) comes in to play.

If virtual or physical tape resources are unavailable to the backup application the user could be automatically alerted to this fact, enabling them to investigate the Data Domain system in the event of a failure.



Figure 26. User alerts

Benefit

DPA analyses provide monitoring of critical systems and intelligent reporting of failures, exceptions, and predicted events. With the addition of integration into standard event management tools DPA removes many layers of complexity in monitoring a backup environment and allows the equivalent of a highly skilled and vigilant operator on a 24x7 basis.

Conclusion

EMC Data Protection Advisor (DPA) collects, monitors, analyzes, and reports on information from your customer's entire data protection infrastructure, providing a unified data protection management interface. Data Domain system support in DPA gathers information about the configuration, status, and performance of Data Domain system components, via SNMP from the Data Domain MIB. DPA can present this information in multiple formats, arranging assets in logical groups (such as business unit or geography), greatly enhancing readability, and giving the user the ability to perform advanced reporting, troubleshooting, performance management, and capacity planning operations. Through heterogeneous support, DPA reduces the cost and risk to manage a data protection environment, enabling our customers to get more from their existing investments.

References

The following can provide additional information and can be found on Powerlink[®], EMC's password-protected customer- and partner-only extranet.

- *EMC Data Protection Advisor Version 5.6 Architecture Overview*
- *EMC Data Protection Advisor Version 5.6 Compatibility Matrix*
- *EMC Data Protection Advisor Version 5.6 Installation Guide*
- *EMC Data Protection Advisor Version 5.6 Administration Guide*
- *EMC Data Protection Advisor Version 5.6 Reference Guide*
- *EMC Data Protection Advisor Version 5.6 Release Notes*
- *EMC Data Protection Advisor Version 5.6 User Guide*