

EMC Data Protection Advisor 6.0

Applied Technology

Abstract

EMC® Data Protection Advisor provides a comprehensive set of features to reduce the complexity of managing data protection environments, improve compliance with business and regulatory requirements, and reduce the risk of data loss. This white paper outlines the technology behind DPA 6.0 and provides an overview of its features.

March 2013

Copyright © 2013 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.

Part Number H11363

Table of Contents

Executive summary	4
Audience.....	4
Functionality Overview	5
Flexible Data Gathering Architecture.....	5
Broad Data Collection.....	5
Advanced Analysis	6
Powerful Reporting	7
Wide Range of Reports	8
Drill Down Reports.....	8
Chargeback Reports	8
Custom Reports.....	8
Forecast Reports.....	8
New Functionality in DPA 6.0	9
Highly Scalable Architecture.....	9
Embedded DPA Datastore.....	10
Enhanced Analysis Engine.....	11
Redesigned User Interface.....	11
Enhanced VMware Support	12
Industry-standard REST API	12
Conclusion	13

Executive summary

EMC® Data Protection Advisor (DPA) is a comprehensive monitoring, analytics, alerting and reporting platform that provides businesses and service providers with full visibility into the utilization, effectiveness and compliance status of their data protection environment. It performs this by monitoring infrastructure used to store and protect data, including backup software and devices, servers, storage arrays, databases and virtual infrastructure.

DPA 6.0 has been engineered to help organizations respond to current and future data protection challenges and to assist with IT transformation to an IT-as-a-Service model. It has a new scalable architecture that can grow in line with the constant increase in data usage and the associated challenges in protecting that data, as well as an enhanced analysis engine that provides real-time processing, correlation and alerting on events across data protection technologies. The architecture and features of DPA 6.0 supports easy implementation and integration into Service Provider environments.

The new dashboard driven user interface has been designed to be intuitive and to simplify common tasks with wizard-style processes. Deep integration with portals and other applications is simplified through the use of a new industry standard REST API.

Audience

This white paper is intended for new and existing DPA users seeking an understanding of the product and information on the latest technology and features in DPA 6.0.

Functionality Overview

EMC® Data Protection Advisor (DPA) is a comprehensive monitoring, analytics, alerting and reporting platform that supports a wide range of data protection products and infrastructure.

Significantly more than a reporting tool, DPA provides comprehensive visibility into an organizations data protection environment, enabling a unified view of data protection status. DPA gathers information from multiple sources, analyses the information in real time, and generates automated alerts on user-defined rules. The DPA user interface gives access to consolidated dashboards, as well as chargeback and show-back, status, trend and forecast reports.

DPA has been designed to enable quick access to the required level of data protection status information. The DPA dashboards and consolidated summary reports give assurance to IT management and auditors that data is protected according to required organizational protection policies, while detailed drill-down reports and proactive alerts enable operational staff to rapidly identify, investigate and resolve data protection issues.

As more organizations transform to IT as a Service operational environment or outsource aspects of their data protection to outside service providers, DPA is able to scale to cloud-level architecture designs and provide the required proof of service information and charge-back reports that these environments require.

Flexible Data Gathering Architecture

DPA has been designed to be rapidly self-installed and configured so that data collection and reporting can occur within minutes.

A simple wizard interface enables configuration of data gathering through either the DPA Collection Agent installed with the DPA Server, or with an agent installed on, or near, data protection infrastructure.

For example, a DPA agent can be installed on an EMC® NetWorker server to gather NetWorker and operating system data locally, or a remote DPA Agent could gather NetWorker data without a local agent being installed. In either scenario, within minutes the DPA Discovery Wizard could be used to add the NetWorker server to DPA and configuration, job and performance data would be written to the DPA datastore and be accessible through DPA's vast collection of built-in reports.

The DPA Application Server and Datastore Server can be installed on a number of 64-bit server platforms, including Windows, Solaris and Linux.

Broad Data Collection

DPA is able to gather data from a wide range of data protection infrastructure, including backup products, backup appliances, switches and replication solutions. DPA data collection support includes:

- Backup servers (such as EMC® Avamar, EMC® NetWorker as well as products from IBM, Symantec and others)
- Storage devices (such as EMC® Data Domain)
- Tape libraries from various vendors (including IBM, HP, Quantum and Oracle StorageTek)
- Fibre Channel switches (from vendors such as Cisco, Brocade and McData)
- VMware infrastructure
- Server operating systems (including Windows, RHEL, SUSE Linux, Solaris, AIX, HP-UX)
- EMC Storage Array replication (EMC Symmetrix, EMC VMAX, EMC CLARiiON, EMC VNX and EMC VPLEX)
- Software replication solutions (EMC RecoverPoint)
- Databases (Microsoft SQL Server, Oracle and PostgreSQL)
- Application replication analysis (Microsoft Exchange, MS-SQL Server, Oracle)



Figure 1: Broad data protection infrastructure support

A full list of supported products is available in the EMC Data Protection Advisor 6.0 Software Compatibility Guide

Advanced Analysis

DPA enables users to create advanced analysis policies and automatically generate email / SNMP alerts when the rules that make up a policy are triggered. Analysis policies can be either:

- Event based, meaning that an alert is triggered in response to data streaming into the DPA Server, providing real-time event processing and alerting
- Schedule based, meaning that DPA analyzes data in the datastore on a scheduled basis and triggers an alert when appropriate

An analysis policy is made up of a set of rules, which are instructions to DPA to analyze data, compute and correlate results across domains and determine if a certain condition, or multiple conditions, are met.

DPA 6.0 includes a number of default analysis rules in various categories, such as data protection, change configuration, capacity planning, recoverability and performance. DPA also provides the ability to customize existing rules and to create new rules. An analysis policy can be created with a combination of rules according to requirements.

Some examples of analyses policies that could be created are:

- Email a backup administrator if a backup job has not completed successfully for more than 48 hours
- Email a backup administrator if a NetWorker bootstrap has not been generated within the past 24 hours
- Email a storage administrator if a SAN replication job did not complete in a specified window
- Email a network administrator if the WAN usage on a RecoverPoint RPA is higher than a specified threshold
- Generate a SNMP trap if DPA forecasts that a file system on a Tier 1 system will reach 90% utilization within the next 3 weeks.

Analysis Policies						
Analysis Policy Library	Rule Templates	Applied Analysis Policies	Protection Policies	Chargeback Policies	Data Collection Policies	
Rule Template Name	Description	Category	Rule Type	Object Types	Alert	F
Backup Not Occurred For Many Days	Alert if a backup not occurred for many days	DataProtection		BackupClient		
Filesystem file utilization high	Alert if a filesystem file utilization is higher than a given threshold	ResourceUtilization		HostFileSystem		
No NetWorker bootstrap generated	Alert if no NetWorker bootstrap has been generated for a period of t..	Status		NetWorker		
Fibre Channel port reporting more than x% errors	Alert if a Fibre Channel port is reporting errors on more than a given..	Troubleshooting		FibreChannelPort		
Filesystem utilization high	Alert if a filesystem utilization is higher than a given threshold	ResourceUtilization		HostFileSystem		
Too many backups without a full	Alert if there have been too many backups without a full	DataProtection		BackupClient		
Recover point journal utilization high	Alerts if the Journal utilization on a RecoverPoint RPA is higher the...	ResourceUtilization		RecoverPointCo...		

1 Selected

Figure 2: DPA Rules

Powerful Reporting

DPA contains a powerful and flexible reporting engine, enabling on-demand reports to be run in the GUI as well as the ability to email reports on a scheduled basis. DPA also includes a set of customizable dashboards giving administrators and IT management an immediate visual overview of Key Performance Indicators.

Reports can be run against a specific system or aspect of data protection, or consolidated reports can be generated for custom groups of systems. For example, consolidated reports could be run for the following custom groups:

- Group based on business unit (such as finance department)

- Group based on location (such as New York based servers)
- Group based on custom defined attributes (such as a customer name in a service provider environment)

Wide Range of Reports

DPA includes hundreds of out of the box reports; from summary reports designed to provide managers with a high-level overview of data protection status, to operational reports with details on system configuration, job status, and specific errors. Report categories include data protection, chargeback, capacity planning, regulatory and configuration, and reports may also include data from other sources such as an external database or text file.

Drill Down Reports

Many DPA reports include ‘drill-down’ ability, where a user can review a high level summary of data protection status and drill-down multiple levels to get detailed information. For example, an administrator can review a summary report for last night’s backups or replication jobs, and then drill-down to get further information on failures, down to a report listing errors for a specific server or job.

Chargeback Reports

DPA includes a backup and replication chargeback reporting mechanism, with the ability for an administrator to create multiple charge back policies (a different cost amount per GB replicated or backup taken for example). These charge back policies can be applied to different groups or cost centers, such as one cost basis for a Gold Level service offering and a different cost basis for a Silver Level offering.

Custom Reports

In addition to built-in reports, DPA also includes a powerful custom report editor for creation of new reports tailored to specific environments and use cases. These custom reports can be configured to use external inputs such as a database or a CSV file generated by another application, in order to facilitate Asset Management integration.

Forecast Reports

DPA is able to generate forecast reports where a trend line is used to extrapolate into the future, such as a report forecasting Avamar Server capacity utilization for the next 3 months. DPA uses historical information stored in the datastore to perform trend analysis.

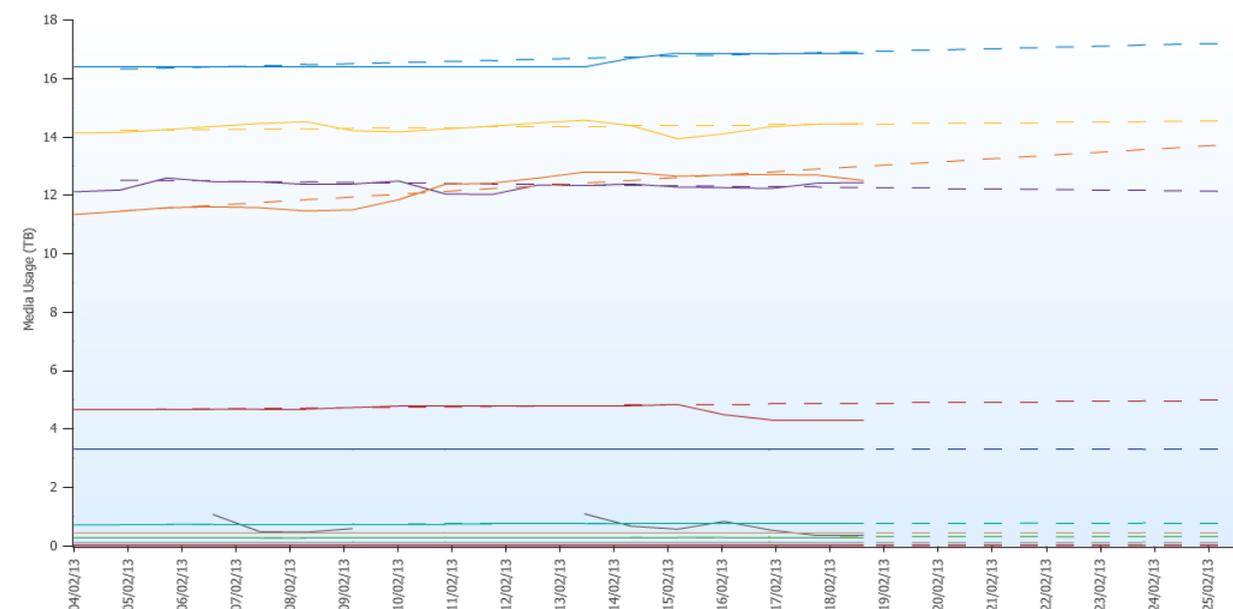


Figure 3: DPA 6 Forecast Report

New Functionality in DPA 6.0

DPA 6.0 offers significant product enhancements to support customer confidence in data protection, including a new highly scalable architecture that is purpose-built for the cloud; a complete refresh of the user interface; a new analysis engine to drive real-time visibility and predictability in cross-domain correlation analysis and trending; and a standards based API mechanism.

Highly Scalable Architecture

The separate Controller, Reporter, Listener, Publisher and Analysis Engine processes from DPA 5.x are replaced in DPA 6.0 with a new unified engine. This new architecture offers increased flexibility, performance and scaling capabilities in large environments.

DPA 6.0 flexible deployment options include:

Combined Application and Datastore: Not recommended for production environments but this architecture allows for simple deployment in test environments.

Single Application Server and Remote Datastore: This architecture is recommended for most small to mid-size production environments and separates the DPA Application Service on one host from the DPA datastore on a separate host.

Multiple Application Servers and Single Datastore: This architecture is recommended for larger environments and cloud deployments and provides for multiple DPA Application Servers working with a single DPA datastore. With this deployment architecture a customer supplied load balancing switch may be placed in front of the

DPA Application Servers to distribute user load over Application Servers, or manual load balancing without a load balancing switch can be used to assign groups of users across different Application Servers. In virtual environments DPA Application Servers may be brought on and offline as demand requires, such as bringing additional servers online to handle increased load during peak monthly reporting periods.

One of the key design specifications for the new architecture was the ability to offer increased flexibility and scale-out options for large customers and multi-tenanted service provider environments. A result of this is the ability to partition DPA Application Servers into separate clusters so that the load placed on DPA by one group of users will not affect the load on a separate cluster for a different set of users. This new functionality also enables a design architecture that separates DPA Application Servers used for processing DPA Collector Agent data from DPA Application Servers used by groups of users to run reports.

See Figure 4 for an example of DPA 6.0 Architecture Design in a Service Provider environment.

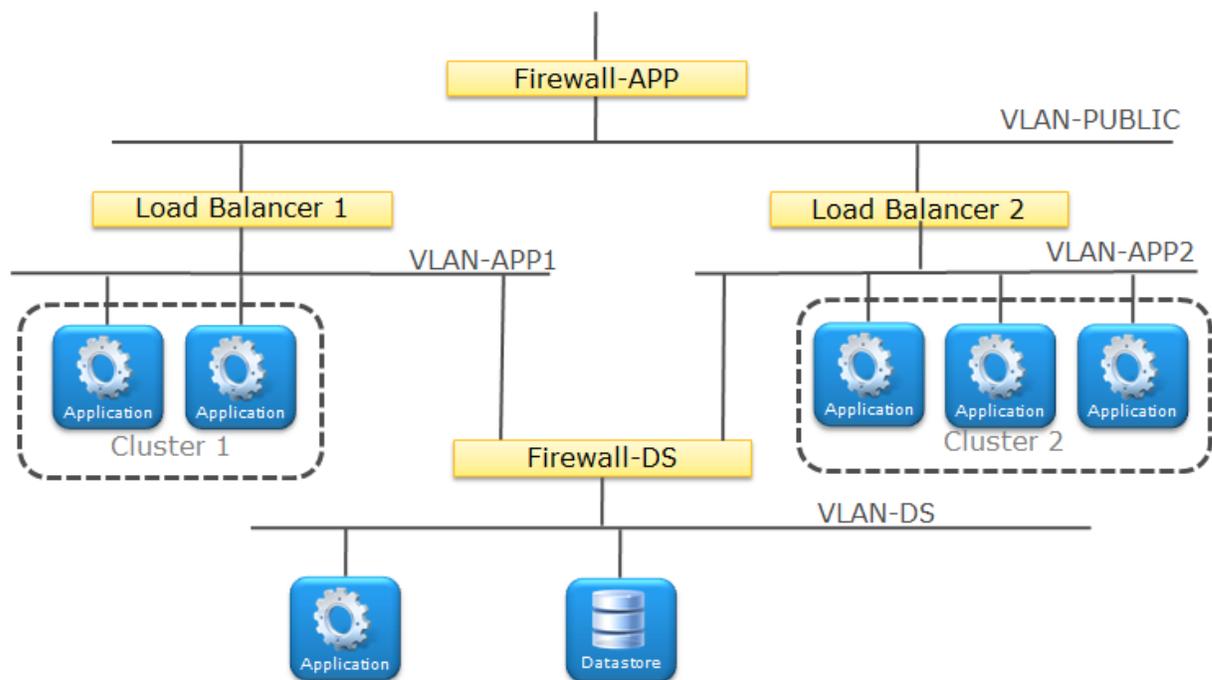


Figure 4: Example Architecture for a Service Provider environment

Embedded DPA Datastore

DPA 6.0 includes an embedded datastore, offering performance and support benefits due to closer alignment of the DPA application with its database. This architecture change removes the need for customers to supply and manage a separate database instance, external to DPA, leading to lower deployment and management costs.

A migration tool included with DPA 6.0 enables customers to migrate DPA v5.5.1 (and later) databases to the new embedded datastore. This powerful migration tool provides flexibility by enabling users to import all DPA 5.x data or to limit data

migration by type and age of data. In addition to data migration, the migration tool also migrates the DPA 5.x configuration, including custom reports, monitored objects, users and existing licenses.

Enhanced Analysis Engine

DPA includes a new analysis engine with the ability to monitor and process events as data is received from collector agents. In addition to the schedule based rules in DPA v5.x, administrators now have the ability to create event based rules which are able to analyze incoming data within a few seconds of the data arriving at the DPA server. This functionality enables faster notification of missed events and compromised or erroneously trending Key Performance Indicators.

The advanced analysis engine in DPA allows for complex cross-domain rules to be created, such as rules that are built from a combination of backup job information, switch status and system performance metrics.

The DPA 6.0 analysis engine is based on Complex Event Processing (CEP) concepts. CEP supports processing many events happening across all the layers of an organization, identifying the most meaningful events within the event cloud, analyzing their impact, and alerting as data is streamed into the DPA Server.

Redesigned User Interface

The DPA 6.0 user interface has been completely redesigned to offer an intuitive navigation model and simplified administration, while aligning with the familiar look and feel of other EMC software. The new UI works with any browser that has the Adobe Flash plugin installed.

One of the key features of the new UI is the use of dashboards, designed to provide at-a-glance insight into the overall state of the data protection environment. These dashboards are updated on a scheduled basis and present key information when a user logs into DPA without the need to first find and run relevant reports.

For example, a backup administrator logging into DPA can be presented with a dashboard showing summary information such as clients not backed up, client backup success rate, and failed backups and restores. Each DPA user is able to customize the dashboard they see on login from a selection of Key Performance Indicators.

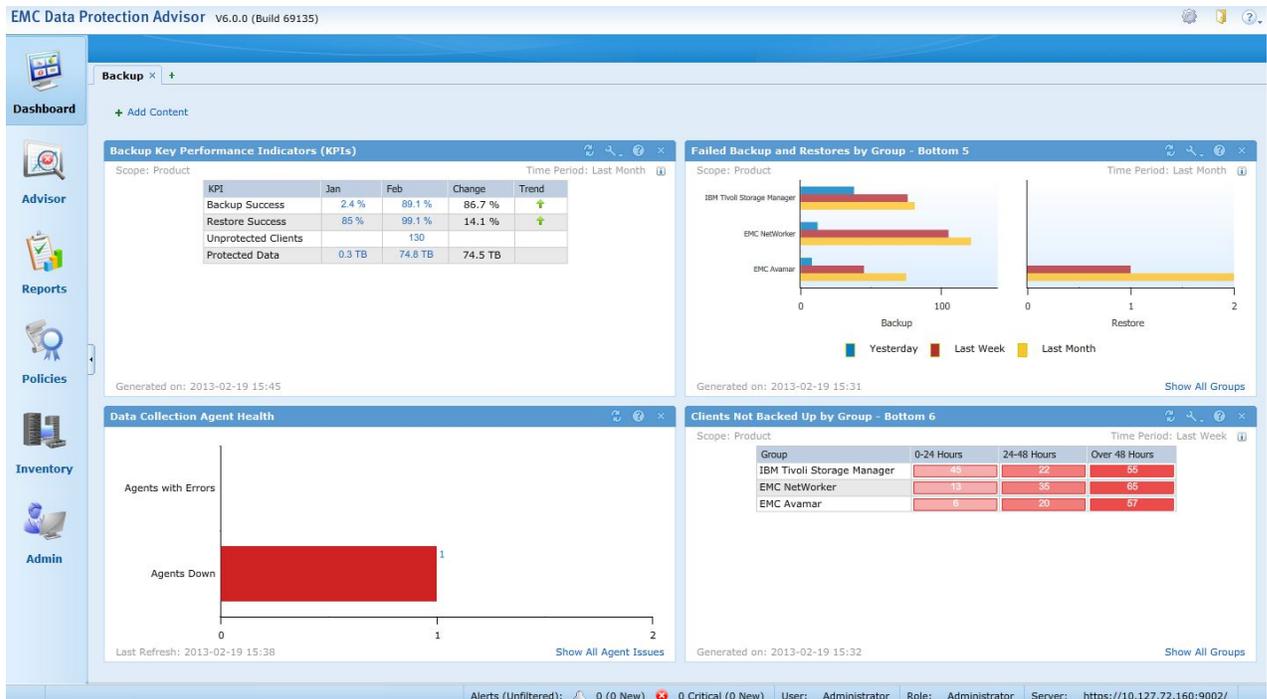


Figure 5: Dashboards in the redesigned DPA 6.0 User Interface

Enhanced VMware Support

While DPA 5.x supported monitoring of vSphere infrastructure, including the ability to validate that all virtual systems on an ESX server are being protected, this functionality has been further enhanced.

DPA 6.0 introduces support for monitoring virtual machines that are migrated between ESX servers using VMware vMotion technology. This enhancement helps organizations to keep track of the protection status of virtual machines as they are dynamically moved around the virtual infrastructure.

Industry-standard REST API

DPA 6.0 implements a new Application Programming Interface (API) based on the industry-standard REST API architecture.

This new API implementation allows access to DPA functions through standard HTTP(S)/XML protocols, providing the ability to seamlessly integrate DPA with external portals and third party applications, enabling new web-based composite applications and dashboards.

Demonstrating the scope and power of the new API, the DPA 6.0 user interface functions entirely by making REST API based calls to the DPA engine.

The DPA REST API is versioned to ensure future advances are compatible for customers utilizing the API.

Conclusion

EMC Data Protection Advisor 6.0 ensures that DPA can scale with an organizations data growth and journey to the cloud, increases the ability to proactively monitor data protection infrastructure, enables deeper integration with portals and other applications, and presents a redesigned user interface aligned with other EMC software.

As a powerful tool for monitoring, analyzing, alerting and reporting across the data protection environment, DPA provides a fast return on investment by:

- Greatly reducing the effort required to assure that data is protected according to agreed policies and contracted Service Level Agreements
- Generating chargeback and show-back reports for data protection billing or cost recovery
- Rapidly alerting on issues in the data protection environment
- Helping to plan for future growth with capacity planning reporting
- Allowing fast access to reports required for regulatory requirements

With a new architecture and analysis engine, DPA 6.0 is designed to meet the complex data protection management requirements of organizations of all sizes, from small single site businesses to large enterprises transitioning to an IT as a Service model as well as service providers with geographically dispersed data centers.