



Veritas Resiliency Platform

Disaster Recovery Orchestration
for Amazon Web Services

VERITAS™

The truth in information.

Contents

INTRODUCTION	4
RESILIENCY PLATFORM COMPONENTS	4
DEPLOYMENT OPTIONS	5
SOLUTION VALUE	6
SIZING GUIDANCE	8
BEST PRACTICES AND RECOMMENDATIONS	9
CONCLUSION	9
APPENDIX	9

REVISION HISTORY

Version	Date	Changes	Author
1.00	2018-Dec	Initial Version	Ryan Behiel
1.01	2018-Dec	Feedback, additional data	Ryan Behiel
1.02	2019-Jan	Update deployment options, AWS Icons	Ryan Behiel

INTRODUCTION

EXECUTIVE SUMMARY

Veritas Technologies is a leader in developing data resiliency solutions that focus on the protection and management of digital assets critical for a company's success and business continuity. One of our flagship products, Veritas Resiliency Platform, is designed to enable high availability and disaster recovery (HA/DR) for data centers, hybrid and multi-cloud environments. Adding to the Veritas portfolio and legacy of creating stable solutions customers have trusted and relied on, Resiliency Platform is an enterprise-class solution designed to address the HA/DR needs of organizations using multiple platforms, including Amazon Web Services (AWS). Resiliency Platform acts as an orchestration engine that can manage a wide range of data center workloads and enable failover, fallback, migration and testing of workloads, as required.

TARGET AUDIENCE








This document is for customers, partners and Veritas field personnel interested in learning more about using Resiliency Platform as a solution to provide HA/DR with a combination of traditional VMware data center infrastructure and systems in AWS.

SCOPE

The purpose of this document is to provide technical details to assist in understanding Resiliency Platform as a solution for HA/DR between an on-site VMware environment and AWS. It describes the components of this solution, its value, sizing guidance and some best practices. Although this document provides some deployment examples, we advise you to refer to the product documentation for installation, configuration and administration information. We update this documentation periodically, and you can download the latest version from this [link](#).

RESILIENCY PLATFORM COMPONENTS

The following tables will outline the names and descriptions of the components involved in the Resiliency Platform setup, configuration and operational process.

Component		Description
VMware ESXi Server		A purpose-built bare-metal hypervisor that installs directly onto a physical server. Resiliency Platform installs a Managed Host Package on the VM guests running on ESXi hosts. The Managed Host Package acts as an I/O tap that relays data from the VMs to the local Resiliency Platform Replication Gateway.
VMware vCenter Server		A centralized management application that allows users to manage VMware virtual machines (VMs) and ESXi hosts centrally. You can use vCenter Server to install the Resiliency Platform components by using the 'Deploy OVF Template' option for the Resiliency Platform Open Virtual Appliance (OVA) files.
Resiliency Platform Resiliency Manager		The Resiliency Platform component that provides the services required for protecting assets (i.e., VMs) within the logical scope of a Resiliency Platform deployment (known as a Resiliency Domain). The Resiliency Manager discovers and manages information about data center assets from the Infrastructure Management Server. You only need a Resiliency Manager in the on-site data center when using Resiliency Platform to orchestrate a failover operation from AWS to the on-site data center. The Resiliency Manager is deployed as a virtual software appliance.
Resiliency Platform Infrastructure Management Server		The Resiliency Platform component that discovers and monitors assets within a data center and enables management operations on assets (i.e., starting or stopping a VM). The Infrastructure Management Server scales horizontally and is deployed as a virtual software appliance.
Resiliency Platform Replication Gateway		The Resiliency Platform component that manages replication across sites and across hypervisors. The Replication Gateway scales horizontally and is deployed as a virtual software appliance. Replication of data between an on-site VMware environment and an AWS environment does not require a data format conversion. The Replication Gateway receives data from the in-guest I/O tap module on the VM guests and creates update sets using this data. The update sets are replicated every 2 minutes or 500mb, whichever is sooner. These parameters are configurable.
Resiliency Platform Data Gateway		The Resiliency Platform component that is deployed in the AWS environment to enable replication using Amazon S3 object storage. This component is optional because it is not a requirement to use object storage within the AWS environment for Resiliency Platform-managed replication. The Data Gateway provides additional resiliency for the replication process and can also help reduce the cost of infrastructure required in the AWS environment. The Data Gateway deploys some additional components within the AWS environment described here . You need to deploy the Data Gateway in an AWS region where FIFO (first in first out) queue support is available.
Amazon CloudWatch		Amazon CloudWatch is a component of AWS that provides monitoring for AWS resources and the customer applications running on the Amazon infrastructure. Installing the Resiliency Platform Data Gateway also creates various alarms, and CloudWatch is only implemented and required as part of this process.

Amazon DynamoDB		A fully managed NoSQL database provided by AWS that is used by Resiliency Platform to maintain consistent ordering of data being replicated into and out of the AWS environment. DynamoDB is only implemented and required when you install a Resiliency Platform Data Gateway is installed.
Amazon EBS (Elastic Block Storage)		A durable, block-level storage device that is attached to an EC2 instance. EBS volumes can be provisioned using SSD-backed storage and HDD-backed storage. The Resiliency Platform components within AWS use EBS volumes for local storage.
Amazon EC2 (Elastic Cloud Compute)		A virtual computer system (known as an instance) that is provided by Amazon on which users run their own computer applications. EC2 instances are used to run the Resiliency Platform components.
Amazon Machine Image (AMI)		A master image used for the creation of Amazon EC2 instances. You can create AMIs created from the OVA files available for the Resiliency Platform components. Doing so requires an Amazon S3 bucket, which is used to store the Resiliency Platform OVA files.
Amazon S3 Storage		Object storage provided by AWS that is optionally used by Resiliency Platform as a target for storing replicated block-level data from on-site VMware systems. You can also use S3 for replication in the reverse direction, where it can store block-level data from cloud-based systems being replicated to an on-site VMware environment.
Amazon SNS (Simple Notification Service)		Amazon Simple Notification Service is a notification service provided as part of Amazon Web Services. As part of the Resiliency Platform deployment, a topic called VeritasDGWMonitoringTopic is created. Amazon SNS is only implemented and required when a Resiliency Platform Data Gateway is installed.
Amazon SQS (Simple Queue Service)		A fully managed distributed message queuing service provided by Amazon that is used by Resiliency Platform to queue requests. Amazon SQS is only implemented and required when a Resiliency Platform Data Gateway is installed.
AWS API Gateway		A fully managed service that enables Resiliency Platform to move data in and out of Amazon S3 using API calls made by the Resiliency Platform Data Gateway.
AWS CloudFormation		The option within AWS that enables the creation of templates for provisioning services or applications within AWS. You can use a CloudFormation template to easily deploy the Resiliency Platform components within an AWS environment.
AWS Identity and Access Management (IAM)		AWS IAM enables users to manage secure access to AWS services and resources. Using IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources. It is necessary to ensure that the IAM user who deploys the Resiliency Platform Data Gateway has the required permissions for deploying the Data Gateway in AWS.
AWS Lambda		The Amazon computing service that runs in response to events. When data replicated by Resiliency Platform from an on-site data center lands in Amazon S3 storage, AWS Lambda triggers a notification. The AWS Lambda functions are only required and implemented as part of the install process for the Resiliency Platform Data Gateway.
AWS Marketplace		The online store where AWS customers can discover, purchase, migrate and deploy software applications within AWS. Resiliency Platform is available in the AWS Marketplace and you can deploy it using a CloudFormation template.
AWS Virtual Private Cloud (VPC)		Amazon VPC is a service that allows users to provision a logically isolated section of the AWS Cloud where they can launch AWS resources in a virtual network they define. Resiliency Platform is provisioned inside a VPC.
VPN Gateway		This is the anchor on the AWS side that links a data center or network to an Amazon VPC. A customer VPN gateway (physical or software appliance) is also required on the other side of the connection.

Table 1) Resiliency Platform and AWS Component Descriptions

DEPLOYMENT OPTIONS

As listed in Table 2, the Resiliency Platform deployment consists of components that run in both the on-site VMware environment and in the AWS environment. Some components are not required or applicable within both environments. Table 3 will describe the deployment options available for the Resiliency Platform solution.

Component	On-Site Data Center	AWS
VMware ESXi Server	■	
VMware vCenter Server	■	
Resiliency Platform Resiliency Manager	■	■
Resiliency Platform Infrastructure Management Server	■	■

Resiliency Platform Replication Gateway		
Resiliency Platform Data Gateway		■
Amazon CloudWatch		■
Amazon DynamoDB		■
Amazon EBS		■
Amazon EC2		■
Amazon Machine Image (AMI)		■
Amazon S3 Storage		■
Amazon SNS		■
Amazon SQS (Simple Queue Service)		■
AWS API Gateway		■
AWS CloudFormation		■
AWS Lambda		■
AWS Marketplace		■
AWS VPC		■
VPN Gateway		■

Table 2) Resiliency Platform and AWS Component Requirements

■ Required
■ Optional

Option	VMware On-Site	AWS
OVA File Import	OVA files are used to install all the Resiliency Platform components: Resiliency Manager, Infrastructure Management Server and Replication Gateway. You can import them into VMware using the process described in this link .	N/A
Amazon Machine Image (AMI)	N/A	AMIs are available for all the Resiliency Platform components and you can use them to deploy Resiliency Manager, Infrastructure Manager (IMS) and Replication Gateway in AWS.
Express Install	An option that bundles the Resiliency Platform Infrastructure Management Server and Replication Gateway appliances into a single package (vApp) that you can install using a single process.	This is a CloudFormation template available in the AWS Marketplace that you can use to quickly deploy the full Resiliency Platform component stack within AWS. The process is described in this link .
Replication Gateway Install	You can install Replication Gateway in a VMware environment by importing the corresponding OVA file and configuring it within the Resiliency Platform domain. At least one Replication Gateway is required in the VMware environment to replicate data to the AWS environment.	This is a CloudFormation template available in the AWS Marketplace to install a Replication Gateway only. This option assumes you have already provisioned the Express Install option or the required Resiliency Platform AMIs within the AWS environment.

Table 3) Resiliency Platform Deployment Options

LICENSING

The Resiliency Platform stack available in the AWS Marketplace is provided as a Bring Your Own License (BYOL) model. Please review the product overview sheet and pricing information prior to deploying Resiliency Platform in an AWS environment.

Resiliency Platform installs with an embedded 60-day trial license. Once this trial expires, the product will no longer function.

SOLUTION VALUE

Creating an HA/DR solution between an on-site data center and an AWS environment poses some challenges that in most cases cannot be completely resolved with native AWS configuration options. Resiliency Platform is designed to integrate with AWS to address these challenges by providing additional functionality and automation of the DR process.

Common AWS DR Scenarios

The following scenarios outlined by Amazon describe some common DR configurations you can use to achieve DR capability between an on-site data center and an AWS environment.

- **Backup & Restore:** You can deploy an AWS Storage Gateway on-site that enables you to create snapshots of on-site data volumes that can then be transparently copied into Amazon S3 storage for backup purposes. You can subsequently create local volumes or Amazon EBS volumes from these snapshots that you could use in a DR scenario.
- **Pilot Light:** This is a term used to describe a DR scenario in which a minimal version of an on-site environment is always running in the cloud. To provision the remainder of the infrastructure and restore business-critical services, you would typically have some preconfigured servers bundled as AMIs, which are ready to be started up at a moment's notice.
- **Warm Standby:** This is a DR scenario in which a scaled-down version of a fully functional on-site environment is always running in the cloud. This option differs from the Pilot Light scenario in that the systems running in the cloud are fully provisioned and already running.

Resiliency Platform Managed DR with AWS

The common AWS DR scenarios described above all require some degree of configuration and user intervention to be successfully executed. Here's how Resiliency Platform can help simplify and automate these DR scenarios and reduce the need for manual user intervention:

- **Backup & Restore:** The AWS Storage Gateway requires you to manually create snapshots of on-site data that can then be copied into Amazon S3 storage. Resiliency Platform can automate and improve the data movement process by filtering and replicating on-site data to the cloud as it's created/changed. Resiliency Platform ensures the data is kept in-sync between the on-site systems and the cloud systems, allowing you to start applications instantly in the event of a failover. Resiliency Platform eliminates the need for you to manually create and copy snapshots and then create usable volumes from these snapshots when needed in a DR scenario.
- **Pilot Light:** Resiliency Platform provides replication of data between sites, automates DNS updates, manages network mappings between sites and starts applications inside AWS in the event of a failover—all with a single click. Resiliency Platform can also run a DR rehearsal on an isolated, non-production network segment within AWS to ensure systems in the cloud are working properly prior to a full DR failover event. You can run DR rehearsals by using snapshots of production data that are then attached to temporarily provisioned systems used for testing purposes. Resiliency Platform also manages the cleanup of the rehearsal environment when it's no longer needed.
- **Warm Standby:** Similar to the Pilot Light scenario, in Warm Standby Resiliency Platform provides replication of data between sites, automates DNS updates, manages network mappings between sites and starts applications inside AWS with a single click. The DR rehearsal option described in the Pilot Light scenario is also available in a Warm Standby configuration.

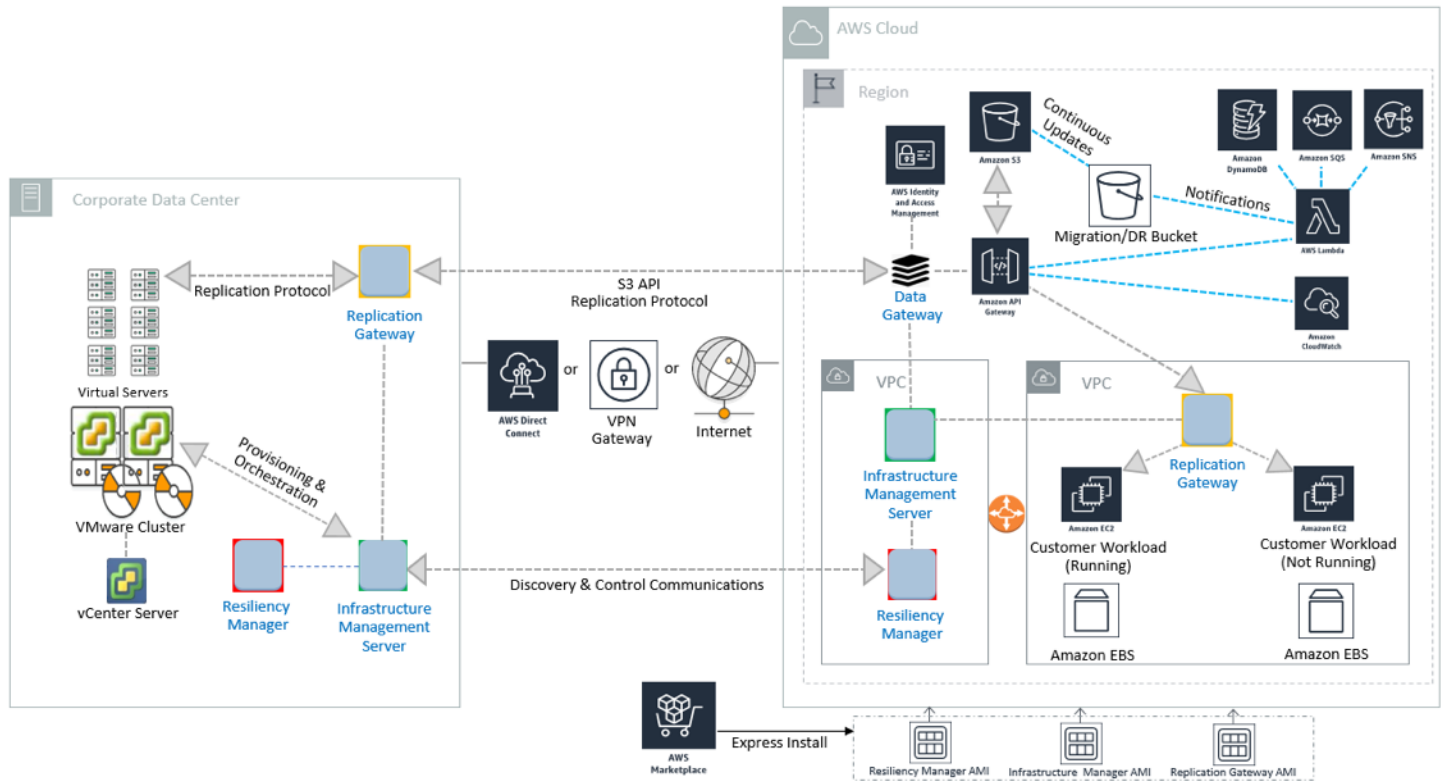
Resiliency Platform can further facilitate DR orchestration using an integrated option where application tiers can be grouped together in a way that represents the entire business service the application provides. This option is known as a **Virtual Business Service (VBS)**. A VBS represents a multi-tier application as a single, consolidated entity and can build on the HA/DR provided for the individual tiers by integrating with products such as Veritas Cluster Server and Veritas ApplicationHA.

Additional orchestration options are available with Resiliency Platform:

- **Resiliency Plans:** Provide the ability to create a custom automated workflow consisting of a specific set of tasks. This workflow can include tasks such as starting, stopping, migrating and taking over a resiliency group or VBS. You can also include DR rehearsals as part of a resiliency plan.
- **Evacuation Plans:** Provide the ability to takeover a resiliency group or VBS from one data center (either on-site or in the cloud) in the event of a site evacuation.

Resiliency Platform also simplifies management by providing a modern and intuitive user interface you can use to manage all the components required to fully orchestrate the DR process. The Resiliency Platform user interface provides a clear view into all operations and can proactively notify administrators of potential risks (i.e., replication lag) that may exist within the environment that could prevent the successful execution of a DR plan.

SOLUTION ARCHITECTURE



SIZING GUIDANCE

The following system resources are required for Resiliency Platform appliances. System resource use may vary based on factors such as your organization's environment size, performance requirements and usage patterns.

VMware vSphere system requirements:

- **Resiliency Platform Resiliency Manager:** 8 vCPUs and 32 GB RAM. Minimum of 60 GB disk space. Resiliency Manager is optional for the on-site data center. It is only required for providing failover orchestration from the AWS environment to the on-site VMware environment.
- **Resiliency Platform Infrastructure Management Server:** 8 vCPUs and 16 GB RAM. Minimum of 60 GB disk space.
- **Resiliency Platform Replication Gateway:** 8 vCPUs and 16 GB RAM. Minimum of 40 GB disk space. An additional (thick) data disk with a minimum of 50 GB is required and each protected VM has a disk space requirement that will vary depending on the configuration of the update set parameters. As such, the data disk may need more than 50 GB, depending on the number of protected VMs. You can find additional capacity planning information for Replication Gateway [here](#).

AWS system requirements:

- **Resiliency Platform Resiliency Manager:** 8 vCPUs and 32 GB RAM. Minimum of 150 GB disk space.
 - **AWS instance type:** m4.2xlarge or better (must satisfy minimum system requirements).
- **Resiliency Platform Infrastructure Management Server:** 8 vCPUs and 16 GB RAM. Minimum of 60 GB disk space.
 - **AWS instance type:** m4.xlarge or better (must satisfy minimum system requirements).
- **Resiliency Platform Replication Gateway:** 8 vCPUs and 16 GB RAM. Minimum of 40 GB disk space. An additional (thick) data disk with a minimum of 50 GB is required, and each protected VM has a disk space requirement that will vary depending on the configuration of the update set parameters. As such, the data disk may need more than 50 GB depending on the number of protected VMs.
 - **AWS instance type:** m4.xlarge or better (must satisfy minimum system requirements).

You can find current Resiliency Platform User Guides and product documentation on the Veritas Services and Operations Readiness Tools (SORT) website at this [link](#).

BEST PRACTICES AND RECOMMENDATIONS

- **Deploy Resiliency Platform in AWS using Express Install.** To do so, go to the AWS Marketplace and find the CloudFormation template that has been created as a deployment option. There is a CloudFormation template for both the Express Install and the Replication Gateway install. The AWS Marketplace deployment offers a much more automated process compared to using AMIs and deploying manually. This option also eliminates the need to manually provide bootstrap inputs for Resiliency Platform appliances.
- **Use the Resiliency Platform Data Gateway for additional resiliency and scale in the replication process.** The Data Gateway stores replicated data in Amazon S3 prior to its being written to EBS volumes attached to the target systems running in AWS. This approach provides additional resiliency by ensuring data is not lost in the event of a failure or service interruption with the Resiliency Platform Replication Gateway appliance.
- **Use Amazon EC2 Reserved Instances for the Resiliency Platform components in AWS.** Using Reserved Instances in AWS will result in a significant discount compared to EC2 On-Demand instance pricing. Reserved Instances also provide a capacity reservation, which is ideal for Resiliency Platform and provides additional confidence in your ability to launch instances when they are needed.

CONCLUSION

Resiliency Platform has been designed to integrate with an evolving IT landscape in a way that helps you achieve resiliency for environments that consist of both on-site and AWS cloud infrastructure and services. You can realize some key benefits when using Resiliency Platform to provide an HA/DR orchestration solution for IT environments consisting of both an on-site data center and systems/services within AWS:

- **Scalability:** By integrating with S3 storage, Resiliency Platform automatically scales according to your requirements by using AWS services to achieve scalability. Using S3 storage to facilitate data movement between an on-site environment and AWS can provide additional resiliency and help you reduce operating costs in the cloud.
- **Simplified management:** Reduce the need for manual processes that are time-consuming and error-prone. Virtual Business Services further simplify management by logically representing complex, multi-tier applications as a single entity you can migrate between sites with a single click.
- **Increased visibility and control:** Resiliency Platform provides a visual representation and single management console for the entire HA/DR domain that helps eliminate complexity and increases your confidence in often-unpredictable situations.
- **Increased confidence:** Resiliency Platform's integrated, non-disruptive DR rehearsals preserve production uptime and increase your confidence in rolling out new technology.

Meeting uptime service-level objectives (SLOs) across an on-site data center and AWS with multiple point tools can be complicated and costly. Resiliency Platform helps you proactively ensure application resiliency across constantly evolving hybrid-cloud environments with a single solution.

APPENDIX

- **Amazon S3-IA:** An Amazon S3 storage class for data that you access less frequently.
- **API:** Application Programming Interface. A set of routines, protocols and tools for building software applications.
- **FIFO Queue:** First in first out queue. In the context of Amazon S3, FIFO queues are designed to enhance messaging between applications when the order of operations and events is critical or where duplicates can't be tolerated.
- **OVA:** Open Virtual Appliance. A preconfigured virtual machine image, ready to run on a hypervisor.
- **Reserved Instance:** An Amazon EC2 instance type that enables you to commit to usage parameters at the time of purchase to achieve a lower hourly rate.
- **Stack:** A collection of AWS resources that you can manage as a single unit.
- **Update Set:** A set of workload I/Os collected over a time by the I/O tap drivers in the systems being managed by Resiliency Platform.

- **Veritas ApplicationHA:** A software application that provides high availability for business-critical applications through application visibility and control across virtual environments.
- **Veritas Cluster Server:** High-availability cluster software for Unix, Linux and Microsoft Windows computer systems typically used for systems with a low RTO/RPO.

DISCLAIMER

THIS PUBLICATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS PUBLICATION. THE INFORMATION CONTAINED HEREIN IS SUBJECT TO CHANGE WITHOUT NOTICE.

No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

ABOUT VERITAS TECHNOLOGIES LLC

Veritas Technologies empowers businesses of all sizes to discover the truth in information—their most important digital asset. Using the Veritas platform, customers can accelerate their digital transformation and solve pressing IT and business challenges including multi-cloud data management, data protection, storage optimization, compliance readiness and workload portability—with no cloud vendor lock-in. Eighty-six percent of Fortune 500 companies rely on Veritas today to reveal data insights that drive competitive advantage. Learn more at www.veritas.com or follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054 USA
+1 (866) 837 4827
veritas.com

For specific country offices and contact numbers,
please visit our website.
veritas.com/about/contact

VERITAS[™]
The truth in information.

V0837 02/19