

A large, intricate red scribble graphic on the left side of the page, composed of many overlapping, flowing lines that form a dense, organic shape. It extends from the top left towards the center of the page.

# **Veritas Access**

## **7.4.2.301 Release Update**

---

Linux

**VERITAS™**  
The truth in information.



**Contents**

About the release..... 3

Layout changes to Veritas Data Deduplication storage..... 4

Changes to log locations for episodic replication ..... 5

Changes to episodic replication alerts ..... 5

Changes to S3 certificates ..... 5

Improvements to the support debuginfo command ..... 6

Deprecated functionality ..... 6

Advance notice for deprecating commands ..... 7

Known issues..... 10

Fixed issues in this release ..... 16



## About the release

This update contains critical updates and enhancements for the Veritas Access 7.4.2 release. This patch can be installed only on Veritas Access 7.4.2 release with Red Hat Enterprise Linux (RHEL) version 7.5. If you are on an earlier version of RHEL, you must upgrade to RHEL version 7.5 before installing the patch.

After upgrading to 7.4.2.301 version, Veritas recommends that you upgrade to RHEL version 7.7.

**Note:** If you have upgraded from an earlier version of RHEL to RHEL version 7.5 or 7.7, then ensure that the faulted or stale repositories have been removed from the `/etc/yum.repos.d/` directory before you perform rolling upgrade.

For upgrading Veritas Access, refer to the *Veritas Access Installation Guide*.

**Warning:** Rolling upgrade is not supported if you have deployed Veritas Access using a shared nothing storage configuration.

Contact [<DL-VTAS-AA742-300-Support@veritas.com>](mailto:DL-VTAS-AA742-300-Support@veritas.com) to get assistance with the upgrade.

### To determine if your environment uses a shared nothing storage configuration:

You can use the following command to determine the disk IDs of the shared disks available in your configuration:

```
# ls /dev/disk/by-id/scsi-* 2>/dev/null | grep -v part | uniq | awk -F'scsi-' '{print $2}'
```

If this command does not return any shared disks, it means that you have deployed Veritas Access using a shared nothing storage configuration.

For upgrading to Veritas Access 7.4.2.301, refer to the following technote:

[https://www.veritas.com/support/en\\_US/article.100049254](https://www.veritas.com/support/en_US/article.100049254)

### New features and enhancements included in this release:

- The storage layout of Veritas Data Deduplication is changed so that multiple file systems of 100 TB are created when you grow or configure a new deduplication server.
- Performance improvements have been done for optimized duplication.
- Default size of the stripe unit of the deduplication file system has been changed to 64K.
- New AutoSupport alerts for episodic replication.
- Changes to S3 certificates.
- Changes have been made to the log location of episodic replication.



## Layout changes to Veritas Data Deduplication storage

Starting from this release, changes have been done to layout of the Veritas Data Deduplication storage. If you have Veritas Access 7.4.2.301 installed in your system and you configure Veritas Data Deduplication, multiple file systems of 100TB and a separate catalog file system of size 5TB are created.

The **dedupe grow** command creates a new file system for every 100TB.

If you upgrade to Veritas Access 7.4.2.301 and you had previously configured Veritas Data Deduplication, then the old layout of Veritas Data Deduplication is preserved. If a separate catalog file system existed before upgrade, the same catalog file system is present after the upgrade also. If the size of the data file systems is more than 100 TB, then multiple directories are created in proportion to the size considering 100 TB as the threshold. If multiple directories are already present, then they are not created again.

**Note:** Veritas Data Deduplication configuration requires 5 TB for creating the catalog file system. This is over and above the size specified for the deduplication pool in the GUI.

Reconfiguration is a two-step process involving unconfiguring the system followed by configuration. After an upgrade to Veritas Access 7.4.2.301, if you want to reconfigure the Veritas Data Deduplication service, the following checks are performed:

- If a separate catalog file system existed before reconfiguration, then the same catalog file system is passed to the **dedupe config** command during reconfiguration.
- If a separate catalog file system did not exist before reconfiguration, then the catalog file system cannot be passed to the **dedupe config** command during reconfiguration.
- Configuration fails if there is a directory mismatch before and after configuration.

You can use the **dedupe config** command to configure the deduplication server.

```
Dedupe> config `` *filesystem1* [, *filesystem2*,...] *ip* *username*  
[*catalogfilesystem*]
```

This command configures the deduplication server using the given set of data file systems, IP address, username and catalog file system. The data and catalog file systems are used to store the deduplicated data and its metadata respectively.

Reconfiguration of Veritas Data Deduplication can be done only using CLISH.

### Steps to reconfigure Veritas Data Deduplication

1. Run the **dedupe show** command and make a note of the file system names for both the data and catalog file system.
2. Unconfigure Veritas Data Deduplication using the **dedupe unconfig** command.
3. Run the following command to configure Veritas Data Deduplication:  
# dedupe config <filesystem\_list> <Virtual\_Ip> <username> <catfs\_name>  
where the <catfs\_name> depends on the configuration.  
The credentials should be the same that you used when you previously configured Veritas Data Deduplication.



The catalog file system needs to be mentioned only if the previous Veritas Data Deduplication configuration had a catalog file system.

You can use the **dedupe catgrow <size>** command to grow the catalog file system to the specified size. The default **dedupe grow** command does not modify the size of the catalog file system.

If there are multiple data file systems, specify the file system names in the same order in which you specified them in the **dedupe show** command previously.

From this release, the default size of the stripe unit of the deduplication file system created is 64K.

As a result of the creation of catalog file system and VxFS optimizations, performance improvements can be seen in optimized duplication and restore.

## Changes to log locations for episodic replication

The following changes have been made to the episodic replication log locations:

- The log directory for episodic replication is changed to `/log/VRTSvxfs/replication/log/`.
- The `/var/VRTSvxfs/replication/log` directory is now a symbolic link to the `/log/VRTSvxfs/replication/log/` location.
- The log files in `/var/VRTSvxfs/replication/log` directory will be rotated daily if the file size has reached a size of 1 MB.

## Changes to episodic replication alerts

A new AutoSupport alert for job failure has been added for episodic replication. An AutoSupport alert is raised if a job fails, and the alert is resolved on completion of the job.

## Changes to S3 certificates

Starting from this release, if data endpoints are set, then the first data endpoint of the server is used as the common name (CN) for the certificate. Earlier the CN for the certificate was `s3.<clustername>`.

On setting the data endpoint or admin endpoint, the server certificate is renewed.

For example:

```
Objectaccess> set data_endpoints acc1.sunday,acc2.sunday,acc3.sunday
ACCESS ObjectAccess SUCCESS V-493-10-0 Data endpoints set successfully. Server
certificate renewed.
Objectaccess> set admin_endpoints accadmin1.sunday,accadmin2.sunday
ACCESS ObjectAccess SUCCESS V-493-10-0 Admin endpoints set successfully. Server
certificate renewed
```



## Improvements to the support debuginfo command

In this release, some improvements have been made to the **support debuginfo** command.

- MSDP option has been renamed to VDD.
- A new "default" option has been introduced which does not collect SOS report. Now, SOS report is collected only when reports are collected for all the modules or when the corresponding module (os) is run. This enhancement reduces the time taken during log collection.
- Logging has been introduced for the **debuginfo** command and logs are placed at */log/VRTSnas/log/support\_debuginfo\_\*.log*

## Deprecated functionality

The Server Message Block(SMB) version 1 protocol is deprecated in version 7.4.2.300. SMBv1 clients can no longer connect to CIFS shares. If you use SMBv1, you need to upgrade to SMBv2 or SMBv3 protocol.



## Advance notice for deprecating commands

Veritas plans to deprecate the following commands in future versions. The commands mentioned below are currently supported but no longer recommended as these might become obsolete in future releases.

### User authentication to Active Directory for accessing CIFS shares

The following commands, which are used for authenticating to Active Directory will be deprecated in the next release. An alternative method to configure authentication to Active Directory will be provided.

```
cifs set domain
cifs set domaincontroller
cifs set domainuser
cifs set workgroup
```

### Scale-out file system

The scale-out file system will no longer be supported in the next release, and the following commands, which are used to create and manage the scale-out file system will be deprecated. Starting with next release, only the Cluster File System (CFS) will be supported and guidelines to migrate from a scale-out file system to CFS will be introduced.

```
storage fs create largefs
storage fs create largefs simple
storage fs create largefs mirrored
storage fs create largefs striped
storage fs create largefs mirrored-stripe
storage fs create largefs striped-mirror
objectaccess set fs_type largefs
objectaccess set fs_type largefs ecoded
objectaccess set fs_type largefs simple
objectaccess set fs_type largefs mirrored
objectaccess set fs_type largefs striped
objectaccess set fs_type largefs mirrored-stripe
objectaccess set fs_type largefs striped-mirror
objectaccess group set fs_type largefs
objectaccess group set fs_type largefs ecoded
objectaccess group set fs_type largefs simple
objectaccess group set fs_type largefs mirrored
objectaccess group set fs_type largefs striped
objectaccess group set fs_type largefs mirrored-stripe
objectaccess group set fs_type largefs striped-mirror
```

### Erasure-coded file systems

File systems with erasure-coded layout will no longer be supported in the next release, and the following commands, which are used to create and manage the erasure-coded file systems will be deprecated. The guidelines to migrate erasure-coded file systems to other file system types will be provided in the next release.

```
objectaccess set fs_type ecoded
objectaccess group set fs_type ecoded
storage fs create ecoded
storage fs rebalance
storage fs relayout
```

### Cloud as a tier for scale-out file systems

Configuring cloud as a tier for scale-out file systems will no longer be supported in the next release because the





scale-out file system will be deprecated. Starting with next release, only the Cluster File System (CFS) will be supported. If you have configured cloud as a tier for scale-out file systems, the guidelines to migrate will be provided in the next release.

The following commands, which manage the storage tier operations and collect statistics on data usage in the cloud tier will be deprecated. You can use the alternative commands mentioned below.

**Note:** The commands will be deprecated for both on-premises storage as well as cloud storage.

Command to be deprecated	Recommended command
<code>storage tier addcolumn</code>	<code>storage fs addcolumn</code>
<code>storage tier rmcolumn fs_name</code>	<code>storage fs rmcolumn</code>
<code>storage tier addmirror</code>	<code>storage fs addmirror</code>
<code>storage tier rmmirror</code>	<code>storage fs rmmirror</code>
<code>storage tier setfastresync</code>	<code>storage fs setfastresync</code>
<code>storage tier unsetfastresync</code>	<code>storage fs unsetfastresync</code>
<code>storage tier stats show</code>	Unavailable
<code>storage tier stats monitor</code>	Unavailable
<code>storage tier stats reset</code>	Unavailable

### File system deduplication

File system deduplication will no longer be supported in the next release, and the following commands, which are used to configure and manage file system deduplication will be deprecated:

```
Storage dedup enable
Storage dedup disable
Storage dedup list
Storage dedup start
Storage dedup stop
Storage dedup status
Storage dedup set cpu
Storage dedup set memory
Storage dedup set priority
Storage dedup schedule set
Storage dedup schedule modify
Storage dedup schedule delete
Storage dedup dryrun
Storage dedup remove
```

### Network and support commands

The **tethereal** commands that are used dump and analyze the network traffic will no longer be supported in the next release, and the following commands will be deprecated:

```
Support tethereal export
Support tethereal show
```

The **ntp enable** command internally synchronization of date with the NTP server. Hence, the **ntp sync** command which is used to sync the date with NTP server on all the nodes in the cluster will no longer be supported in the next release and will be deprecated:

```
System ntp sync
```





---

### **NFS-Ganesha server support**

Veritas plans to deprecate the NFS server support to the NFS-Ganesha server (GNFS) in the next release. Steps to switch to the Kernel-based NFS server (KNFS) will be provided in the next release.



## Known issues

### The provisioning for Veritas Data Deduplication is displayed as failed in GUI (IA-23055)

If there are any subtasks which are in running state, the provision for Veritas Data Deduplication fails.

#### Workaround:

Check the task details to make sure that no subtask is in running state and check the reason for the failure. Wait for all the subtasks to get completed before retrying the operation.

### Stale file handle error is displayed during rolling upgrade (IA-27274)

When upgrading from an earlier version to 7.4.2.301, read and write operations to the node where NFS shares are mounted might fail. Some of the NFS shares might be unmounted from the system during the rolling upgrade and the applications accessing these shares might face read and write issues.

After the upgrade is complete, the NFS shares are automatically mounted and the file writes are started automatically.

### Replication setting in episodic replication does not work if the folder name has space in its path or name (IA-27557)

In episodic replication, setting replication is not supported if the folder name has a space in its path or name.

#### Workaround:

Ensure that the replication path does not include any spaces.

### One of the CIFS share may become inaccessible after upgrade (IA-28194)

After you perform an upgrade from version 7.4.2 to 7.4.2.x, one of the CIFS shares may becoming inaccessible.

#### Workaround:

Restart the CIFS server.

### During upgrade, CVM goes into a faulted or partial state on some of the existing node(s) in the cluster and into a partial state on the joiner node(s) (IA-28215)

When you perform an upgrade, one node is removed from the cluster and packages are upgraded on that node. After upgrading the packages, the node is added back to the cluster. In some scenarios, when a transaction is committed after a node is removed from the cluster for upgrade, a vxconfigd core dump is seen on some of the nodes in the cluster. The upgraded node gets added back to the cluster but is stuck in the 'vxconfigd in join' stage due to the vxconfigd core dump on some of the nodes. The nodes where vxconfigd is not running cause the join reconfiguration to hang.

#### Workaround:

1. Get the cluster view using the **vxclustadm nidmap** command from any one of the nodes in cluster (apart from the joiner nodes).



2. Disable the VCS service on all the nodes.  
`# systemctl disable vcs`
3. Bring down the whole cluster by restarting the nodes in cluster.
4. After restart, run the **hastart** command on the master node first. Verify that it has formed a cluster and is the master.
5. Run the **hastart** command on all the other nodes which were part of the cluster, apart from the joiner nodes.
6. Wait for these nodes to join the cluster.
7. Run the **hastart** command on the joiner nodes and check that they have also joined the cluster successfully.
8. Enable the VCS service on all the nodes of the cluster.  
`# systemctl enable vcs`

#### **GNFS fails to start if port 875 is in use by rpcbind process (IA-28384)**

Port 875 is required by the rquota service of GNFS. If rpcbind is listening on port 875, GNFS fails to start and the following error message is displayed:

*Bind\_sockets\_V6 :DISP :WARN :Cannot bind RQUOTA udp6 socket, error 98 (Address already in use)*

#### **Workaround:**

1. Verify if the process using port 875 is rpcbind.  
Run the following command on the node where GNFS failed to start:  
`netstat -anlp | grep 875`
2. Restart rpcbind service.  
Run the following command on the node where GNFS failed to start:  
`systemctl restart rpcbind.service`
3. Login to CLISH and run the autofix command.  
`/opt/VRTSnas/clish/bin/clish -u admin`  
Support> services autofix

#### **During upgrade, the CVM and CFS agents are not stopped (IA-28386)**

When CIFS is configured in the cluster and you perform an upgrade to version 7.4.2.301, the upgrade process may fail to stop the running services. This occurs due to a resource dependency violation.

#### **Workaround:**

1. Run the following command on the cluster node only if the other node got evacuated during the upgrade process  
`/opt/VRTS/install/installaccess -start <node_ip>`
2. Run the following command on the cluster:  
Support> service autofix
3. Stop the CIFS server.



```
CIFS> server stop
```

4. Restart the CIFS server.

```
CIFS> server start
```

### The upgrade operation to version 7.4.2.301 fails if VVR is configured (IA-28466)

If VVR is configured and you perform an upgrade to version 7.4.2.301, the VVR service groups try to stop the vradmin service in order to go offline while some other procedures try to start the vradmin service. The VVR service groups are not able to go offline and this causes the upgrade operation to hang.

#### Workaround:

1. Login to Access CLISH either using the console IP address (ssh master@<consoleIP>) or by using the **su – master** command.
2. Execute the following command before starting the upgrade  

```
replication> continuous service stop
```
3. Start upgrade and wait for the upgrade operation to complete.
4. After the upgrade is complete, login to the Access CLISH and execute the following command:  

```
replication> continuous service start
```

### Rolling upgrade fails when the cluster has space-optimized rollback in online state (IA-28502)

Space-optimized rollbacks are always created with a simple file system layout irrespective of the underlying file system layout. Hence, there is a possibility that the CVM/CFS service group may go into a faulted state during rolling upgrade. If the CVM/CFS service group are in faulted state, then the rolling upgrade fails.

#### Workaround:

Bring all the space-optimized rollbacks to offline state before performing rolling upgrade.

To list the rollbacks:

```
Storage> rollback list
```

To bring all the rollbacks to offline state:

```
Storage> offline <rollback_name>
```

### System may panic with a kernel stack during upgrade operation (IA-28647)

During upgrade, if the I/O workload is in-progress on a bucket backed by a scale-out file system and the S3 server is stopped, the system may panic due to a race condition.

#### Workaround:

Unmount the all scale-out filesystems from the node which is being upgraded.

1. Use the following command to list the mounted scale-out filesystems:



```
/opt/VRTSnas/bin/tfsmntadm display <nodename>
```

2. Unmount all the scale-out filesystems one-by-one using the following command:

```
/opt/VRTSnas/bin/tfsmntadm umount <mount point> <nodename>
```

### **VVR sync does not work after you restart a node (IA-28648)**

If you restart a node, the VVR sync does not work as the ports required for VVR (8989,4145,8199) are not added in the iptables.

#### **Workaround:**

1. Log in to the Veritas Access CLISH.

```
/opt/VRTSnas/clish/bin/clish
```

2. Run the following commands:

```
replication> continuous service stop
```

```
replication> continuous service start
```

3. If replication does not start after performing the above step, run the following command from the primary cluster:

```
replication> continuous start fs_name
```

where fs\_name is the name of the file system.

### **After an upgrade from version 7.4.2 to 7.4.2.301, LDAP and NIS are in disabled state (IA-28673)**

After you perform an upgrade from version 7.4.2 to 7.4.2.301, the LDAP and NIS are in disabled state.

#### **Workaround:**

Execute the following command to enable LDAP and NIS:

```
network> nis enable
```

```
network> ldap enable
```

### **During upgrade, CVM goes into a faulted or partial state on some of the existing node(s) in the cluster and into a partial state on the joiner node(s) (IA-28958)**

During upgrade, one node is removed from the cluster and packages are upgraded on that node. After upgrading the packages, the node is added back to the cluster. In some scenarios, when a transaction is committed after a node is removed from the cluster for upgrading, a vxconfigd core dump is seen on some of the nodes in the cluster. The upgraded node tries to rejoin the cluster, but its join is stuck in the 'vxconfigd in join' stage due to the vxconfigd core dump on some of the nodes. The nodes where vxconfigd is not running cause the join reconfiguration to hang.

#### **Workaround:**

1. Get the cluster view using the **vxclustadm nidmap** command from any one of the nodes in cluster (apart from the joiner nodes).



2. Disable the VCS service on all the nodes.  
`# systemctl disable vcs`
3. Bring down the whole cluster by restarting the nodes in cluster.
4. After restart, run the **hastart** command on the master node first. Verify that it has formed a cluster and is the master.
5. Run the **hastart** command on all the other nodes which were part of the cluster, apart from the joiner nodes.
6. Wait for these nodes to join the cluster.
7. Run the **hastart** command on the joiner nodes and check that they have also joined the cluster successfully.
8. Enable the VCS service on all the nodes of the cluster.  
`# systemctl enable vcs`

**Upgrade fails if VVR is configured and IO processes are ongoing in the filesystem that are configured for VVR (IA-29072)**

VVR upgrade fails if primary and secondary replication links are not in sync or up-to-date.

**Workaround:**

Stop all the IO processes for the file systems on the primary cluster and ensure that both primary and secondary clusters are in sync.

1. Login to Veritas Access CLISH.  
`/opt/VRTSnas/clish/bin/clish -u master`
2. Check if the primary and secondary replication links are up-to-date.  
`replication> continuous status <fs_name>`

The output of the above command should display the data status as "up-to-date".

*Secondary Site Info:*

*Host name 192.168.10.20*

*Configured mode synchronous-override*

*Data status consistent, up-to-date*

If the data status is not "up-to-date", wait for the sync to complete. If both primary and secondary cluster are in-sync, continue the upgrade.

**After upgrading from version 7.4.2 to version 7.4.2.301, the default route entry in the ip rule table on one of the cluster nodes is missing (IA-29632)**

After upgrading to version 7.4.2.301, the default route entry is not present in the ip rule table on one of the cluster nodes. After installing the patch, you need to wait until the restart is complete to view the ip rules in the table. Some of the cluster nodes become accessible after installing the patch. However, the gateway setting is deleted on one of the cluster nodes because of which the node becomes inaccessible.

**Workaround:**



1. Run the below script on the management console node:

```
/opt/VRTSnas/scripts/lib/rpm_install_post.sh 1
```

2. Run the following commands on the node where problem is observed:

```
# hstop -local  
# hstart
```

**Note:** After completing the above steps if the Access Storage Unit (STU) is not reachable from NetBackup, restart the deduplication services using the **dedupe stop** and **dedupe start** commands.

### After upgrade, the UI discovery and the storage fs list <fs name> command may fail (IA-30263)

After the rolling upgrade is complete, the storage fs list <fs name> command fails if the vxdclid process is not in running state. Hence, the UI discovery fails, and no data appears in the UI.

The following error message may appear when you run UI discovery:

```
/opt/VRTSnas/pysnas/bin/isagui_cluster_perf.py --full  
Traceback (most recent call last):  
File "/opt/VRTSnas/pysnas/bin/isagui_cluster_perf.py", line 755, in <module>  
obj.run()  
File "/opt/VRTSnas/pysnas/bin/isagui_cluster_perf.py", line 677, in run  
if (self.init_perf_gather() != 0):  
File "/opt/VRTSnas/pysnas/bin/isagui_cluster_perf.py", line 665, in init_perf_gather  
self.set_up_fs_db()  
File "/opt/VRTSnas/pysnas/bin/isagui_cluster_perf.py", line 519, in set_up_fs_db  
fsData['block_size'] = fs_info['block_size']  
KeyError: 'block_size'
```

### Workaround:

1. Log in to the bash prompt.
2. Manually start the vxdclid service using the following command:  

```
/opt/VRTSsfmh/adm/dclisetup.sh
```





## Fixed issues in this release

This section includes the issues fixed since the last release.

Fixed Issues	Description
IA-17743	The <b>debuginfo</b> command hangs when running the VxExplorer utility.
IA-18206	The <b>debuginfo</b> command fails when logged in as any user other than root user.
IA-23833	After upgrading to 7.4.2.200, Memcached vulnerability fixes are not applied.
IA-24163	Upgrade to 7.4.2.200 fails to bring the cluster online if GNFS is enabled.
IA-24728	Enhancement to Veritas Data Deduplication log collection with Access <b>debuginfo</b> command.
IA-25134	Unable to migrate primary and secondary roles.
IA-25377	Dalloc setting issue during CFS mount.
IA-25591	Storage checkpoints are not excluded from the <code>/etc/updatedb.conf</code> file.
IA-25607	Episodic replication fails with error <code>Connection reset by peer</code> .
IA-25608	If fail over of continuous replication fails, the file system on the primary site is left in an inconsistent state and the fail over destroys all the configuration of the file system.
IA-25611	Full or partial discovery cannot proceed if replication is configured.
IA-25612	No alert generated for replication failure.
IA-26509	SMB server signing issue.
IA-26510	The command <b>storage fs worm set fs-name</b> wrongly prompts for the <b>autocorr</b> parameter when enabling WORM support for a non-WORM file system.
IA-26513	Log rotate does not work as expected.
IA-26903	TLS Vulnerability on WebServer port 14161
IA-27287	Unable to migrate primary and secondary roles.
IA-27571	Virtual IPs are not balanced after all nodes are restarted
IA-28572	During rolling upgrade with NFS v4 using NFS lease, system may experience hang

For information on all the other features, refer to the Veritas Access 7.4.2 documentation, which is available online. The latest version of the product documentation is available on the [SORT](#) website.



---

## 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 [veritas.com](https://veritas.com) or follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

---

Veritas World  
Headquarters 2625  
Augustine Drive  
Santa Clara, CA 95054  
+1 (866) 837 4827  
[www.veritas.com](https://www.veritas.com)

For specific country offices  
and contact numbers,  
please visit our website.

**VERITAS**<sup>™</sup>  
The truth in information.