

Testing DR Readiness

Overview

This paper confronts the problem of ensuring Disaster Recovery (DR) success, and reducing the expense of DR test coverage across the entire enterprise.

Problem

How can you guarantee success when recovering systems?

Natural disasters, terrorist threats, and the normal life cycle of hardware have lately raised the profile of Disaster Recovery. CEOs and CIOs want certainty that the enterprise is covered. DR testing has thus become a crucial part of the enterprise's operations.

But DR testing is expensive and time consuming. A good DR test requires:

- Finding and setting up a test site
- Finding hardware to re-create the backup server and library
- Finding hardware to re-create the targeted machines
- Moving the right set of backup media to the test site
- Recovering the backup server
- Recovering each target machine's operating system
- Recovering each target machine's filesystem data
- Ensuring each target machine starts and functions correctly.

Which machines do you pick for an expensive DR test?

All this takes a lot of time and money, often six figures for a large organization. So a DR manager typically picks a few "mission-critical" machines and does a full DR test on them quarterly or annually.

But the machines *not* considered "mission-critical" probably *are*. Their outage will lose the company time and money, maybe a *lot* of time and money. The DR manager doesn't really know how much damage a given machine's outage will cause. So which machines are worthy of a full DR test?

It's easy to point to five boxes and say, "These are mission-critical!" The unspoken implication is "None of the others are very important." That's not true.

WHITE PAPER

Boosting Coverage

How can one get test coverage for “all” machines in the enterprise?

Ideally, the DR manager would schedule a full DR test on every machine every day. Of course, the resources required for such a setup would be extravagant.

Random sampling augments full DR tests.

A more pragmatic approach augments normal DR tests with statistical sampling of a few randomly selected files from each machine. This sampling technique can

- **raise confidence:** if a *few* files restore successfully, then it's probable that *all* files will.
- **uncover problems:** if a file fails to restore, there is a root cause (for example, missing tape media). Remedying this root cause often prevents similar problems on hundreds of other machines.

The combination of full DR tests for *some* machines deemed “mission-critical” and a statistical sampling of *all* machines gives management greater confidence in the overall preparedness of the enterprise.

Caveats

What can go wrong?

Such a solution must tread lightly. It must:

Testing should not disrupt operations.

- **Limit the size of restores.**
Restoring 500GB of data as a test puts undue demands on the network and the backup server.
- **Do no harm to backup server operations.**
Backup servers, especially TSM servers, get busy at unpredictable times when internal “housekeeping” jobs run. Restores, high-priority by definition, must be careful not to pre-empt these operations.
- **Do no harm to the tested machines.**
Machines in a production environment must obey stringent change control guidelines. Managers may not casually install new software, perform test restores, or do anything else which might impact operations.
- **Maintain confidentiality.**
Randomly chosen files may include confidential information, and should not remain on the testing platform.
- **Address large consolidated servers.**
One machine may host many applications, and each may have a different backup scheme. A successful restore of a file from one application doesn't predict success on another.

WHITE PAPER

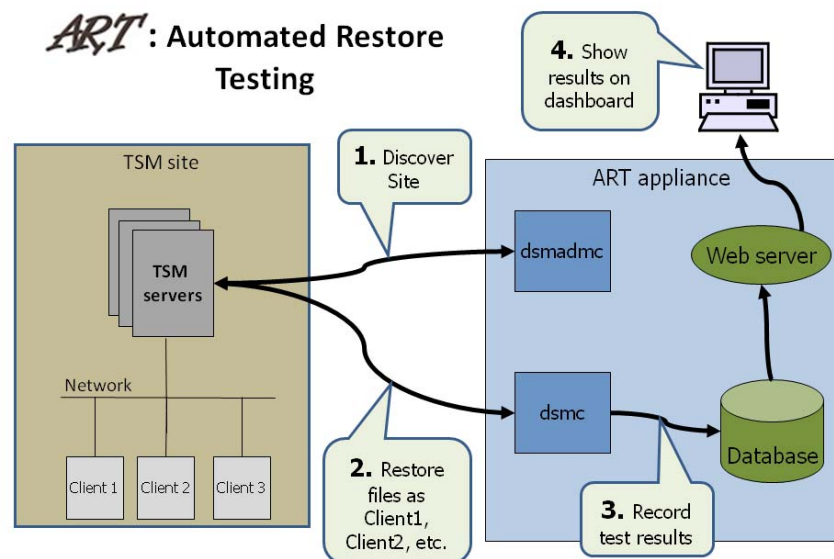
Solution

ART: Automated Restore Testing for Tivoli Storage Manager

ART, TSMworks' Automated Restore Testing appliance, implements such a sampling solution for the Tivoli Storage Manager environment. ART addresses the above concerns.

- **Limit the size of restores.**
The ART Virtual Appliance's default setup limits its samples to 10MB per file, and one file per filesystem.
- **Do no harm to the normal operations.**
ART tests for the availability of system and library resources prior to each test restore.
- **Do no harm to the tested machines.**
ART uses TSM's "grant proxy" system to masquerade as each production machine. The production machine's restored files come to ART. The production machine itself is completely untouched by the restore.
- **Maintain confidentiality.**
ART immediately deletes restored files.
- **Address large consolidated servers.**
ART auto-discovers all nodes on the TSM servers it monitors. Further, it discovers all filesystems on those nodes. The filesystem is the usual separator for applications on consolidated servers: one application may use many filesystems, but rarely would two applications share the same filesystem. Thus a filesystem-by-filesystem test effectively samples each application.

How ART manages to meet these requirements.



WHITE PAPER

Surprising Results

Real-world tests uncover problems.

Since across-the-board testing is usually impossible, managers are forced to pretend that it's not needed. The assumption (often unspoken) becomes "If the backups all succeed, then the restores will succeed too."

But ART has sampled several large TSM sites with results that surprised the administrators. It found:

*Users find
problems they
didn't think they
had.*

- machines that had never been placed on a schedule
- filesystems that had not been backed up for several months
- tape volumes that had not been checked in correctly, thus unavailable
- tape library too busy for restores (without disrupting daily operations)
- restores that took longer than five minutes (user would perceive as error).

Repairing these problems was easy, and is a clear win in terms of enterprise reliability.

Conclusion

As an adjunct to traditional DR testing, statistical sampling of restorability offers great benefits. Because it touches every computer in the enterprise, sampling is effective at uncovering problems that would otherwise go undetected for months, and raises managers' confidence in the robustness of their DR preparedness.

About TSMworks

Based in the Research Triangle area of North Carolina, TSMworks develops leading-edge software that addresses problems specific to the Tivoli Storage Manager market. More information is available at www.tsmworks.com.