



Pre-Testing Disaster Recovery & Business Continuity Plans: The Network Simulation Solution

A Shunra Software White Paper

Executive summary

Effective IT disaster recovery (DR) and business continuity planning is essential for every business. All businesses depend on their IT services for moment-to-moment operations. So they must all take measures to ensure that those services are not disrupted due to a natural or man-made disaster.

However, because IT environments have become so complex – and because IT budgets are stretched so thinly – IT organizations face real challenges as they attempt to safeguard the business against such disasters. These challenges include:

1. **Ensuring that planned DR measures will actually meet the needs of the business if and when they are called upon**
2. **Determining exactly when such measures should, in fact, be activated**
3. **Documenting to senior management, auditors, insurers, and/or regulators that best efforts have been made to protect the business from a full range of potentially disruptive eventualities**

All of these challenges can be addressed by pre-testing DR plans in a simulated network environment. By performing these simulations, DR planners can cost-effectively assess the end-to-end performance of critical applications under any projected condition. Equipped with this insight, DR planners can make more informed decisions about how to best ensure business continuity – and can fully justify those decisions to any appropriate third parties.

Why is IT disaster recovery planning so hard?

While events such as the attacks of 9/11 and Hurricane Katrina have called greater attention to the possibility of catastrophic business interruption, businesses have always had to plan for the worst. IT's role in this planning has grown over the years as IT services have become an increasingly critical component in day-to-day business operations. In fact, for many businesses survival depends almost entirely on the ability to provide at least some minimal number of core IT services to at least some minimal number of end-users – regardless of where those end-users may be.

Several factors make it difficult to plan for the delivery of core IT services to end-users who may or may not be working at their usual locations:

“Live” testing is often impractical. DR measures often include use of “hot sites,” “cold sites,” alternative network service providers, and other resources that can't be tested “live” with any frequency – if at all. This makes it difficult if not impossible to re-assess such measures as applications are added or other types of changes are made to the enterprise computing environment.



Business continuity depends on application performance. It's not enough to simply give end-users in some remote location networked access to a server somewhere. They also have to be able to actually use the application they're accessing. If the application behaves too sluggishly – or network latency causes it not to perform properly at all – the disaster recovery plan will fail.

Application performance can be difficult to predict. Different applications respond in much different ways to bandwidth constraints, added latency, and intermittent connectivity. Without real insight into the idiosyncrasies of each critical application, DR planners may under- or over-provision contingency infrastructure.

There are so many different contingencies to consider. IT has to be prepared for everything from an accidental cable cut to the complete destruction of the corporate data center. The sheer number of potential disaster scenarios can tax the ability of even the largest IT organization to assess them all and formulate appropriate recovery plans.

DR planning resources are limited. Most IT organizations already have their hands full just maintaining the systems they already have in place and supporting strategic, high-ROI technology initiatives. This leaves little in the way of financial and human resources for a “necessary evil” like DR.

It has to work right the first time. There's no room for error when it comes to DR. If IT discovers a problem with a typical system upgrade or new piece of hardware it can quickly roll things back to how they were before. But there's nothing to roll back to in the event of a disaster. So IT has to have a high level of confidence in its contingency plans.

For these reasons and others, IT organizations often struggle to achieve confidence in their ability to fully protect the business from all possible contingencies. These factors can also make it difficult for IT to satisfy the demands of upper management and/or outside auditors for proof that adequate planning and testing has been performed. That's why many IT organizations are actively seeking a better way to perform and manage their DR planning.

The simulation solution

One powerful way to improve DR planning is to pre-test disaster recovery scenarios using network simulation. Network simulation in a lab environment provides a safe and flexible means of observing end-to-end application performance under a full range of possible scenarios. It does this by accurately simulating conditions as they exist in the current production network environment – or as they might exist in a disruption or recovery scenario – so IT staff can observe, analyze, and experiment with applications and infrastructure.



Effective simulation must encompass all factors relevant to application performance across the network, including:

- All network links and their impairments (physical distance and associated latency, bandwidth, jitter, packet loss, CIR, QoS/MPLS classification schemes, etc.)
- The number and distribution of end-users at each remote location
- Application traffic loads
- Specific scenario events, such as the disruption of a specific network link

Ideally, it should be possible to directly import these attributes from the current production network and/or any provisioned DR resources – such as remote “hot” or “cold” site facilities and failover network circuits. It should also be possible to modify environmental attributes at will in the simulated environment in order to assess how changes in conditions – such as increased utilization or intermittent connectivity – will affect application performance.

It’s important to note that testing in a simulated network environment doesn’t just enable IT staff to measure functional attributes such as the utilization of a particular link or the congestion at a certain switch. By connecting the “virtual” simulated network to production servers running live applications, IT staff can directly observe and accurately measure the end-user’s experience at the desktop for each and every remote location. In the case of VoIP, phones can be connected to the testing environment so that testers can actually hear for themselves what a call will sound like given the contingency conditions being simulated.

This ability to accurately assess end-to-end application performance and to create a limitless range of “what-if” scenarios makes simulation-based testing an exceptionally powerful technology for DR planners.

Ensuring the adequacy of DR plans in simulated test environments

The applicability of network simulation to DR planning is fairly obvious. Instead of having to actually “go live” with any given recovery scenario, DR planners can readily simulate those scenarios in the lab. This is far less expensive and much more convenient than activation of DR sites or backup network circuits.



Using this kind of testing environment, DR planners can make a variety of determinations about projected contingency scenarios, including:

- Does the DR plan provide adequate application performance to the minimum acceptable number of end-users?
- Can the planned recovery solution meet the needs of all end-users at all remote locations?
- What is the maximum number of users the solution can support without bringing performance below acceptable levels?
- Are there any specific applications that the planned recovery solution can't adequately support?
- Will additional bandwidth improve performance? Or are performance issues primarily related to other network characteristics, such as distance and latency? Does this mean that the hot site or secondary data center should be re-located?
- Will an alternative architecture, such as the use of Citrix servers, help performance? How many Citrix servers would be required to support a given number of end-users?

Of course, the answers to these questions are likely to change over time. As a business grows, for example, the minimum number of end-users it needs to get through a typical day is likely to increase. The new and/or modified applications that get rolled out may have different performance characteristics – which mean they will have to be tested under a full range of DR conditions. With simulation technology, it is relatively easy to test and re-validate DR plans in light of these changes. Without it, IT faces the choice of either re-running costly production “fire drills” or waiting several months for the next scheduled one – potentially leaving the business vulnerable until then.

From a business perspective, it's also important to note that this kind of testing eliminates the need to over-provision DR resources. With budgets tight, IT organizations can ill afford to spend more on DR infrastructure than they have to. But many businesses do in fact over-spend in order to make sure that their DR infrastructure will hold up under pressure. By performing appropriate testing in a virtual environment, DR planners can effectively “right-size” such infrastructure and avoid this inefficient resource allocation.

Determining when to “pull the trigger”

When a flood or regional power failure forces the abandonment of a primary facility, then there's no question that DR plans have to be activated immediately. But what if a less drastic problem occurs? What if a router starts to experience problems, but doesn't actually fail? At what point should the decision be made to cut over to an alternative network link? Should non-critical applications be turned off to preserve the performance of critical ones? And on what basis should such a decision be made?



These decisions can be difficult to make on the spur of the moment without adequate preparation. Many IT organizations don't even know how these types of less-than-disastrous events will impact their critical IT services until they actually occur. The resulting confusion can wind up creating a business interruption that's nearly as serious as an actual disaster.

Sometimes, such decisions have to be made on the basis of cost as well as application performance. The use of a secondary network provider's infrastructure for failover, for example, may be very expensive. In these cases, a line-of-business manager or other executive may have to make the call – rather than someone from IT.

Simulation technology can be extremely useful for dealing with these kinds of “near-disaster” scenarios by showing exactly how infrastructure impairments will impact critical IT services – and how various alternative tactics can potentially remediate those impairments over the short term.

Simulation testing also helps IT organizations and the businesses they support prepare for such decisions by letting end-users and managers experience impaired application performance first-hand in the lab. It's probably not very useful to ask a non-technical person “Should we cut over to our alternative network service provider if latency on this circuit jumps from 5 milliseconds to 100 milliseconds?” A more reasonable approach would be to bring them to the lab and show them what the response of their primary applications will look like under the impaired condition – and to then ask “Is it worth \$20,000 to avoid this sluggishness on our end-users' desktops for 48 hours?” Such a demonstration prepares managers for making appropriate decisions on short notice, so that the impact of such problems on the business's bottom-line can be minimized.

Documenting DR diligence

IT organizations may be quite diligent in their DR planning. But if they are not able to fully document that diligence to others, they may have difficulty responding to the inquiries of upper management and external auditors. In fact, if they can't provide ready documentation validating their preparedness for multiple disaster scenarios, they may be asked to demonstrate the adequacy of their plans in ways that are both time-consuming and costly.

Documentation of DR diligence may become especially important in the event that a disaster occurs and DR plans yield results that don't fully meet the expectations of executives, board members, or stockholders.

In both the pre- and post-disaster situations, DR planners must be able to show:

- that specific types of disaster scenarios have been considered
- that DR plans for these various scenarios were appropriately tested
- that to-the-desktop performance of all current critical applications was assessed
- that tests were performed based on appropriate utilization levels by an appropriate number of end-users
- that outcomes were approved by authorized managers



This kind of documentation is much easier to generate and manage in the context of a controlled virtual testing environment than it is with conventional “live” production trials. Under a conventional model, for example, the only way to show a business manager how an application will perform on a desktop in Cincinnati, it would be necessary to fly the manager to Cincinnati. If the manager wasn’t satisfied with the result, it would be necessary to modify the DR plan, implement it, and then fly the manager there again.

With network simulation technology, on the other hand, such demonstrations can be performed right at headquarters. If results aren’t satisfactory, parameters such as distance and bandwidth can simply be adjusted in the lab – and an improved demonstration can be shown. Those parameters can then be applied to actual DR provisioning.

Key benefits of Network Simulation for DR planners

Simulation testing supports good DR planning in many ways. It helps ensure the accuracy and adequacy of DR implementations. It helps IT and business managers understand exactly when they may need to put DR plans into action. And it enables full documentation of DR planning efforts.

By taking full advantage of this technology across all DR activities, IT organizations and the businesses they support can reap a variety of significant benefits, including:

Reduced risk of catastrophic business interruption. By enabling DR planners to more readily consider various disaster scenarios and carefully develop effective continuity strategies, simulation testing helps ensure the ongoing operation of the business.

Reduced DR provisioning costs. By enabling DR planners to “right-size” infrastructure such as network connections and servers, simulation testing eliminates the need to over-spend on DR provisioning.

Faster response to change. By allowing DR planners to quickly assess the impact of technological and organizational change on DR plans, simulation testing eliminates potential “windows of vulnerability” that can arise when new applications, new locations, new users, and other modifications are introduced into the IT environment.

Reduced DR planning costs. By making it easier for DR planners to experiment with various scenarios and solutions, simulation testing reduces dependency on production trials and allows IT organizations to avoid associated charges from DR vendors.

Improved defensibility of DR plans. By providing a controlled environment for validating and documenting DR scenarios, simulation testing makes it easier for DR teams to prove their diligence to upper management and external auditors.



As IT environments become increasingly complex – and as businesses become increasingly dependent on the services delivered across those IT environments – DR planning will become an increasingly critical and challenging discipline. DR planners should therefore strongly consider taking advantage of today’s simulation testing solutions. By doing so, they will be able to better protect the business, while at the same time taking cost and hassle out of core DR processes.

About Shunra

Shunra is the pioneer and market leader in predicting how business applications and network services will perform for remote end-users - before deployment. The Shunra VE solutions simulate any production network environment in the pre-deployment lab. It delivers a powerful, flexible and easy way to test the performance of applications or network equipment under a wide variety of network impairments - as if they were running in a real production environment. Through this process users understand the impact that the network and applications have on each other’s performance and on remote end-users’ experiences, and uncover and resolve production-related problems before rollout.

Over 1,700 leading companies worldwide use Shunra’s award-winning network simulation solutions – including 3M, Boeing, Cisco, Dow Chemical, EMC, FedEx, General Motors, JPMorgan Chase, Kelly Services, Merrill Lynch, Motorola, Nestle, Pitney Bowes, and Vodafone. Shunra’s headquarters are located in New York City, and the company has a major R&D center in Israel, and offices in the UK and Sweden. Shunra is also supported through a global network of channel partners. For more information, please visit www.shunra.com.

