A Project Methodology for Disaster Recovery Testing in a Server Based Environment

Russ Neal

Project execution methodologies for conducting disaster recovery testing for a server-based environment at a hotsite are not readily available. If it is the first time an organization is testing the recovery of several applications and a network cutover, devising or locating a methodology for recovery testing at a hotsite can seem overwhelming. The majority of methodologies range from on-site failure and recovery, through mirroring or clustering techniques, to remote recovery through mirroring and or replication, rather than the more predominant tape backup and restoration at the hotsite. Frequently, the methodologies do not detail what steps to take to prepare for, and to execute, applications and network recovery at a hotsite, or how to conduct debriefing of the test results back to the testing team and senior management. This is not to say there are not good methodologies — they are just difficult to find.

The purpose of this article is to present a reasonably detailed methodology for conducting the disaster recovery testing over four phases: Phase 1: Testing Readiness Assessment; Phase 2: Testing Preparation; Phase 3: Testing Execution; and Phase 4: Testing Debriefing. The format for presenting these phases follows this structure: The Methodology of the Testing Project, Results and Analysis of the Testing, and Summary and Recommendations.

The target audience for this article is broad. IT managers, security managers, disaster recovery administrators, facilities managers, physical security managers, product line managers, network engineers, and consultants in any of these areas should gain a reasonable understanding of a sound testing approach.

The methodology is applicable to private-sector, public-sector, and nonprofit organizations. Throughout the discussion, the term "client" is used to represent a given organization of any type. This term does not automatically imply the project manager as consultant; rather, it denotes that the organization is either an internal client for a resident project manager, or an external client for a consulting project manager.

THE METHODOLOGY OF THE TESTING PROJECT

There are several components to the project that are assumed to be in place at its start.

Russ Neal is a Senior Network Systems Consultant, Project Manager, at International Network Services.

First, it is assumed that a project manager has been appointed. Second, despite the management sponsors stating that the client was ready to conduct preparation for and subsequent execution of the test, if a Testing Readiness Assessment shows significant preparation weaknesses, those weaknesses must be corrected prior to the next phase of testing preparation. That a hotsite contract is in place is a given, because scouting for and contracting with a hotsite company is a very large project in itself. With a contract in place, it can be assumed that all equipment required for a test or disaster has been detailed in the contract schedules.

A fourth component is that backup circuits are in place from whichever sites are to be linked to the hotsite during the test and the hotsite itself. Typically, these are major hub sites that connect nationally or internationally to several small sites such that when the backup circuit is activated during testing or a disaster, the hub site is cutover to the hotsite. It is also assumed that management has reserved a testing window (e.g., 48 hours) and the test time reserved falls no earlier than 90 calendar days from the start of the project.

A further assumption is that a Business Impact Analysis (BIA) was conducted prior to the testing project, and that Recovery Time Objectives for the business functions, network and support devices, and the network itself were documented. Moreover, recovery scripts for the applications, the network, and infrastructure should exist, but the larger disaster recovery plan itself is not a requisite for purposes of the testing methodology portrayed here.

PHASE 1: CONDUCTING A TESTING READINESS ASSESSMENT

A Testing Readiness Assessment must be conducted prior to commencing with the formal preparation for the execution of the test itself. This is conducted for several reasons. The foremost reason is to ensure that if certain processes, recovery scripts, network pieces, hardware, applications, or whatever else is required to conduct the test, has X days worth of work to be done by either the client or the hotsite before the test can be conducted, and the test is in 90 days (generally the minimum time required for a testing methodology), and X > 90 days by a significant amount, then X must be either eliminated from the scope of the test or the test dates must be moved out.

If a readiness assessment was not conducted, and the project manager is 20 days into the preparation for the test execution and X is discovered, it becomes problematic. First, if X must be eliminated, then there is a good chance that much of the test is compromised because other functions depend on X. Second, it is not easy to move a test date out because test time is an increasingly scarce and expensive hotsite commodity. Ideally, a Testing Readiness Assessment would have been performed prior to test dates having been reserved. However, this is not always or easily done. The secondary reason for a Testing Readiness Assessment is to identify any major gaps that must be focused on which, despite not having the X > 90 days problem, still require large amounts of group work, and to identify these gaps early in the project.

The methodology to determine testing readiness should entail several steps, the first of which is to develop a structured questionnaire and interview format that enables open-ended responses from management and team members. Examples of such questions are described later in this section. The questions must probe the testing facets as to whether or not the client is ready to move to the formal testing project.

The next step is to conduct structured interviews. These must be held with each of the groups that will participate in the test. When the findings are published, responses should be recorded by each interview question. A representative list of IT groups that would be involved in the hotsite testing — and used as examples in the methodology discussion — of several server-based applications, both NT/Windows 2000 and UNIX based, and a network cutover are given here. For illustrative purposes, the applications,

A Testing Readiness Assessment must be conducted prior to commencing with the formal preparation for the execution of the test itself.

type of network circuits, type of databases, client desktop systems, and Web services will be identified later when addressing the steps involved in executing the test. Results of the interviews are then analyzed, and gaps, trends, and readiness are reported for each group.

- □ *Network Engineering*: the group that designs, troubleshoots, and maintains the corporate WAN and the LAN.
- □ *Applications Development and Support*: a collection of sub-groups that develops and maintains their respective in-house applications programs and third-party vendor programs.
- □ *Desktop Services*: the group that maintains workstations, including images, clients necessary to access server-based software, and company laptops.
- □ *Operations*: the group that manages the network, production control functions, the help desk, and tape backups and rotation.
- □ *Applications Sevices, Windows-Based Servers*: the group that manages and supports NT- and Windows 2000-based servers, including the firewall and infrastructure, as well as the server interfacing to the LAN.
- □ *Applications Services, UNIX-Based Servers*: the group that manages and supports all UNIX-based servers and the server interfacing to the LAN.
- □ *Database Systems*: the group that develops, manages, and supports all database structures which support the various applications, both Oracle and SQL.
- □ *Web Services*: the group that develops and manages the corporate Web site and provides support to the applications subgroups that have applications that must interface to the corporate Web site.

The readiness questions and the implications for the questions are detailed here. This list does not exhaust all of the major issues that might be encountered, but it does present issues that tend to come up during the readiness interview. (Other typical gaps and administrative tasks are addressed later.):

- □ What are the applications that are to be tested?
  - – Are recovery scripts in place, and are they current?
  - – Is the data center undergoing migrations to upgrades to the OS or the applications that must be timed with when backups are taken that will be used for the test?
  - – Are there licensing issues for running the applications at the hotsite?
  - – Are applications not planned to be tested required by applications that are planned to be tested?
  - – Are there internal or external service level agreements (SLAs) stating recovery time objectives that are shorter than the planned recovery time during a test?
  - – Is a change control system in place that will identify changes that affect testing?
  - – How many workstations will be required for recovery of the applications?
  - – Is there an application that is to be retired shortly after the test? If yes, why test it when another application that is not planned on being retired soon could possibly be substituted?
- □ What operating systems are required for each server to be tested?
  - – Can the hotsite equipment support the OS versions?
  - – Can the hotsite do any pre-loading of the OS for the client?
  - – If pre-loading is allowed, will it contaminate any specially required application-OS load sequencing?
  - – How will disk arrays be set up at the hotsite?
  - – Is disk sizing at the hotsite, as reported in the hotsite contract schedules of equipment, based on a before or after RAID basis?
- □ What hardware and related resources are required for the test, are they in the contract schedules, and are they available for the time of the test?

- The hotsite cannot suddenly adjust to a new demand for hardware or other resources made known by the team two weeks prior to the test due to upgrading. Typically, the hotsite needs 60 to 90 days lead-time. If hardware upgrades are planned for an application that also is to be tested, and the upgrade falls close to the recovery test, a forced decision may be required that results in either the upgrade to go into production and not be tested, or the upgrade to not go into production and the current application be tested.
- Will the client have servers that are owned by the client that must be brought to the hotsite or the hotsite company's associated end-user site (often a separate facility that could be thousands of miles away from the hotsite and that is connected by a LAN bridge to the hotsite)? Example: infrastructure composed of WINS, DNS, DHCP, and AD that in a real disaster would be replicated to the hotsite versus recovered at the hotsite — how will this replication be simulated?
- Are data center tape backup devices migrating from, for example, DLT to AIT, such that by the time the test is conducted, tape backups are not inadvertently made from a DLT device and an impossible restore is attempted on an AIT device?

□ Are the network requirements and strategy in place to conduct the test?

- If disaster recovery circuits from business sites or hub sites to the hotsite are not in place, expect an average time of 45 to 60 days to get the circuits installed to a hotsite located in most U.S. regions once the circuit order is placed. Otherwise, if the hotsite is in New York or New Jersey, allow 90 days.
- How deep into the hotsite network will the team test? In the majority of cases, the hotsite cannot use the same private IP addressing scheme as the business or hub site that will connect to the hotsite if the connecting business or hub site is still linked to the production network at the time of testing and the test plan calls for network access not only to the hotsite WAN router, but also into the LAN switch. If the hub site or business site is isolated from the production network, then testing can run from the given client site, into the hotsite LAN, and thereby access the servers; otherwise, the network test must stop at the hotsite WAN router (assuming the WAN router is disconnected from the LAN).
- If a public corporate Web site is to be tested for recovery, have special testing-specific external IP addresses been assigned by the hotsite for use by the client's ISP, along with a testing-specific Web site name? For example, if the regular URL for the client is *client.com*, the test URL might be *clientdr.com*.
- Have some applications used hard-coded IP addresses that become invalid at the time of testing?
- Do the disaster recovery circuits have adequate bandwidth?

□ Are backup procedures in place to support testing?

- Are all the backups required for the test actually getting backed up and rotated offsite (or at least backed up and available for use at the test)?
- How often is the client backing up? If the client is backing up only weekly, then the client will be testing with, at a minimum, week-old tapes. Therefore, with a 100 percent successful restore at the backup site, the best conclusion that can be drawn is that the client is capable of restoring up to where it was one week ago and, therefore, had it been a real recovery, the remaining data must be manually re-entered.
- Are traditional disaster recovery resources to be used in the test and at the time of a disaster being stored offsite and updated? Examples include installation media, vendor manuals,

internal recovery scripts, other forms of documentation, and a CD burner. (There are certainly many more items that should be held offsite outside of tapes, media, and documentation. The list runs from credit cards, to ready-to-be cashed checks and petty cash, to disaster recovery manuals. Our concern here is limited to the items required for testing.)

- □ Are restore procedures in support of disaster recovery testing in place?
  - Do restore scripts exist (or can they be written prior to the time of the test) that describe in detail how to recover the OS, service packs and security patches, the database kernel, the database itself, data restoration, the application, and how to conduct subsequent application testing?
- □ Is an offsite storage facility used, and are procedures in place for pulling the tapes and other resources required for testing?
  - What company is used as the off-site storage facility? Are the tapes required for testing reported to be easily identifiable and clearly labeled? Does the client want to actually use the only set of backup tapes from a certain production date for the test, or should special jobs be run to make extra backups to be used only for the test?
- □ Will customers, business partners, or internal end users be participants in the test?
  - Expectations must be set up and managed very early. A solid technical, systematic approach to the test as managed by the Information Technology Division could be seriously skewed if a high-ranking internal user from a product line division has unchecked decision-making power as to what is to be tested and how the test is to proceed or be managed. Furthermore, these testing partners must be pulled into the test preparation as early as possible.
- □ Is a disaster recovery plan in place?
  - It would seem obvious that a full disaster recovery *plan* must be in place to conduct the testing described here, but that is not true. In reality, recovery teams are drafted as long as recovery *scripts* are in place, or can be in place by the time of the test. The difference is that the plan drives how the first, second, and third levels of contact come about; who the team members are; what they do and in what sequence; how to get to the emergency control center; how to pull the tapes required; and how to activate and get to the hotsite. The scripts describe how to restore the OS, databases, applications, and the network.

So why ask if a plan is in place? To determine if the recovery scripts are part of a master plan that has a unity of purpose to it. If there is no plan, the test project can continue, but there should be a strong thrust toward developing the recovery plan as soon as the test is completed.

The above questions will not pull out all issues the project manager would want to be aware of prior to starting the test preparation phase. Nonetheless, the questions will prevent late discoveries of missing items that could cause the test to be delayed.

As noted above, there are other gaps and administrative tasks that may come about during the subsequent test preparation phase of the project. Representative of these are items such as how turnkey is the hotsite? When the client arrives, will the LAN, its switch, and all servers be ready to go, without the client having to run cables and make adjustments? Which testing participants can bring and use their own company laptops, and which require a hotsite desktop that the client's technicians must image themselves? With respect to equipment, has the detailed listing of all equipment to be recovered as that equipment is in production been cross-indexed to the detailed listing of all equipment that is actually available at the hotsite for testing and recovery?

There are two reasons for the equipment cross-indexing. First, it is very easy for hardware to have been upgraded without the hotsite contract being revised to support the change. Second, understand that there is the equipment used in production, as well as the equipment that the client will get at the hotsite at the time of a disaster or testing. However, there is one qualification here: the equipment the client gets for the test may not be exactly what the hotsite promises the client in the contract for a disaster. In most instances, the client will be fine; but it is not until around two weeks prior to the test, at the hotsite's pre-test conference call, that exactly what equipment the client will be assigned at the hotsite for the test is revealed. The differences may be subtle on the surface, but in reality can have major consequences. Restoring an application on a dual processor server in production may have always gone smoothly, but restoring the same application on the same server with quad processors may not go smoothly.

Other gaps to be aware of include: Will there will be enough ports on the hotsite LAN switch for all the servers and workstations getting restored? Are there are servers that have applications so data-bound that the server requires gigabit CPU-to-disk connectivity? Is the ratio of servers to KVMs (keyboard-video-mouse) too high? This would result in those doing the server recovery stepping over each other and having to unnecessarily wait. Does the firewall at the hotsite have enough NIC cards?

PHASE 2: DISASTER RECOVERY TESTING PREPARATION

The second phase is the heart of the methodology. It is exemplified by four primary activities: (1) weekly status meetings for the application and associated teams, and separate weekly status meetings for the network team; (2) the building of a project administration task plan and the completion of those tasks; (3) the creation, tracking, and resolution of all open issues identified during the Testing Readiness Assessment and during Phase 2 itself; and (4) any required ad hoc sub-group meetings. These testing preparation activities take the form of the steps addressed here.

Step 1: Set the Testing Objectives

To illustrate Phase 2, assume that the fictitious client is an Idaho-headquartered international cosmetics manufacturing company, and its data center is centralized and located on the same campus as the headquarters office. There is one large domestic sales office in San Diego, one domestic hub site in Nevada, and one European hub site in Scotland; and each of these three sites connects to the data center through their respective T1 Frame Relay PVCs.

Several smaller domestic sales offices connect through fractional T1 Frame Relay to the Nevada hub site, and several smaller European sales offices connect through fractional T1 Frame Relay to the Scotland hub site. The client uses an ISP for its corporate Web site. A high-level diagram of this network that is to be recovered is provided in Exhibit 1. With reference to the diagram, the network engineering testing objectives are described here. The testing objectives of the other teams are then presented.

Network Engineering Group Objectives (Team at East Coast Hotsite). The first objective of the network team is to cutover the San Diego sales office to a disaster recovery T1 Frame Relay circuit that links from the sales office to the hotsite through the IXC's (long-haul carrier's) Frame Relay cloud over to a POP on the hotsite company's national network, and then into the main hotsite. The IXC will be scheduled several days in advance to participate in the test. Because the client will be running a 10.1.x.x on the hotsite LAN, which is the same as what is run on the production LAN, the sales office must be temporarily cut off from production during the network portion of the test. Ping the hotsite LAN switch, run trace routes, and then have an end user access the various applications.

EXHIBIT 1 Network Cutover Test

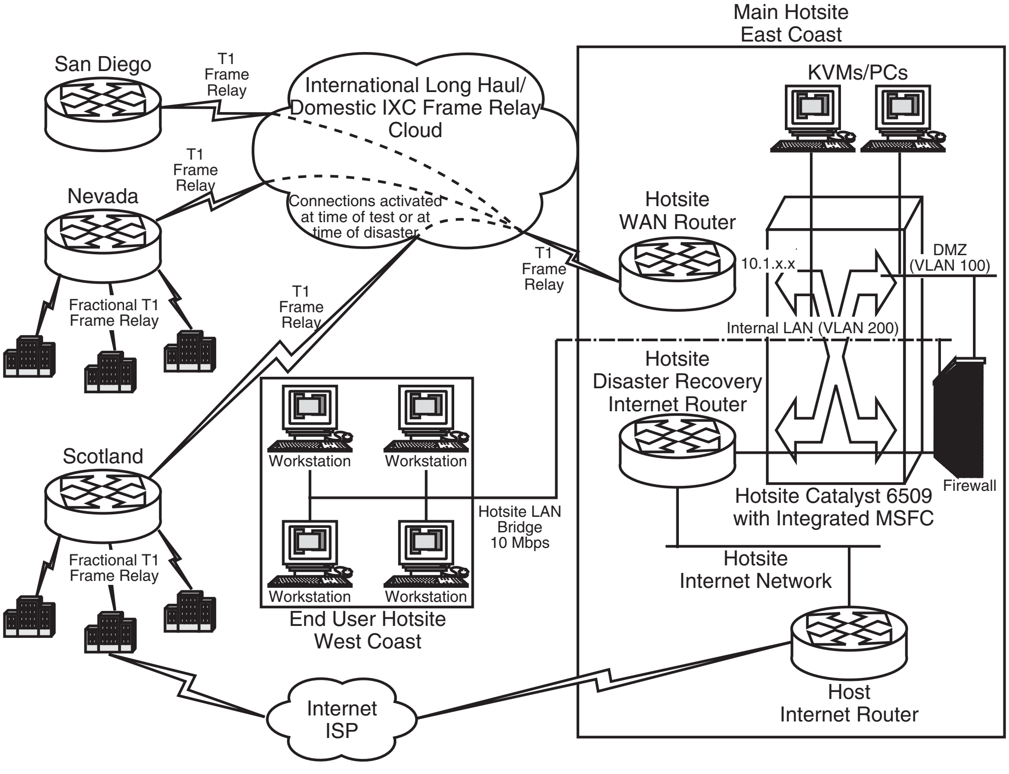
Main Hotsite East Coast
San Diego
T1 Frame Relay
International Long Haul/ Domestic IXC Frame Relay Cloud
KVMs/PCs
Nevada
T1 Frame Relay
Connections activated at time of test or at time of disaster
Hotsite WAN Router
10.1.x.x
DMZ (VLAN 100)
T1 Frame Relay
Fractional T1 Frame Relay
T1 Frame Relay
Internal LAN (VLAN 200)
Hotsite Disaster Recovery Internet Router
Scotland
Workstation
Workstation
Firewall
Hotsite Catalyst 6509 with Integrated MSFC
Hotsite LAN Bridge 10 Mbps
Fractional T1 Frame Relay
Hotsite Internet Network
Workstation
Workstation
End User Hotsite West Coast
Internet ISP
Host Internet Router

A second objective is to cutover the domestic hub site and the European hub site per the San Diego sales office approach. The regional and smaller sales offices gain connectivity to the main hotsite through their regular access into the Nevada hub site. All of these sites will be cut off from the production network. A third objective is to enable end users to access the corporate Web site through any office via the regular ISP connection from the given office, thereby verifying that the backup URL of clientdr.com was recovered. Finally, the network team must ensure that all network components are cutover within the RTO of 48 hours.

Applications Development and Support (Team at West Coast Hotsite). This team's objectives are fourfold. They are to restore the PeopleSoft Customer Relationship Management (CRM) system, the SAP system, the Time and Expense system, and the Datamart system, all within the RTO of 48 hours. The RTOs include the OS portions conducted by the applicable applications services.

Desktop Services (Team at West Coast Hotsite). There is one objective for Desktop Services: image 20 workstations within the RTO of eight hours.

Applications Services: Windows Based Servers (Team at East Coast Hotsite). This group is faced with three major objectives. First, restore the OS, including SPs and security patches, and data on NT- and Windows 2000-based servers, within the above RTO of 48 hours. This applies to the Time and Expense system and the Datamart system. Second, recover the firewall within the RTO of ten hours. Third, prior to the test, pre-build the infrastructure (WINS, DNS, DHCP, and AD) on spare servers by enabling replication from the production system; position the servers at the West

Coast hotsite, and simulate replication by replicating to the East Coast hotsite. Note that during a real disaster, replication would emanate from either the Nevada site or the Scotland site, from infrastructure servers at those locations. Therefore, the simulation is logically valid for recovery testing purposes.

Applications Services: UNIX-Based Servers (Team at East Coast Hotsite). The UNIX team is to restore the OS, including patches, on UNIX-based servers, including restore of the applicable database, within the above RTO of 48 hours. This applies to the PeopleSoft Customer Relationship Management system and the SAP system.

Database Systems (Team at West Coast Hotsite). The Database team's objectives are to recover the database for Time and Expense, Datamart, SAP, and PeopleSoft CRM. These four applications interface with their respective Oracle databases. This team must also recover the database for the corporate Web site. The Web site relies on an SQL database. All of these databases must be recovered within the same RTO of 48 hours set above for those applications.

Web Services. Recover the corporate Web server within the RTO of 48 hours.

Step 2: Develop the Administrative Project Plan and the Gaps List

Administrative Project Plan. Exhibit 2 depicts the administrative project plan that would be developed by the project manager. Although this is an abbreviated version lacking start and end dates, duration, percentage completed, owners, and comments, it still serves as a reasonable starting point for portraying what tasks should be completed by each team in preparation for the test. The project plan and gaps list are the primary instruments for managing the preparation of the test. A third management tool, the Test Execution and Sequencing Wall Chart, which provides timelines of when each function, from firewall to network to each of the applications and database restores, to desktop restores, showing interdependencies for all of these, is also a critical tool for managing the preparation of the test, but its use in Phase 2 is largely limited to Walkthroughs 1 and 2. The primary value of the Wall Chart lies in its use for managing test execution at the hotsites. The Wall Chart is discussed in Phase 3: Testing Execution.

Weekly status meetings should be held, with the meetings broken into two groups: Group 1 — a weekly half-hour status meeting for the Network Team, and Group 2 — a separate weekly one-hour status meeting attended by the remaining teams. Breaking the status meetings into two groups greatly helps each group concentrate on where they have come from, where they are at, and where they are going with assigned tasks and associated issues that are much more specific to each group's test preparation.

The Gaps List. The source of the gaps list consists of all of the issues and other outstanding items that were winnowed out during the Testing Readiness Assessment, and communicated by the various groups to the project manager during the course of test preparation. For a list of gaps typical of a server environment, reference the above Testing Readiness list of issues. As is typical in project management, the gaps list is usually a list of unclosed issues that tend to prevent completion of tasks in the project administration plan. The gaps list should give a description of the gap, who owns it, what test team is affected by the gap, the target resolution date, and a comments section.

Step 3: Develop the Test Execution and Sequencing Wall Chart

Exhibit 3 shows the Test Execution and Sequencing Wall Chart. As noted previously, the chart serves two functions: (1) for review during the two Walkthroughs, and (2) for the management of the test during its execution. This style of horizontal bar chart is very practical for managing and tracking the entire 48 hours of the test for several reasons. One important reason is that it is *perceptu-*

The project plan and gaps list are the primary instruments for managing the preparation of the test.

EXHIBIT 2 Project Task Plan (abbreviated version)

Task #	Description
1	**Project Administration Team**
1.1	Preliminary Testing Readiness: assessment and report.
1.2	PowerPoint presentation for management on Testing Readiness.
1.3	Obtain management approval for scope of applications and network testing.
1.4	Notify hotsite we want to position three pre-built servers at West Coast hotsite to simulate infrastructure replication.
1.5	Update hotsite Equipment Schedules and addenda to reflect production changes.
1.6	Review lessons learned from previous test; incorporate into planning for current test.
1.7	RTOs (Recovery Time Objectives): adjust RTOs by polling senior end users and sponsors.
1.8	Ensure all licensing authorizations are OK for test.
1.9	Contact hotsite and ensure test dates of 10/16 and 10/17 are scheduled.
1.10	Confirm with hotsite that the number of NETWORK test hours matches APPLICATION test hours = 48 hours.
1.11	Confirm with hotsite that the number of APPLICATION test hours matches scheduled hours = 48 hours.
1.12	Obtain hotsite pricing for added or changed equipment.
1.13	Obtain management approval for updated pricing and contract update.
1.14	Determine hotsite's Program Management Services or other Professional Services to be contracted for.
1.15	Finalize updated contract minimum 60 days prior to test to ensure equipment purchased by hotsite or available from hotsite pool.
1.16	Request network engineer to schedule DR testing cutover with carrier.
1.17	Hold kickoff meeting: network and applications combined.
1.18	Schedule all weekly status meetings for applications.
1.19	Schedule all weekly status meetings for network.
1.20	Prepare and maintain gaps/issues list.
1.21	Poll applications teams for the number of (1) private IP addresses required and (2) public IP addresses required.
1.22	Request all existing applications and OS restore/recovery scripts to be updated.
1.23	Request all NEW applications and OS restore/recovery scripts to be written.
1.24	Request all network restore/recovery scripts to be updated.
1.25	Prepare TEST EXECUTION AND SEQUENCING Wall Chart.
1.26	Request all participants to book hotel.
1.27	Request all participants to book flight.
1.28	Set up bridge lines for common communication point, for test weekend.
1.29	Set up test communications plan: who gets paged/notified of what milestone successes and issues at what time.
1.30	Prepare logistics plan: teams and locations; objectives; travel and lodging info; testing work schedule; etc.
1.31	Hold backup tapes pulling strategy meeting; select production date for the pulling of backed up tapes — all players to sign off on date and approach.
1.32	Request participants to review and update contents of off-site DR Recovery Bins prior to shipment to one or both hotsites: media, recovery scripts, CD-burner, vendor manuals, special tapes, etc.
1.33	Purchase several different colors of CAT5 straight-through cables to enable color coding from ports on server to LAN switch *if* hotsite cannot supply.
1.34	Ship DR Recovery Bins to West Coast and East Coast hotsites.
1.35	Ship DR Recovery Tape Bins to East Coast hotsite.
1.36	Hold Walkthrough 1 by holding high-level review and validation of what portions of whose systems restore in what order in relation to each other; modify test scripts accordingly.
1.37	Hold Walkthrough 2 by holding high-level review and validation of what portions of whose systems restore in what order in relation to each other; modify test scripts accordingly.
1.38	Adjust TEST EXECUTION AND SEQUENCING Wall Chart per walkthrough results.
1.39	Set up and hold hotsite pre-test conference call: hotsite takes order of what resources required for the test, customer validates, hotsite reads back the order, customer validates again.
1.40	Conduct the test.
1.41	Set up the white board in NT server room East Coast; use for determining what stage of restore/recovery each application is at.

EXHIBIT 2 Project Task Plan (abbreviated version) (Continued)

1.42 Document test milestones, by function, noting time and event.

1.43 Document significant problems, errors, impeding variables.

1.44 Prepare debriefing report.

1.45 Hold debriefing meeting.

1.46 Prepare client closeout report.

1.47 Hold client closeout meeting.

2 Network Testing Team

2.1 Revise existing network cutover script or write new script.

2.2 Hold preliminary approach internal meeting.

2.3 Review lessons learned from previous test; incorporate into planning for current test.

2.4 Determine sites to be tested.

2.5 Confirm team members.

2.6 Mentor new team members.

2.7 Receive private and public IP address needs from project manager.

2.8 Request public IP addresses from hotsite for forwarding to ISP.

2.9 Notify ISP of Web/Internet testing.

2.10 Confirm that zone files at ISP are set up to point to hotsite; ensure propagation.

2.11 Obtain and confirm IXC/carrier cutover procedures for DR testing of PVCs.

2.12 Update hotsite Equipment Schedules and addenda to reflect production changes.

2.13 Create network test Visio diagrams.

2.14 Define scope of testing: who at what sites will do what when?

2.15 Ensure all gaps for network portion of test are closed.

2.16 Define approach to subnets, firewall, VLAN segments, routers, and switch.

2.17 Book flight.

2.18 Book hotel.

2.19 Hold pre-test meeting with hotsite; detail and confirm network approach and resources.

2.20 Review and update contents of off-site DR Recovery Bins prior to shipment to one or both hotsites: media, recovery scripts, CD-burner, vendor manuals, special tapes, etc.

2.21 Perform any lab testing if required.

2.22 Attend Walkthrough 1 by holding high-level review and validation of what portions of whose systems restore in what order in relation to each other; modify test scripts accordingly.

2.23 Attend Walkthrough 2 by holding high-level review and validation of what portions of whose systems restore in what order in relation to each other; modify test scripts accordingly.

2.24 Attend hotsite pre-test conference call: hotsite takes order of what resources required for the test, customer validates, hotsite reads back the order, customer validates again.

2.25 Incorporate any changes from Walkthroughs into recovery script.

2.26 Conduct test.

3 Applications Development and Support, Applications Services (Windows and UNIX), Database Services, and Web Services Teams

3.1 Revise existing restore script or write new script.

3.2 Assign team members.

3.3 Forward desktop image requirements to Desktop Services.

3.4 Identify any gaps/issues and forward to project manager.

3.5 Update hotsite Equipment Schedules and addenda to reflect production changes.

3.6 Attend special network configuration meeting to determine the setup of servers, firewall, and other hardware in relation to VLANs, IP addressing, and subnets.

3.7 Ensure all tape backups required for test are being made and stored offsite.

3.8 Mutually decide with other groups what tape backup date will be used for the test.

3.9 Identify and correct any production issues prior to the test to ensure they are not manifested in the test.

3.10 Review and update contents of off-site DR Recovery Bins prior to shipment to one or both hotsites: media, recovery scripts, CD-burner, vendor manuals, special tapes, etc.

3.11 Conduct any pilot work/lab work if required.

3.12 Attend Walkthrough 1 by holding high-level review and validation of what portions of whose systems restore in what order in relation to each other; modify test scripts accordingly.

3.13 Attend Walkthrough 2 by holding high-level review and validation of what portions of whose systems restore in what order in relation to each other; modify test scripts accordingly.

3.14 Attend hotsite pre-test conference call: hotsite takes order of what resources required for the test, customer validates, hotsite reads back the order, customer validates again.

3.15 Book flight.

3.16 Book hotel.

3.17 Conduct test.

EXHIBIT 2 Project Task Plan (abbreviated version) (Continued)

4 Desktop Services Team

4.1 Revise existing restore script or write new script.

4.2 Assign team members.

4.3 Review lessons learned from previous test debriefing report.

4.4 Collect desktop image requirements from those assigned to West Coast hotsite.

4.5 Identify any gaps/issues and forward to project manager.

4.6 Update hotsite Equipment Schedules and addenda to reflect production changes.

4.7 Ensure all installation CD and tape backups required for test are being made and stored offsite.

4.8 Identify and correct any desktop issues prior to test to ensure they are not manifested in test.

4.9 Review and update contents of offsite DR Recovery Bins prior to shipment to one or both hotsites: media, recovery scripts, CD-burner, vendor manuals, special tapes, etc.

4.10 Ensure all desktop licensing authorizations are OK for test.

4.11 Conduct any pilot work/lab work if required.

4.12 Contact West Coast hotsite to see if working space configuration has changed since previous test; obtain new configuration if necessary.

4.13 Prepare workstation seating assignment and forward to project manager.

4.14 Attend Walkthrough 1 by holding high-level review and validation of what portions of whose systems restore in what order in relation to each other; modify test scripts accordingly.

4.15 Attend Walkthrough 2 by holding high-level review and validation of what portions of whose systems restore in what order in relation to each other; modify test scripts accordingly.

4.16 Attend hotsite pre-test conference call: hotsite takes order of what resources required for the test, customer validates, hotsite reads back the order, customer validates again.

4.17 Book flight.

4.18 Book hotel.

4.19 Conduct test.

5 Operations/Production Control Team

5.1 Mutually decide with other groups what tape backup date will be used for test.

5.2 Hold DR Recovery Bins on-site one week if necessary to allow DR team users to review contents and update/replace items.

5.3 Receive special job runs from DR team leads for creation of extra backup tapes to be used for the test (do not use production backup tapes).

5.4 Submit backup tape creation requests/jobs to operators.

5.5 Operators run backup jobs, create special tape listings identifying the tape volume numbers.

5.6 Forward DR Tape Bins and DR Recovery Bins (media, CDs, vendor manuals, etc.) to project manager for shipping to the hotsites.

ally easy to use. Often, a project manager will use project planning software and attempt to view the restore sequencing and recovery from the software's PERT perspective. However, this is frequently difficult because of how the sequencing is portrayed over numerous pages. Alternately, a "launch schedule" type of narrative is often used. This is usually a multi-page table listing step number, start time, end time, and description. This is not necessarily bad, but a table-style document reads vertically, and it is not easy to depict what the dependencies are in a manner that enables instantaneous visual recognition; rather, the dependencies are presented cognitively and must be waded through.

A second reason for use of the Wall Chart approach is that it is *conceptually easy to use*. It captures all of the restore activity for both the network and the servers on one or two pages by printing the Wall Chart on 34 × 44-inch plotting paper or the smaller poster-sized paper. Draft copies can be reviewed by the teams at the Walkthroughs, where they can be edited to become the final Wall Chart to be used at the test. Further, the chart enables the team to instantly determine the degree of delay to, for example, Server A that is caused by the troublesome recovery of Server B that sits upstream on the critical path of dependent Server A.

The project manager composes the Wall Chart using Microsoft's Excel. The arrows are drawn using the Block Arrows selected

EXHIBIT 3 Test Execution and Sequencing

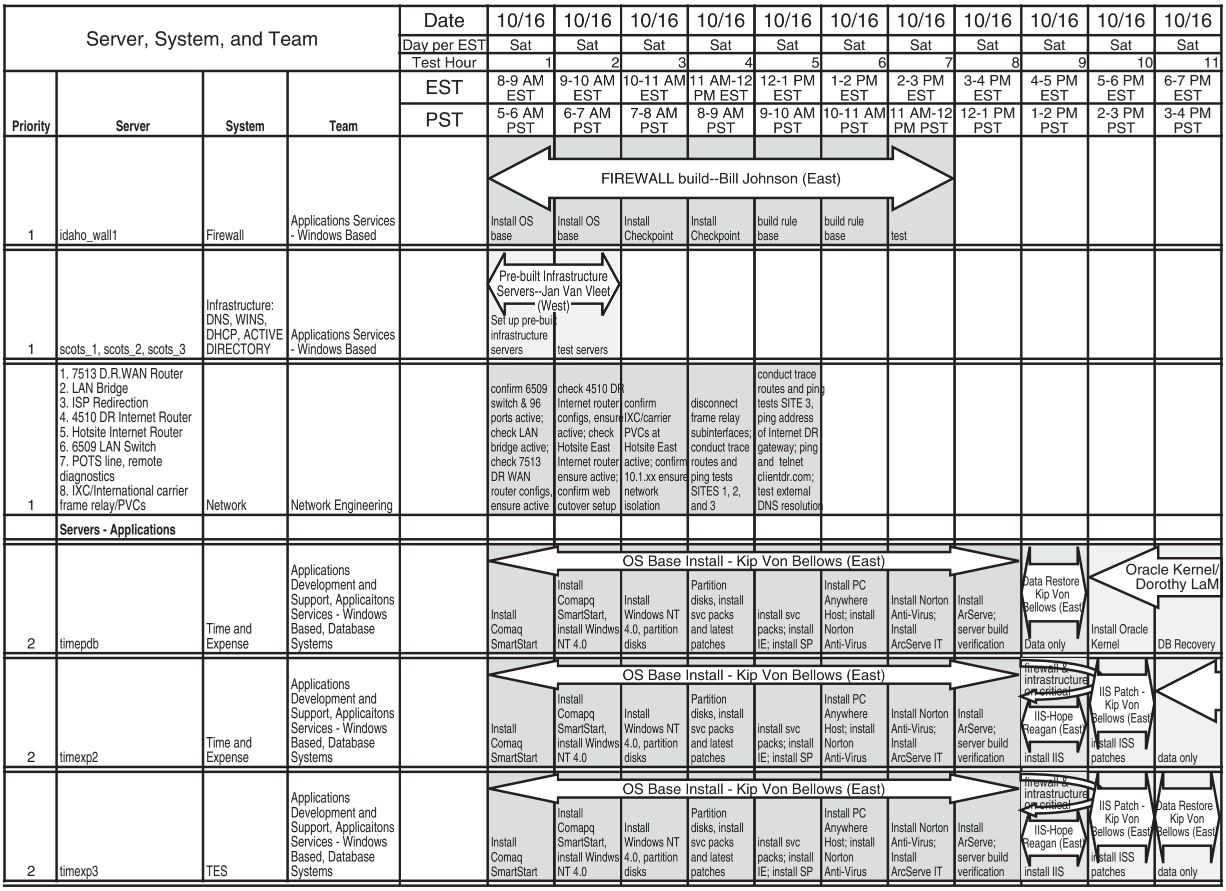

Server, System, and Team				Date	10/16	10/16	10/16	10/16	10/16	10/16	10/16	10/16	10/16	10/16	10/16
				Day per EST	Sat	Sat	Sat	Sat	Sat	Sat	Sat	Sat	Sat	Sat	Sat
				Test Hour	1	2	3	4	5	6	7	8	9	10	11
				EST	8-9 AM EST	9-10 AM EST	10-11 AM EST	11 AM-12 PM EST	12-1 PM EST	1-2 PM EST	2-3 PM EST	3-4 PM EST	4-5 PM EST	5-6 PM EST	6-7 PM EST
Priority	Server	System	Team	PST	5-6 AM PST	6-7 AM PST	7-8 AM PST	8-9 AM PST	9-10 AM PST	10-11 AM PST	11 AM-12 PM PST	12-1 PM PST	1-2 PM PST	2-3 PM PST	3-4 PM PST
1	idaho_wall1	Firewall	Applications Services - Windows Based		FIREWALL build--Bill Johnson (East) Install OS base	Install OS base	Install Checkpoint	Install Checkpoint	build rule base	build rule base	test				
1	scots_1, scots_2, scots_3	Infrastructure: DNS, WINS, DHCP, ACTIVE DIRECTORY	Applications Services - Windows Based		Pre-built Infrastructure Servers--Jan Van Vleet (West) Set up pre-built infrastructure servers	test servers									
1	1. 7513 D.R.WAN Router 2. LAN Bridge 3. ISP Redirection 4. 4510 DR Internet Router 5. Hotsite Internet Router 6. 6509 LAN Switch 7. POTS line, remote diagnostics 8. IXC/International carrier frame relay/PVCs	Network	Network Engineering		confirm 6509 switch & 96 ports active; check LAN bridge active; check 7513 DR WAN router configs, ensure active	check 4510 DR Internet router configs, ensure active; check Hotsite East Internet router ensure active; confirm web cutover setup	confirm IXC/carrier PVCs at Hotsite East active; confirm 10.1.xx ensure network isolation	disconnect frame relay subinterfaces; conduct trace routes and ping tests SITES 1, 2, and 3	conduct trace routes and ping tests SITE 3, ping address of Internet DR gateway; ping and telnet clientdr.com; test external DNS resolution						
	Servers - Applications														
2	timepdb	Time and Expense	Applications Development and Support, Applicaitons Services - Windows Based, Database Systems		OS Base Install - Kip Von Bellows (East) Install Comaq SmartStart	Install Comapq SmartStart, install Windws NT 4.0	Install Windows NT 4.0, partition disks	Partition disks, install svc packs and latest patches	install svc packs; install IE; install SP	Install PC Anywhere Host; install Norton Anti-Virus	Install Norton Anti-Virus; Install ArcServe IT	Install ArServe; server build verification	Data Restore Kip Von Bellows (East) Data only	Oracle Kernel/ Dorothy LaM Install Oracle Kernel	DB Recovery
2	timexp2	Time and Expense	Applications Development and Support, Applicaitons Services - Windows Based, Database Systems		OS Base Install - Kip Von Bellows (East) Install Comaq SmartStart	Install Comapq SmartStart, install Windws NT 4.0	Install Windows NT 4.0, partition disks	Partition disks, install svc packs and latest patches	install svc packs; install IE; install SP	Install PC Anywhere Host; install Norton Anti-Virus	Install Norton Anti-Virus; Install ArcServe IT	Install ArServe; server build verification	firewall & intrastructure on critical IIS-Hope Reagan (East) install IIS	IIS Patch - Kip Von Bellows (East) install ISS patches	data only
2	timexp3	TES	Applications Development and Support, Applicaitons Services - Windows Based, Database Systems		OS Base Install - Kip Von Bellows (East) Install Comaq SmartStart	Install Comapq SmartStart, install Windws NT 4.0	Install Windows NT 4.0, partition disks	Partition disks, install svc packs and latest patches	install svc packs; install IE; install SP	Install PC Anywhere Host; install Norton Anti-Virus	Install Norton Anti-Virus; Install ArcServe IT	Install ArServe; server build verification	firewall & intrastructure on critical IIS-Hope Reagan (East) install IIS	IIS Patch - Kip Von Bellows (East) install ISS patches	Data Restore Kip Von Bellows (East) data only

from the Autoshapes of Excel's Draw facility. For space reasons, Exhibit 3 shows a very small section of the test, but the idea should still be clear by studying the composition of the Wall Chart example.

Step 4: Hold Walkthroughs 1 and 2

At each of the Walkthroughs, all test participants should meet and review the chart. Out of the Walkthroughs should come several deliverables. One is an understanding by each test participant of what process must be restored first before their own process can be restored, whether the process is the entire process or a sub-process for a server or other function. A second deliverable is an understanding by each participant of two key pieces of time: (1) the amount of time their portion of a server or network recovery is in relation to the recovery time required for the entire recovery of that server or network, and (2) the amount of time required for the entire recovery of that server or network in relation to the allocated test time.

Additional deliverables from the Walkthroughs include an understanding of the main critical path that runs through the entire recovery testing period, including the slack time, as well as an understanding of how a major delay in up-front processes could have serious consequences for later processes, to the point of running out of time for completing the test. Finally, there is the expectation that the team members will gain an understanding of some type of triage that must be used if a given process encounters trouble. Is the process a requirement for all systems, or can it be discarded as not recoverable during the test, with a knowledge of what remaining processes can still be recovered?

Step 5: Prepare the Logistics Kit and Distribute to All Participants

A Logistics Kit is essentially the roadmap of the "who, what, when, where, and how" for the team members participating in the test from the point the flights leave and arrive at the hotsite destinations, through what happens during the test, and getting back home. This kit should be distributed approximately one week prior to the test.

Representative contents of the Logistics Kit would include an overview of what the kit contains; a list of team members, roles, and locations (e.g., East Coast or West Coast as shown in our examples); and the test calendar, which at a high level provides key events that must occur on what dates. There should be a section for travel that lists for each participant their complete itinerary, a brief description of orientation to be provided by the hotsite on the day prior to or day of the test, and the test objectives for each system to be recovered and for the network. Also included should be a section that lists several restaurants in the area and if food is to be catered at the hotsite.

The kit should have a section describing the parameters and assumptions for the test. This part should state what times, on what dates, what IOS will be loaded in the routers, any preloaded OS provided, what servers will be tested (refer the reader to the Wall Chart), IP addressing for the subnets, static and dynamic IP addressing, and public IP addressing. Parameters would also include what circuits will be activated or will automatically swing over to the hotsite, circuit capacity, the redirecting of a gateway to ISP services, and what type of platforms are provided for the workstations.

Also contained in the kit are two important sections: (1) a listing from the hotsite, and as validated by the joint customer-hotsite pre-test conference call, of every service and item of equipment the hotsite is providing for the test, and (2) a section that addresses team communications. This latter section is extremely important for managing test execution. It describes how each participant, some on opposite ends of the country at different hotsites, is to know where the other participants are in their recovery; how interested executive managers wanting to remotely observe the recovery test are to be notified of each milestone accomplished; and how all or a sub-set of test participants, spread out over the country, discuss a recovery problem ad hoc recovery strategy. The communications plan provides the hotsite room phone, what bridge phone number to jump on for diagnostics, and the use of a recorded message line that provides milestone updates to test participants and managers wanting general knowledge of test progress.

A diagram section should also be included. The project manager should provide all network diagrams and the seating chart diagrams with associated desk phone numbers for those in the workstation recovery area. Finally, there should be a section that contains a fold-out copy of the Wall Chart. For the actual recording of the test events and progress on an hour-by-hour basis, an electronic version of the chart is used that has a column added for notes. The extent of the notes from the project management perspective does not need not be any more detailed than the following example, for a firewall server. The tracking notes, plus post-test gathered notes from the participants, make up the main body of information for the debriefing report.

SUNDAY 8:30 a.m.: recovery started
SUNDAY 2:45 p.m.: incorrect application version installed; must rebuild
SUNDAY 5: 10 p.m.: behind schedule; estimated time for completion is 10 p.m.
SUNDAY 6:33 p.m.: running into problems, apparently due to NIC cards
SUNDAY 8:10 p.m.: problems with new application version continuing; might be sequence of security patch loading
SUNDAY 11:27 p.m.: server fixed; source of problem regarding NIC cards was a bad cable

SUNDAY 12:14 p.m.: server fully recovered

Finally, the kit should have a section for hotel and transportation information, and hotsite location information, which give the hotel and hotsite addresses and phone numbers and the driving instructions for getting to each.

PHASE 3: TESTING EXECUTION

By now, all of the dominoes should have been set up, and it is seemingly a matter of just lightly pushing on the first domino to get all others to fall in place. However, no test goes perfectly smoothly, and many tests end up in the rough very early. Here is a suggested approach for avoiding serious problems. Additional advisories for testing are listed in the "Summary and Recommendations" section of this article.

There are several checks that must be conducted within the first one to two hours upon arrival at the hotsite. First, the network engineer of the test team must link up with the hotsite technicians to validate the network. This assumes network cutover has already occurred. The network engineer's first task should be the inspection of all LAN and WAN hardware, all configurations, hardware for getting out to the Internet, subnet configurations, the LAN bridge between the two hotsites, POTS line availability for remote diagnostics, analog lines for dialing out, and all aspects of Internet connectivity.

Second, the network engineer should inspect all communication schemes: bridge lines, status recording lines, and the capability of dialing to any of the workstation areas.

Very early in the test, ensure the infrastructure is working — DNS, DHCP, WINS, and AD. It should be up within one to two hours of arriving at the hotsite (assuming replication emulation from prebuilt servers). During the first hours of restoration, ensure that as soon as an OS is recovered, the server image is ghosted.

Additionally, make sure the Wall Chart is posted in the recovery room and in the hotsite conference room from where the test is managed. Pay careful attention to slack time in the critical path for the recovery as a whole. A team restoring the OS can easily get two hours behind, but still think it has sufficient time in light of the entire slack time, when in reality it is not considering what problems the subsequent database or application restores might have. This results in groups stating they are doing well when they are not doing well.

As the test progresses, do not over-report. Excess reporting of test progress requires over-observation and that results in constant "where are we?" question-asking. There is nothing wrong with asking each participant, each hour, where he (she) is in the recovery (actuals), in relation to the Wall Chart (planned). In the Notes section of the electronic Wall Chart, record the main milestones, as described above, and any major problems encountered. This is the raw data and the foundation of the debriefing report.

For server engineers and network engineers, there is relatively little else to do at the hotsite other than following their restore scripts, logging observations, engaging in bridge calls, working through problems, calling vendors, interfacing with the hotsite technicians during troubleshooting, and reporting hourly to the project manager where they are in the recovery process. For the project manager, the majority of time is spent tracking and documenting test progress, running bridge calls, updating the recorded message line, driving troubleshooting, and ensuring that the troubleshooting does not get bogged down.

It is essential that a debriefing report be prepared and distributed to all team members and management sponsors.

TESTING DEBRIEFING

Based on the tracking and documenting during testing, it is essential that a debriefing report be prepared and distributed to all team members and management sponsors. This is the primary mode of closing the feedback loop. From the debriefing report, recovery scripts are modified, and test tactics and strategy get changed. A full description of a debriefing report goes far beyond the scope of the write-up. Nonetheless, major reporting areas are essential to even a modest report.

The first major reporting area is descriptive statistics. For each server or other equipment or circuit recovered, present the mean hours ahead or behind scheduled completion time, the actual recovery time (actual end time minus actual start time, including lulls), mean recovery time for all workstations, and standard deviations for these averages. The statistics portion should conclude with problems, observations, and recommendations in general. For each server, detail the following: server/function name, recovery priority, the team, planned start, actual start, planned finish, actual finish, finish variance, RTO and the hour recovered in, and the number of hours ahead or behind schedule.

The second major section should contain the chronology of events, what major recovery events were accomplished, lessons learned, whether one or more recovery scripts related to the function need to be changed, and recommendations specific to the server or network piece.

SUMMARY AND RECOMMENDATIONS

Locating or creating methodologies for disaster recovery testing at a hotsite for a server-based environment is challenging. The purpose of this article was to present a relatively detailed methodology for hotsite testing in a server-based environment. The importance of conducting a Testing Readiness Assessment was described, an abbreviated project task plan was presented, and typical gaps encountered during the assessment and the course of testing preparation were listed. Further, the criticality of the use of the Testing Execution and Sequencing Wall Chart for the Walkthroughs and for the management of the test was discussed.

Recommendations based on the project management of several recovery tests follow.

General Recommendations

□ Conduct a Testing Readiness Assessment to ensure that there are no missing items with implementation times greater than the 90 days required for testing preparation.

□ Use a structured methodology, and always adapt it to the needs of the client.

□ At the earliest point in the project, focus hard on all items in the project with the longest implementation times (e.g., hardware ordering, circuit orders, network cutover approach, isolated versus non-isolated testing network, and hotsite contract negotiations for equipment upgrades or modifications).

□ Understand the crux of recovery rests on the recovery scripts, *not* the DR plan.

□ Concentrate early on the IP addressing schema, including what severs will link to what VLANs using what ports; public and private IP addressing; and obtaining from and providing information to the client's ISP.

□ Be assertive with the hotsite company. The project manager, *not* the hotsite company, should be running the testing project within the classical administrative, budgetary, political, legal, regulatory, physical, and technical constraints.

□ Fight off silly approaches to testing that are rarely beneficial. For example, a client wanting to march into the data center with plane tickets and a surprise test that begins that night as a means of "really discovering if we would have recovered." That is meaningless.

□ As a follow-up to silly approaches, discontinue the use of the very concept of "testing" and that it is only a test so that if you do not recover some systems, that this is allowable because all will "learn something." It is not really testing — it is *recovery practice*. The visibility of a test (still using the term because it is so engrained) is far too political to return home to report to senior management that half of the servers were not recovered and the network could not be cutover. Testing is a scrimmage. Does any reader recall a scrimmage that was acceptable to lose?

□ Early in the project, delineate what is to be tested, stick to these targets, and move forward. If control is not gained early, there will be hopeless confusion and frustration.

□ Alternate status meetings by concentrating on gaps one week, then tasks completion the next week. For a given week's meeting, attend to only those tasks or gaps that are due, or that will be due in the next two or three weeks. Holding endless status meetings is interminable for participants and results in alienation and project slow-down.

Specific Recommendations

These recommendations may seem somewhat far-fetched. However, they are based on actual tests. Obviously, there are scores of reasons why a server is not getting restored, but consider the following recommendations as coming from those returning from the front:

□ Be excruciatingly careful about what tapes are pulled for the test.
□ The Project Manager must maintain hyper-vigilance for team sleep-deprivation. Incorrect OS versions or application versions will get loaded, and diagnostics begin to fail.
□ Check to see if a physical IP address, instead of the logical IP address, is inadvertently being used if a server is not recovering.
□ Ghost the image after recovery.
□ If cables are not color coded, it is difficult to determine which port on a server is connected to which VLAN, and diagnostics become troublesome.
□ If a server is rebooting continuously, check to see if the boot was set to the NIC instead of the CD-ROM.
□ Consider the use of a hotnode firewall that is pre-built and ready to go.
□ Verify that the hardware at the hotsite is of sufficient capacity.
□ Despite having scheduled the test with the carrier, just prior to the test, double-check to make sure that the carrier really knows that the test is to occur.
□ Look for externally attached storage early in the restores to avoid scrambling for SCSI adapter drivers.
□ Be very careful about licensing keys.
□ If the hotsite was trying to do you a favor by providing eight processors on a server when your NT system in production calls for four, then demand from the hotsite technicians that the server be brought down to four processors, if the overage prevents recovery.
□ On UNIX servers, ensure that all previous testing customer junk is cleaned off.
□ Check all system names and static IP addresses prior to the test/recovery.

Again, general and specific recommendations are limitless, and as with all technical experiences, there is variation according to what is encountered. Successful recovery tests will always be a reflection of a successful production environment, and will always be a function of a sound and intrepid team, structured project management, adequate time and resources, and management support. ■